

# On the Optimal Succinctness and Efficiency of Functional Encryption and Attribute-Based Encryption

Aayush Jain<sup>1</sup>

Huijia Lin<sup>2</sup>

Ji Luo<sup>2</sup> 

<sup>1</sup> Carnegie Mellon University  
aayushja@andrew.cmu.edu

<sup>2</sup> University of Washington  
{rachel, luoji}@cs.washington.edu

October 2022

## Abstract

In this work we investigate the asymptotically best-possible succinctness and efficiency of functional encryption (FE) and attribute-based encryption (ABE), focusing on simultaneously minimizing the sizes of secret keys and ciphertexts and the decryption time. To this end, we consider the notion of partially hiding functional encryption (PHFE) that captures both FE and ABE, and the most efficient computation model of random access machine (RAM). A PHFE secret key  $sk_f$  is tied to a function  $f$ , whereas a ciphertext  $ct_x(y)$  is tied to a public input  $x$  (e.g., ABE attribute) and encrypts a private input  $y$ . Decryption reveals  $f(x, y)$  and nothing else about  $y$ .

We present the first PHFE scheme for RAMs based on the necessary assumption of FE for circuits. It achieves nearly optimal succinctness and efficiency:

- The secret keys  $sk_f$  are of (optimal) *constant size*, independent of the description size  $|f|$  of the function tied to it.
- The ciphertexts  $ct_x(y)$  have (nearly optimal) *rate-2* dependency on the private input length  $|y|$  and (optimal) independency of the public input length  $|x|$ .
- Decryption is efficient, running in time linear in the instance running time  $T$  of the RAM computation, in addition to the input and function sizes, i.e.,  $T_{\text{Dec}} = (T + |f| + |x| + |y|) \text{poly}(\lambda)$ .

Our construction significantly improves upon the asymptotic efficiency of prior schemes. As a corollary, we obtain the first ABE scheme with both constant-size keys and constant-size ciphertexts, and the best-possible decryption time matching an existing lower bound.

We show barriers to further improvement on the asymptotic efficiency of (PH-)FE. We prove the first unconditional space-time trade-off for (PH-)FE. No secure (PH-)FE scheme can have both key size and decryption time sublinear in the function size  $|f|$ , and no secure PHFE scheme can have both ciphertext size and decryption time sublinear in the public input length  $|x|$ . These space-time trade-offs apply even in the simplest selective 1-key 1-ciphertext secret-key setting. Furthermore, we show a conditional barrier towards achieving the optimal decryption time  $T_{\text{Dec}} = T \text{poly}(\lambda)$  — any such (PH-)FE scheme implies a primitive called secret-key doubly efficient private information retrieval (SK-DE-PIR), for which so far the only known candidates rely on new and non-standard hardness conjectures.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Technical Overview . . . . .	6
1.2	Related Works . . . . .	13
<b>2</b>	<b>Preliminaries</b>	<b>15</b>
2.1	Multi-Tape Random Access Machine . . . . .	15
2.2	Laconic Garbled RAM . . . . .	17
2.3	Partially Hiding Functional Encryption and FE for Circuits . . . . .	19
2.4	Universal RAM and PHFE for RAM . . . . .	21
2.5	Indistinguishability Obfuscation . . . . .	22
2.6	Laconic Oblivious Transfer . . . . .	23
2.7	Garbled Circuits . . . . .	25
2.8	Puncturable Pseudorandom Function . . . . .	26
2.9	Secret-Key Encryption . . . . .	26
2.10	Oblivious RAM . . . . .	27
<b>3</b>	<b>Efficiency Trade-Offs of PHFE for RAM</b>	<b>29</b>
3.1	Contention Between Storage Overhead and Decryption Time . . . . .	29
3.2	Barrier to Time Optimality . . . . .	33
<b>4</b>	<b>Bounded LGRAM with Fixed-Memory Security</b>	<b>39</b>
4.1	Construction . . . . .	40
4.2	Security . . . . .	43
<b>5</b>	<b>Transformations of LGRAM</b>	<b>53</b>
5.1	Fixed-Memory to Fixed-Address . . . . .	53
5.2	Fixed-Address to Full Security . . . . .	57
5.3	Bounded to Unbounded . . . . .	60
<b>6</b>	<b>PHFE for RAM</b>	<b>63</b>
6.1	Bounded Private Input . . . . .	63
6.2	Full-Fledged PHFE for RAM . . . . .	72
	<b>References</b>	<b>75</b>

# 1 Introduction

Functional encryption (FE) [BSW11,O’N10] is a powerful enhancement of public-key encryption enabling creation of *functional secret keys* that allow for computing specific functions on encrypted data and reveal only the output. A closely related concept is attribute-based encryption (ABE) [SW05,GPSW06], which supports fine-grained access control. An ABE secret key is tied to a predicate and can only decrypt ciphertexts that are tied to attributes satisfying the predicate. Both FE and ABE have been extensively studied in the literature and have numerous applications.

In this work, we consider the most general notion of partially hiding functional encryption (PHFE) [GVW15,AJL<sup>+</sup>19,JLMS19] that captures both FE and ABE. A secret key  $sk_f$  is tied to a function  $f$ , and a ciphertext  $ct_x(y)$  is tied to a public input  $x$  in addition to encrypting a private input  $y$ , so that decryption produces the computation output  $z = f(x, y)$ . Collusion-resistant (indistinguishability-based) security ensures that given unboundedly (polynomially) many secret keys  $\{sk_{f_q}\}_q$  tied to different functions, ciphertexts  $ct_x(y_0)$  and  $ct_x(y_1)$  encrypting different private inputs w.r.t. the same public input remain indistinguishable so long as none of the secret keys separate the ciphertexts through outputs, i.e.,  $f_q(x, y_0) = f_q(x, y_1)$ . Put simply, the only information revealed about the private input  $y$  is the outputs  $\{f_q(x, y)\}_q$ . Thanks to its generality and flexibility, advances in PHFE construction immediately imply those in FE, ABE, and their special cases such as IPFE, broadcast encryption, IBE, etc.

Over the past decade, significant progress has been made in establishing the feasibility of FE and PHFE, for various classes of computation, with different levels of efficiency, and from different assumptions. However, we still yet to understand the asymptotic optimality and theoretical limits of the efficiency of (PH-)FE. We ask

*Can we construct PHFE with best-possible asymptotic efficiency?  
Can we show matching lower bounds?*

We make progress towards answering the above questions. For the lower bounds, we show inherent trade-offs between the sizes of keys/ciphertexts and the decryption time, and show barriers towards achieving asymptotically optimal decryption time. From the constructive side, we present the first collusion-resistant PHFE for RAM based on the necessary assumption of (polynomially secure) collusion-resistant FE for circuits, which in turn can be based on well-studied assumptions [JLS21,JLS22]. Our scheme has nearly minimally sized keys and ciphertexts, and the best-possible decryption time matching our lower bounds. As a corollary, we obtain the first ABE with both constant-size keys and constant-size ciphertexts, and the best-possible decryption time matching the recently discovered lower bound [Luo22]. We emphasize again that this work focuses on asymptotic efficiency rather than concrete efficiency.

**Dream Efficiency.** Before describing our results, we first picture the dream efficiency w.r.t. three important dimensions. Each dimension has been a consistent research theme across many primitives in cryptography.

- *Efficient Computation Model.* Functions should be represented by random access machines (RAM), the most efficient computation model subsuming both circuits and Turing machines.

*RAM Computation.* It would be flexible to represent a RAM computation by a constant-size RAM machine  $U^1$  with access to multiple tapes, a function tape containing  $f$ , a public input tape for  $x$ , and a private input tape for  $y$ . Moreover, the computation may have *arbitrarily long outputs*.

This model flexibly captures many natural scenarios, e.g., binary search where the database could be a part of  $f$ ,  $x$ , or  $y$ . It can emulate the evaluation of a circuit  $C$  on input  $(x, y)$  by including the circuit description in the function tape. In ciphertext-policy ABE, the ciphertext is tied to a predicate  $C$ , which can be captured by putting the description of  $C$  on the public input tape. These examples tell us that any of  $f$ ,  $x$ , and  $y$  could have a long description and we want to optimize the efficiency dependency on their lengths.

- *Succinctness.* Having small communication- and storage- overhead means having small-sized master public key  $\text{mpk}$ , secret keys  $\text{sk}_f$ , and ciphertext  $\text{ct}_x(y)$ . At the most basic level, the size of each component should be *polynomial* in the length of the information it encodes —  $|\text{mpk}| = O(1)$ ,  $|\text{sk}_f| = \text{poly}(|f|)$ , and  $|\text{ct}| = \text{poly}(|x|, |y|)$ , where  $|f|, |x|, |y|$  are the description lengths of  $f, x, y$ , respectively — referred to as *polynomial efficiency*.<sup>2</sup> (In this introduction,  $O(\cdot)$  hides  $\text{poly}(\lambda)$  factors.) However, there is much to be desired beyond this basic level of efficiency. For instance, linear efficiency means  $|\text{sk}_f| = O(|f|)$  and  $|\text{ct}| = O(|x| + |y|)$ , and rate-1 efficiency means  $|\text{sk}_f| = |f| + O(1)$  and  $|\text{ct}| = |x| + |y| + O(1)$ .

In fact, even smaller parameters are possible. Since (PH-)FE does not guarantee the privacy of the function  $f$ , it is possible to give a description of  $f$  in the clear in the secret key, and the *non-trivial* part of the secret key may be shorter than  $f$ . In this case, the right measure of efficiency should be the size of the non-trivial part (i.e., the overhead), which we now view as *the* secret key. We can now aim for secret keys of size *independent* of that of the function — i.e.,  $|\text{sk}_f| = O(1)$  — referred to as *constant-size* keys. The same observation applies to the public input  $x$  tied to the ciphertext and we can hope for ciphertexts of size *independent* of  $|x|$  while having optimal, rate-1 dependency on the length  $|y|$  of the private input — i.e.,  $|\text{ct}_x(y)| = |y| + O(1)$ . In summary,

**Optimal Succinctness:**  $|\text{mpk}| = O(1), \quad |\text{sk}_f| = O(1), \quad |\text{ct}_x(y)| = |y| + O(1)$  .

Note that the ideal component sizes are completely independent of the running time or output length of the computation.

- *Decryption Time.* Decryption is also a RAM computation,  $\text{Dec}^{f,x,\text{sk}_f,\text{ct}_x(y)}(\text{mpk})$ , which on input  $\text{mpk}$  and with random access to  $f, x, \text{sk}_f, \text{ct}_x(y)$ , computes the output  $U^{x,y,f}()$ . We want decryption to be efficient, ideally taking time linear in the instance running time  $T$  of the RAM computation with  $(f, x, y)$ , that is,

**Optimal Decryption Time:**  $T_{\text{Dec}} = O(T)$  .

<sup>1</sup>With loss of generality, one can think of  $U$  as the universal RAM machine.

<sup>2</sup>It may appear that polynomial efficiency is the bare minimal. However, it is possible to consider components whose sizes depend on an upper bound on the length of the information *not* encoded in them. Many schemes are as such, e.g., the FE scheme of [GGH<sup>+</sup>13] has  $|\text{mpk}| = \text{poly}(\max |y|)$ , and the sizes of master public keys and ciphertexts of the celebrated ABE scheme by [BGG<sup>+</sup>14] grow polynomially with the maximal depth of the computation. When a scheme requires knowing an upper bound on parameter  $Z$  (e.g., input length, depth, or size), it is referred to as  $Z$ -bounded FE/ABE.

**Table 1.** Comparison among some FE schemes.  $T$  is the instance running time of the Turing machine computation. 1-key sls FE stands for 1-key sublinearly succinct FE. Sub-exp iO stands for sub-exponentially secure iO.

citation	functionality	$ \text{sk} $	$ \text{ct} $	$T_{\text{Dec}}$	adaptive	assuming
[GGH <sup>+</sup> 13]	Circuit-FE	$\text{poly}( C )$	$\text{poly}( y )$	$\text{poly}( C )$	no	iO
[KNTY19]	Circuit-FE	$\text{poly}( C )$	$\text{poly}( y )$	$\text{poly}( C )$	yes	1-key sls FE
[GWZ22]	Circuit-FE	$\text{poly}( C )$	$ y  + O(1)$	$\text{poly}( C )$	no	iO
[AS16]	TM-FE	$\text{poly}( f )$	$\text{poly}( y )$	$\text{poly}( f ,  y )T$	yes	iO
[AJS17]	TM-FE	$c \cdot  f  + O(1)$	$c \cdot  y  + O(1)$	$\text{poly}( f ,  y )T$	yes	sub-exp iO
[AM18]	TM-FE	$\text{poly}( f )$	$O( y )$	$\text{poly}( f )T$	yes	Circuit-FE
[KNTY19]	TM-FE	$\text{poly}( f )$	$\text{poly}( y )$	$\text{poly}( f ,  y )T$	no	1-key sls FE
This work	RAM-PHFE	$O(1)$	$2 y  + O(1)$	$O(T +  f  +  x  +  y )$	yes	Circuit-FE

**Our Results.** The question is whether the dream efficiency is simultaneously attainable in all above three dimensions. Towards understanding this, we present both new construction and lower bounds.

*New PHFE for RAM with Almost Optimal Succinctness.* Starting from polynomially secure FE for circuits with only polynomial efficiency (i.e., all of  $|\text{mpk}|, |\text{sk}_f|, |\text{ct}(y)|$  are just  $\text{poly}(|f|, |y|)$ ), which in turn can be constructed from three well-studied assumptions [JLS21, JLS22], we construct PHFE for RAMs with almost optimal succinctness.

**Theorem 23** (informal). *Assume polynomially hard selectively secure FE for circuits. There is an adaptively secure PHFE scheme for RAM with arbitrarily long outputs, satisfying*

$$\begin{aligned} \text{Efficiency of Our PHFE: } \quad & |\text{mpk}| = O(1), \quad |\text{sk}_f| = O(1), \quad |\text{ct}_x(y)| = 2|y| + O(1), \\ & T_{\text{Dec}} = O(T + |f| + |x| + |y|). \end{aligned}$$

Our construction gives the first collusion-resistant FE for RAM, and also the first FE for any model of computation (e.g., circuit or TMs) with almost optimal succinctness. The only gap to optimality is that the ciphertext has rate-2 dependency on  $|y|$  instead of rate-1. Prior constructions work with either circuits [GGH<sup>+</sup>13, JLS21, JLS22] or Turing machines [AS16, AM18, KNTY19], and achieve only polynomial efficiency as summarized in Table 1 (more discussion on related works in Section 1.2).

As a corollary, we obtain the first ABE for RAM from falsifiable assumptions, and the first ABE for any model of computation with both constant-size keys and constant-size ciphertexts. The only prior construction of ABE for RAM by [GKP<sup>+</sup>13] relies on non-falsifiable assumptions like SNARK and differing input iO. In terms of efficiency, known schemes achieve either constant-size keys or constant-size ciphertexts [ALdP11, YAHK14, Tak14, Att16, ZGT<sup>+</sup>16, AT20, LL20, LLL22]. Achieving constant-size keys and ciphertexts simultaneously has been an important theoretical open question in the area of ABE (see discussion in [LLL22]). The state-of-the-art is summarized in Table 2.

**Corollary** (informal). *Assume polynomially secure FE for circuits. There is an adaptively secure key-policy ABE scheme for RAM, and an adaptively secure ciphertext-policy ABE scheme*

**Table 2.** Comparison among some key-policy ABE schemes.  $d$  is maximal depth of the circuit. GGM stands for Generic Group Model, ABP for Arithmetic Branching Program, diO for differing input iO. In [GKP<sup>+</sup>13], the challenge attribute can be chosen after seeing mpk and keys for adaptively chosen TMs, but no keys can be published after seeing the challenge ciphertext.

citation	functionality	sk	ct	$T_{\text{Dec}}$	adaptive	assuming
[BGG <sup>+</sup> 14]	Circuit-ABE	$\text{poly}(d)$	$ x  \text{poly}(d)$	$ C  \text{poly}(d)$	no	LWE
[LL20]	ABP-ABE	$O( C  \cdot  x )$	$O(1)$	$O( C  \cdot  x )$	yes	k-lin
[LLL22]	Circuit-ABE	$O(1)$	$\text{poly}(d)$	$ C  \text{poly}(d)$	yes	LWE & GGM
[GKP <sup>+</sup> 13]	RAM-ABE	$O(1)$	$\text{poly}( x )$	$O(T +  f  +  x )$	no	SNARK & diO
This work	RAM-ABE	$O(1)$	$O(1)$	$O(T +  f  +  x )$	yes	Circuit-FE

for RAM, satisfying

$$\begin{aligned} \text{Efficiency of Our KP-ABE:} \quad & |\text{mpk}| = O(1), \quad |\text{sk}_f| = O(1), \quad |\text{ct}(x, \perp)| = O(1), \\ & T_{\text{Dec}} = O(T + |f| + |x|). \end{aligned}$$

$$\begin{aligned} \text{Efficiency of Our CP-ABE:} \quad & |\text{mpk}| = O(1), \quad |\text{sk}_x| = O(1), \quad |\text{ct}(f, \perp)| = O(1), \\ & T_{\text{Dec}} = O(T + |f| + |x|). \end{aligned}$$

It appears that the decryption time of our PHFE (or ABE) scheme is sub-optimal: Besides from the linear dependency on the RAM instance running time  $T$ , it additionally depends linearly on  $|f|$ ,  $|x|$ , and  $|y|$ . It turns out that, *optimal succinctness and optimal decryption time are at conflict*: We show that conditioned on optimal succinctness, the linear dependency of decryption time on  $|f|$ ,  $|x|$  is inherent. We also show barriers towards removing the dependency on  $|y|$ , or  $|f|$ , or  $|x|$ , irrespective of succinctness. Hence, our PHFE scheme matches the lower bounds and barriers. For ABE, our lower bounds and barriers do not apply. Nevertheless, our ABE scheme matches an existing lower bound by [Luo22] showing that any ABE must satisfy  $|\text{ct}|T_{\text{Dec}} = \Omega(|x|)$  and  $|\text{sk}|T_{\text{Dec}} = \Omega(|f|)^3$ . Given that our ABE has constant size keys and constant-size ciphertexts, its decryption time matches the lower bound of [Luo22] and is optimal.

Trade-Offs Between Succinctness and Fast Decryption. We now describe our lower bounds in more detail. First, consider the efficiency dependency on the length of public information  $f$  and  $x$ . We show that unconditionally, it is impossible to simultaneously have key size sublinear in  $|f|$ , and decryption time with sublinear dependency on  $|f|$ . Similarly, it is impossible to simultaneously have ciphertext size sublinear in  $|x|$ , and decryption time sublinear in  $|x|$ . In fact, these trade-offs exist for the simplest, secret key, 1-key, 1-ciphertext, selectively secure PHFE, and the first trade-off w.r.t.  $|f|$  applies to plain FE. More precisely,

**Theorem 4** (informal). *For a secret-key, 1-key, 1-ciphertext, selectively secure PHFE with the following efficiency, it must hold that  $\alpha \geq 1$  or  $\beta \geq 1$ .*

$$|\text{sk}| = O(|f|^\alpha) \quad \text{and} \quad T_{\text{Dec}} = O(T + |f|^\beta + |y|) \text{poly}(|x|)$$

The same holds for FE where  $|x| = 0$ .

<sup>3</sup>Theorem 14 in [Luo22] proved the first trade-off w.r.t. ciphertext size. Essentially the same proof also gives the second trade-off w.r.t. key size.

**Theorem 5** (informal). *For a secret-key, 1-key, 1-ciphertext, selectively secure PHFE with the following efficiency, it must hold that  $\alpha \geq 1$  or  $\beta \geq 1$ .*

$$|\text{ct}| = O(|x|^\alpha) \text{ poly}(|y|) \quad \text{and} \quad T_{\text{Dec}} = O(T + |f| + |x|^\beta) \text{ poly}(|y|)$$

Our PHFE schemes achieve one profile of optimality,  $\alpha = 0$  and  $\beta = 1$ . A natural question that our work leaves open is whether the other profile of optimality,  $\alpha = 1$  and  $\beta < 1$  or even  $\beta = 0$ , is attainable. Another question is whether decryption time must depend on the length of the private input  $y$ .

Barriers to Decryption Time Optimality. We illustrate barriers to positive answers to the above two questions. We show that a PHFE scheme with decryption time independent of  $|y|$ , or  $|x|$ , or  $|f|$ , implies a primitive called secret-key doubly efficient private information retrieval (SK-DE-PIR), which currently have no construction based on standard assumptions.

**Theorem 9** (informal). *If there is a secret-key, 1-ciphertext, many-key, selectively secure PHFE scheme with following decryption time, there is a secret-key doubly efficient PIR scheme.*

$$T_{\text{Dec}} = O(T^{e_T} + |f|^{e_f} + |x|^{e_x} + |y|^{e_y}), \text{ where } e_x = 0 \text{ or } e_y = 0$$

*If there is a secret-key, many-ciphertext, 1-key, selectively secure PHFE scheme with the following decryption time, there is a secret-key doubly efficient PIR scheme.*

$$T_{\text{Dec}} = O(T^{e_T} + |f|^{e_f} + |x|^{e_x} + |y|^{e_y}), \text{ where } e_f = 0$$

*The same holds for FE where  $|x| = 0$ , w.r.t.  $e_y$  and  $e_f$ .*

SK-DE-PIR, introduced by [BIPW17,CHR17], allows a client to privately encode a database  $D$  into  $\tilde{D}$  while keeping a secret key  $k$ . Later, client can outsource the encoded database  $\tilde{D}$  to a remote storage server, and *obliviously* query the database using  $k$  hiding the logical access pattern. Different from ORAM, the server never updates the encoded database nor keeps any additional state. Different from PIR, SK-DE-PIR allows the database to be privately encoded, in exchange for *double efficiency* – for each query, the complexity of both the client and server is independent of the database size  $|D|$ , whereas PIR necessarily have server complexity  $\Omega(|D|)$ . The double efficiency of SK-DE-PIR makes it highly desirable. However, so far, there are only candidates [BIPW17,CHR17] based on a new conjecture that permuted local-decoding queries (for a Reed-Muller code with suitable parameters) are computationally indistinguishable from uniformly random sets of points. More recently, a simple “toy conjecture” inspired by (but formally unrelated to) these SK-DE-PIR schemes has been broken [BHMW21]. Constructing SK-DE-PIR based on well-studied assumptions (including primitives like  $i\mathcal{O}/\text{FE}$ ) has been left open. Our theorem shows that achieving optimal decryption time in PHFE based on standard assumptions entails resolving this open problem. In fact, even improving the decryption time to be sublinear in  $|f|$ , or  $|y|$ , or  $|x|$  would imply SK-DE-PIR where the server complexity per query is sublinear in database size  $|D|$ . Moreover, having a ciphertext of size  $c|y|$  implies a SK-DE-PIR where the encoded database size is  $|\tilde{D}| = c|D|$  for any  $c$ , a even more challenging task.

Asymptotically Efficient Succinct Garbled RAM, and Constant-Overhead iO. The main tool in our construction of efficient RAM-FE is succinct Garbled RAM (GRAM). Initiated by [KLW15,BGL<sup>+</sup>15,CHJV15], a sequence of works have constructed succinct garbled RAM based on sub-exponentially hard circuit-FE [CH16,CCC<sup>+</sup>16,ACC<sup>+</sup>16,CCHR16] and succinct garbled Turing machines based on (polynomially-hard) circuit-FE [KLW15,AJS17,AL18,GS18b].

In this work, we formulate a new notion of succinct GRAM (informally described in the technical overview and formal definition in Section 2.2) geared for building efficient RAM-FE, and construct it based on circuit-FE. Our construction has two consequences: 1) we obtain the first succinct GRAM (under our or the standard notion) based on polynomial hardness, as opposed to sub-exponential hardness as in prior constructions, and 2) using iO for circuits, we obtain iO for RAM with constant-overhead – the size of the obfuscated program is  $2|M| + \text{poly}(\lambda, l)$ , where  $M$  is the original RAM and  $l$  is the input length. Constant overhead iO was only known for Turing machines before [AJS17].

## 1.1 Technical Overview

We start with an overview of our negative results.

**Unconditional Lower Bound.** As introduced earlier, we show that it is impossible for a secure PHFE to enjoy sublinear dependency on  $|f|$  (resp.  $|x|$ ) simultaneously for  $|\text{sk}_f|$  (resp.  $|\text{ct}_x(y)|$ ) and  $T_{\text{Dec}}$  when  $T_{\text{Dec}}$  is linear in  $T, |x|, |y|$  (resp.  $T, |f|, |y|$ ). We illustrate our ideas of proving the contention between

$$|\text{sk}_f| = O(|f|^\alpha) \quad \text{and} \quad T_{\text{Dec}} = O(T + |f|^\beta + |x| + |y|) \quad \text{for } \alpha < 1 \text{ and } \beta < 1$$

by exhibiting an efficient adversary breaking the security of PHFE for RAM (polynomial factors in the security parameter are ignored).

Adversarial Function and Strategy. The adversary will selectively request one secret key and one ciphertext. Let  $n < N$  be determined later.

- The function  $f$  is described by a string  $R \in \{0, 1\}^N$ .
- There is no public input, so  $x = \perp$ .
- The private input  $y$  is either  $(I \subseteq [N], w \in \{0, 1\}^n)$  or  $z \in \{0, 1\}^n$ , where  $I$  is a set containing  $n$  indexes and  $w$  is a one-time pad.

The functionality is simply *reading and XORing* or *outputting as-is*, i.e.,

$$f(x, y) = \begin{cases} R[I] \oplus w, & \text{if } y = (I, w); \\ z, & \text{if } y = z; \end{cases}$$

where  $R[I]$  means the  $n$  bits of  $R$  at the indices in  $I$ . Clearly,

$$|f| = O(N), \quad |\text{sk}| = O(N^\alpha), \quad |x| = O(1), \quad |y| = O(n), \quad T = O(n), \quad T_{\text{Dec}} = O(n + N^\beta).$$

More precisely,  $|y| = O(n \log n)$ , but the  $\log n$  factor is absorbed by the  $\text{poly}(\lambda)$  factor hidden in the big-O notation.



The adversary chooses

key query  $f$  with  $R \xleftarrow{\$} \{0, 1\}^N$ ,  
 challenge  $x \leftarrow \perp$ ,  $y_0 \xleftarrow{\$} \text{random}(I, w)$ ,  $y_1 \leftarrow z = R[I] \oplus w$ .

By our choice,  $f(x, y_0) = R[I] \oplus w = z = f(x, y_1)$ , so the challenge is well-formed. Upon receiving sk and ct, the adversary simply runs the decryption algorithm on (sk, ct) with random access to the function, i.e.,  $R$ , in the clear. Let  $L \subseteq [N]$  be the set of indexes in  $R$  that is read during decryption.

$$R[I] \oplus w \leftarrow \text{Dec}^{R, x=\perp}(\text{sk}, \text{ct}), \text{ where Dec reads } R[L]$$

The adversary determines that

$$\text{ct is an encryption of } \begin{cases} y_0 = (I, w), & \text{if } |L \cap I| \text{ is large;} \\ y_1 = z, & \text{if } |L \cap I| \text{ is small;} \end{cases}$$

where the threshold for *large* and *small* is described below.

Intuition and Toy Proof. The intuition behind the adversary's strategy is simple. Let  $L_b$  be the index set  $L$  that decryption accessed when decrypting ct encrypting  $y_b$ .

- When ct encrypts  $y_0$ , and decryption must correctly recover  $R[I]$  (as the adversary knows  $w$ ). Decryption can only obtain information of  $R[I]$  either from sk or via accesses to the tape  $R$ . Since  $R[I]$  contains  $n$  bits of information and by setting  $|\text{sk}| = O(N^\alpha) \ll n$ , decryption must read a *large* portion of information of  $R[I]$  from the tape  $R$ , implying  $|L \cap I|$  is large  $\Omega(n)$ .
- In contrast, when ct encrypts  $y_1$ , observe that the joint distribution of  $(R, \text{ct}, \text{sk})$  is independent of  $I$ , as  $w$  is a one-time pad and completely hides  $I$  in  $y_1 = R[I] \oplus w$ . Thus the behavior of, Dec, is independent of  $I$ . Since Dec runs in a short time  $O(n + N^\beta) \ll N$ , without knowing  $I$ , where it reads in  $R$  cannot overlap with  $I$  for a large portion. Therefore,  $|L \cap I|$  is likely to be *small*.

It remains to analyze how large and small  $|L \cap I|$  is in the above two cases. Let us first consider a toy proof, under the hypothesis that sk contains no information about  $R[I]$  at all. We will remove this hypothesis below. By this hypothesis, when ct encrypts  $y_0$ , the decryption algorithm must read the entire  $R[I]$  from the tape  $R$  and hence  $|L \cap I| = n$ . When ct encrypts  $y_1$ , since the indexes  $L_1$  that the decryption algorithm reads from  $R$  is independent of  $I$ ,  $|L_1 \cap I|$  follows a hypergeometric distribution, and hence

$$\mathbb{E}[|L_1 \cap I|] \leq \frac{T_{\text{Dec}} \cdot n}{N} \ll n \quad (1)$$

This means the adversary can distinguish when ct encrypts  $y_0$  or  $y_1$  with good probability, and contradicts the security of PHFE.

Removing the Hypothesis. The hypothesis that  $\text{sk}$  contains no information about  $R[I]$  at all may well be false. When it is removed, we can no longer argue that  $I \subseteq L_0$ , as the adversary may obtain some information of  $R[I]$  from  $\text{sk}$ . Our intuition is that  $|L_0 \cap I| \geq |I| - |\text{sk}|$ , but proving this formally is not trivial as  $\text{sk}$  may contain arbitrary information of  $R[I]$ , such as, its parity.

We employ a compression argument. The basic rationale behind a compression argument is that there is no pair of encoding and decoding algorithms (Encode, Decode), with arbitrarily long shared randomness  $s$ , is able to transmit a  $n$ -bit random string  $u$  (independent of  $s$ ) from one end to the other, via an encoding  $v$  containing less than  $n$  bits. Informally,

$$\text{if } \Pr \left[ \text{Decode}(s, v) = u \mid \begin{array}{l} s \leftarrow \mathcal{D}_s \\ u \leftarrow \{0, 1\}^n \\ v \leftarrow \text{Encode}(s, u) \end{array} \right] = 1, \text{ then } |v| \geq |u|$$

(Formal statement by [DTT10] is in Lemma 8.) We show that suppose  $|L_0 \cap I| < |I| - |\text{sk}|$  then we can design a pair of (Encode, Decode) violating the above information theoretic bound.

- the shared randomness  $s$  contains the PHFE randomness and  $I, w, R[[N] \setminus I]$ .
- To encode  $u \in \{0, 1\}^n$ , Encode sets  $R[I] = u$ . Using  $s$ , it then generates a PHFE key  $\text{sk}$  for  $R$  and a ciphertext  $\text{ct}$  encrypting  $y_0 = (I, w)$ , runs Dec to obtain the locations  $L_0$  in  $R$  that decryption reads. Finally, it sets the codeword to be  $v = (\text{sk}, R[L_0 \cap I])$ .
- To decode, Decode regenerates  $\text{ct}$  using  $s$ , and runs Dec to obtain the output  $z = R[I] \oplus w$  and recover  $u = z \oplus w$ . During decryption, Dec queries for locations  $L_0$  in  $R$ . Every query is in either  $R[[N] \setminus I]$  or  $R[L_0 \cap I]$ ; the former can be answered by finding the right element in  $s$  and the latter in  $v$ .

Suppose  $|L_0 \cap I| < |I| - |\text{sk}|$ , we have  $|v|$  is less than  $|I| = |u|$ , which contradicts the incompressibility of  $u$ . (In the formal proof, we make  $v$  fixed-length and suffer from incorrect decoding, hence the statements are probabilistic. See Section 3.1 for more details.)

Stepping back, the compression argument shows that  $|L_0 \cap I| \geq |I| - |\text{sk}|$ . In contrast, by Equation (1),  $|L_1 \cap I| \leq n/2$  with high probability. To show that the adversary can distinguish  $\text{ct}$  encrypting  $y_0$  or  $y_1$ , we can set, e.g.,  $n = N^{(\alpha+1)/2}$ , so that  $|\text{sk}| = O(N^\alpha) \ll n$ , and  $|L_0 \cap I| \geq |I| - |\text{sk}| \geq n/2$ . (In the formal proof,  $N$  itself is a large poly( $\lambda$ ) to overwhelm poly( $\lambda$ ) factors, which is ignored in this overview.) In summary, any PHFE with both  $|\text{sk}|$  sublinear in  $|f|$  and  $T_{\text{Dec}}$  sublinear in  $|f|$  (and linear in  $T, |x|, |y|$ ) is insecure.

**Computational Implausibility for Fast Decryption.** As described earlier, we show implausibility results against constructing a PHFE scheme with fast decryption. Consider a PHFE scheme where the decryption time is  $O(T_{\varphi(f,x,y)}^{\beta_T} + |f|^{\beta_f} + |x|^{\beta_x} + |y|^{\beta_y})$  for some constants  $\beta_T, \beta_f, \beta_x, \beta_y$ . We show that even if any of  $\beta_x, \beta_y$  or  $\beta_f$  is zero, then such a scheme implies SKDEPIR protocols (an informal description of SKDEPIR is in the introduction and formal definition in Section 3.2).

To illustrate our main idea, we start by describing this transformation for the case when the decryption is efficient in the length of the public input  $x$ , namely when  $\beta_x = 0$ . The ideas naturally extend to the other cases.

Since the decryption is efficient in  $|x|$ , as a first attempt, we set  $DB \in \{0, 1\}^n$  as  $x$ , and  $y$  as empty. The client processes the database  $DB$  by first sampling  $(\text{mpk}, \text{msk})$  for the PHFE scheme and then encodes the database into  $\widetilde{DB} = (x, \text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y)))$  and sends it over to the server. To look up a location  $DB_{i_j}$ , the client can compute a PHFE function key  $\text{sk}_{f_j}$  for the program  $f_j$  that looks up and outputs  $DB_{i_j}$  and sends the key as the query to the Server.

$$\widetilde{DB} = (DB, \text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (DB, \perp))) , \quad \text{query} = \text{sk}_{f_j} \text{ where } f_j^{\text{DB}} = DB_{i_j}$$

The server responds to the query by decrypting  $\text{ct}_{\text{PHFE}}$  in  $\widetilde{DB}$  using the key  $\text{sk}_{f_j}$  and with random access to  $DB$ , and learns  $DB_{i_j}$ . Note that double efficiency requirement is already satisfied. Client only needs to compute a function key  $\text{sk}_{f_j}$  that can be computed in time polynomial in the description length of  $f_j$ , and hence polynomial in  $\lambda$  and  $\log n$ . On the other hand, due to the supposed efficiency of decryption, the decryption time is polynomial in  $T_{f_j(x,y)}, |f_j|, |y|$ , which are also polynomial in  $\lambda$  and  $\log n$ .

While this idea solves the core issue, we have completely missed one important aspect. The scheme reveals the indices  $\{i_j\}$  to the Server as the function key  $\{\text{sk}_{f_j}\}$  is not guaranteed to hide the function description  $\{f_j\}$ . To resolve this issue, we observe that if we had a function hiding PHFE scheme, we would have been done. To enable this, we will use similar techniques as used to convert any functional encryption scheme to a function hiding functional encryption scheme [BS15]. Namely, we will compute a symmetric key encryption SKE of the index  $i$  (denoted as  $\text{ctk}_1$ ). We will hardwire  $\text{ctk}_1$  in the function secret key instead of the index  $i$ . The corresponding secret key  $\text{SKE.sk}_1$  will be put in the private component  $y$ , which will be used to decrypt  $\text{ctk}_1$  to learn index  $i$ . While this might seem to be enough, we face yet another issue. Learning  $DB_{i_j}$  in the clear upon decryption can reveal information about the index  $i_j$  to the Server. To fix this, the decryption will output an encryption of  $DB_{i_j}$  under another secret key  $\text{SKE.sk}_2$  of the secret-key encryption scheme. We will put this key in the private input  $y$  along with a PRF key to derive randomness to compute the encryption.

$$\begin{aligned} \widetilde{DB} &= (DB, \text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (DB, (\text{SKE.sk}_1, \text{SKE.sk}_2, \text{PRF.k})))) , \\ \text{query} &= \text{sk}_{f_j} \text{ where } f_j[\text{ctk}_1(i_j), \$]_{\text{DB}, (\text{SKE.sk}_1, \text{SKE.sk}_2, \text{PRF.k})} = \text{SKE}(\text{SKE.sk}_2, DB_{i_j} ; \text{PRF}(\text{PRF.k}, \$)) \end{aligned}$$

Observe how double efficiency is still preserved. We have increased the complexity of  $f_j$  multiplicatively by a polynomial amount (in  $\lambda$  and  $\log n$ ), similarly the size of  $y$  is also polynomial in  $\lambda$  to store secret keys of SKE and a PRF key. There are few more subtleties : to make the proof go through, we need to use the trojan method in the FE literature [DIJ<sup>+</sup>13] which requires another encryption key  $\text{SKE.sk}_3$  and additional programming.

**Overview of Our FE for RAM.** At a very high level, we use a *succinct garbled RAM (GRAM) scheme* to lift a FE for circuits to a FE for RAMs. This former can be viewed as a 1-key, 1-ciphertext, secret-key FE for RAMs, where succinctness implies that the running time of key generation and encryption is independent of the running time of the RAM computation. The (collusion resistant) FE for circuits is then used to lift the one-time security to many-time security. This high-level approach first appeared in [AS16] for building FE for TMs. In this work, towards nearly optimally efficient FE for RAMs, we first observe that existing definitions and constructions of succinct GRAM [BGL<sup>+</sup>15,

[CHJV15, CH16, CCC<sup>+</sup>16, ACC<sup>+</sup>16, CCHR16] are insufficient. Therefore, we first formulate a new variant of succinct GRAM, termed *laconic GRAM*, and then construct it using polynomially hard FE for circuit. Along the way, we weaken the assumption underlying succinct GRAM schemes from iO, which requires sub-exponentially hard FE for circuit, to polynomially-hard FE for circuits.

Let us first review the syntax and security of standard GRAM schemes. They consist of the following algorithms. The encoding algorithm encodes a database  $D$  into  $\hat{D}$  and outputs a private state  $\tau$ . The garbling algorithm uses  $\tau$  to garble a RAM  $M$  into  $\hat{M}$  and outputs a collection of input labels  $\{L_{i,b}\}_{i,b}$ . The evaluation algorithm given the garbled RAM  $\hat{M}$ , a subset of labels  $L_k$  corresponding to an input  $k$ , and random access to  $\hat{D}$ , returns the output  $M^D(k)$  of the RAM computation. Simulation based security ensures that  $\hat{D}, \hat{M}, L_k = \{L_{i,k_i}\}_i$  can be simulated using only the output  $M^D(k)$ . The efficiency of different algorithms is described below.

$$\begin{aligned} (\hat{D}, \tau) &\leftarrow \text{Encode}(D), & \hat{M}, \{L_{i,b}\}_{i,b} &\leftarrow \text{Garble}(M, \tau), & M^D(k) &= \text{Eval}^{\hat{D}}(\hat{M}, L_k) \\ |\hat{D}| &= |D| \text{ poly}(\lambda), & |\hat{M}| &= \text{poly}(\lambda), & T_{\text{Eval}} &= T \text{ poly}(|M|, \lambda) \end{aligned}$$

We now describe why the standard notion falls short for our purpose of building very efficient FE for RAMs and how to address these issues. An informal definition of our succinct GRAM is provided in Figure 1.

- *many-tape v.s. single-tape*: To start with, we consider RAM computation with multiple tapes  $U^{x,y,f}(1)$  instead of a single tape  $M^D(k)$ .
- *public-tape v.s. private-tape*: Some of the tapes we consider, such as  $x$  and  $f$ , are public. But standard GRAM only provides a mechanism for encoding private tape, and the encoding is necessarily at least as long as the tape itself. However, optimal succinctness requires the FE ciphertext- and key-size to be independent of  $|x|$  and  $|f|$ . Hence we cannot afford to encode  $x, f$  as in standard GRAM. Instead, our new notion of succinct GRAM has a Compress algorithm that compresses the public tapes into hashes/digests  $h_x$  and  $h_f$ ; the Garble algorithm “ties” these hashes to the garbled program  $\hat{U}$ ; and finally Eval makes random access to  $x$  and  $f$  in the clear directly (just as the decryption algorithm of RAM-FE does).
- *rate-2 encoding v.s. rate-poly( $\lambda$ ) encoding of private tape*: Our setting also has private tape, namely  $y$ . But optimal succinctness requires concretely efficient, rate-1 or rate-2, encoding of  $y$ , whereas standard GRAM allows much worse rate poly( $\lambda$ ). To achieve concretely efficient encoding, we can only encrypt  $y$  using a rate-1 encryption scheme. To bind the encryption  $\hat{y}$  with a garbled program, we simply treat  $\hat{y}$  as yet another public tape just like  $x, f$ . In other words, we consider the modified RAM computation  $\tilde{U}^{x,\hat{y},f}(k) = U^{x,y,f}(1)$ , where  $k$  is the secret key of the rate-1 encryption. As such, our succinct GRAM only need to handle public tapes.

In our construction of FE, additional care needs to be taken to ensure that our GRAM security implies that  $y$  is hidden. To achieve this, We rely on existing techniques [NY15], which requires two (rate-1) encryption of  $y$  with independent keys so that different security hybrids can invoke the semantic security of different encryption. This is why our FE has rate-2 dependency on  $|y|$ , instead of rate-1. We omit details in this overview.

- *reusable digests v.s. one-time encoding*: In standard GRAM, the database encoding  $\hat{D}$  can only be used, *once*, by a single garbled program  $\hat{M}$  generated using the right state  $\tau$ . The technical cause of the one-time security of  $\hat{D}$  is due to the use of ORAM in order to hide the access pattern of  $M$  to  $D$ ; the same ORAM storage  $\hat{D}$  cannot be used by multiple garbled programs<sup>4</sup>. The one-time security means that when using succinct GRAM to construct RAM-FE, the decryption of every ciphertext and key pair must generate fresh encoding  $\hat{D}$  (and  $\hat{M}$ ). Such fresh encoding can only be generated using the underlying FE for circuit, by encoding  $D$  in its key or ciphertext, which would lead to large polynomial dependency on  $|D|$ .

Our notion of succinct GRAM compresses the public tapes into hashes  $h_x, h_f, h_y$ . For the same reasons, we cannot afford to generate fresh hashes at decryption of every pair of ciphertext and key. Instead, our hashes are *reusable* – they can be tied to multiple garbled programs; this is ensured by the fact that our Garble algorithm does not take any private state from the Compress algorithm. A technical issue we must resolve is how to hide access pattern to the public tapes  $x, \hat{y}, f$  since they are not encoded using ORAM, which we discuss later.

- *Difference in decryption time*: The reusability comes at a cost. In our new notion, evaluation time is  $(T + |x| + |y| + |f|) \text{poly}(\lambda)$  whereas standard GRAM has evaluation time  $T \text{poly}(|M|, \lambda)$  independent of tape size  $|D|$ . Nevertheless, our lower bound for RAM-FE implies that the dependency on  $|x|, |y|, |f|$  is hard to get around (as our succinct GRAM implies RAM-FE with the same decryption time).
- *RAM with long outputs v.s. single-bit output*: Standard succinct GRAM handles RAM computation with a single-bit output. To handle RAM with  $m$ -bit output, it reduces to creating  $m$  instances of garbled RAM, one for each output bit. Under simulation security, the size of the garbled RAM necessarily grows linearly with the output length  $m$ .

In our notion, we require garbling RAM with arbitrarily long outputs, without efficiency degradation in the output length. To do so, we switch to indistinguishability based security instead of simulation security.

Putting the above pieces together, we formulate succinct GRAM as in Figure 1.

*Our Construction of Succinct GRAM.* One approach towards constructing succinct GRAM for TMs or RAMs is starting from a non-succinct GRAM for TMs or RAMs where the size of the garble program scales with the worst-case time complexity of the TM/RAM  $U$ , into one that is succinct. First introduced in [BGL<sup>+</sup>15], this approach uses iO to obfuscate a circuit that on input an index  $t$ , outputs the  $t$ 'th component in the non-succinct GRAM. If every component of the non-succinct GRAM can be locally generated using a small circuit of size  $\text{poly}(|U|, \lambda)$ , then the obfuscated circuit also has size  $\text{poly}(|U|, \lambda)$  and can be viewed as the succinct garbled program. To prove security, [BGL<sup>+</sup>15] identified that the non-succinct garbling scheme must satisfy another property,

---

<sup>4</sup>This should be distinguished from the scenario of GRAM with persistent database where a sequence of garbled program  $(\hat{M}_1, \hat{M}_2, \dots)^{\hat{D}}$  are executed sequentially with  $\hat{D}$ . The difference lies in that in sequential execution, each garbled RAM can modify  $\hat{D}$  and the changes are kept persistently to the next computation. Here, we are considering the scenario where the unmodified  $\hat{D}$  is used by multiple garbled program, which breaks ORAM security.

### Our Notion of Succinct Garbled RAM

- $\text{Compress}(\tau, D_\tau)$  compresses the  $\tau$ 'th public tape  $D_\tau$  into a short hash digest $_\tau$  of length  $\text{poly}(\lambda)$ . It runs in time  $O(|D_\tau|)$ .
- $\text{Garble}(U, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]})$  outputs a garbled program  $\widehat{U}$  tied with hashes of the public tapes, and pairs of labels  $\{L_{i,b}\}_{i,b}$ . It runs in time  $\text{poly}(\lambda)$  ( $U$  has constant-size).
- $\text{Eval}^{D_1, \dots, D_\mathcal{T}}(U, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}, \widehat{U}, L_k)$  returns the (long) output of RAM computation  $U^{D_1, \dots, D_\mathcal{T}}(k)$  if  $L_k = \{L_{i,k_i}\}_i$ . It runs in time  $(T + \sum_\tau |D_\tau|) \text{poly}(\lambda)$ .

Security guarantees that for two computations  $U^{D_1, \dots, D_\mathcal{T}}(k_0)$  and  $U^{D_1, \dots, D_\mathcal{T}}(k_1)$  with different inputs but identical outputs and running time, the distributions of  $(\widehat{U}, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}, L_{k_0})$  and  $(\widehat{U}, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}, L_{k_1})$  are indistinguishable. This holds when the tapes  $\{D_i\}$  are chosen adaptively dependent on the hashes of previously chosen tapes, before the program  $U$  and inputs  $k_0, k_1$  are chosen.

**Figure 1.** Our notion of succinct GRAM.

articulated later by [AL18], called *local simulation*. Informally, it requires that the non-succinct garbled scheme is proven secure via a sequence of hybrids, where components of every hybrid garbled program can be locally generated using a small circuit, and in neighboring hybrids, only a few components changes. Beyond succinct garbling, local simulation has also found application in achieving adaptive security [BHR12] of garbling schemes. A sequence of works developed local simulation strategies for garbled circuit [HJO<sup>+</sup>16, GS18a], TM [GS18b, AL18], and RAM [GOS18]. Most notably, the work by Garg and Srinivasan [GS18a] introduced a beautiful pebbling technique for realizing local simulation.

Our construction of succinct GRAM proceeds in steps. First, we use the Garg–Srinivasan [GS18a] pebbling technique to obtain a non-succinct GRAM that has a local simulation proof for, however, weak indistinguishability security called fixed memory security. Indistinguishability only holds when the two RAM computations have not only identical outputs and running time, but also identical memory access pattern and content. Then by the same approach of [BGL<sup>+</sup>15, GS18b, AL18], we turn it into a succinct GRAM, still with only *fixed memory security*, relying on iO for polynomial-sized domain which is implied by polynomially-hard FE for circuits. (In the body, we merge these two steps together in Section 4.) Many details need to be ironed out in order to realize our new notion of succinct GRAM. For example, prior works [GS18a, GS18b, AL18] deal with single-bit output RAM and can build intermediate security hybrid where the suffix (i.e., the last certain number of steps) of a computation is simulated by hard-wiring the single-bit output. In contrast, we directly garble RAM with arbitrarily long outputs. Hard-wiring the long output would compromise the local simulation property (since the hybrid garbled program can no longer be locally generated by a small circuit). To avoid this, we build a hybrid GRAM that runs with one input  $k_0$  in the prefix of the computation and with another input  $k_1$  in the suffix (recall that these two inputs produce identical memory). This ensures that the output is always correctly computed, while keeping local simulation. Similar hybrids appeared in [GOS18] for different reasons.

Finally, we lift *fixed memory security* to full security using punctured PRF and ORAM to protect the memory content and access pattern. One issue is that in our succinct GRAM, the public input tapes  $D_1, \dots, D_T$  are not encoded using ORAM, and evaluation is given random access to them in the clear. Yet, to ensure security, evaluation must access these tapes in an oblivious way, independent of the input  $k_0$  or  $k_1$ . To solve this issue, we consider a modified RAM program  $U'$ , which has random access to  $D_1, \dots, D_T$  and additionally a work-tape that contains an ORAM storage that initially contains no elements. The program  $U'$  starts with linearly scanning every tape  $D_\tau$  and inserting every element into the ORAM storage. Only after that, it emulates the execution of the original RAM program  $U$ ; every time  $U$  read from/write to a location in tape  $D_\tau$ ,  $U'$  accesses the corresponding location on its work-tape through ORAM, which hides the access pattern of  $U$ . The intuition is that since the access pattern of  $U'$  is independent of the input, it suffices to garble it using GRAM with fixed memory security. Clearly, the running time of  $U'$  scales linearly with the total length of all tapes  $\sum_\tau |D_\tau|$ . This is why the evaluation time of our succinct GRAM is linear in  $\sum_\tau |D_\tau|$ . Nevertheless, our lower bound shows that this dependency is hard to circumvent. Lastly, we mention that to prove security, one must ensure that the use of ORAM does not hurt local simulation. Fortunately, the techniques by [CH16] provide a solution.

## 1.2 Related Works

Our new construction significantly improve upon the efficiency of prior FE and ABE schemes. The state-of-the-art is summarized in Table 1 and 2; below, we compare with prior works in more detail.

FE for Circuits. The first construction of collusion resistant FE for polynomial-sized circuits is by [GGH<sup>+</sup>13] and based on iO, which in turn relies on subexponential hardness. Later works [GS16, LM16, KNT18, KNTY19] improved the assumption from iO to 1-key FE with sublinearly compact ciphertext,  $|\text{ct}(y)| = \text{poly}(|y|)|T_f|^{1-\varepsilon}$ , where  $\varepsilon$  is a positive constant and  $T_f$  is the circuit complexity of  $f$ . The latter has been recently constructed by [JLS21, JLS22] from the polynomial hardness of three well-studied assumptions. However, these circuit-FE schemes have polynomial efficiency. The only exception is the recent construction by [GWZ22], which has rate-1 ciphertext  $|\text{ct}(y)| = |y| + O(1)$  but still large secret keys  $|\text{sk}_f| = \text{poly}(T_f)$ .

FE for Turing Machines. Several works constructed FE for Turing machines with arbitrary-length inputs, first from the assumption of iO [AS16], then from FE for circuits [AM18], and more recently from 1-key sublinearly compact FE [KNTY19] (which implies collusion resistant FE for circuits). The construction of [AS16] relies on a 1-key 1-ciphertext secret-key FE for TMs, which is essentially a succinct garbling scheme for TMs with indistinguishability security. They constructed it by modifying the succinct garbling for TM of [KLW15]. Later, the works of [AL18, GS18b] improved and simplified the construction of succinct garbling for TM. Following these works, [KNTY19] improved the assumption to the existence of 1-key sublinearly succinct FE, and showed that their garbling scheme can be combined with [AS16] to obtain FE for TMs. On the other hand, [AM18] presented an alternative direct approach to FE for TMs from FE for circuits, that does not use succinct garbling for TM. Prior constructions of FE for TMs [AS16, AM18, KNTY19] focus more on weakening the underlying assumptions, and only show

polynomial efficiency. Examining their schemes, we think they achieve efficiency listed in Table 1.

FE for Bounded-Input RAM. A line of research obtained *bounded-input* iO for Turing machines [KLW15,AJS17,GS18b] and for RAMs [BGL<sup>+</sup>15,CHJV15,CH16,CCC<sup>+</sup>16,ACC<sup>+</sup>16,CCHR16]. Plugging these iO schemes into the [GGH<sup>+</sup>13] construction yields *bounded-input* FE for TMs and RAMs. Unfortunately, these schemes are not full-fledged FE for TMs and RAMs for the following reason: existing iO only handles bounded-input TMs and RAMs in the sense that the obfuscator needs to know the maximal input length  $\max(|y|)$  to the TM/RAM  $f$  being obfuscated. (Constructing iO for unbounded input TMs/RAMs remains an major open question.) Plugging this into the FE construction of [GGH<sup>+</sup>13] gives a scheme where the key generation algorithm needs to know the maximal input length  $\max(|y|)$ , despite that the TM/RAM  $f$  could process arbitrarily long inputs. Such FE is called bounded-input FE. In terms of efficiency, the secret key contains an obfuscated program of size  $\text{poly}(|f|, \max(|y|))$  when using RAM-iO of [CHJV15,CH16], and  $\text{poly}(|f|, \max(|y|), S)$  where  $S$  is the space complexity of  $f$  when using RAM-iO of [BGL<sup>+</sup>15].

In summary, our construction gives the first full-fledged (PH-)FE scheme for RAM computation (with arbitrarily long inputs and outputs), significantly improves the efficiency of prior FE schemes, and matches lower bounds.

ABE for Circuits and Turing Machines. Since FE implies ABE, the above mentioned FE schemes immediately imply ABE with the same level of efficiency. The literature on ABE focuses on constructing ABE from weaker assumptions, or achieving better efficiency, etc. The celebrated works of [GVW13,BGG<sup>+</sup>14] showed that ABE for *bounded-depth* circuits can be constructed from Learning With Errors (LWE). Parameters of these schemes however depends polynomially in the maximal depth  $d$  of the circuits supported, namely,  $|\text{mpk}| = \text{poly}(d)$ ,  $|\text{sk}_f| = \text{poly}(d)$ ,  $|\text{ct}(x, m)| = \text{poly}(d)|x|$ , where the encrypted message  $m$  is a bit, and decryption time is  $T_{\text{Dec}} = \text{poly}(d)T$ . A recent work [LLL22] improved this construction to obtain constant-size keys while keeping the sizes of master public key and ciphertext, but at the cost of additionally relying on the Generic Group Model (GGM). ABE for low-depth computation such as  $\text{NC}^1$  or (arithmetic) Branching programs can be constructed using pairing groups, where several schemes have either constant-size keys *or* constant-size ciphertexts, but never both [ALdP11,YAHK14,Tak14,Att16,ZGT<sup>+</sup>16,AT20,LL20].

The work of [GKP<sup>+</sup>13] constructed ABE for Turing machines and RAMs with constant-size secret keys  $|\text{sk}_f| = O(1)$ , but still large ciphertexts  $|\text{ct}(x, m)| = \text{poly}(|x|)$ . However, their scheme uses SNARK and differing-input iO, which cannot be based on falsifiable assumptions. Another work [AFS19] tries to construct ABE for RAMs from LWE, at the cost of making master public keys, secret keys, and ciphertexts all grow polynomially with the maximal running time of RAMs supported; in other words, this is an ABE for bounded-time RAMs.

In summary, we give the first full-fledged ABE for RAMs from falsifiable assumptions, which has simultaneously constant-size secret keys and constant-size ciphertexts, and best-possible decryption time matching known lower bound [Luo22].



## 2 Preliminaries

Throughout the paper, we denote the security parameter by  $\lambda$ . An ensemble  $X$  of objects is a sequence  $\{X_\lambda\}_{\lambda \in \mathbb{N}}$  of such objects indexed by  $\lambda$ . Except in definitions, we omit  $\lambda$  and (ab-)use the symbol  $X$  of an ensemble to represent the  $\lambda^{\text{th}}$  object  $X_\lambda$ . Let  $S$  be a set and  $n$  a natural number, then  $S^{\leq n}$  is the set of *non-empty* strings over alphabet  $S$  of length at most  $n$ . For a string  $x$ , its  $i^{\text{th}}$  symbol is denoted by  $x[i]$ . As an example, the third *bit* in some sufficiently long  $x \in (\{0, 1\}^2)^{\leq 3}$  is  $x[2][1]$ . For two strings  $x, y$ , we write  $x||y$  for their concatenation. We denote by  $C[x_1]$  the circuit  $C$  with  $x_1$  hardwired as the first portion of input so that  $C[x_1](x_2) = C(x_1, x_2)$ .

### 2.1 Multi-Tape Random Access Machine

We use a model of *multi-tape* random access machines, with several read-only input tapes and one read/write working tape. In addition to the tapes, the machine also has a short input and maintains a (small) state.<sup>5</sup> The behavior of a RAM is described by its step circuit:

$$(\ell_1, \dots, \ell_{\mathcal{T}}, w, \text{oldst}, \text{rdata}) \xrightarrow{\text{CPU}} (\text{done}, \text{newst}, \tau, i, \text{wdata}, \text{out}).$$

The step circuit takes as input *i*) the input tape lengths  $\ell_1, \dots, \ell_{\mathcal{T}}$  and a short input  $w$ , which stay the same throughout the computation, and *ii*) the previous state  $\text{oldst}$  and the last string  $\text{rdata}$  read from the tapes, which change from step to step during the computation. It outputs whether the machine should halt in the flag  $\text{done}$ . If the machine does not halt, it outputs the next state  $\text{newst}$ , the next tape  $\tau$  and address  $i$  to read from. Additionally, if  $\tau$  specifies the working tape, then the step circuit also outputs a string  $\text{wdata}$ , which gets written at address  $i$  of the working tape. Each step also optionally generates  $\text{out}$ , one bit of output.

The execution of a machine  $M$  with  $D_1, \dots, D_{\mathcal{T}}$  on the input tapes and  $w$  as the short input begins by zero-initializing<sup>6</sup> the working tape, the state, and the last-read bit. Throughout the process, the short input stays unchanged and the state is updated by the machine. For each step, the requested location (determined by the output of the previous step) is read, whose value is passed into the last-read bit of the current step, before that location is overwritten. For simplicity, we insist that overwriting happen if and only if the requested tape is the working tape (the input tapes are read-only). The whole sequence of  $\text{out}$ 's, including the timing information of at which step each output bit is produced, is the output of the execution. This process is denoted by  $M^{D_1, \dots, D_{\mathcal{T}}}(w)$ . We now give the formal definition.

**Definition 1** (multi-tape RAM). Let  $\mathcal{T}$  be a natural number. A  $\mathcal{T}$ -tape RAM  $M$  is specified by its step circuit CPU,

$$\text{CPU} : (\ell_1, \dots, \ell_{\mathcal{T}}, w, \text{oldst}, \text{rdata}) \mapsto (\text{done}, \text{newst}, \tau, i, \text{wdata}, \text{out}),$$

which is subject to the constraints below. The set of  $\mathcal{T}$ -tape RAMs is denoted by  $\text{RAM}_{\mathcal{T}}$ .

<sup>5</sup>The distinction between input and state is arbitrary, and many works formulate them as a single entity. We choose to separate them for clearer semantics.

<sup>6</sup>We do not consider persistent memory in this work.

**Addresses and Tapes.** An (*input-tape*) *address* is an  $\ell_{\text{addr}}$ -bit string indexing a *cell*. Each input tape consists of at most  $2^{\ell_{\text{addr}}}$  cells, and each cell is  $\ell_{\text{cell}}$ -bit long, i.e., the content  $D$  of each input tape is in  $(\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}$ . For technical convenience, we allow the working tape to have different address and cell lengths, denoted by  $\ell_{\text{ADDR}}$  and  $\ell_{\text{CELL}}$ , than the input tapes. Without loss of generality (efficiency- and security-wise), we assume  $\ell_{\text{ADDR}} \geq \ell_{\text{addr}}$  and  $\ell_{\text{CELL}} \geq \ell_{\text{cell}}$ . Conceptually, the working tape has exactly  $2^{\ell_{\text{ADDR}}}$  cells.

**Inputs and Outputs of CPU.** The inputs include:

- $\ell_\tau \in [2^{\ell_{\text{addr}}}]$  is the length (in cells) of the  $\tau^{\text{th}}$  input tape.
- $w \in \{0, 1\}^{\ell_{\text{in}}}$  is the short input of length  $\ell_{\text{in}}$ .
- $\text{oldst} \in \{0, 1\}^{\ell_{\text{st}}}$  is the state of length  $\ell_{\text{st}}$ .
- $\text{rdata} \in \{0, 1\}^{\ell_{\text{CELL}}}$  is the string that was read.

Among the outputs, done is a flag indicating whether the machine should halt. If done is set, all of  $\text{newst}, \tau, i, \text{wdata}, \text{out}$  should be  $\perp$ . Otherwise:

- $\text{newst} \in \{0, 1\}^{\ell_{\text{st}}}$  is the new state.
- $\tau \in [\mathcal{T}] \cup \{\text{work}\}$  is the tape to read from (and possibly write to).
  - If  $\tau \in [\mathcal{T}]$ , then  $i \in [\ell_\tau]$  is the address to read from on the  $\tau^{\text{th}}$  input tape, and  $\text{wdata} = \perp$ .
  - If  $\tau = \text{work}$ , then  $i \in [2^{\ell_{\text{ADDR}}}]$  is the address to read from and write to on the working tape, and  $\text{wdata} \in \{0, 1\}^{\ell_{\text{CELL}}}$  is the string to be written.
- $\text{out} \in \{\perp, 0, 1\}$  is an optional output bit.

**Executing RAM.** We now formally present the steps involved in executing RAM. Let  $D_1, \dots, D_{\mathcal{T}} \in (\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}$  be the input tape contents and  $w \in \{0, 1\}^{\ell_{\text{in}}}$  the short input. To execute  $M^{D_1, \dots, D_{\mathcal{T}}}(w)$ :

1. Let  $\text{st}_0 \leftarrow 0^{\ell_{\text{st}}}$ ,  $\text{rdata}_0 \leftarrow 0^{\ell_{\text{CELL}}}$ ,  $D_{\text{work},0} \leftarrow 0^{\ell_{\text{CELL}} 2^{\ell_{\text{ADDR}}}}$ ,  $t \leftarrow 1$ .
2. Let  $(\text{done}_t, \text{st}_t, \tau_t, i_t, \text{wdata}_t, \text{out}_t) \leftarrow \text{CPU}(|D_1|, \dots, |D_{\mathcal{T}}|, w, \text{st}_{t-1}, \text{rdata}_{t-1})$ .
3. If  $\text{done}_t$  is set, the execution halts.
4. If  $\tau_t \in [\mathcal{T}]$ , let  $\text{rdata}_t \leftarrow D_{\tau_t}[i_t] \| 0^{\ell_{\text{CELL}} - \ell_{\text{cell}}}$  and  $D_{\text{work},t} \leftarrow D_{\text{work},t-1}$ .
5. Otherwise,  $\tau_t = \text{work}$ , let  $\text{rdata}_t \leftarrow D_{\text{work},t-1}[i_t]$  and let  $D_{\text{work},t}$  be  $D_{\text{work},t-1}$  with  $D_{\text{work},t}[i_t]$  replaced by  $\text{wdata}_t$ .
6. Let  $t \leftarrow t + 1$  and go back to step 2.

We assume, without loss of generality, that all constraints of RAM are respected during all executions.<sup>7</sup> For a halting execution,

<sup>7</sup>For example, the memory access should never be out of range. Given a step circuit, it is efficient to wrap it inside another circuit checking all the constraints and correcting any error (e.g., by halting immediately when a constraint is violated).

- its *running time*  $T = \text{time}(M, D_1, \dots, D_{\mathcal{T}}, w)$  is the value of  $t$  when it halts;<sup>8</sup>
- its *running space* is  $\text{space}(\dots) = \max_{t < T, \tau_t = \text{work}} i_t$ ;<sup>9</sup>
- its *state sequence* is  $\text{stS}(\dots) = (\text{st}_1, \dots, \text{st}_{T-1})$ ;
- its *address sequence* is  $\text{addrS}(\dots) = (\tau_1, i_1, \dots, \tau_{T-1}, i_{T-1})$ ;
- its *write sequence* is  $\text{writeS}(\dots) = (\text{wdata}_1, \dots, \text{wdata}_{T-1})$ ;
- its *output sequence* is  $\text{outS}(\dots) = (\text{out}_1, \dots, \text{out}_{T-1})$ .

It is worth noting that the machine does *not* necessarily produce non- $\perp$  outputs at the end or even consecutively. For most part of this work, we do not care whether the output sequence is in some particular format. When defining secure evaluation of RAM, we simply require that all information be hidden except the output sequence, which implies that the running time and the time steps when each non- $\perp$  output is produced are not supposed to be hidden as such information is incorporated in the output sequence.

## 2.2 Laconic Garbled RAM

Our notion of garbling RAM laconically involves two steps. First, a short digest is created for each input tape. The digest has length independent of that of the tape itself, and the digest must be computable in linear time. Second, the RAM and the short digests are put together to produce a garbled program and the labels. This procedure runs in time polylogarithmic in the RAM running time. Given a garbled program and one set of labels (selected by the bits of the short input), the evaluation procedure computes the output sequence in time linear in the RAM running time.

We consider indistinguishability-based security for the short input. The input tape contents can be chosen adaptively, but the short input cannot depend on the garbled program (i.e., selectiveness). We also consider several weaker notions, where the running space is bounded and more information about the execution is allowed to be revealed.

**Definition 2 (LGRAM).** Let  $\mathcal{T}$  be a natural number. A *laconic garbling scheme* for  $\mathcal{T}$ -tape RAM consists of 3 algorithms:

- $\text{Compress}(1^\lambda, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, \tau, D_\tau)$  takes as input a cell length  $1^{\ell_{\text{cell}}}$ , an address length  $1^{\ell_{\text{addr}}}$ , an input tape index  $\tau \in [\mathcal{T}]$ , and its content  $D_\tau \in (\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}$ . It outputs a short digest  $\text{digest}_\tau$ . The algorithm runs in time  $|D_\tau| \text{poly}(\lambda, \ell_{\text{cell}}, \ell_{\text{addr}})$ , and its output length is  $\text{poly}(\lambda, \ell_{\text{cell}}, \ell_{\text{addr}})$ .
- $\text{Garble}(1^\lambda, T_{\max}, M, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]})$  takes as input a time bound  $T_{\max} \in \mathbb{N}^*$ , a  $\mathcal{T}$ -tape RAM  $M$ , and  $\mathcal{T}$  input tape digests. It outputs a garbled program  $\widehat{M}$  and  $\ell_{\text{in}}$  pairs of labels  $\{L_{i,b}\}_{i \in [\ell_{\text{in}}], b \in \{0,1\}}$ . The algorithm runs in polynomial time.
- $\text{Eval}^{D_1, \dots, D_{\mathcal{T}}}(1^\lambda, T_{\max}, M, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}, \widehat{M}, \{L_i\}_{i \in [\ell_{\text{in}}]})$  takes as input the time bound  $T_{\max}$ , the  $\mathcal{T}$ -tape RAM  $M$ , the input tape digests, the garbled program  $\widehat{M}$ , and one

<sup>8</sup>If the execution does not halt, we say  $\text{time}(\dots) = +\infty$ .

<sup>9</sup>To be complete,  $\text{space}(\dots)$  is defined to be 1 if the working tape is never accessed.

set of labels. Given random access to the input tapes, it is supposed to compute the output sequence. The algorithm runs in time

$$\left( \min\{T_{\max}, \text{time}(M, D_1, \dots, D_{\mathcal{T}}, w)\} + \sum_{i=1}^{\mathcal{T}} |D_{\tau}| \right) \text{poly}(\lambda, |M|, \log T_{\max}),$$

where  $w$  is the short input corresponding to the labels.

The scheme must be *correct*, i.e., for all  $\lambda, \ell_{\text{cell}}, \ell_{\text{addr}}, \ell_{\text{in}} \in \mathbb{N}, T_{\max} \in \mathbb{N}^*, D_1, \dots, D_{\mathcal{T}} \in (\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}, w \in \{0, 1\}^{\ell_{\text{in}}}$ ,  $\mathcal{T}$ -tape RAM  $M$ , if  $M^{D_1, \dots, D_{\mathcal{T}}}(w)$  halts in time at most  $T_{\max}$ , then

$$\Pr \left[ \begin{array}{l} \text{digest}_{\tau} \stackrel{\$}{\leftarrow} \text{Compress}(1^{\lambda}, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, \tau, D_{\tau}) \quad \forall \tau \in [\mathcal{T}] \\ (\widehat{M}, \{L_{i,b}\}_{i \in [\ell_{\text{in}}], b \in \{0,1\}}) \stackrel{\$}{\leftarrow} \text{Garble}(1^{\lambda}, T_{\max}, M, \{\text{digest}_{\tau}\}_{\tau \in [\mathcal{T}]}) \\ : \\ \text{Eval}^{D_1, \dots, D_{\mathcal{T}}}(1^{\lambda}, T_{\max}, M, \{\text{digest}_{\tau}\}_{\tau \in [\mathcal{T}]}, \widehat{M}, \{L_{i,w[i]}\}_{i \in [\ell_{\text{in}}]}) \\ = \text{outS}(M, D_1, \dots, D_{\mathcal{T}}, w) \end{array} \right] = 1.$$

**Definition 3** (LGRAM security). An LGRAM scheme (Definition 2) is (*tape-adaptively, indistinguishability-based*) *secure* if  $\text{Exp}_{\text{LGRAM}}^0 \approx \text{Exp}_{\text{LGRAM}}^1$ , where  $\text{Exp}_{\text{LGRAM}}^b(1^{\lambda})$  with adversary  $\mathcal{A}$  proceeds as follows:

- **Setup.** Launch  $\mathcal{A}(1^{\lambda})$  and receive from it the cell length  $1^{\ell_{\text{cell}}}$  and the address length  $1^{\ell_{\text{addr}}}$ .
- **Tape Choices.** Repeat this for  $\mathcal{T}$  rounds. In each round,  $\mathcal{A}$  chooses  $\tau \in [\mathcal{T}]$  and  $D_{\tau} \in (\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}$ . Upon receiving such choice, run

$$\text{digest}_{\tau} \stackrel{\$}{\leftarrow} \text{Compress}(1^{\lambda}, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, \tau, D_{\tau})$$

and send  $\text{digest}_{\tau}$  to  $\mathcal{A}$ .

- **Challenge.**  $\mathcal{A}$  chooses the instance running time bound  $1^{\overline{T}}$  (in unary), the time bound  $T_{\max}$  (in binary), a  $\mathcal{T}$ -tape RAM  $M$ , and two inputs  $w_0, w_1$ . Run

$$(\widehat{M}, \{L_{i,b}\}_{i \in [\ell_{\text{in}}], b \in \{0,1\}}) \stackrel{\$}{\leftarrow} \text{Garble}(1^{\lambda}, T_{\max}, M, \{\text{digest}_{\tau}\}_{\tau \in [\mathcal{T}]})$$

and send  $(\widehat{M}, \{L_{i,w_b[i]}\}_{i \in [\ell_{\text{in}}]})$  to  $\mathcal{A}$ .

- **Guess.**  $\mathcal{A}$  outputs a bit  $b'$ . The output of the experiment is  $b'$  if all of the following conditions hold:
  - The  $\tau$ 's in all rounds of **Tape Choices** are distinct.
  - Both  $M^{D_1, \dots, D_{\mathcal{T}}}(w_0)$  and  $M^{D_1, \dots, D_{\mathcal{T}}}(w_1)$  halt in time  $T \leq \overline{T} \leq T_{\max}$  with identical output sequences  $\text{outS}(\dots)$ .

Otherwise, the output is set to 0.

Weaker security notions are obtained by strengthening the second condition in **Guess**.

- *Fixed-address security.* "... with identical  $\text{outS}(\dots)$  and  $\text{addrS}(\dots)$ ".
- *Fixed-memory security.* "... with identical  $\text{outS}(\dots)$ ,  $\text{addrS}(\dots)$ , and  $\text{writeS}(\dots)$ ".

We emphasize that the scheme (Definition 2) must be efficient when  $T_{\max}$  is written in binary. However, security only has to hold for polynomially large instance running time, which is captured by the requirement that the adversary must produce  $1^{\overline{T}}$ , an upper bound of the instance running time in unary.

**Bounded LGRAM.** As an intermediate primitive, we consider *bounded LGRAM*:

**Definition 4** (bounded LGRAM and security). The notion of a *bounded LGRAM scheme* is obtained by modifying Definition 2 as follows:

- The evaluation procedure runs in time

$$\left(T_{\max} + \sum_{i=1}^{\tau} |D_{\tau}| \right) \text{poly}(\lambda, |M|, \log T_{\max}).$$

- Correctness is required only if  $M^{D_1, \dots, D_{\tau}}(w)$  halts using space at most  $T_{\max}$  in time at most  $T_{\max}$ .

Its security notions are obtained by modifying Definition 3 as follows:

- In **Challenge**, the adversary chooses  $1^{T_{\max}}$  (instead of  $1^{\bar{T}}$  and  $T_{\max}$ ).
- In **Guess**, the second condition is strengthened to “.. halt using space  $S \leq T_{\max}$  in time  $T \leq T_{\max}$  with identical ...”

### 2.3 Partially Hiding Functional Encryption and FE for Circuits

We define partially hiding functional encryption (PHFE) with respect to functions of the form

$$\varphi : F \times X \times Y \rightarrow \{\perp\} \cup (\mathbb{N}_+ \times Z),$$

where  $F, X, Y, Z$  are the sets of function description, public input, private input, and output, respectively. Each key is associated with some  $f \in F$ , and each ciphertext encrypts some private input  $y \in Y$  and is associated with some public input  $x \in X$ . The decryptor should be able to recover  $z$  if  $\varphi(f, x, y) = (T, z)$ , in which case  $T$  is the time to compute  $z$  from  $f, x, y$  in the clear and serves as a baseline for decryption efficiency. For security, we only consider  $f, x, y$  for which  $T$  is polynomially bounded. On the other hand, when  $\varphi(f, x, y) = \perp$ , we require neither correctness nor security. This can be used to exclude non-halting computation.

**Definition 5** (PHFE). Let  $\Phi = \{\Phi_{\lambda}\}_{\lambda \in \mathbb{N}}$  be a sequence of functionality families such that each  $\varphi \in \Phi_{\lambda}$  is a function

$$\varphi : F_{\varphi} \times X_{\varphi} \times Y_{\varphi} \rightarrow \{\perp\} \cup (\mathbb{N}_+ \times Z_{\varphi}).$$

A *partially hiding functional encryption scheme* for  $\Phi$  consists of 4 algorithms, with efficiency defined in Definition 6:

- $\text{Setup}(1^{\lambda}, \varphi)$  takes a functionality  $\varphi \in \Phi_{\lambda}$  as input, and outputs a pair of master keys  $(\text{mpk}, \text{msk})$ .
- $\text{KeyGen}(1^{\lambda}, \text{msk}, f)$  takes as input the master secret key  $\text{msk}$  and a function description  $f \in F_{\varphi}$ . It outputs a secret key  $\text{sk}_f$ .
- $\text{Enc}(1^{\lambda}, \text{mpk}, x, y)$  takes as input the master public key  $\text{mpk}$ , a public input  $x \in X_{\varphi}$ , and a private input  $y \in Y_{\varphi}$ . It outputs a ciphertext  $\text{ct}_x$ .

- $\text{Dec}^{f,x,\text{sk}_f,\text{ct}_x}(1^\lambda, \text{mpk})$  takes the master public key  $\text{mpk}$  as input and is given random access to the function description  $f$ , the public input  $x$ , the secret key  $\text{sk}_f$ , and the ciphertext  $\text{ct}_x$ . It is supposed to compute  $z$  in  $\varphi(f, x, y) = (T, z)$  efficiently.

The scheme is *correct* if for all  $\lambda \in \mathbb{N}$ ,  $\varphi \in \Phi_\lambda$ ,  $f \in F_\varphi$ ,  $x \in X_\varphi$ ,  $y \in Y_\varphi$  such that  $\varphi(f, x, y) = (T, z) \neq \perp$ , it holds that

$$\Pr \left[ \begin{array}{l} (\text{mpk}, \text{msk}) \xleftarrow{\$} \text{Setup}(1^\lambda, \varphi) \\ \text{sk}_f \xleftarrow{\$} \text{KeyGen}(1^\lambda, \text{msk}, f) : \text{Dec}^{f,x,\text{sk}_f,\text{ct}_x}(1^\lambda, \text{mpk}) = z \\ \text{ct}_x \xleftarrow{\$} \text{Enc}(1^\lambda, \text{mpk}, x, y) \end{array} \right] = 1.$$

**Definition 6** (PHFE efficiency). The basic efficiency requirements of a PHFE scheme (Definition 5) are as follows:

- Setup, KeyGen, Enc are polynomial-time.
- Dec runs in time  $\text{poly}(\lambda, |\varphi|, |f|, |x|, |y|, T)$ , where  $\varphi(M, f, x, y) = (T, z) \neq \perp$ .

The following time-efficiency properties are considered for a PHFE:

- It has *linear-time* KeyGen [resp. Enc] if KeyGen [resp. Enc] runs in time  $|f| \text{poly}(\lambda, |\varphi|)$  [resp.  $(|x| + |y|) \text{poly}(\lambda, |\varphi|)$ ];
- It has  $(T^{e_T} + |f|^{e_f} + |x|^{e_x} + |y|^{e_y})$ -time Dec (for constants  $e_T, e_f, e_x, e_y$ ) if Dec runs in time

$$(T^{e_T} + |f|^{e_f} + |x|^{e_x} + |y|^{e_y}) \text{poly}(\lambda, |\varphi|),$$

where  $\varphi(M, f, x, y) = (T, z) \neq \perp$ . Furthermore, the scheme has *f-fast* [resp. *x-fast*, *y-fast*] Dec if it has  $(T^{e_T} + |f|^{e_f} + |x|^{e_x} + |y|^{e_y})$ -time Dec with  $e_f = 0$  [resp.  $e_x = 0, e_y = 0$ ].

The following size-efficiency properties are considered for PHFE:

- It is *f-succinct* if  $|\text{sk}_f| = \text{poly}(\lambda, |\varphi|)$ , independent of  $|f|$ .
- It is *x-succinct* if  $|\text{ct}_x| = \text{poly}(\lambda, |\varphi|, |y|)$ , independent of  $|x|$ . Furthermore, it has *rate-c ciphertext* for some constant  $c$  if  $|\text{ct}_x| = \text{poly}(\lambda, |\varphi|) + c|y|$ ;

We consider adaptive IND-CPA for polynomially bounded  $T$ :

**Definition 7** (PHFE security). A PHFE scheme (Definition 5) is (*adaptively IND-CPA*) secure if  $\text{Exp}_{\text{PHFE}}^0 \approx \text{Exp}_{\text{PHFE}}^1$ , where  $\text{Exp}_{\text{PHFE}}^b(1^\lambda)$  with adversary  $\mathcal{A}$  proceeds as follows:

- **Setup.** Launch  $\mathcal{A}(1^\lambda)$  and receive from it a functionality  $\varphi \in \Phi_\lambda$  and a time bound  $1^{\bar{T}}$ . Run

$$(\text{mpk}, \text{msk}) \xleftarrow{\$} \text{Setup}(1^\lambda, \varphi)$$

and send  $\text{mpk}$  to  $\mathcal{A}$ .

- **Query I.** Repeat the following for arbitrarily many rounds determined by  $\mathcal{A}$ . In each round,  $\mathcal{A}$  submits a function description  $f_q \in F_\varphi$ . Upon receiving the query, run

$$\text{sk}_q \xleftarrow{\$} \text{KeyGen}(1^\lambda, \text{msk}, f_q)$$

and send  $\text{sk}_q$  to  $\mathcal{A}$ .

- **Challenge.**  $\mathcal{A}$  submits  $x \in X_\varphi, y_0, y_1 \in Y_\varphi$ . Upon the challenge, run

$$\text{ct} \stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, \text{mpk}, x, y_b)$$

and send ct to  $\mathcal{A}$ .

- **Query II.** Same as Query I.
- **Guess.**  $\mathcal{A}$  outputs a bit  $b'$ . The outcome of the experiment is  $b'$  if

$$\begin{aligned} & |y_0| = |y_1|, \\ & \text{and } \varphi(f_q, x, y_0) = \varphi(f_q, x, y_1) = (T_q, z_q) \neq \perp && \text{for all } q, \\ & \text{and } T_q \leq \bar{T} && \text{for all } q. \end{aligned}$$

Otherwise, the outcome is set to 0.

We will use FE for circuits as a building block:

**Definition 8** (FE for circuits). An *FE scheme for circuits* is a PHFE scheme (Definition 5) for

$$\begin{aligned} \Phi &= \{\Phi_\lambda\}_{\lambda \in \mathbb{N}}, & \Phi_\lambda &= \{\varphi_{\ell,s}\}_{\ell,s \in \mathbb{N}_+}, \\ \varphi_{\ell,s} &: F_{\ell,s} \times X \times Y_\ell \rightarrow \{\perp\} \cup (\mathbb{N}_+ \times Z), \\ F_{\ell,s} &= \{\text{circuits of input length } \ell \text{ and size at most } s\}, \\ X &= \{\perp\}, & Y_\ell &= \{0, 1\}^\ell, & Z &= \{0, 1\}^*, \\ \varphi_{\ell,s}(f, \perp, y) &= (1, f(y)), \end{aligned}$$

where the functionality  $\varphi_{\ell,s}$  is represented by  $(1^\ell, 1^s)$ .

We remark that the first output of  $\varphi_{\ell,s}$  is just a dummy value and all efficiency parameters (Definition 6) are always allowed arbitrary polynomial dependency on  $\lambda, \ell, s$  by our choice of representing  $\varphi_{\ell,s}$  by  $(1^\ell, 1^s)$ . This is intended as we use FE for circuits as a building block and do not wish to start with too strong a scheme. In fact, thanks to many beautiful prior works in the literature, we can further weaken the assumption all the way down to an “obfuscation-minimum FE” with only polynomial security.

**Lemma 1** ([KNTY19]). *Assuming the existence of “weakly selective secure, 1-key, sublinearly succinct FE for circuits” (per definitions in [KNTY19]), i.e., so-called “obfuscation-minimum FE with only polynomial security”, there exists an FE scheme for circuits (Definition 8) with adaptive IND-CPA security (Definition 7).*

## 2.4 Universal RAM and PHFE for RAM

In this section, we define PHFE for RAM after explaining some rationales.

We will employ the standard transformation [QWW18] of using FE for circuits to compute LGRAM to obtain PHFE for RAM. However, in LGRAM (Definition 2), the encryption time, as well as the ciphertext size, could depend on the size of the machine  $M$ . This dependency is inherited by the key size of the resultant PHFE for RAM if we associate each key with a RAM. To make the scheme  $f$ -succinct, we fix some universal RAM  $U$  upon setting up the scheme, and associate with each key a piece of

code interpreted by  $U$ . The size of  $U$  is some fixed  $\text{poly}(\lambda)$ ,<sup>10</sup> and so is the LGRAM garbling time. This translates to  $f$ -succinctness in the constructed PHFE for RAM.

The other issue is that LGRAM puts an upper bound on the running time and its incorrectness in case of exceeding time limit is propagated to the PHFE scheme. We avoid it<sup>11</sup> by defining  $\varphi = \perp$  if the running time exceeds some super-polynomial value prescribed upon set-up.

The above explains the intended usage of our PHFE for RAM, yet we define it for general RAM. Moreover, as a stepping stone, we will first consider PHFE for RAM with *bounded private input*, where the private input is simply the short input to the RAM:

**Definition 9** (PHFE for RAM with bounded private input). A *PHFE scheme for RAM with bounded private input* is a PHFE scheme (Definition 5) for

$$\begin{aligned} \Phi &= \{\Phi_\lambda\}_{\lambda \in \mathbb{N}}, & \Phi_\lambda &= \{\varphi_{M, T_{\max}}\}_{M \in \text{RAM}_2, T_{\max} \in \mathbb{N}^*}, \\ \varphi_{M, T_{\max}} &: F_M \times X_M \times Y_M \rightarrow \{\perp\} \cup (\mathbb{N}_+ \times Z), \\ F_M = X_M &= (\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}, & Y_M &= \{0, 1\}^{\ell_{\text{in}}}, & Z &= \{\perp, 0, 1\}^*, \\ \varphi_{M, T_{\max}}(f, x, y) &= \begin{cases} (T, \text{outS}(M, f, x, y)), & \text{if } \text{time}(M, f, x, y) = T \leq T_{\max}; \\ \perp, & \text{otherwise;} \end{cases} \end{aligned}$$

where  $\varphi_{M, T_{\max}}$  is represented by  $(M, T_{\max})$ .

In a full-fledged PHFE for RAM, the RAM no longer has a short input, and the private input is encoded on a tape:

**Definition 10** (full-fledged PHFE for RAM). A *full-fledged PHFE scheme for RAM* is a PHFE scheme (Definition 5) for

$$\begin{aligned} \Phi &= \{\Phi_\lambda\}_{\lambda \in \mathbb{N}}, & \Phi_\lambda &= \{\varphi_{M, T_{\max}}\}_{M \in \text{RAM}_2 \text{ (with } \ell_{\text{in}}=0) \text{ and } T_{\max} \in \mathbb{N}^*}, \\ \varphi_{M, T_{\max}} &: F_M \times X_M \times Y_M \rightarrow \{\perp\} \cup (\mathbb{N}_+ \times Z), \\ F_M = X_M = Y_M &= (\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}, & Z &= \{\perp, 0, 1\}^*, \\ \varphi_{M, T_{\max}}(f, x, y) &= \begin{cases} (T, \text{outS}(M, f, x || y, \varepsilon)), & \\ \quad \text{if } |x| + |y| \leq 2^{\ell_{\text{addr}}} \text{ and } \text{time}(M, f, x || y, \varepsilon) = T \leq T_{\max}; \\ \perp, & \text{otherwise;} \end{cases} \end{aligned}$$

where  $\varepsilon$  is the empty string and  $\varphi_{M, T_{\max}}$  is represented by  $(M, T_{\max})$ .

## 2.5 Indistinguishability Obfuscation

We will use indistinguishability obfuscation for circuits *with polynomial-sized domains* as a building block.

<sup>10</sup> $U$  is not a fixed RAM — its input address length should be  $\omega(\log \lambda)$  to accommodate all polynomially long input.

<sup>11</sup>An alternative solution is to blatantly reveal everything if the running time is too large so that correctness in that case can be implemented by executing the machine in the clear. Security is not affected because the adversary is not allowed to choose keys and ciphertexts with super-polynomial instance running time. However, non-halting computation still needs to be handled separately.



**Definition 11** ( $i\mathcal{O}$ ). An *indistinguishability obfuscator* is an efficient algorithm  $i\mathcal{O}(1^\lambda, C)$  taking a circuit  $C$  as input. It outputs an obfuscated circuit  $\widetilde{C}$ .

The obfuscator must be *correct*, i.e., for all  $\lambda \in \mathbb{N}$ , circuit  $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ , input  $x \in \{0, 1\}^n$ , it holds that

$$\Pr[i\mathcal{O}(1^\lambda, C)(x) = C(x)] = 1.$$

The obfuscator is *secure for polynomial-sized domains* if for all polynomial-sized  $(1^{2^n}, C_0, C_1)$  such that  $C_0, C_1 : \{0, 1\}^n \rightarrow \{0, 1\}^m$  have identical truth tables and sizes, it holds that

$$\{1^\lambda, C_0, C_1, i\mathcal{O}(1^\lambda, C_0)\} \approx \{1^\lambda, C_0, C_1, i\mathcal{O}(1^\lambda, C_1)\}.$$

We remark that the above definition does not allow the obfuscator to work by simply outputting the truth table, as the constraint of having polynomial-sized domains only applies to security, not efficiency nor correctness. Such an obfuscator is implied by the existence of FE for circuits (with polynomial security loss), either via a tight security reduction of FE to  $i\mathcal{O}$  transformation [LT17] or via decomposable obfuscation [LZ17]:

**Lemma 2** ([LT17, LZ17]). *If there exists a secure FE for circuits (Definition 8), there exists an indistinguishability obfuscator for circuits that is secure for polynomial-sized domains.*

## 2.6 Laconic Oblivious Transfer

**Definition 12** (LOT [CDG<sup>+</sup>17]). A *laconic oblivious transfer scheme* consists of 4 algorithms:

- $\text{HashGen}(1^\lambda, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}})$  takes as input the cell length  $1^{\ell_{\text{cell}}}$  and the address length  $1^{\ell_{\text{addr}}}$ , and outputs a hash key  $\text{hk}$ . It runs in polynomial time.
- $\text{Hash}(1^\lambda, \text{hk}, D)$  takes as input the hash key  $\text{hk}$  and a database  $D \in (\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}$ , and *deterministically* outputs a hash  $h$  and a processed database  $\widehat{D}$ . It runs in time  $|D| \text{poly}(\lambda, \ell_{\text{cell}}, \ell_{\text{addr}})$ , and the hash length is  $\ell_{\text{hash}} = \text{poly}(\lambda, \ell_{\text{cell}}, \ell_{\text{addr}})$ .
- $\text{SendRead}(1^\lambda, \text{hk}, h, i, \{m_j^b\}_{j \in [\ell_{\text{cell}}]}^{b \in \{0, 1\}})$  takes as input the hash key  $\text{hk}$ , a hash  $h$ , an address  $i \in [2^{\ell_{\text{addr}}}]$ , and  $\ell_{\text{cell}}$  pairs of messages of identical length. It outputs a read ciphertext  $\text{rct}$ . The algorithm runs in polynomial time.
- $\text{RecvRead}^{\widehat{D}}(1^\lambda, \text{hk}, h, i, \text{rct})$  takes as input the hash key  $\text{hk}$ , the hash  $h$ , the address  $i$ , and the read ciphertext  $\text{rct}$ . Given random access to the processed database  $\widehat{D}$ , it is supposed to recover  $\{m_j^{D[i][j]}\}_{j \in [\ell_{\text{cell}}]}$ . The algorithm runs in time  $\text{poly}(\lambda, \ell_{\text{cell}}, \ell_{\text{addr}}, |\text{rct}|)$ .

The scheme must be *correct*, i.e., for all  $\lambda, \ell_{\text{cell}}, \ell_{\text{addr}} \in \mathbb{N}$ ,  $D \in (\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}$ ,  $i \in [|D|]$ ,  $\{m_j^b\}_{j \in [\ell_{\text{cell}}]}^{b \in \{0, 1\}}$  of identical length, it holds that

$$\Pr \left[ \begin{array}{l} \text{hk} \xleftarrow{\$} \text{HashGen}(1^\lambda, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}) \\ (h, \widehat{D}) \leftarrow \text{Hash}(1^\lambda, \text{hk}, D) \\ \text{rct} \xleftarrow{\$} \text{SendRead}(1^\lambda, \text{hk}, h, i, \{m_j^b\}_{j \in [\ell_{\text{cell}}]}^{b \in \{0, 1\}}) \\ \quad : \text{RecvRead}^{\widehat{D}}(1^\lambda, \text{hk}, h, i, \text{rct}) = \{m_j^{D[i][j]}\}_{j \in [\ell_{\text{cell}}]} \end{array} \right] = 1.$$

An *updatable* LOT has 2 additional algorithms:

- $\text{SendWrite}(1^\lambda, \text{hk}, h, i, \text{wdata}, \{m_j^b\}_{j \in [\ell_{\text{hash}}]}^{b \in \{0,1\}})$  takes as input the hash key  $\text{hk}$ , the hash  $h$ , the address  $i$ , an overwriting string  $\text{wdata} \in \{0,1\}^{\ell_{\text{cell}}}$ , and  $\ell_{\text{hash}}$  pairs of messages of identical length. It outputs a write ciphertext  $\text{wct}$ . The algorithm runs in polynomial time.
- $\text{RecvWrite}^{\widehat{D}}(1^\lambda, \text{hk}, h, i, \text{wdata}, \text{wct})$  takes as input the hash key  $\text{hk}$ , the hash  $h$ , the address  $i$ , the overwriting string  $\text{wdata}$ , and the write ciphertext  $\text{wct}$ . Given random access to the processed database  $\widehat{D}$ , it is supposed to update  $\widehat{D}$  to  $\widehat{D}'$  and recover  $\{m_j^{h'[j]}\}_{j \in [\ell_{\text{hash}}]}$ , where  $D'$  is  $D$  with  $D'[i]$  replaced by  $\text{wdata}$  and  $h'$  is the hash of  $D'$ . The algorithm runs in time  $\text{poly}(\lambda, \ell_{\text{cell}}, \ell_{\text{addr}}, |\text{wct}|)$ .

It is required that for all  $\lambda, \ell_{\text{cell}}, \ell_{\text{addr}} \in \mathbb{N}$ ,  $D \in (\{0,1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}$ ,  $i \in [|D|]$ ,  $\text{wdata} \in \{0,1\}^{\ell_{\text{cell}}}$ ,  $\{m_j^b\}_{j \in [\ell_{\text{hash}}]}^{b \in \{0,1\}}$  of identical length, it holds that

$$\Pr \left[ \begin{array}{l} \text{hk} \xleftarrow{\$} \text{HashGen}(1^\lambda, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}) \\ (h, \widehat{D}) \leftarrow \text{Hash}(1^\lambda, \text{hk}, D), \quad (h', \widehat{D}') \leftarrow \text{Hash}(1^\lambda, \text{hk}, D') \\ \text{wct} \xleftarrow{\$} \text{SendWrite}(1^\lambda, \text{hk}, h, i, \text{wdata}, \{m_j^b\}_{j \in [\ell_{\text{hash}}]}^{b \in \{0,1\}}) \\ \quad : \text{RecvWrite}^{\widehat{D}}(1^\lambda, \text{hk}, h, i, \text{wdata}, \text{wct}) = \{m_j^{h'[j]}\}_{j \in [\ell_{\text{hash}}]} \\ \text{and } \widehat{D} \text{ is updated to } \widehat{D}' \text{ by RecvWrite} \end{array} \right] = 1,$$

where  $D'$  is  $D$  with  $D'[i]$  replaced by  $\text{wdata}$ .

Part of our LGRAM garbling procedure involves hashing the all-zero string (the initial working tape) under a newly generated LOT hash key. It must be done very efficiently:

**Definition 13** (LOT fast initialization). An LOT scheme (Definition 12) has *fast initialization* if there is an efficient algorithm  $\text{Hash0s}(1^\lambda, \text{hk}, S)$  computing the hash of  $0^{\ell_{\text{cell}}S}$  given the hash key and the length (in binary). More precisely, for all  $\lambda, \ell_{\text{cell}}, \ell_{\text{addr}} \in \mathbb{N}$ ,  $S \in [2^{\ell_{\text{addr}}}]$ , it holds that

$$\Pr \left[ \begin{array}{l} \text{hk} \xleftarrow{\$} \text{HashGen}(1^\lambda, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}) \\ (h, \widehat{D}) \leftarrow \text{Hash}(1^\lambda, \text{hk}, 0^{\ell_{\text{cell}}S}) \end{array} : \text{Hash0s}(1^\lambda, \text{hk}, S) = h \right] = 1.$$

The generic bootstrapping procedure [CDG<sup>+</sup>17, AL18, KNTY19] for LOT using Merkle tree always yields a scheme with fast initialization, because the Merkle hash of an all-zero string of length  $S$  can be computed in time  $\text{poly}(\lambda, \log S)$ .

**Security.** Following [KNTY19], we consider database-selective security for LOT.

**Definition 14** (LOT read security [CDG<sup>+</sup>17]). An LOT scheme (Definition 12) is *read-secure* if there exists an efficient simulator  $\text{SimRead}$  such that for all polynomial-sized  $1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}$ ,  $D \in (\{0,1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}$ ,  $i \in [|D|]$ ,  $\{m_j^b\}_{j \in [\ell_{\text{cell}}]}^{b \in \{0,1\}}$  of identical length, it holds that

$$\begin{aligned} \{1^\lambda, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, \text{hk}, D, i, \{m_j^b\}_{j \in [\ell_{\text{cell}}]}^{b \in \{0,1\}}, \text{rct}\} &\approx \{\dots, \widetilde{\text{rct}}\}, \text{ where} \\ \text{hk} &\xleftarrow{\$} \text{HashGen}(1^\lambda, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}), \quad (h, \widehat{D}) \leftarrow \text{Hash}(1^\lambda, \text{hk}, D), \\ \text{rct} &\xleftarrow{\$} \text{SendRead}(1^\lambda, \text{hk}, h, i, \{m_j^b\}_{j \in [\ell_{\text{cell}}]}^{b \in \{0,1\}}), \\ \widetilde{\text{rct}} &\xleftarrow{\$} \text{SimRead}(1^\lambda, \text{hk}, D, i, \{m_j^{D[i][j]}\}_{j \in [\ell_{\text{cell}}]}). \end{aligned}$$

**Definition 15** (LOT write security [CDG<sup>+</sup>17]). An updatable LOT scheme (Definition 12) is *write-secure* if there exists an efficient simulator  $\text{SimWrite}$  such that for all polynomial-sized  $1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, D \in (\{0, 1\}^{\ell_{\text{cell}}})^{\leq 2^{\ell_{\text{addr}}}}, i \in [|D|], \text{wdata} \in \{0, 1\}^{\ell_{\text{cell}}}, \{m_j^b\}_{j \in [\ell_{\text{hash}}]}^{b \in \{0, 1\}}$  of identical length, letting  $D'$  be  $D$  with  $D'[i]$  replaced by  $\text{wdata}$ , it holds that

$$\begin{aligned} \{1^\lambda, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, \text{hk}, D, i, \text{wdata}, \{m_j^b\}_{j \in [\ell_{\text{hash}}]}^{b \in \{0, 1\}}, \text{wct}\} &\approx \{\dots, \widetilde{\text{wct}}\}, \text{ where} \\ \text{hk} &\stackrel{\$}{\leftarrow} \text{HashGen}(1^\lambda, 1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}), \\ (h, \widehat{D}) &\leftarrow \text{Hash}(1^\lambda, \text{hk}, D), \quad (h', \widehat{D}') \leftarrow \text{Hash}(1^\lambda, \text{hk}, D'), \\ \text{wct} &\stackrel{\$}{\leftarrow} \text{SendWrite}(1^\lambda, \text{hk}, h, i, \text{wdata}, \{m_j^b\}_{j \in [\ell_{\text{hash}}]}^{b \in \{0, 1\}}), \\ \widetilde{\text{wct}} &\stackrel{\$}{\leftarrow} \text{SimWrite}(1^\lambda, \text{hk}, D, i, \text{wdata}, \{m_j^{h'[j]}\}_{j \in [\ell_{\text{hash}}]}). \end{aligned}$$

We remark that the above definitions is also index-selective and message-selective. Without loss of generality, we may assume that security holds even if the adversary is allowed to choose  $i$  and  $m_j^b$ 's adaptively (depending on  $\text{hk}$ ). Adaptive security with respect to  $i$  can be obtained by a standard guessing argument as noted in [GOS18], and that with respect to  $m_j^b$ 's by using the scheme as a key encapsulation mechanism (or by encrypting bit by bit and following a standard hybrid argument).

**Lemma 3** ([LZ17, CDG<sup>+</sup>17, AL18, KNTY19]). *Assuming the existence of FE for circuits (Definition 8), there exists an updatable LOT with fast initialization that is both read-secure and write-secure.*

## 2.7 Garbled Circuits

Following the formulation in [GS18a], we make the garbling procedure to take the labels as input instead of letting it generate them. However, their formulation cannot be perfectly correct, which we fix by incorporating the point-and-permute [BMR90] technique:

**Definition 16** (garbled circuits [Yao82]). A *circuit garbling scheme* (with label length  $\ell_L(\lambda)$ ) consists of 2 efficient algorithms:

- $\text{Garble}(1^\lambda, C, \pi, \{L_{i,b}\}_{i \in [n], b \in \{0, 1\}})$  takes as input a circuit  $C$  of input length  $n$ , a point-and-permute string  $\pi \in \{0, 1\}^n$ , and  $n$  pairs of labels, each of length  $\ell_L(\lambda)$ . It outputs a garbled circuit  $\widehat{C}$ .
- $\text{Eval}(1^\lambda, \widehat{C}, x \oplus \pi, \{L_i\}_{i \in [n]})$  takes as input the garbled circuit  $\widehat{C}$ , the permuted input, and  $n$  labels. It is supposed to compute  $C(x)$ .

The scheme must be *correct*, i.e., for all  $\lambda, n \in \mathbb{N}$ , circuit  $C$  of input length  $n$ , input  $x \in \{0, 1\}^n$ , it holds that

$$\Pr \left[ \begin{array}{l} \pi \stackrel{\$}{\leftarrow} \{0, 1\}^n \\ L_{i,b} \stackrel{\$}{\leftarrow} \{0, 1\}^{\ell_L(\lambda)} \text{ for } i \in [n], b \in \{0, 1\} \\ \widehat{C} \stackrel{\$}{\leftarrow} \text{Garble}(1^\lambda, C, \pi, \{L_{i,b}\}_{i \in [n], b \in \{0, 1\}}) \\ y \leftarrow \text{Eval}(1^\lambda, \widehat{C}, x \oplus \pi, \{L_{i,x[i]}\}_{i \in [n]}) \end{array} : y = C(x) \right] = 1.$$

The scheme is *secure* if there exists a simulator  $\widetilde{\text{Garble}}$  such that for all polynomial-sized  $C, x$ , the following distributions are indistinguishable:

$$\begin{aligned} & \{1^\lambda, C, x, \text{Garble}(1^\lambda, C, \pi, \{L_{i,b}\}_{i \in [n], b \in \{0,1\}}), x \oplus \pi, \{L_{i,x[i]}\}_{i \in [n]}\} \\ & \approx \{1^\lambda, C, x, \widetilde{\text{Garble}}(1^\lambda, 1^{|C|}, C(x), x \oplus \pi, \{L_{i,x[i]}\}_{i \in [n]}), x \oplus \pi, \{L_{i,x[i]}\}_{i \in [n]}\}, \end{aligned}$$

where  $L_{i,b} \xleftarrow{\$} \{0,1\}^{\ell_L(\lambda)}$  for all  $i \in [n], b \in \{0,1\}$  and  $\pi \xleftarrow{\$} \{0,1\}^n$ .

## 2.8 Puncturable Pseudorandom Function

We need puncturable pseudorandom functions as a building block.<sup>12</sup>

**Definition 17** (puncturable PRF [BW13]). A *puncturable pseudorandom function* with key (resp. input, output) length  $\ell_{\text{key}}(\lambda)$  (resp.  $\ell_{\text{in}}(\lambda), \ell_{\text{out}}(\lambda)$ ); all polynomial in  $\lambda$ ) consists of 2 efficient algorithms:

- $\text{Eval}(1^\lambda, k, x)$  takes as input a (punctured or not) key  $k$ , an input  $x \in \{0,1\}^{\ell_{\text{in}}(\lambda)}$ , and *deterministically* outputs a bit-string of length  $\ell_{\text{out}}(\lambda)$ .
- $\text{Puncture}(1^\lambda, k, S)$  takes as input a non-punctured key  $k \in \{0,1\}^{\ell_{\text{key}}(\lambda)}$  and a set  $S \subseteq \{0,1\}^{\ell_{\text{in}}(\lambda)}$ . It outputs a punctured key  $\mathring{k}_S$ .

The scheme must be *correct*, i.e., for all  $\lambda \in \mathbb{N}$ ,  $k \in \{0,1\}^{\ell_{\text{key}}(\lambda)}$ ,  $S \subseteq \{0,1\}^{\ell_{\text{in}}(\lambda)}$ ,  $x \in \{0,1\}^{\ell_{\text{in}}(\lambda)} \setminus S$ , it holds that

$$\Pr[\text{Eval}(1^\lambda, \text{Puncture}(1^\lambda, k, S), x) = \text{Eval}(1^\lambda, k, x)] = 1.$$

The scheme is *secure* if for all polynomial-sized  $S \subseteq \{0,1\}^{\ell_{\text{in}}(\lambda)}$ , it holds that

$$\begin{aligned} \{1^\lambda, S, \mathring{k}_S, \{\text{Eval}(1^\lambda, k, s)\}_{s \in S}\} & \approx \{1^\lambda, S, \mathring{k}_S, \{r_s\}_{s \in S}\}, \text{ where} \\ k & \xleftarrow{\$} \{0,1\}^{\ell_{\text{key}}(\lambda)}, \quad \mathring{k}_S \xleftarrow{\$} \text{Puncture}(1^\lambda, k, S), \\ r_s & \xleftarrow{\$} \{0,1\}^{\ell_{\text{out}}(\lambda)} \text{ for all } s \in S. \end{aligned}$$

## 2.9 Secret-Key Encryption

We need secret-key encryption as a building block.

**Definition 18** (SKE). A *secret-key encryption scheme* consists of 3 efficient algorithms

- $\text{Gen}(1^\lambda)$  outputs a key  $k$ .
- $\text{Enc}(1^\lambda, k, m)$  takes as input a key  $k$  and a message  $m$  of arbitrary length, and outputs a ciphertext  $c$ .
- $\text{Dec}(1^\lambda, k, c)$  takes as input a key  $k$  and a ciphertext  $c$ . It is supposed to recover the message.

<sup>12</sup>We also use selectively secure usual PRF, yet omit it here as it is implied by PPRF.

The scheme must be *correct*, i.e., for all  $\lambda \in \mathbb{N}$  and  $m \in \{0, 1\}^*$ , it holds that

$$\Pr \left[ k \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda) : \text{Dec}(1^\lambda, \text{Enc}(1^\lambda, k, m)) = m \right] = 1.$$

It has *pseudorandom ciphertexts* if the ciphertext length is a function of  $\lambda$  and the message length (independent of key generation and encryption randomness) and for all  $\lambda \in \mathbb{N}$  and polynomially long sequence of messages  $\{m_q\}_{q \in [Q]}$ , it holds that

$$\{1^\lambda, 1^Q, \{m_q, c_q\}_{q \in [Q]}\} \approx \{1^\lambda, 1^Q, \{m_q, r_q\}_{q \in [Q]}\}, \text{ where}$$

$$k \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda), \quad c_q \stackrel{\$}{\leftarrow} \text{Enc}(1^\lambda, k, m_q), \quad r_q \stackrel{\$}{\leftarrow} \{0, 1\}^{|c_q|} \quad \text{for all } q \in [Q].$$

## 2.10 Oblivious RAM

We define ORAM as a mechanism to translate logical accesses into physical accesses. It separates the logic of protecting memory access from the computation performed. Compared to defining it as an algorithm transforming an underlying RAM to an ORAM'ed RAM, we avoid having to incorporate randomness<sup>13</sup> into the definition of RAM. We also add error (overflow) checking into our definition so that our LGRAM and PHFE can be made perfectly correct.

**Definition 19** (ORAM). An *oblivious RAM scheme* is an efficient deterministic algorithm

$$\text{MakeORAM}(1^\lambda, 1^{\ell_{\text{CELL}}}, 1^{\ell_{\text{ADDR}}}, T_{\text{max}})$$

taking as input the logical cell length  $1^{\ell_{\text{CELL}}}$ , the logical address length  $1^{\ell_{\text{ADDR}}}$ , and a time bound  $T_{\text{max}} \in \mathbb{N}$ . It outputs a physical space bound  $S'_{\text{max}}$ , a physical cell length  $1^{\ell'_{\text{CELL}}}$ , an ORAM randomness length  $1^{\ell_r}$ , an ORAM state length  $1^{\ell_{\text{ost}}}$ , and a sequence of circuits  $\{\text{ORW}_{t_0}\}_{t_0 \in [T_0]}$  satisfying the following conditions:

- $\ell'_{\text{CELL}} \leq \text{poly}(\lambda, \ell_{\text{CELL}}, \ell_{\text{ADDR}}, \log T_{\text{max}})$ .
- $S'_{\text{max}} \leq T_{\text{max}} \text{poly}(\lambda, \ell_{\text{CELL}}, \ell_{\text{ADDR}}, \log T_{\text{max}})$ .
- $T_0 \leq \text{poly}(\lambda, \ell_{\text{CELL}}, \ell_{\text{ADDR}}, \log T_{\text{max}})$  is the number of physical steps to perform one logical access.
- The circuits have the following input/output syntax:

$$\begin{aligned} \text{ORW}_1 & : & (i, \text{wdata}, r_1) & \mapsto (\text{ost}_1, i'_1, \text{wdata}'_1) & \text{ or } \perp, \\ \text{ORW}_{t_0} & : & (\text{ost}_{t_0-1}, \text{rdata}'_{t_0-1}, r_{t_0}) & \mapsto (\text{ost}_{t_0+1}, i'_{t_0+1}, \text{wdata}'_{t_0+1}) & \text{ or } \perp, & \text{ for } 1 < t_0 < T_0, \\ \text{ORW}_{T_0} & : & (\text{ost}_{T_0-1}, \text{rdata}'_{T_0-1}, r_{T_0}) & \mapsto (i'_{T_0}, \text{wdata}'_{T_0}, \text{rdata}) & \text{ or } \perp. \end{aligned}$$

Here,

- $i \in [2^{\ell_{\text{ADDR}}}]$  is the logical address to read from and write to,
- $\text{rdata} \in \{0, 1\}^{\ell_{\text{CELL}}}$  is the logical string that was read,

<sup>13</sup>Our LGRAM only works with deterministic RAM and dealing with garbling of randomized computation is cumbersome. An ORAM'ed machine is probabilistic and cannot be directly fed into LGRAM with lesser security to obtain LGRAM with better security, so the notion of ORAM'ed RAM does not simplify formalism and might lead to confusion (it bears the name of RAM yet cannot be used as an input to LGRAM). We choose to avoid this notion and define ORAM as a memory access translation mechanism.

- $wdata \in \{0, 1\}^{\ell_{\text{CELL}}}$  is the logical string to write,
- $r_1, \dots, r_{T_0} \in \{0, 1\}^{\ell_r}$  are the ORAM randomness,
- $ost_1, \dots, ost_{T_0-1} \in \{0, 1\}^{\ell_{\text{ost}}}$  are the ORAM states,
- $i'_1, \dots, i'_{T_0} \in [S'_{\text{max}}]$  are the physical addresses to read from and write to,
- $wdata'_1, \dots, wdata'_{T_0} \in \{0, 1\}^{\ell'_{\text{CELL}}}$  are the physical string to write.

The scheme must be *correct*, i.e., for all  $\lambda, \ell_{\text{CELL}}, \ell_{\text{ADDR}}, T_{\text{max}} \in \mathbb{N}$  and all sequence  $\{(i_t, wdata_t)\}_{t \in [T_{\text{max}}]}$  of logical accesses, let

$$(S'_{\text{max}}, 1^{\ell'_{\text{CELL}}}, 1^{\ell_{\text{ost}}}, \{\text{orW}_{t_0}\}_{t_0 \in [T_0]}) \leftarrow \text{MakeORAM}(1^\lambda, 1^{\ell_{\text{CELL}}}, 1^{\ell_{\text{ADDR}}}, T_{\text{max}}),$$

$$r_{t,t_0} \stackrel{\$}{\leftarrow} \{0, 1\}^{\ell_r} \quad \text{for } t \in [T_{\text{max}}], t_0 \in [T_0],$$

and consider the following process:

- Let  $D_0 \leftarrow 0^{\ell_{\text{ADDR}}}$  and  $D'_{0,0} \leftarrow 0^{S'_{\text{max}}}$ .
- For  $t = 1, \dots, T_{\text{max}}$ :
  - (Logical step.) Let  $rdata_t \leftarrow D_{t-1}[i_t]$  and let  $D_t$  be  $D_{t-1}$  with  $D_t[i_t]$  replaced by  $wdata_t$ .
  - Perform  $T_0$  physical steps by

$$(\text{ost}_1, i'_{t,1}, wdata'_{t,1}) \leftarrow \text{orW}_1(i_t, wdata_t, r_{t,1}),$$

$$(\text{ost}_{t_0}, i'_{t,t_0}, wdata'_{t,t_0}) \leftarrow \text{orW}_{t_0}(\text{ost}_{t_0-1}, rdata'_{t,t_0-1}, r_{t,t_0}) \quad \text{for } 1 < t_0 < T_0,$$

$$(i'_{t,T_0}, wdata'_{t,T_0}, \widetilde{rdata}_t) \leftarrow \text{orW}_{T_0}(\text{ost}_{t,T_0-1}, rdata'_{t,T_0-1}, r_{t,T_0}), \quad \text{where}$$

$$D'_{t+1,0} \leftarrow D'_{t,T_0}, \quad D'_{t,t_0} \leftarrow D'_{t,t_0-1} \text{ with } D'_{t,t_0}[i'_{t,t_0}] \text{ replaced by } wdata'_{t,t_0},$$

$$rdata'_{t,t_0} \leftarrow D'_{t,t_0}[i'_{t,t_0}].$$

It is required that (over the random choices of  $r$ 's)

$$\Pr[\text{any of } \text{orW}_{t,t_0} \text{ outputs } \perp] \leq 2^{-\lambda}, \quad \text{and}$$

$$\Pr[(\text{some of } \text{orW}_{t,t_0} \text{ outputs } \perp) \vee (\widetilde{rdata}_t = rdata_t \text{ for all } t \in [T_{\text{max}}])] = 1.$$

In our definition, error is indicated by  $\perp$  from  $\text{orW}_{t,t_0}$ 's. The correctness requirements postulates that it is negligibly likely that any error is reported and that every logical access is perfectly fulfilled until the first error is reported. The latter ensures that errors are always caught before they can corrupt computation.

**Security.** Following [CH16], we need an ORAM with localized randomness.

**Definition 20** (ORAM localized randomness). An ORAM scheme (Definition 19) has *localized randomness* if there exist efficient deterministic algorithms  $\text{PartRnd}$ ,  $\text{SimORAM}$  such that for all  $\lambda, \ell_{\text{CELL}}, \ell_{\text{ADDR}}, T_{\text{max}} \in \mathbb{N}$  and all sequence  $\{(i_t, wdata_t)\}_{t \in [T]}$  of logical accesses with  $T \in [T_{\text{max}}]$ ,

$$\text{PartRnd}(1^\lambda, 1^{\ell_{\text{CELL}}}, 1^{\ell_{\text{ADDR}}}, T_{\text{max}}, \{(i_t, wdata_t)\}_{t \in [T]}) \rightarrow \{R_t\}_{t \in [T]}$$

satisfies

$$R_t \subseteq [T] \times [T_0] \text{ for all } t \in [T], \quad \max_{t \in [T]} |R_t| \leq \text{poly}(\lambda, \ell_{\text{CELL}}, \ell_{\text{ADDR}}, \log T_{\text{max}}),$$

$$\Pr \left[ \begin{array}{l} \text{SimORAM}(1^\lambda, 1^{\ell_{\text{CELL}}}, 1^{\ell_{\text{ADDR}}}, T_{\text{max}}, t, t_0, \{r_i\}_{i \in R_t}) \neq i'_{t, t_0} \\ \text{for some } t \in [T], t_0 \in [T_0] \end{array} \right] \leq 2^{-\lambda},$$

where the probability is over  $r$ 's and the notations follow those in Definition 19.

### 3 Efficiency Trade-Offs of PHFE for RAM

#### 3.1 Contention Between Storage Overhead and Decryption Time

In this section, we show that it is impossible to achieve

$$|\text{sk}| = O(|f|^\alpha) \quad \text{and} \quad T_{\text{Dec}} = O(T + |f|^\beta + |x| + |y|)$$

simultaneously for a secure PHFE for RAM when  $\alpha, \beta < 1$ , where polynomial factors in the security parameter are ignored. This leaves us with two candidate optima:

- $\alpha = 0$  and  $\beta = 1$  for succinct keys; or
- $\alpha = 1$  and  $\beta = 0$  for  $f$ -fast decryption.

Similarly, it is impossible to achieve

$$|\text{ct}| = O(|x|^\alpha) \text{ poly}(|y|) \quad \text{and} \quad T_{\text{Dec}} = O(T + |f| + |x|^\beta + |y|)$$

simultaneously if  $\alpha, \beta < 1$ , which implies a contention between succinct ciphertexts and  $x$ -fast decryption.

Formally, our theorems are slightly stronger than the discussion above:

**Theorem 4** (contention of  $|f|$ -dependency between  $|\text{sk}|$  and  $T_{\text{Dec}}$ ; ¶). *For a secure full-fledged PHFE for RAM (Definitions 5, 7, and 10), if*

$$|\text{sk}| \leq |f|^\alpha (\lambda + |\varphi|)^C \quad \text{and} \quad T_{\text{Dec}} \leq (T + |f|^\beta + |y|) (\lambda + |\varphi| + |x|)^C$$

for infinitely many  $\lambda$ , where  $\alpha, \beta, C$  are constants, then  $\alpha \geq 1$  or  $\beta \geq 1$ .

**Theorem 5** (contention of  $|x|$ -dependency between  $|\text{ct}|$  and  $T_{\text{Dec}}$ ). *For a secure full-fledged PHFE for RAM (Definitions 5, 7, and 10), if*

$$|\text{ct}| \leq |x|^\alpha (\lambda + |\varphi| + |y|)^C \quad \text{and} \quad T_{\text{Dec}} \leq (T + |f| + |x|^\beta) (\lambda + |\varphi| + |y|)^C$$

for infinitely many  $\lambda$ , where  $\alpha, \beta, C$  are constants, then  $\alpha \geq 1$  or  $\beta \geq 1$ .

We will only prove Theorem 4. The proof of Theorem 5 is similar.

*Proof* (Theorem 4). Let (Setup, KeyGen, Enc, Dec) be a secure PHFE for RAM. Suppose for contradiction that  $\alpha, \beta < 1 - 5\varepsilon$  for some  $\varepsilon > 0$ . By enlarging  $C$  as needed, we could assume  $|\varphi| \leq \lambda^C - \lambda - 1$  for all sufficiently large  $\lambda$ , where

$$\varphi = (M_\lambda, 2^\lambda), \quad f = R \in \{0, 1\}^{\leq 2^\lambda}, \quad x = \perp,$$

$$y = \begin{cases} (I, w) = (i_1, w[1], \dots, i_n, w[n]) \in ([2^\lambda] \times \{0, 1\})^{\leq 2^\lambda}; \\ z = (\perp, z[1], \dots, \perp, z[n]) \in (\{\perp\} \times \{0, 1\})^{\leq 2^\lambda}; \end{cases}$$

$$M^{f,x||y}() = \begin{cases} (R[i_1] \oplus w[1], \dots, R[i_n] \oplus w[n]), & \text{if } y = (I, w); \\ (z[1], \dots, z[n]), & \text{if } y = z. \end{cases}$$

Under appropriate encoding and step circuit design,  $y$  has exactly  $n$  cells and  $M$  halts in exactly  $(n + 1)$  steps.

We focus on the values of  $\lambda$  (hereafter, “ $\lambda$  with efficiency”) such that

$$|\text{sk}| \leq |f|^\alpha (\lambda + |\varphi|)^C \quad \text{and} \quad T_{\text{Dec}} \leq (T + |f|^\beta + |y|)(\lambda + |\varphi| + |x|)^C$$

By setting

$$|R| = N = \lceil \lambda^{(C^2+1)/\varepsilon} \rceil, \quad n = \lfloor N^{1-3\varepsilon} \rfloor,$$

we would have  $n < N < 2^\lambda$  for sufficiently large  $\lambda$ . Consider the following adversary  $\mathcal{A}$  (Definition 7):

- Upon launching, it computes  $\varphi$  define above and  $N, n$ , sets up the PHFE scheme for  $\varphi$ , and submits  $1^{n+1}$  as the time bound.
- It samples  $R \xleftarrow{\$} \{0, 1\}^N$  and requests a key  $\text{sk}$  for  $f = R$ .
- It samples  $w \xleftarrow{\$} \{0, 1\}^n$  and a list  $I$  of  $n$  distinct random elements from  $[N]$ , sets

$$z = (R[i_1] \oplus w[1], \dots, R[i_n] \oplus w[n]).$$

It challenges with

$$x = \perp, \quad y_0 = (I, w), \quad y_1 = z,$$

and obtains a ciphertext  $\text{ct}$  encrypting either  $y_0$  or  $y_1$ .

- It runs  $\text{Dec}^{f,x,\text{sk},\text{ct}}(\text{mpk})$  and notes down the list  $L$  of indices into  $R = f$  where it is read during decryption.  $\mathcal{A}$  outputs 1 if and only if

$$|L \cap I| > N^{1-4\varepsilon},$$

where  $L$  and  $I$  are regarded as sets (unordered and deduplicated) for the intersection operation.

Clearly,  $\mathcal{A}$  would be efficient and its challenge would satisfy the constraints of PHFE security for sufficiently large  $\lambda$ . We claim:

*Claim 6 (♣).* For sufficiently large  $\lambda$  with efficiency,

$$\Pr[|L \cap I| > N^{1-4\varepsilon} \text{ in } \text{Exp}_{\text{PHFE}}^0] \geq \frac{3}{4}.$$

*Claim 7 (♣).* For sufficiently large  $\lambda$  with efficiency,

$$\Pr[|L \cap I| > N^{1-4\varepsilon} \text{ in } \text{Exp}_{\text{PHFE}}^1] \leq \frac{1}{4}.$$



The two claims together would contradict the security of PHFE, as the advantage of  $\mathcal{A}$  would be at least  $\frac{1}{2}$  for infinitely many  $\lambda$ . Therefore,  $\alpha \geq 1$  or  $\beta \geq 1$ .  $\square$

To prove Claim 6, we need the following lemma about incompressibility of information:

**Lemma 8** ([DTT10]). *Suppose  $E : S \times U \rightarrow V$  and  $D : S \times V \rightarrow U$  are functions and  $S$  is a distribution over  $S$ , then*

$$|V| \geq |U| \cdot \Pr_{\substack{s \leftarrow S \\ u \leftarrow U}} [D(s, E(s, u)) = u].$$

*Proof* (Claim 6). We use the PHFE scheme to compress a string  $u$  of length  $n$ . To encode, we embed  $u$  into a string  $R$  of length  $N$  at random locations (i.e.,  $I$ ) and generate a PHFE key for  $R$ . The encoding is the key plus some bits in  $R$  used during decryption. To decode, run the decryption algorithm. Lemma 8 will generate the following inequality equivalent to the desired one:

$$\Pr[|L \cap I| \leq \lfloor N^{1-4\epsilon} \rfloor \text{ in } \text{Exp}_{\text{PHFE}}^0] \leq \frac{1}{4}.$$

Formally, let

$$\mathcal{S} = \left\{ \left( \begin{array}{l} \text{mpk, msk, } I, w, R', \\ r_{\text{KeyGen}}, r_{\text{Enc}}, r_{\text{Dec}} \end{array} \right) : \begin{array}{l} (\text{mpk, msk}) \leftarrow \text{Setup}(\varphi) \\ (I, w) \text{ as how } \mathcal{A} \text{ samples it} \\ R'[i] \leftarrow \{0, 1\} \text{ for } i \in [N] \setminus I \\ r_{\text{KeyGen}}, r_{\text{Enc}}, r_{\text{Dec}} \leftarrow \text{algorithm randomness} \end{array} \right\},$$

$$U = \{0, 1\}^n, \quad V = \{0, 1\}^{\lfloor N^{1-4\epsilon} \rfloor} \times \{0, 1\}^{\lfloor N^{1-4\epsilon} \rfloor}.$$

The encoding procedure  $E(s, u)$  works as follows:

- Parse  $I = (i_1, \dots, i_n)$  and set

$$R[i] = \begin{cases} R'[i], & \text{if } i \in [N] \setminus I; \\ u[j], & \text{if } i = i_j. \end{cases}$$

- Run

$$\begin{aligned} \text{sk} &\leftarrow \text{KeyGen}(\text{msk}, R; r_{\text{KeyGen}}), \\ \text{ct} &\leftarrow \text{Enc}(\text{mpk}, \perp, (I, w); r_{\text{Enc}}), \\ u \oplus w &\leftarrow \text{Dec}^{R, \perp, \text{sk}, \text{ct}}(\text{mpk}; r_{\text{Dec}}), \end{aligned}$$

and note down the list  $L = (\ell_1, \dots)$  of indices into  $R$  read by Dec.

- Output  $v = (v_1, v_2)$  with  $v_1, v_2 \in \{0, 1\}^{\lfloor N^{1-4\epsilon} \rfloor}$  and

$$\begin{aligned} v_1 &= 0^{\lfloor N^{1-4\epsilon} \rfloor - |\text{sk}| - 1} \parallel \text{sk}, \\ v_2[i] &= \begin{cases} R[\ell_j], & \text{if } |\{\ell_1, \dots, \ell_{j-1}\} \cap I| = i - 1 \text{ and } |\{\ell_1, \dots, \ell_{j-1}, \ell_j\} \cap I| = i; \\ 0, & \text{if no such } j \text{ exists.} \end{cases} \end{aligned}$$

Here,  $v_1$  is a fixed-length encoding of  $\text{sk}$  and is indeed well-defined since

$$|\text{sk}| \leq |f|^\alpha (\lambda + |\varphi|)^C \leq N^{1-5\epsilon} (\lambda + (\lambda^C - \lambda - 1))^C \leq N^{1-5\epsilon} \lambda^{C^2} < \lfloor N^{1-4\epsilon} \rfloor - 1$$

for sufficiently large  $\lambda$  with efficiency. The string  $v_2$  records, *sequentially*, the bits in  $R$  at each *distinct* index read by Dec that are part of  $u$  and not known from  $R'$ , for at most  $\lfloor N^{1-4\epsilon} \rfloor$  bits.

The decoding procedure  $D(s, v)$  works as follows:

- Run  $\text{ct} \leftarrow \text{Enc}(\text{mpk}, \perp, (I, w); r_{\text{Enc}})$ .
- Parse  $v = (v_1, v_2)$  and recover  $\text{sk}$  from  $v_1$  as specified in  $E$ .
- Initialize  $j$ , an index into  $v_2$ , by  $j \leftarrow 0$ , and initialize  $R$  by

$$R[i] = \begin{cases} R'[i], & \text{if } i \in [N] \setminus I; \\ \perp, & \text{if } i \in I. \end{cases}$$

Run  $z \leftarrow \text{Dec}^{R, \perp, \text{sk}, \text{ct}}(\text{mpk}; r_{\text{Dec}})$  with  $R$  filled on the fly. When Dec reads  $R[i]$ :

- if  $R[i] = \perp$  and  $j < \lfloor N^{1-4\epsilon} \rfloor$ , then let  $j \leftarrow j + 1$  and set  $R[i] \leftarrow v_2[j]$ ;
- if  $R[i] = \perp$  and  $j = \lfloor N^{1-4\epsilon} \rfloor$ , then abort by outputting  $0^n$ ;
- otherwise,  $R[i] \neq \perp$ , then just proceed without aborting;

and return  $R[i]$  to Dec if not aborting.

- Output  $z \oplus w$ .

$D$  will fill  $v_2$  into the correct indices of  $R$  since Enc and Dec are derandomized with the same randomness as in  $E$ .

The sampling of  $s, u$  and the setting of  $R$  in  $E(s, u)$  simulate  $\mathcal{A}$  in  $\text{Exp}_{\text{PHFE}}^0$ . If  $s$  and  $u$  are such that  $|L \cap I| \leq \lfloor N^{1-4\epsilon} \rfloor$  in  $E(s, u)$ , then  $D$  will successfully recover  $u$ . By Lemma 8,

$$\begin{aligned} \Pr[|L \cap I| \leq \lfloor N^{1-4\epsilon} \rfloor \text{ in } \text{Exp}_{\text{PHFE}}^0] &= \Pr_{\substack{s \leftarrow \mathcal{S} \\ u \leftarrow \mathcal{U}}} [ |L \cap I| \leq \lfloor N^{1-4\epsilon} \rfloor \text{ in } E(s, u) ] \\ &\leq \Pr_{\substack{s \leftarrow \mathcal{S} \\ u \leftarrow \mathcal{U}}} [ D(s, E(s, u)) = u ] \\ &\leq \frac{|\mathcal{V}|}{|\mathcal{U}|} = \frac{2^{2\lfloor N^{1-4\epsilon} \rfloor}}{2^n} = 2^{2\lfloor N^{1-4\epsilon} \rfloor - \lfloor N^{1-3\epsilon} \rfloor} \leq \frac{1}{4} \end{aligned}$$

for sufficiently large  $\lambda$  with efficiency. □

*Proof* (Claim 7). For sufficiently large  $\lambda$  with efficiency,

$$\begin{aligned} |L| \leq T_{\text{Dec}} &\leq (T + |f|^\beta + |y|)(\lambda + |\varphi| + |x|)^C \\ &\leq ((n+1) + N^{1-5\epsilon} + n)(\lambda + (\lambda^C - \lambda - 1) + 1)^C \\ &\leq (2N^{1-3\epsilon} + N^{1-5\epsilon} + 1)\lambda^{C^2} \leq N^{1-2\epsilon}. \end{aligned}$$

In  $\text{Exp}_{\text{PHFE}}^1$ , the input to Dec is independent of  $I$ , which only symbolically appears in ct as

$$y_1 = z = (R[i_1] \oplus w[1], \dots, R[i_n] \oplus w[n])$$

and is fully hidden by the one-time pad  $w$ . Therefore, the list of indices into  $R$  read by Dec (i.e.,  $L$ ) is independent of  $I$ . Conditioned on  $L$ , the intersection size  $|L \cap I|$  follows a hypergeometric distribution. By the law of total expectation,

$$\mathbb{E}[|L \cap I|] = \mathbb{E}\left[\mathbb{E}[|L \cap I| \mid L]\right] = \mathbb{E}\left[\frac{|I| \cdot |L|}{N}\right] \leq \frac{N^{1-3\epsilon} \cdot N^{1-2\epsilon}}{N} = N^{1-5\epsilon}$$

for sufficiently large  $\lambda$  with efficiency, which implies, by Markov's inequality,

$$\Pr[|L \cap I| > N^{1-4\epsilon} \text{ in } \text{Exp}_{\text{PHFE}}^1] \leq \frac{\mathbb{E}[|L \cap I|]}{N^{1-4\epsilon}} \leq \frac{N^{1-5\epsilon}}{N^{1-4\epsilon}} = N^{-\epsilon} \leq \frac{1}{4}. \quad \square$$

### 3.2 Barrier to Time Optimality

In this section, we show that if we build PHFE schemes that have very efficient decryption, i.e. ones that are  $f$ -fast, or  $x$ -fast or  $y$ -fast naturally imply doubly efficient secret-key PIR protocols (SKDEPIR) which are currently known only from a very specific assumption [BIPW17,CHR17].

We start by defining syntax of a SKDEPIR, along with the correctness, efficiency and security notions we consider. These definitions are suitably adapted from [BIPW17,CHR17].

**Definition 21.** (Syntax of SKDEPIR) An SKDEPIR scheme consists of five probabilistic polynomial-time algorithms (Keygen, Process, Query, Resp, Dec) with the following specification:

- Keygen takes as input the security parameter  $1^\lambda$  and samples a key  $k$ .
- Process takes as input a key  $k$ , and a database  $\text{DB} \in \{0,1\}^n$  and outputs a processed database  $\widetilde{\text{DB}}$ .
- Query takes as input the key  $k$ , an index  $i \in [n]$ , and outputs a query  $\sigma$  (and possibly a state  $\text{st}$ ).
- Resp is a deterministic algorithm that takes as input  $\sigma$  and  $\widetilde{\text{DB}}$  and outputs a server response  $\rho$ .
- Dec is a deterministic algorithm that takes as input  $(\sigma, \rho, k, \text{st})$  and outputs a data element in  $\{0,1\}$ .

Next, we define the correctness and the efficiency property of SKDEPIR.

**Definition 22.** (Correctness/Efficiency of SKDEPIR) Let  $c > 0$  be a constant,  $\lambda \in \mathbb{N}$  be the security parameter and  $n = \lambda^c$ . Let  $\text{DB} \in \{0,1\}^n$ ,  $i \in [n]$  and consider the following process:

- Run  $k \leftarrow \text{Keygen}$ ,  $\widetilde{\text{DB}} \leftarrow \text{Process}(k, \text{DB})$  and  $(\sigma, \text{st}) \leftarrow \text{Query}(k, i)$ ,

- Compute  $\rho = \text{Resp}(\sigma, \widetilde{\text{DB}})$  and  $m \leftarrow \text{Dec}(\sigma, \rho, k, \text{st})$ .

Then, we say that SKDEPIR is correct if  $m = \text{DB}_i$  with probability at least  $1 - \text{negl}(\lambda)$  for some negligible  $\text{negl}$  where the probability is over the coins of the process. Further, we say that it is doubly-efficient if the running time of  $\text{Resp}$  with RAM access to  $\widetilde{\text{DB}}$  is polynomial in  $\lambda$  and logarithm of  $n$ .

We now define the security notion we consider. For strengthened notions of security see [BIPW17,CHR17].

**Definition 23.** (Security of SKDEPIR) We say that SKDEPIR is secure if there exists a stateful p.p.t. simulator  $\mathcal{S}$  such that for any stateful p.p.t. adversary  $\mathcal{A}$  the following distributions are computationally indistinguishable. The first distribution is:

- $\mathcal{A}(1^\lambda)$  outputs  $1^n$ , a database  $\text{DB} \in \{0, 1\}^n$  and  $i_1, \dots, i_\ell$  where every  $i_j \in [\ell]$ .
- Run  $k \leftarrow \text{Keygen}$ ,  $\widetilde{\text{DB}} \leftarrow \text{Process}(k, \text{DB})$  and  $(\sigma_j, \text{st}_j) \leftarrow \text{Query}(k, i_j)$  for  $j \in [\ell]$ ,
- Output  $(\text{DB}, i_1, \dots, i_\ell, \widetilde{\text{DB}}, \sigma_1, \dots, \sigma_\ell)$

The second distribution is:

- $\mathcal{A}(1^\lambda)$  outputs  $1^n$ , a database  $\text{DB} \in \{0, 1\}^n$  and  $i_1, \dots, i_\ell$  where every  $i_j \in [\ell]$ .
- The simulator outputs  $(\widetilde{\text{DB}}, \sigma_1, \dots, \sigma_\ell) \leftarrow \mathcal{S}(\text{DB}, \ell)$ .
- Output  $(\text{DB}, i_1, \dots, i_\ell, \widetilde{\text{DB}}, \sigma_1, \dots, \sigma_\ell)$

One could consider a stronger security definition where an adversary can ask for index queries adaptively. As of currently, this weaker definition is not known to be any easier to achieve. We will show that any PHFE scheme satisfying a weak form indistinguishability security, where the functions and the challenge messages are declared in the beginning, that is additionally  $f$ -fast/ $x$ -fast/ $y$ -fast implies such an SKDEPIR scheme. We now give overview of three (very similar) constructions. The first construction converts an  $x$ -fast PHFE scheme to an SKDEPIR scheme. The second, converts an  $f$ -fast PHFE scheme to an SKDEPIR scheme. Finally, the third scheme converts an  $y$ -fast PHFE scheme to an SKDEPIR scheme. Our main result is:

**Theorem 9.** *If there exists a PHFE scheme that is either  $x$ -fast or  $y$ -fast or  $f$ -fast as per definition 6, then, then there exist a secret-key doubly efficient PIR scheme as per Definitions 23 and 22.*

Overview We aim to show that an  $x$ -fast/ $f$ -fast/ $y$ -fast PHFE scheme imply a SKDEPIR scheme. We start by describing the main idea assuming  $x$ -fast PHFE scheme, and then we will sketch the ideas for the other cases. In a  $x$ -fast PHFE scheme the time it takes to compute decryption in the RAM model is proportional to  $(T_{\varphi(f,x,y)}^{\beta_T} + |f|^{\beta_f} + |y|^{\beta_y})$  for some constants  $\beta_T, \beta_f, \beta_y$  upto polynomial factors in  $\lambda$ . We will exploit the fact that there is no dependence on  $|x|$  to turn this into a SKDEPIR scheme.

As a first attempt, we set  $\text{DB}$  as  $x$ , and  $y$  as empty. Then  $\text{Process}$  algorithm computes  $\text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y))$  as a preprocessing of the database. To look up  $\text{DB}_i$ , the  $\text{Query}$  algorithm computes a PHFE function key  $\text{sk}_{f_j}$  for the program that looks up and outputs  $\text{DB}_i$  and sends the key to the Server. Observe that  $\text{Query}$  algorithm runs in time polynomial in  $\lambda$  and  $\log n$  as  $\text{PHFE.KeyGen}$  is a polynomial time in the description size of  $f$  and the security parameter.

Further, the Resp algorithm simply decrypts  $ct_{\text{PHFE}}$  using  $sk_f$ . Observe that due to  $x$ -fastness, the time it takes to run Resp is also polynomial in  $\lambda$  and  $\log n$ . While this solves the core issue, we have completely missed one aspect. The scheme completely reveals the indices  $i$  to the Server as the function key  $sk_f$  is not guaranteed to hide the function description  $f$ . To resolve this issue, we observe that if we had a function hiding PHFE scheme, we would have been done. To enable this, we will use similar techniques as used to convert any functional encryption scheme to a function hiding functional encryption scheme [BS15]. Namely, we will compute a symmetric key encryption SKE of the index  $i$  (denoted as  $ct_{k_1}$ ). We will hardwire  $ct_{k_1}$  in the function secret key instead of the index  $i$ . The corresponding secret key  $\text{SKE.sk}_1$  will be put in the private component  $y$ , which will be used to decrypt  $ct_{k_1}$  to learn index  $i$ .

This might seem to be enough, but we face yet another issue. Learning  $\text{DB}_i$  in the clear upon decryption can reveal information about the index  $i$  to the Server. To fix this, the decryption will output the encryption of  $\text{DB}_i$  under another secret key  $\text{SKE.sk}_2$  of the secret-key encryption scheme. We will put this key in  $y$  component of the message along with a PRF key to derive randomness to compute the encryption. To make the proof go through, we need to use another encryption key  $\text{SKE.sk}_3$  and compute another ciphertext that will be hardwired in the functional secret key, along with a mode variable and a separate slot in  $y$  used in encryption. This is done in a standard fashion inspired by [BS15].

While these are the ideas to convert an  $x$ -fast PHFE scheme to a SKDEPIR, they immediately extend to converting a  $y$ -fast PHFE scheme to a SKDEPIR scheme. The idea is that while encrypting one could set  $x = \perp$  and just use the hidden slot  $y$  alone. Now, the database  $\text{DB}$  will also be a part of the hidden input.

Conversion of an  $f$ -fast PHFE to SKDEPIR scheme is a slightly different. The idea is yet again a natural one. Instead of now encrypting  $\text{DB}$  inside the ciphertext, the client will process it as a part of a function key  $sk_{\text{DB}}$ . To look up an index  $i_j$ , the client will generate ciphertexts  $ct_{\text{PHFE},j}$ . These ciphertexts will encrypt a public component  $x_j = \perp$  and a secret component  $y_j = (\text{SKE.sk}, r_j, i_j, \perp)$  where  $\text{SKE.sk}$  is a secret key of a secret key encryption scheme chosen once and for all by the Process algorithm and  $r_j$  is a randomness of length polynomial in  $\lambda$ , and  $\perp$  represents an additional slot used in the proof. The function key is set so that upon decryption we get  $\text{SKE.Enc}(\text{SKE.sk}, \text{DB}_{i_j}; r_j)$ . Now using  $\text{SKE.sk}$  client can learn  $\text{DB}_{i_j}$ . Observe that the function corresponding to the PHFE secret key on the input encrypted inside the ciphertexts takes time  $T$  which is polynomial in  $\lambda$  and  $\log n$  in the RAM model. The length  $|y_j|$  is also polynomial in  $\lambda$  and  $\log n$ . Therefore, the scheme satisfies the desired efficiency properties. Owing to similarities, we omit a formal treatment of this scheme. Below we give a formal description of how to convert an  $x$ -fast PHFE to a SKDEPIR scheme.

Ingredients for the constructions Our first ingredient is a PHFE scheme  $\text{PHFE} = (\text{Setup}, \text{Enc}, \text{Keygen}, \text{Dec})$  that is additionally  $x$ -fast. We also use a secret-key encryption scheme  $\text{SKE} = (\text{Setup}, \text{Enc}, \text{Dec})$ .

Construction:  $x$ -fast PHFE to SKDEPIR

- Keygen : On input  $1^\lambda$ , we run  $(\text{mpk}, \text{msk}) \leftarrow \text{PHFE.Setup}(1^\lambda)$ . We also sample three secret keys for a secret-key encryption scheme  $\{\text{SKE.sk}_i \leftarrow \text{SKE.Setup}(1^\lambda)\}_{i \in [3]}$ . We additionally sample a PRF key  $k_{\text{PRF}}$  for a PRF  $= (\text{Setup}, \text{Eval})$  with input

<b>Computation</b> $\varphi(f, x, y)$	
<b>Hardwired:</b>	Ciphertexts $\text{ctk}_1, \text{ctk}_3$ and $t \in \{0, 1\}^\lambda$
<b>Input:</b>	$x = \text{DB}$ and $y = (\text{mode}, \alpha_1, \alpha_2, \alpha_3, \alpha_4)$
<b>Output:</b>	$\rho$
<ul style="list-style-type: none"> <li>• If <math>\text{mode} = 0</math>, parse <math>\alpha_1 = \text{SKE.sk}_1</math>, <math>\alpha_2 = \text{SKE.sk}_2</math>, <math>\alpha_3 = k_{\text{PRF}}</math> otherwise parse <math>\alpha_4 = \text{SKE.sk}_3</math>.</li> <li>• <b>Case</b> <math>\text{mode} = 0</math>: Compute <math>i \leftarrow \text{SKE.Dec}(\text{SKE.sk}_1, \text{ctk}_1)</math> and output <math>\rho = \text{SKE.Enc}(\text{SKE.sk}_2, \text{DB}_i; \text{PRF.Eval}(k_{\text{PRF}}, t))</math>.</li> <li>• <b>Case</b> <math>\text{mode} = 1</math>: Output <math>\rho = \text{SKE.Dec}(\text{SKE.sk}_3, \text{ctk}_3)</math></li> </ul>	

**Figure 2.** The computation  $\varphi(f, x, y)$ .

and output lengths implicitly defined. The output of this algorithm is  $k = (\text{mpk}, \text{msk}, \text{SKE.sk}_1, \text{SKE.sk}_2, \text{SKE.sk}_3, k_{\text{PRF}})$ .

- **Process** : On input  $k$  and  $\text{DB} \in \{0, 1\}^n$ , we set the public-input  $x = \text{DB}$  and  $y = (\text{mode} = 0, \alpha_1 = \text{SKE.sk}_1, \alpha_2 = \text{SKE.sk}_2, \alpha_3 = k_{\text{PRF}}, \alpha_4 = \perp)$  (where the mode is a bit that can take two values in  $\{0, 1\}$  and  $\perp$  will be used in the proof. The algorithm outputs  $\widetilde{\text{DB}} = (x = \text{DB}, \text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y)))$ .
- **Query**( $k, i$ ) : We sample a random string  $t \leftarrow \{0, 1\}^\lambda$  along with the following quantities:
  - Compute  $\text{ctk}_1 \leftarrow \text{SKE.Enc}(\text{SKE.sk}_1, i)$ ,
  - Compute  $\text{ctk}_2 \leftarrow \text{SKE.Enc}(\text{SKE.sk}_2, 0)$  and  $\text{ctk}_3 = \text{SKE.Enc}(\text{SKE.sk}_3, \text{ctk}_2)$ ,
  - Compute  $\text{sk}_f \leftarrow \text{PHFE.Keygen}(\text{msk}, f)$  for a RAM program  $f$  where we describe the computation  $\varphi(f, x, y)$  in Figure 2. Output of the algorithm is  $\sigma = (f, \text{sk}_f)$  and  $\text{st} = \perp$ .
- **Resp**( $\sigma = (f, \text{sk}_f), \widetilde{\text{DB}} = \text{DB}, \text{ct}_{\text{PHFE}}$ ) : Output  $\rho = \text{PHFE.Dec}^{f, \text{DB}, \text{sk}_f, \text{ct}_{\text{PHFE}}}(\perp)$ .
- **Dec**( $\sigma = (f, \text{sk}_f), \rho, k = (\text{mpk}, \text{msk}, \text{SKE.sk}_1, \text{SKE.sk}_2, \text{SKE.sk}_3, k_{\text{PRF}}), \perp$ ) : Output  $\text{SKE.Dec}(\text{SKE.sk}_2, \rho)$ .

We now prove various properties associated below.

Correctness: The correctness follows from the correctness of the SKE, and PHFE schemes. Note that Process encodes DB by computing a PHFE ciphertext  $\text{ct}_{\text{PHFE}}$  encrypting DB in the public component and SKE secret keys ( $\text{SKE.sk}_1, \text{SKE.sk}_2$ ) and a PRF key in the secret component. Query on input  $i$ , produces a PHFE function secret key that upon decrypting the ciphertext produces an SKE encryption of for any index  $i$ , under  $\text{SKE.sk}_2$ . Resp then simply performs this decryption to come up with  $\rho$ , which is a SKE ciphertext encrypting  $\text{DB}_i$  using key  $\text{SKE.sk}_2$ . Finally, Dec uses  $\text{SKE.sk}_2$  to learn  $\text{DB}_i$ .

Efficiency: We now bound the running time of Query and Resp.

- **Running time of Query:** Observe that Query computes a PHFE secret key  $sk_f$  for a function  $f$  that has a size polynomial in  $\lambda$  and  $\log n$ . This step takes time polynomial in  $\lambda$  and  $\log n$  as KeyGen algorithms of PHFE takes polynomial time. Further, to construct  $f$ , the Query algorithm constructs  $ctk_1$  and  $ctk_3$ , which also take time polynomial in  $\lambda$  and  $\log n$ .
- **Running time of Resp:** Resp simply performs  $\text{PHFE.Dec}^{f, \text{DB}, sk_f, ct_{\text{PHFE}}}(1^\lambda)$ . Since PHFE is  $x$ -fast the running time of decryption is  $(T_{\varphi(f,x,y)}^{\beta_T} + |f|^{\beta_f} + |y|^{\alpha_y}) \text{poly}(\lambda, |\varphi|)$  for some constants  $\beta_T, \beta_f, \beta_y$ , since PHFE is  $x$ -fast. Observe that  $T_{\varphi(f,x,y)}$  is bounded by polynomial in  $\lambda$  and  $\log n$  as it consists of a constant number of SKE decryptions/encryption, a PRF computation, and a single lookup of DB at an index  $i$ . Furthermore,  $y$  is of size polynomial in  $\lambda$  as it consists of a PRF key along with SKE secret keys. Finally,  $f$  is of size polynomial in  $\lambda$  and  $\log n$  as argued above. As a result, total time complexity of Resp in the RAM model is polynomial in  $\lambda$  and  $\log n$ .

We now prove security for the constructed SKDEPIR scheme.

**Theorem 10.** *Assuming PRF is a secure pseudorandom function, SKE is secure secret-key encryption scheme and PHFE is a secure partially hiding functional encryption scheme satisfying security Definition 7, the SKDEPIR scheme described above satisfies the security Definition 23.*

We prove this by providing a set of indistinguishable hybrids. The first hybrid corresponds to the first distribution of the security game, whereas the last hybrid corresponds to the second distribution. The simulator is implicit from the distribution in the last hybrid. We then formally describe our simulator. We use properties of various primitives involve to prove that the intermediate hybrids are indistinguishable.

- $H_1$  is the first distribution of the SKDEPIR game, where the adversary receives  $(\text{DB}, \{i_j\}_{j \in [\ell]}, \widetilde{\text{DB}} = (\text{DB}, ct_{\text{PHFE}}), \{\sigma_j = (f_j, sk_{f_j})\}_{j \in [\ell]})$  where  $ct_{\text{PHFE}}, sk_{f_j}$  are computed as follows:

$$\begin{aligned}
& (\text{mpk}, \text{msk}) \leftarrow \text{PHFE.Setup}(1^\lambda), \quad \{\text{SKE.sk}_i \leftarrow \text{SKE.Setup}(1^\lambda)\}_{i \in [3]}, \quad k_{\text{PRF}} \leftarrow \text{PRF.Setup}(1^\lambda) \\
& x = \text{DB}, \quad y = (\text{mode} = 0, \text{SKE.sk}_1, \text{SKE.sk}_2, k_{\text{PRF}}, \perp) \\
& ct_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y)) \\
& ctk_{1,j} = \text{SKE.Enc}(\text{SKE.sk}_1, i_j), \quad ctk_{2,j} = \text{SKE.Enc}(\text{SKE.sk}_2, 0), \\
& ctk_{3,j} = \text{SKE.Enc}(\text{SKE.sk}_3, ctk_{2,j}) \\
& t_j \leftarrow \{0, 1\}^\lambda \quad sk_{f_j} = \text{PHFE.KeyGen}(\text{msk}, f_j = f[ctk_{1,j}, ctk_{3,j}, t_j])
\end{aligned}$$

- $H_2$  In this hybrid, the only change is that  $ctk_{3,j}$  is now computed as an encryption of  $ctk_{2,j}$  that is computed differently.  $ctk_{2,j}$  is now computed as an SKE encryption

of  $DB_i$  using randomness derived from  $k_{\text{PRF}}$ :

$(\text{mpk}, \text{msk}) \leftarrow \text{PHFE.Setup}(1^\lambda), \quad \{\text{SKE.sk}_i \leftarrow \text{SKE.Setup}(1^\lambda)\}_{i \in [3]}, \quad k_{\text{PRF}} \leftarrow \text{PRF.Setup}(1^\lambda)$   
 $x = \text{DB}, \quad y = (\text{mode} = 0, \text{SKE.sk}_1, \text{SKE.sk}_2, k_{\text{PRF}}, \perp)$   
 $\text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y))$   
 $\text{ctk}_{1,j} = \text{SKE.Enc}(\text{SKE.sk}_1, i_j), \quad \boxed{\text{ctk}_{2,j} = \text{SKE.Enc}(\text{SKE.sk}_2, \text{DB}_{i_j}; \text{PRF.Eval}(k_{\text{PRF}}, t_j))},$   
 $\boxed{\text{ctk}_{3,j} = \text{SKE.Enc}(\text{SKE.sk}_3, \text{ctk}_{2,j})}$   
 $t_j \leftarrow \{0, 1\}^\lambda \quad \text{sk}_{f_j} = \text{PHFE.KeyGen}(\text{msk}, f_j = f[\text{ctk}_{1,j}, \text{ctk}_{3,j}, t_j])$

Observe that  $H_1$  and  $H_2$  are indistinguishable due to the security of SKE. The key  $\text{SKE.sk}_3$  is hidden from the adversary.

- $H_3$  In this hybrid, the only change is that we now compute  $\text{ct}_{\text{PHFE}}$  differently. The private component  $y$  is now designed so that it sets  $\text{mode} = 1$ . The secret keys  $\text{SKE.sk}_1, \text{SKE.sk}_2$  and the PRF key  $k_{\text{PRF}}$  is removed and they are replaced with  $\text{SKE.sk}_3$ :

$(\text{mpk}, \text{msk}) \leftarrow \text{PHFE.Setup}(1^\lambda), \quad \{\text{SKE.sk}_i \leftarrow \text{SKE.Setup}(1^\lambda)\}_{i \in [3]}, \quad k_{\text{PRF}} \leftarrow \text{PRF.Setup}(1^\lambda)$   
 $x = \text{DB}, \quad \boxed{y = (\text{mode} = 1, \perp, \perp, \perp, \text{SKE.sk}_3)}$   
 $\boxed{\text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y))}$   
 $\text{ctk}_{1,j} = \text{SKE.Enc}(\text{SKE.sk}_1, i_j), \quad \text{ctk}_{2,j} = \text{SKE.Enc}(\text{SKE.sk}_2, \text{DB}_{i_j}; \text{PRF.Eval}(k_{\text{PRF}}, t_j)),$   
 $\text{ctk}_{3,j} = \text{SKE.Enc}(\text{SKE.sk}_3, \text{ctk}_{2,j})$   
 $t_j \leftarrow \{0, 1\}^\lambda \quad \text{sk}_{f_j} = \text{PHFE.KeyGen}(\text{msk}, f_j = f[\text{ctk}_{1,j}, \text{ctk}_{3,j}, t_j])$

Observe that  $H_2$  and  $H_3$  are indistinguishable due to the security of PHFE. Observe that in both the cases  $\varphi(f_j, x, y)$  produce the same outcome  $\text{ctk}_{2,j}$ . Therefore due to the security of the PHFE scheme these hybrids are computationally indistinguishable.

- $H_4$  In this hybrid, the only change is that now we encrypt  $\text{ctk}_{1,j}$  as an encryption of 0:

$(\text{mpk}, \text{msk}) \leftarrow \text{PHFE.Setup}(1^\lambda), \quad \{\text{SKE.sk}_i \leftarrow \text{SKE.Setup}(1^\lambda)\}_{i \in [3]}, \quad k_{\text{PRF}} \leftarrow \text{PRF.Setup}(1^\lambda)$   
 $x = \text{DB}, \quad y = (\text{mode} = 0, \perp, \perp, \perp, \text{SKE.sk}_3)$   
 $\text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y))$   
 $\boxed{\text{ctk}_{1,j} = \text{SKE.Enc}(\text{SKE.sk}_1, 0)}, \quad \text{ctk}_{2,j} = \text{SKE.Enc}(\text{SKE.sk}_2, \text{DB}_{i_j}; \text{PRF.Eval}(k_{\text{PRF}}, t_j)),$   
 $\text{ctk}_{3,j} = \text{SKE.Enc}(\text{SKE.sk}_3, \text{ctk}_{2,j})$   
 $t_j \leftarrow \{0, 1\}^\lambda \quad \text{sk}_{f_j} = \text{PHFE.KeyGen}(\text{msk}, f_j = f[\text{ctk}_{1,j}, \text{ctk}_{3,j}, t_j])$

Observe that  $H_3$  and  $H_4$  are indistinguishable due to the security of SKE. The key  $\text{SKE.sk}_1$  is hidden from the adversary.

- $H_5$  In this hybrid, the only change is that to compute  $\text{ctk}_{2,j}$  instead of deriving



randomness using the PRF key  $k_{\text{PRF}}$ :

$$\begin{aligned} & (\text{mpk}, \text{msk}) \leftarrow \text{PHFE.Setup}(1^\lambda), \quad \{\text{SKE.sk}_i \leftarrow \text{SKE.Setup}(1^\lambda)\}_{i \in [3]}, \quad k_{\text{PRF}} \leftarrow \text{PRF.Setup}(1^\lambda) \\ & x = \text{DB}, \quad y = (\text{mode} = 0, \perp, \perp, \perp, \text{SKE.sk}_3) \\ & \text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y)) \\ & \text{ctk}_{1,j} = \text{SKE.Enc}(\text{SKE.sk}_1, 0), \quad \boxed{\text{ctk}_{2,j} = \text{SKE.Enc}(\text{SKE.sk}_2, \text{DB}_{i_j})}, \\ & \text{ctk}_{3,j} = \text{SKE.Enc}(\text{SKE.sk}_3, \text{ctk}_{2,j}) \\ & t_j \leftarrow \{0, 1\}^\lambda \quad \text{sk}_{f_j} = \text{PHFE.KeyGen}(\text{msk}, f_j = f[\text{ctk}_{1,j}, \text{ctk}_{3,j}, t_j]) \end{aligned}$$

Observe that  $H_4$  and  $H_5$  are indistinguishable due to the security of the PRF. The key  $k_{\text{PRF}}$  is hidden from the adversary and appears only in the form of values  $\text{PRF.Eval}(k_{\text{PRF}}, t_j)$ .

- $H_6$  In this hybrid, the only change is that  $\text{ctk}_{2,j}$  are now computed as an encryption of 0:

$$\begin{aligned} & (\text{mpk}, \text{msk}) \leftarrow \text{PHFE.Setup}(1^\lambda), \quad \{\text{SKE.sk}_i \leftarrow \text{SKE.Setup}(1^\lambda)\}_{i \in [3]}, \quad k_{\text{PRF}} \leftarrow \text{PRF.Setup}(1^\lambda) \\ & x = \text{DB}, \quad y = (\text{mode} = 0, \perp, \perp, \perp, \text{SKE.sk}_3) \\ & \text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y)) \\ & \text{ctk}_{1,j} = \text{SKE.Enc}(\text{SKE.sk}_1, 0), \quad \boxed{\text{ctk}_{2,j} = \text{SKE.Enc}(\text{SKE.sk}_2, 0)}, \\ & \text{ctk}_{3,j} = \text{SKE.Enc}(\text{SKE.sk}_3, \text{ctk}_{2,j}) \\ & t_j \leftarrow \{0, 1\}^\lambda \quad \text{sk}_{f_j} = \text{PHFE.KeyGen}(\text{msk}, f_j = f[\text{ctk}_{1,j}, \text{ctk}_{3,j}, t_j]) \end{aligned}$$

Observe that  $H_5$  and  $H_6$  are indistinguishable due to the security of SKE. The key  $\text{SKE.sk}_2$  is hidden from the adversary.

From the description of the hybrid above, we now define the simulator  $\mathcal{S}$  below:

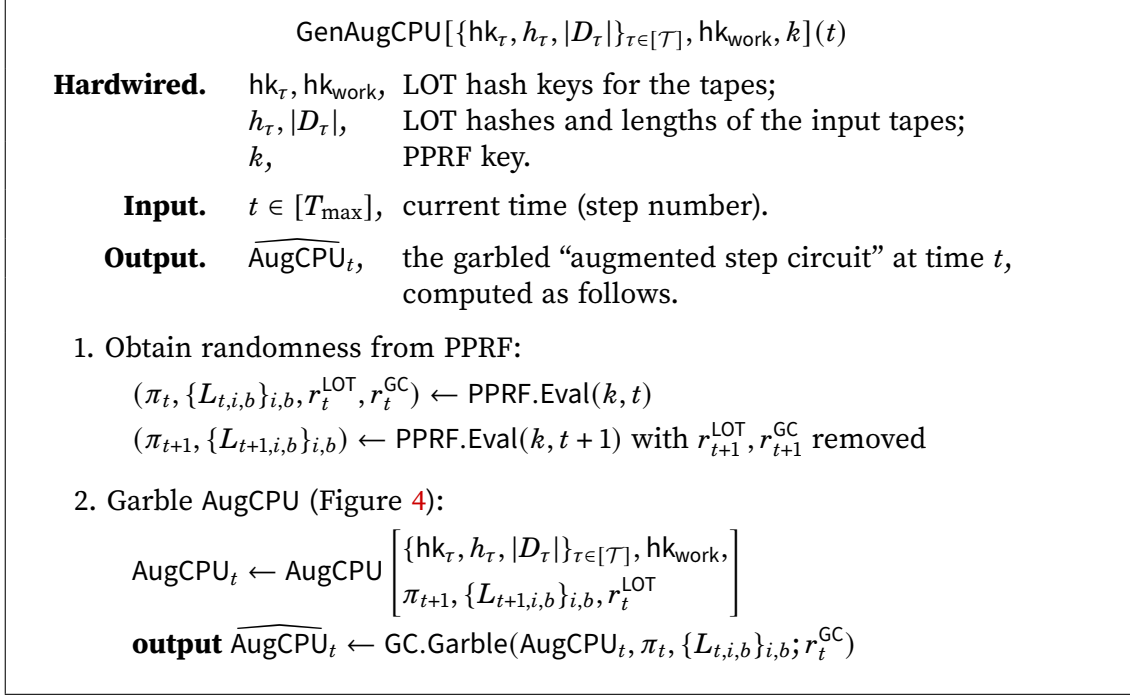
Simulator  $\mathcal{S}$ :

- **Simulating  $\widetilde{\text{DB}}$ :** We compute  $\text{ct}_{\text{PHFE}} = \text{PHFE.Enc}(\text{mpk}, (x, y))$ . where  $x = \text{DB}$  and  $y = (\text{mode} = 1, \perp, \perp, \perp, \text{SKE.sk}_3)$ .
- **Simulating  $\text{sk}_{f_j}$ :** We sample  $t_j \leftarrow \{0, 1\}^\lambda$ ,  $\text{ctk}_{1,j} = \text{SKE.Enc}(\text{SKE.sk}_1, 0)$  and  $\text{ctk}_{2,j} = \text{SKE.Enc}(\text{SKE.sk}_2, 0)$  and  $\text{ctk}_{3,j} = \text{SKE.Enc}(\text{SKE.sk}_3, \text{ctk}_{2,j})$ . Then, we set  $\text{sk}_{f_j} = \text{PHFE.KeyGen}(\text{msk}, f_j = f[\text{ctk}_{1,j}, \text{ctk}_{3,j}, t_j])$ .

## 4 Bounded LGRAM with Fixed-Memory Security

In this section, we present our LGRAM with fixed-memory security. The construction is based on the works of [GS18a, GOS18, AL18, KNTY19].

Roughly speaking, [GS18a] constructs a circuit garbling scheme from *adaptively* secure LOT with *local* proof of security and lifts it to an adaptively secure scheme using somewhere equivocal encryption, which is further developed to obtain a garbled RAM scheme with fixed-memory security (or unprotected memory access) [GOS18]. The work of [AL18] modularizes the construction and shows how to obtain a succinct garbling scheme using obfuscation for circuits with polynomial-sized domains, which is in turn



**Figure 3.** The circuit GenAugCPU in Construction 1.

implied by FE for circuits. The work of [KNTY19] shows that *selectively* secure LOT suffices and that such an LOT is implied by FE for circuits.

All of the works consider garbling schemes with no public input, and the security notions for their final products are simulation-based. Consequently, the length of the garbling necessarily [AIKW13] grows linearly with the input and output lengths. In contrast, our notion of garbling has a short private input and we are interested in indistinguishability and succinctness.

## 4.1 Construction

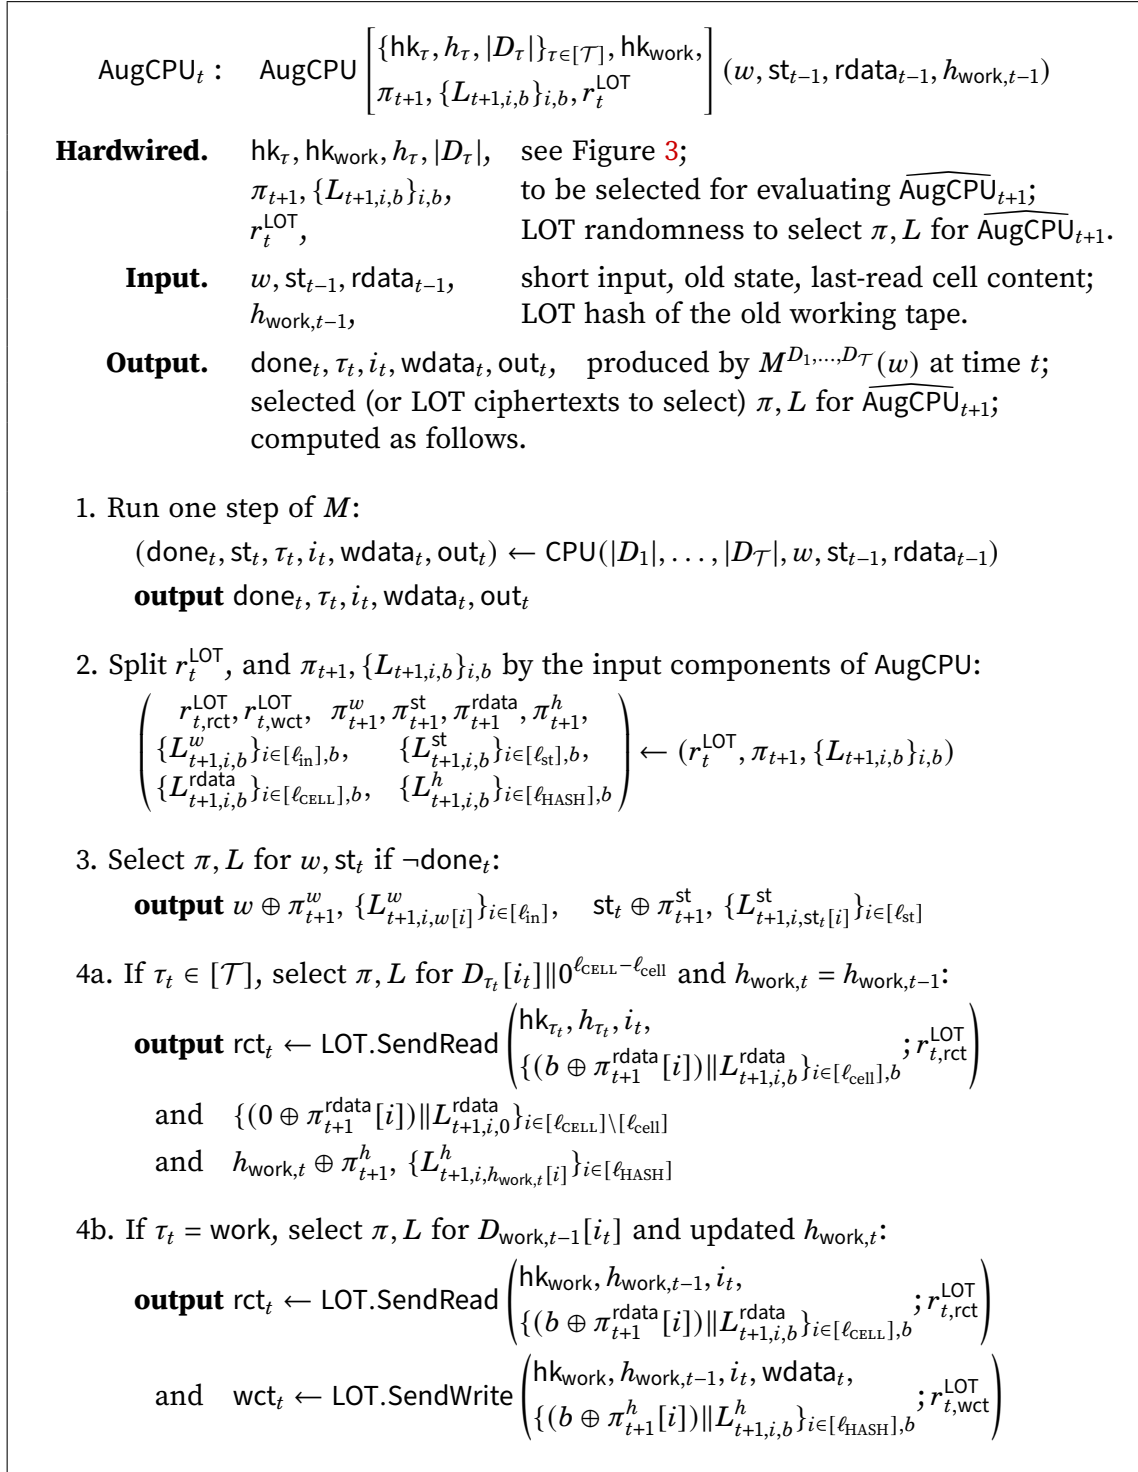
**Ingredients of Construction 1.** Let

- $i\mathcal{O}$  be a circuit obfuscator,
- LOT an updatable LOT with fast initialization,

$$\text{LOT} = (\text{LOT.HashGen}, \text{LOT.Hash}, \text{LOT.SendRead}, \text{LOT.RecvRead}, \\ \text{LOT.SendWrite}, \text{LOT.RecvWrite}, \text{LOT.Hash0s}),$$

- GC = (GC.Garble, GC.Eval) a circuit garbling scheme, and
- PPRF = (PPRF.Puncture, PPRF.Eval) a puncturable PRF.

**Construction 1** (bounded LGRAM with fixed-memory security). Our bounded LGRAM works as follows:



**Figure 4.** The circuit  $\text{AugCPU}$  in Construction 1.

- $\text{Compress}(1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, \tau, D_\tau)$  takes as input the cell length, the address length, the tape index, and the tape content. It runs

$$\text{hk}_\tau \xleftarrow{\$} \text{LOT.HashGen}(1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}), \quad (h_\tau, \widehat{D}_\tau) \leftarrow \text{LOT.Hash}(\text{hk}_\tau, D_\tau),$$

and outputs  $\text{digest}_\tau = (\text{hk}_\tau, h_\tau, |D_\tau|)$ .

- $\text{Garble}(T_{\text{max}}, M, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]})$  takes as input an upper bound of the running time, the machine, and the input tape digests. It runs

$$\text{hk}_{\text{work}} \xleftarrow{\$} \text{LOT.HashGen}(1^{\ell_{\text{CELL}}}, 1^{\ell_{\text{ADDR}}}), \quad h_{\text{work},0} \leftarrow \text{LOT.Hash0s}(\text{hk}_{\text{work}}, T_{\text{max}}),$$

samples PPRF key  $k$ , prepares  $\text{GenAugCPU}$  using  $M$  (Figure 3), and runs

$$\widehat{\text{GenAugCPU}} \xleftarrow{\$} i\mathcal{O}(\text{GenAugCPU}[k, \{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}]).$$

$\widehat{\text{GenAugCPU}}$  is padded to size  $\text{poly}(\lambda, |M|, \log T_{\text{max}})$  for the security proof to work. The algorithm computes using  $k$  and splits  $\pi, L$  for  $\widehat{\text{AugCPU}}_1$  as defined in Figure 4,

$$\begin{aligned} \pi_1 : & \pi_1^w, \pi_1^{\text{st}}, \pi_1^{\text{rdata}}, \pi_1^h, \\ \{L_{1,i,b}\}_{i,b} : & \{L_{1,i,b}^w\}_{i \in [\ell_{\text{in}}], b}, \{L_{1,i,b}^{\text{st}}\}_{i \in [\ell_{\text{st}}], b}, \{L_{1,i,b}^{\text{rdata}}\}_{i \in [\ell_{\text{CELL}}], b}, \{L_{1,i,b}^h\}_{i \in [\ell_{\text{HASH}}], b}, \end{aligned}$$

sets  $\text{st}_0 = 0^{\ell_{\text{st}}}$ ,  $\text{rdata}_0 = 0^{\ell_{\text{CELL}}}$ , and outputs

$$\begin{aligned} \widehat{M} = & \left( \text{hk}_{\text{work}}, \widehat{\text{GenAugCPU}}, \text{st}_0 \oplus \pi_1^{\text{st}}, \text{rdata}_0 \oplus \pi_1^{\text{rdata}}, h_{\text{work},0} \oplus \pi_1^h, \right. \\ & \left. \{L_{1,i,\text{st}_0[i]}^{\text{st}}\}_{i \in [\ell_{\text{st}}]}, \{L_{1,i,\text{rdata}_0[i]}^{\text{rdata}}\}_{i \in [\ell_{\text{CELL}}]}, \{L_{1,i,h_{\text{work},0}[i]}^h\}_{i \in [\ell_{\text{HASH}}]}, \right) \\ & \{L_{i,b} = (b \oplus \pi_1^w[i]) \| L_{1,i,b}^w\}_{i \in [\ell_{\text{in}}], b \in \{0,1\}}. \end{aligned}$$

- $\text{Eval}^{D_1, \dots, D_\tau}(T_{\text{max}}, M, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}, \widehat{M}, \{L_i\}_{i \in [\ell_{\text{in}}]})$  takes as input the upper bound of the running time, the machine, the input tape digests, the garbled machine, and a set of labels. It parses  $\widehat{M}$  as specified in Garble, and does the following:

1. Initialize by

$$\begin{aligned} (h_\tau, \widehat{D}_\tau) & \leftarrow \text{LOT.Hash}(\text{hk}_\tau, D_\tau) \quad \text{for } \tau \in [\mathcal{T}], \\ (h_{\text{work},0}, \widehat{D}_{\text{work},0}) & \leftarrow \text{LOT.Hash}(\text{hk}_{\text{work}}, 0^{T_{\text{max}}}), \\ \text{st}_0 & \leftarrow 0^{\ell_{\text{st}}}, \quad \text{rdata}_0 \leftarrow 0^{\ell_{\text{CELL}}}, \quad t \leftarrow 1. \end{aligned}$$

2. Writing

$$\begin{aligned} X_t & = w \| \text{st}_{t-1} \| \text{rdata}_{t-1} \| h_{\text{work},t-1}, \\ Y_t & \text{ is parts of } (X_{t+1} \oplus \pi_{t+1}) \text{ and } L_{t+1} \text{ and } \text{rct}_t, \text{wct}_t, \end{aligned}$$

run<sup>14</sup>

$$\begin{aligned} \widehat{\text{AugCPU}}_t & \leftarrow \widehat{\text{GenAugCPU}}(t), \\ (\text{done}_t, \tau_t, i_t, \text{wdata}_t, \text{out}_t, Y_t) & \leftarrow \text{GC.Eval}(\widehat{\text{AugCPU}}_t, X_t \oplus \pi_t, \{L_{t,i,X_t[i]}\}_i), \end{aligned}$$

and halt if  $\text{done}_t$  is set. Otherwise, output  $\text{out}_t$  and continue.

<sup>14</sup> $X_t$  itself is not needed for evaluation and is not supposed to be efficiently computable. The values of  $(X_t \oplus \pi_t)$  and  $\{L_{t,i,X_t[i]}\}_i$  for  $t = 1$  are read from  $\widehat{M}$ , and those for  $t > 1$  are computed by evaluating  $\widehat{\text{AugCPU}}_{t-1}$  as explained later.

3. If  $\tau_t \in [\mathcal{T}]$ , pick  $\text{rct}_t$  from  $Y_t$  and run

$$\{(\text{rdata}_t[i] \oplus \pi_{t+1,i}^{\text{rdata}}) \| L_{t+1,i,\text{rdata}_t[i]}^{\text{rdata}}\}_{i \in [\ell_{\text{cell}}]} \leftarrow \text{LOT.RecvRead}^{\widehat{D}_{\tau_t}}(\text{hk}_{\tau_t}, h_{\tau_t}, i_t, \text{rct}_t).$$

Combine it with parts of  $(X_{t+1} \oplus \pi_{t+1})$  and  $L_{t+1}$  already in  $Y_t$  to obtain  $(X_{t+1} \oplus \pi_{t+1})$  and  $\{L_{t+1,i,X_{t+1}[i]}\}_i$ .

4. Otherwise,  $\tau_t = \text{work}$ , then pick  $\text{rct}_t, \text{wct}_t$  from  $Y_t$ , run

$$\{(\text{rdata}_t[i] \oplus \pi_{t+1,i}^{\text{rdata}}) \| L_{t+1,i,\text{rdata}_t[i]}^{\text{rdata}}\}_{i \in [\ell_{\text{CELL}}]} \leftarrow \text{LOT.RecvRead}^{\widehat{D}_{\text{work},t-1}}(\text{hk}_{\text{work}}, h_{\text{work},t-1}, i_t, \text{rct}_t),$$

$$\{(h_{\text{work},t} \oplus \pi_{t+1,i}^h) \| L_{t+1,i,h_{\text{work},t}[i]}^h\}_{i \in [\ell_{\text{HASH}}]} \leftarrow \text{LOT.RecvWrite}^{\widehat{D}_{\text{work},t-1}} \left( \begin{array}{l} \text{hk}_{\text{work}}, h_{\text{work},t-1}, \\ i_t, \text{wdata}_t, \\ \text{wct}_t \end{array} \right),$$

updating  $\widehat{D}_{\text{work},t-1}$  into  $\widehat{D}_{\text{work},t}$ . Again, combine it with parts of  $(X_{t+1} \oplus \pi_{t+1})$  and  $L_{t+1}$  already in  $Y_t$  to obtain  $(X_{t+1} \oplus \pi_{t+1})$  and  $\{L_{t+1,i,X_{t+1}[i]}\}_i$ .

5. Let  $t \leftarrow t + 1$  and go back to step 2.

**Correctness and Efficiency.** They follow from those of all the ingredients and the invariant that  $\text{AugCPU}_t$  is evaluated on  $w, \text{st}_{t-1}, \text{rdata}_{t-1}, h_{\text{work},t-1}$  coinciding with those from the execution of  $M$ . We remark that the evaluation time is *not* instance-specific, because processing the empty working tape for LOT already takes time  $T_{\max} \text{poly}(\lambda, |M|)$ .

## 4.2 Security

**Theorem 11** (¶). *Suppose in Construction 1,  $i\mathcal{O}$  is secure for polynomial-sized domains (Definition 11), LOT is read- and write-secure (Definitions 14 and 15), and GC, PPRF are secure (Definition 16 and 17), then the constructed scheme is fixed-memory secure (Definition 4).*

The proof follows the pebbling strategy in [GS18a]. Recall that in  $\text{Exp}_{\text{LGRAM}}^b$ , the LGRAM labels are selected for  $w_b$ , and each  $\text{AugCPU}_t$  performs a step of  $M^{D_1, \dots, D_{\mathcal{T}}}(w_b)$ , for which we write  $M^{\dots}(w_b)$  hereafter in this proof. Along the hybrids from  $\text{Exp}_{\text{LGRAM}}^0$  to  $\text{Exp}_{\text{LGRAM}}^1$ , we gradually switch the trailing steps from those for  $M^{\dots}(w_0)$  to those for  $M^{\dots}(w_1)$ .

**Specification of Hybrids.** Each hybrid  $H_{t,s}$  is indexed by a natural number  $0 \leq t \leq T_{\max} + 1$  and a set  $s \subseteq [T_{\max}]$ . The indices  $t, s$  specify how  $\text{AugCPU}_t$  is garbled and what it does:

- when  $t < t$ , it runs the  $t^{\text{th}}$  step of  $M^{\dots}(w_0)$ ;
  - if  $t \notin s$ , the step is garbled normally;
  - if  $t \in s$ , the step is simulated with true randomness for labels and LOT encryption that selects the labels for  $\text{AugCPU}_{t+1}$ ;
- when  $t = t$ , it is simulated and outputs the labels so that  $\text{AugCPU}_{t+1}$  runs the  $(t+1)^{\text{st}}$  step of  $M^{\dots}(w_1)$ , i.e., it switches the execution from  $M^{\dots}(w_0)$  to  $M^{\dots}(w_1)$ ;
- when  $t > t$ , it is garbled normally but runs the  $t^{\text{th}}$  step of  $M^{\dots}(w_1)$ , due to the behavior of  $\text{AugCPU}_t$ .

There are two special cases:

- when  $t = 0$ , we give the LGRAM labels for  $w_1$  (think of them as the output of the zeroth step circuit) to the adversary so that  $\text{AugCPU}_1$  runs the first step of  $M^{\dots}(w_1)$  when the LGRAM is evaluated normally;
- when  $t > T$  for  $t \in \{t\} \cup \mathfrak{s}$ , neither execution has a  $t^{\text{th}}$  step, and we simulate this step as if it were the last step of the execution, whose output would not select any label for the  $(t + 1)^{\text{st}}$  step.<sup>15</sup>

In the parlance of [GS18a], consider a line graph where vertex  $t$  represents step  $t$ ,

- a gray pebble is placed on  $t \in \mathfrak{s}$  if  $t < t$  (the step is being simulated), and
- a black pebble is placed on  $t \geq t$  (the step is done for good).

However, the first black pebble (on  $t$ ) is special for us as step  $t$  facilitates the transition from  $M^{\dots}(w_0)$  to  $M^{\dots}(w_1)$ . We will use an optimized pebbling sequence [GS18a] to connect the hybrids.

Formally, let

$$\begin{aligned} \text{stS}(M, D_1, \dots, D_T, w_0) &= (\text{st}_{0,1}, \dots, \text{st}_{0,T-1}), \\ \text{stS}(M, D_1, \dots, D_T, w_1) &= (\text{st}_{1,1}, \dots, \text{st}_{1,T-1}), \\ \text{addrS}(\dots, w_0) = \text{addrS}(\dots, w_1) &= (\tau_1, i_1, \dots, \tau_{T-1}, i_{T-1}), \\ \text{writeS}(\dots, w_0) = \text{writeS}(\dots, w_1) &= (\text{wdata}_1, \dots, \text{wdata}_{T-1}), \\ \text{outS}(\dots, w_0) = \text{outS}(\dots, w_1) &= (\text{out}_1, \dots, \text{out}_{T-1}). \end{aligned}$$

We also write  $D_{\text{work},t}$  for the working tape content,  $\text{rdata}_t$  for the read cell content, and  $h_t, h_{\text{work},t}$  for the LOT hashes — due to the constraint of fixed-memory security, the values of them in the two executions coincide. Define

$$X_{b,t} = w_b \parallel \text{st}_{b,t-1} \parallel \text{rdata}_{t-1} \parallel h_{\text{work},t-1}.$$

In  $H_{t,\mathfrak{s}}$ , we change  $\widehat{\text{GenAugCPU}}$  in  $\widehat{M}$  to

$$i\mathcal{O} \left( \widehat{\text{GenAugCPU}}' \left[ \left\{ \text{hk}_t, h_t, |D_t| \right\}_{t \in [T]}, \text{hk}_{\text{work}}, \right. \right. \\ \left. \left. t, \mathfrak{s}, \mathring{k}_{\{t\} \cup \mathfrak{s}}, \{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t \}_{t \in \{t\} \cup \mathfrak{s}} \right] \right),$$

where

- $\widehat{\text{GenAugCPU}}'$  is shown in Figure 5,
- $\mathring{k}_{\{t\} \cup \mathfrak{s}}$  is a PPRF key punctured at  $\{t\} \cup \mathfrak{s}$ ,
- $\pi_t$ 's and  $L_{t,i,b}$ 's are random strings, and
- $\widehat{\text{AugCPU}}_t$ 's are the simulated garbled step circuits.

$$\text{GenAugCPU}' \left[ \begin{array}{l} \{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \\ t, \mathfrak{s}, \mathring{k}_{\{t\} \cup \mathfrak{s}}, \{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t\} \cup \mathfrak{s}} \end{array} \right] (t)$$

**Hardwired.**  $\text{hk}_\tau, \text{hk}_{\text{work}}, h_\tau, |D_\tau|$ , see Figure 3;  
 $t, \mathfrak{s}$ , switching time, hardwired step numbers;  
 $\mathring{k}_{\{t\} \cup \mathfrak{s}}$ , punctured PPRF key;  
 $\{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t\} \cup \mathfrak{s}}$ , hardwired randomness and garbled steps.  
**Input.**  $t \in [T_{\max}]$ , current time (step number).  
**Output.**  $\widehat{\text{AugCPU}}_t$ , computed as follows.

1. Check for hardwired steps, if  $t \in \{t\} \cup \mathfrak{s}$ :

**output** hardwired  $\widehat{\text{AugCPU}}_t$  and **terminate**

2. Otherwise, obtain randomness, either hardwired or from PPRF:

$$(\pi_t, \{L_{t,i,b}\}_{i,b}, r_t^{\text{LOT}}, r_t^{\text{GC}}) \leftarrow \text{PPRF.Eval}(\mathring{k}_{\{t\} \cup \mathfrak{s}}, t)$$

$$(\pi_{t+1}, \{L_{t+1,i,b}\}_{i,b}) \leftarrow \begin{cases} \text{hardwired,} & \text{if } t+1 \in \{t\} \cup \mathfrak{s}; \\ \text{PPRF.Eval}(\mathring{k}_{\{t\} \cup \mathfrak{s}}, t+1) \text{ with } r_{t+1}^{\text{LOT}}, r_{t+1}^{\text{GC}} \text{ removed,} & \\ \text{if } t+1 \notin \{t\} \cup \mathfrak{s}; \end{cases}$$

3. Garble AugCPU (Figure 4) as done in GenAugCPU (Figure 3):

$$\text{AugCPU}_t \leftarrow \text{AugCPU} \left[ \begin{array}{l} \{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \\ \pi_{t+1}, \{L_{t+1,i,b}\}_{i,b}, r_t^{\text{LOT}} \end{array} \right]$$

$$\text{output } \widehat{\text{AugCPU}}_t \leftarrow \text{GC.Garble}(\text{AugCPU}_t, \pi_t, \{L_{t,i,b}\}_{i,b}; r_t^{\text{GC}})$$

**Figure 5.** The circuit GenAugCPU' in the proof of Theorem 11.

We still use  $\pi_t, L_{t,i,b}$  for part of the PPRF output when  $t \notin \{t\} \cup \mathfrak{s}$ .

$\widetilde{\text{AugCPU}}_t$ 's are generated by

$$\text{GC.Garble}(1^{|\text{AugCPU}|}, Z_{b,t}, X_{0,t} \oplus \pi_t, \{L_{t,i,X_{0,t}[i]}\}_i) \text{ with } b = \begin{cases} 0, & \text{if } t < t \text{ and } t \in \mathfrak{s}; \\ 1, & \text{if } t = t; \end{cases}$$

where  $\text{GC.Garble}$  is a simulator for GC, and  $Z_{b,t}$  for  $b \in \{0,1\}$  consists of  $(\text{done}_t, \tau_t, i_t, \text{wdata}_t, \text{out}_t)$ , parts of  $(X_{b,t+1} \oplus \pi_{t+1})$  and  $L_{t+1}$ , and  $\widetilde{\text{rct}}_t, \widetilde{\text{wct}}_t$  (cf. Figure 4):

- If  $\neg \text{done}_t$ , then  $Z_{b,t}$  contains (dependent on  $b$ )

$$w_b \oplus \pi_{t+1}^w, \{L_{t+1,i,w_b[i]}^w\}_{i \in [\ell_{\text{in}}]}, \quad \text{st}_{b,t} \oplus \pi_{t+1}^w, \{L_{t+1,i,\text{st}_{b,t}[i]}^{\text{st}}\}_{i \in [\ell_{\text{st}}]}.$$

- If  $\tau_t \in [\mathcal{T}]$ , then  $Z_{b,t}$  contains (independent of  $b$ )

$$\begin{aligned} \widetilde{\text{rct}}_t &\stackrel{\mathfrak{s}}{\leftarrow} \text{LOT.SimRead} \left( \begin{array}{l} \text{hk}_{\tau_t}, D_{\tau_t}, i_t, \\ \{(rdata_t[i] \oplus \pi_{t+1}^{\text{rdata}}[i]) \| L_{t+1,i,rdata_t[i]}^{\text{rdata}}\}_{i \in [\ell_{\text{cell}}]} \end{array} \right), \\ &\{(0 \oplus \pi_{t+1}^{\text{rdata}}[i]) \| L_{t+1,i,0}^{\text{rdata}}\}_{i \in [\ell_{\text{CELL}}] \setminus [\ell_{\text{cell}}]}, \\ &h_{\text{work},t} \oplus \pi_{t+1}^h, \{L_{t+1,i,h_{\text{work},t}[i]}^h\}_{i \in [\ell_{\text{HASH}}]}. \end{aligned}$$

- If  $\tau_t = \text{work}$ , then  $Z_{b,t}$  contains (independent of  $b$ )

$$\begin{aligned} \widetilde{\text{rct}}_t &\stackrel{\mathfrak{s}}{\leftarrow} \text{LOT.SimRead} \left( \begin{array}{l} \text{hk}_{\text{work}}, D_{\text{work},t-1}, i_t, \\ \{(rdata_t[i] \oplus \pi_{t+1}^{\text{rdata}}[i]) \| L_{t+1,i,rdata_t[i]}^{\text{rdata}}\}_{i \in [\ell_{\text{CELL}}]} \end{array} \right), \\ \widetilde{\text{wct}}_t &\stackrel{\mathfrak{s}}{\leftarrow} \text{LOT.SimWrite} \left( \begin{array}{l} \text{hk}_{\text{work}}, D_{\text{work},t-1}, i_t, \text{wdata}_t, \\ \{(h_{\text{work},t}[i] \oplus \pi_{t+1}^h[i]) \| L_{t+1,i,h_{\text{work},t}[i]}^h\}_{i \in [\ell_{\text{HASH}}]} \end{array} \right). \end{aligned}$$

**Connecting the Hybrids.** We start with the experiments in Definition 4:

**Claim 12** (¶).  $\text{Exp}_{\text{LGRAM}}^0 \approx \text{H}_{T_{\max}+1, \emptyset}$  and  $\text{Exp}_{\text{LGRAM}}^1 \approx \text{H}_{0, \emptyset}$ .

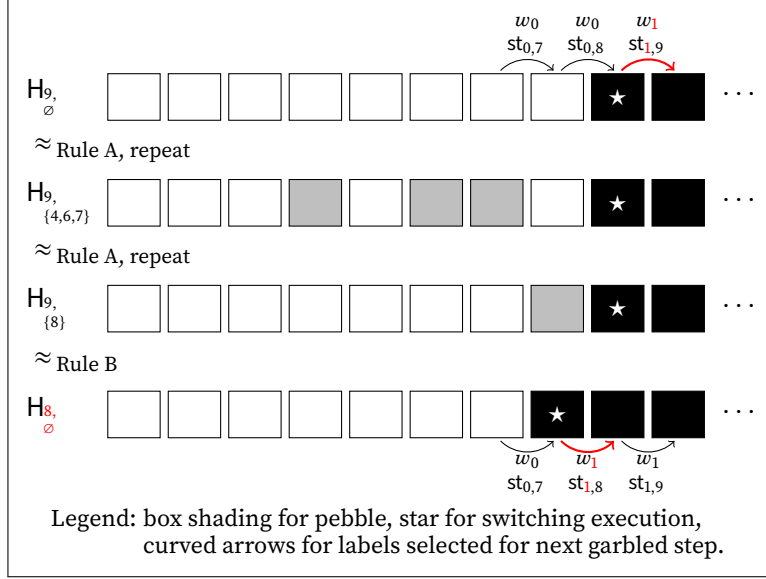
It remains to connect  $\text{H}_{T_{\max}+1, \emptyset}$  and  $\text{H}_{0, \emptyset}$ , for which we employ the pebbling strategy of [GOS18]. We state our claims in parallel with the pebbling rules:

**Claim 13** (Rule A; ¶). *A gray pebble can be placed on or removed from  $t^*$  if  $t^* = 1$  or a gray pebble is placed on  $(t^* - 1)$ . Formally,  $\text{H}_{t,\mathfrak{s}} \approx \text{H}_{t, \{t^*\} \cup \mathfrak{s}}$  for  $t^* \notin \mathfrak{s}$  and  $t^* < t$  if  $t^* = 1$  or  $t^* - 1 \in \mathfrak{s}$ .*

**Claim 14** (Rule B; ¶). *A gray pebble on  $t^*$  can be replaced by a black pebble if black pebbles are placed on all vertices greater than  $t^*$ . Formally,  $\text{H}_{t,\mathfrak{s}} \approx \text{H}_{t-1, \mathfrak{s} \setminus \{t-1\}}$  if  $t-1 \in \mathfrak{s}$  (here,  $t^* = t-1$ ). Moreover,  $\text{H}_{1, \emptyset} \approx \text{H}_{0, \emptyset}$ .*

<sup>15</sup>This pretention trick unifies the usage of garbled circuit security. Since an out-of-range  $t^{\text{th}}$  step is pretended to repeat the last existent step, simulating it removes all the labels for the  $(t+1)^{\text{st}}$  step, and consequently (and inductively), we can pretend that the set of labels for the  $(t+1)^{\text{st}}$  step corresponding to the input to the last existent step were revealed, so that the  $(t+1)^{\text{st}}$  step can also be pretended to repeat the last step. Without pretending, we will need another security notion of garbled circuits, namely that it can be simulated without any output if no input labels are given. This notion holds for many known constructions, but is not implied by the usual version.





**Figure 6.** Typical hybrid transitions in the proof of Theorem 11.

Claim 14 is different from rule B in [GOS18] (no need for a gray pebble on  $(t^* - 1)$  for us) because the step corresponding to the first black pebble is always simulated.

It is helpful to consider a virtual vertex 0 (corresponding to  $\tilde{M}$  and the LGRAM labels), where a gray pebble is initially placed, and put 0 into  $\mathfrak{s}$ . With the virtual vertex, the separate case  $t^* = 1$  in rule A is just a special case of  $t^* - 1 = 0 \in \mathfrak{s}$ , and the separate case  $H_{1, \emptyset} \approx H_{0, \emptyset}$  in rule B becomes a special case of the general rule (which would read  $H_{1, \{0\}} \approx H_{0, \emptyset}$  instead). However, for consistency with existing literature, we will go without the virtual vertex in the text below.

The hybrid sequence can be built using an optimized pebbling strategy:

**Lemma 15** ([Ben89, GS18a]). *There exists an efficient algorithm that on input  $1^{T_{\max}}$  computes a pebbling sequence for a line graph of size  $T_{\max}$  with  $O(\log T_{\max})$  gray pebbles at any time during the  $\text{poly}(T_{\max})$  moves using rule A or B, starting with no pebbles and ending with black pebbles on all vertices. As a corollary, the efficient algorithm can compute*

$$t_0 = T_{\max} + 1, \mathfrak{s}_0 = \emptyset, \quad c_1, t_1, \mathfrak{s}_1, \quad c_2, t_2, \mathfrak{s}_2, \quad \dots,$$

$$c_{\text{poly}(T_{\max})-1}, t_{\text{poly}(T_{\max})-1} = 1, \mathfrak{s}_{\text{poly}(T_{\max})-1} = \emptyset, \quad c_{\text{poly}(T_{\max})}, t_{\text{poly}(T_{\max})} = 0, \mathfrak{s}_{\text{poly}(T_{\max})} = \emptyset,$$

such that  $\max_j |s_j| = O(\log T_{\max})$  and  $H_{t_{j-1}, s_{j-1}} \approx H_{t_j, s_j}$  by Claim  $c_j \in \{13, 14\}$ .

A typical sequence of hybrid transitions is depicted in Figure 6.

*Proof* (Theorem 11). By Claims 12, 13, and 14, and Lemma 15,  $\text{Exp}_{\text{LGRAM}}^0 \approx \text{Exp}_{\text{LGRAM}}^1$ .  $\square$

We remark that for full rigor,  $T_{\max}$  is a random variable, not a fixed number for each  $\lambda$ . It cannot be derandomized even in the non-uniform setting, because  $\mathcal{A}$ , given  $\text{hk}$ 's returned for the input tapes, can choose  $T_{\max}$  adaptively, whose randomness originates from that of the LGRAM scheme. This phenomenon is not uncommon and can be solved by either guessing  $T_{\max}$  (or targeting a specific  $T_{\max}$  in the non-uniform setting), or considering a polynomial upper bound of  $T_{\max}$  and supplying appropriate definitions for hybrids with out-of-range indices.

*Proof* (Claim 12). Both follow from the security of  $i\mathcal{O}$  for polynomial-sized domains.  $\square$

*Proof* (Claim 13). Let  $t, s, t^*$  be such that  $t^* \notin s$  and  $t^* < t$  and either  $t^* = 1$  or  $t^* - 1 \in s$ . The hybrids below are mostly identical to  $H_{t,s}$  except the contents hardwired into  $\text{GenAugCPU}'$  are modified. We specify them:

- $G_0$ . This is just  $H_{t,s}$ , where the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, t, s, \mathring{k}_{\{t\} \cup s},$$

$$\{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t\} \cup s}.$$

- $G_1$ . In this hybrid, the PPRF key  $k$  is punctured at  $\{t^*, t\} \cup s$  instead of  $\{t\} \cup s$ , and the  $(t^*)^{\text{th}}$  garbled step (but not its randomness) is hardwired into  $\text{GenAugCPU}'$ , i.e., the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, t, \boxed{\{t^*\} \cup s}, \boxed{\mathring{k}_{\{t,t^*\} \cup s}},$$

$$\{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t\} \cup s}, \boxed{\widehat{\text{AugCPU}}_{t^*}},$$

where the newly hardwired information is computed using (Figure 5 Step 3)

$$(\pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}, r_{t^*}^{\text{LOT}}, r_{t^*}^{\text{GC}}) \leftarrow \text{PPRF.Eval}(k, t^*),$$

$$\widehat{\text{AugCPU}}_{t^*} \leftarrow \text{AugCPU} \left[ \begin{array}{l} \{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \\ \pi_{t^*+1}, \{L_{t^*+1,i,b}\}_{i,b}, r_{t^*}^{\text{LOT}} \end{array} \right],$$

$$\widehat{\text{AugCPU}}_{t^*} \leftarrow \text{GC.Garble}(\widehat{\text{AugCPU}}_{t^*}, \pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}; r_{t^*}^{\text{GC}}).$$

This is different from  $\widehat{\text{AugCPU}}_t$ 's for  $t \in \{t\} \cup s$ , which are simulated. The absence of  $\pi_{t^*}$  and  $\{L_{t^*,i,b}\}_{i,b}$  in  $\text{GenAugCPU}'$  is fine, because they are only used in Step 2 (Figure 5) for  $t = t^* - 1$  and  $t = t^*$ . In  $G_1$ , that branch is not taken for those two values of  $t$  as both of them are handled by Step 1. Therefore, the two versions of  $\text{GenAugCPU}'$  in  $G_0$  and  $G_1$  have identical truth tables (and are padded appropriately to the same size), and  $G_0 \approx G_1$  follows from the security of  $i\mathcal{O}$  for polynomial-sized domains.

- $G_2$ . In this hybrid,  $\text{PPRF.Eval}(k, t^*)$  is replaced by a uniformly random string, i.e., the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, t, \{t^*\} \cup s, \mathring{k}_{\{t,t^*\} \cup s},$$

$$\{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t\} \cup s}, \widehat{\text{AugCPU}}_{t^*},$$

where

$$(\pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}, r_{t^*}^{\text{LOT}}, r_{t^*}^{\text{GC}}) \leftarrow \boxed{\text{uniformly random}},$$

$$\widehat{\text{AugCPU}}_{t^*} \leftarrow \text{AugCPU} \left[ \begin{array}{l} \{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \\ \pi_{t^*+1}, \{L_{t^*+1,i,b}\}_{i,b}, r_{t^*}^{\text{LOT}} \end{array} \right],$$

$$\widehat{\text{AugCPU}}_{t^*} \leftarrow \text{GC.Garble}(\widehat{\text{AugCPU}}_{t^*}, \pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}; r_{t^*}^{\text{GC}}).$$

$G_1 \approx G_2$  follows from the security of PPRF.

- $G_3$ . In this hybrid,  $\widehat{\text{AugCPU}}_{t^*}$  is simulated instead of being normally generated, i.e., the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t, \quad \{t^*\} \cup \mathfrak{s}, \quad \mathring{k}_{\{t, t^*\} \cup \mathfrak{s}}, \\ \{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t\} \cup \mathfrak{s}}, \quad \widehat{\text{AugCPU}}_{t^*},$$

where

$$(\pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}, r_{t^*}^{\text{LOT}}, r_{t^*}^{\text{GC}}) \leftarrow \text{uniformly random}, \\ \widehat{\text{AugCPU}}_{t^*} \stackrel{\$}{\leftarrow} \text{GC.Garble}(1^{|\text{AugCPU}|}, Z_{0,t^*}, X_{0,t^*} \oplus \pi_{t^*}, \{L_{t^*,i,X_{0,t^*}[i]}\}_i), \\ Z_{0,t^*} : \pi_{t^*+1}, \{L_{t^*+1,i,b}\}_{i,b} \begin{cases} \text{in the clear,} \\ \text{in } \widetilde{\text{rct}}_{t^*} \leftarrow \text{LOT.SendRead}(\dots; r_{t^*,\text{rct}}^{\text{LOT}}), \\ \text{in } \widetilde{\text{wct}}_{t^*} \leftarrow \text{LOT.SendWrite}(\dots; r_{t^*,\text{wct}}^{\text{LOT}}). \end{cases}$$

This is different from  $Z_{b,t}$ 's for  $t \in \{t\} \cup \mathfrak{s}$ , which might contain simulated  $\widetilde{\text{rct}}_t$ 's and  $\widetilde{\text{wct}}_t$ 's. In  $G_2$ , the augmented step circuit  $\widehat{\text{AugCPU}}_{t^*}$  is garbled with truly random  $\pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}, r_{t^*}^{\text{GC}}$ . If  $t^* = 1$ , then  $\pi_1, \{L_{1,i,b}\}_{i,b}$  only additional appear in  $\widetilde{M}$  and the LGRAM labels given to the adversary as  $(X_{0,1} \oplus \pi_1)$  and  $\{L_{1,i,X_{0,1}[i]}\}_i$ . If  $t^* - 1 \in \mathfrak{s}$ , then they are only additionally used to simulate

$$\widehat{\text{AugCPU}}_{t^*-1} \stackrel{\$}{\leftarrow} \text{GC.Garble}(\dots, Z_{0,t^*-1}, \dots),$$

where  $Z_{0,t^*-1}$  contains only  $(X_{0,t^*} \oplus \pi_{t^*})$  and  $\{L_{t^*,i,X_{0,t^*}[i]}\}_i$ , potentially in the clear, in  $\widetilde{\text{rct}}_{t^*-1}$ , in  $\widetilde{\text{wct}}_{t^*-1}$ , and by imagination.<sup>16</sup> Also,  $\widehat{\text{AugCPU}}_t(X_{0,t^*}) = Z_{0,t^*}$ . It follows from the security of GC that  $G_2 \approx G_3$ .

- $G_4$ . In this hybrid,  $Z_{0,t^*}$  contains  $\widetilde{\text{rct}}_{t^*}, \widetilde{\text{wct}}_{t^*}$  simulated using  $\text{LOT.SimRead}$  and  $\text{LOT.SimWrite}$  as needed, instead of normally generated ones, i.e., the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t, \quad \{t^*\} \cup \mathfrak{s}, \quad \mathring{k}_{\{t, t^*\} \cup \mathfrak{s}}, \\ \{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t\} \cup \mathfrak{s}}, \quad \widehat{\text{AugCPU}}_{t^*},$$

where

$$(\pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}, r_{t^*}^{\text{LOT}}) \leftarrow \text{uniformly random}, \\ \widehat{\text{AugCPU}}_{t^*} \stackrel{\$}{\leftarrow} \text{GC.Garble}(1^{|\text{AugCPU}|}, Z_{0,t^*}, X_{0,t^*} \oplus \pi_{t^*}, \{L_{t^*,i,X_{0,t^*}[i]}\}_i), \\ Z_{0,t^*} : \begin{cases} \text{in the clear,} \\ \text{in } \widetilde{\text{rct}}_{t^*} \stackrel{\$}{\leftarrow} \text{LOT.SimRead}(\dots), \\ \text{in } \widetilde{\text{wct}}_{t^*} \stackrel{\$}{\leftarrow} \text{LOT.SimWrite}(\dots). \end{cases}$$

The LOT ciphertexts  $\widetilde{\text{rct}}_{t^*}, \widetilde{\text{wct}}_{t^*}$  are encrypted using truly random  $r_{t^*}^{\text{LOT}}$  in  $G_3$ . Each database is known before its LOT hash key is provided to the adversary: for an

<sup>16</sup>When  $t^* > T$ , none of  $\pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}$  appears in  $Z_{0,t^*-1}$ . We pretend  $X_{0,t^*} = X_{0,T}$  and that  $(X_{0,t^*} \oplus \pi_{t^*})$  and  $\{t^*, i, X_{0,t^*}[i]\}_i$  are revealed, imagining that the  $(t^*)^{\text{th}}$  step repeated the last existent step.

input tape,  $hk_\tau$  is provided after  $D_\tau$  is chosen; for the working tape,  $D_{\text{work},t^*-1}$  and  $i_{t^*}, wdata_{t^*}$  are known once the adversary commits to the challenge  $M, w_0, w_1$ , after which  $hk_{\text{work}}$  is provided. Therefore, it follows from the read/write security of LOT that  $G_3 \approx G_4$ .

- $G_5$ . In this hybrid, we hardwire  $\pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}$  into GenAugCPU', i.e., the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t, \quad \{t^*\} \cup \mathfrak{s}, \quad \mathring{k}_{\{t,t^*\} \cup \mathfrak{s}}, \\ \{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t\} \cup \mathfrak{s}}, \quad \boxed{\pi_{t^*}, \{L_{t^*,i,b}\}_{i,b}}, \widehat{\text{AugCPU}}_{t^*}.$$

$G_4 \approx G_5$  follows from the security of  $i\mathcal{O}$  for polynomial-sized domains.

By inspection,  $G_5$  is just  $H_{t,\{t^*\} \cup \mathfrak{s}}$ . Therefore,  $H_{t,\mathfrak{s}} \equiv G_0 \approx G_5 \equiv H_{t,\{t^*\} \cup \mathfrak{s}}$ .  $\square$

*Proof (Claim 14).* Let  $t, \mathfrak{s}$  be such that  $t = 1$  or  $t - 1 \in \mathfrak{s}$ . The hybrids are mostly identical to  $H_{t,\mathfrak{s}}$  except for the contents hardwired into GenAugCPU':

- $G_0$ . This is just  $H_{t,\mathfrak{s}}$ , where the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t, \quad \mathfrak{s}, \quad \mathring{k}_{\{t\} \cup \mathfrak{s}}, \\ \{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t\} \cup \mathfrak{s}}.$$

- $G_1$ . In this hybrid, we no longer directly hardwire  $\pi_t, \{L_{t,i,b}\}_{i,b}$  into GenAugCPU' for  $t = t$ , i.e., the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t, \quad \mathfrak{s}, \quad \mathring{k}_{\{t\} \cup \mathfrak{s}}, \\ \{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{\boxed{t \in \mathfrak{s}}}, \quad \boxed{\pi_t, \{L_{t,i,b}\}_{i,b}}, \widehat{\text{AugCPU}}_t,$$

where

$$\begin{aligned} & (\pi_t, \{L_{t,i,b}\}_{i,b}) \stackrel{\mathfrak{s}}{\leftarrow} \text{uniformly random;} \\ & \left\{ \begin{array}{ll} \widehat{M}, \text{ LGRAM labels} \leftarrow X_{0,1} \oplus \pi_1, \{L_{1,i,X_{0,1}[i]}\}_i, \dots, & \text{if } t = 1; \\ \widehat{\text{AugCPU}}_{t-1} \stackrel{\mathfrak{s}}{\leftarrow} \text{GC.Garble} \left( \begin{array}{c} 1^{|\text{AugCPU}|}, Z_{0,t-1}, \\ X_{0,t-1} \oplus \pi_{t-1}, \{L_{t-1,i,X_{0,t-1}[i]}\}_i \end{array} \right), \\ Z_{0,t-1} : X_{0,t} \oplus \pi_t, \{L_{t,i,X_{0,t}[i]}\}_i \text{ in the clear, } \widetilde{\text{rct}}_{t-1}, \widetilde{\text{wct}}_{t-1}, & \text{if } t - 1 \in \mathfrak{s}; \\ \widehat{\text{AugCPU}}_t \stackrel{\mathfrak{s}}{\leftarrow} \text{GC.Garble} \left( \begin{array}{c} 1^{|\text{AugCPU}|}, Z_{1,t}, \\ X_{0,t} \oplus \pi_t, \{L_{t,i,X_{0,t}[i]}\}_i \end{array} \right), \\ Z_{1,t} : X_{1,t+1} \oplus \pi_{t+1}, \{L_{t+1,i,X_{1,t+1}[i]}\}_i \text{ in the clear, } \widetilde{\text{rct}}_t, \widetilde{\text{wct}}_t, & \text{if } t \leq T_{\max}. \end{array} \right. \end{aligned}$$

Note that GenAugCPU' still indirectly contains  $\pi_t, \{L_{t,i,b}\}_{i,b}$  via  $\widehat{M}$  and LGRAM labels, or  $Z_{0,t-1}$ . Removal of their direct hardwiring does not alter the truth table of GenAugCPU' as they are only used in Step 2 (Figure 5) for  $t = t$  and  $t = t - 1$ , a branch never taken for those values of  $t$  because they are handled by Step 1. By the security of  $i\mathcal{O}$  for polynomial-sized domains,  $G_0 \approx G_1$ .

- $G_2$ . In this hybrid, we change the input of  $\widehat{\text{AugCPU}}_t$  from  $X_{0,t}$  to  $X_{1,t}$  by modifying the permuted input and the labels selected for  $\widehat{\text{AugCPU}}_t$ , i.e., the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t, \quad \mathfrak{s}, \quad \mathring{k}_{\{t\} \cup \mathfrak{s}},$$

$$\{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \mathfrak{s}}, \quad \widehat{\text{AugCPU}}_t,$$

where

$$(\pi_t, \{L_{t,i,b}\}_{i,b}) \stackrel{\mathfrak{s}}{\leftarrow} \text{uniformly random};$$

$$\left\{ \begin{array}{l} \widehat{M}, \text{ LGRAM labels} \leftarrow \boxed{X_{1,1} \oplus \pi_1, \{L_{1,i,X_{1,1}[i]}\}_i}, \dots, \quad \text{if } t = 1; \\ \widehat{\text{AugCPU}}_{t-1} \stackrel{\mathfrak{s}}{\leftarrow} \text{GC.Garble} \left( \begin{array}{c} 1^{|\text{AugCPU}|}, Z_{0,t-1}, \\ X_{0,t-1} \oplus \pi_{t-1}, \{L_{t-1,i,X_{0,t-1}[i]}\}_i \end{array} \right), \\ Z_{0,t-1} : \boxed{X_{1,t} \oplus \pi_1, \{L_{t,i,X_{1,t}[i]}\}_i} \text{ in the clear, } \widetilde{\text{rct}}_{t-1}, \widetilde{\text{wct}}_{t-1}, \quad \text{if } t-1 \in \mathfrak{s}; \\ \widehat{\text{AugCPU}}_t \stackrel{\mathfrak{s}}{\leftarrow} \text{GC.Garble} \left( \begin{array}{c} 1^{|\text{AugCPU}|}, Z_{1,t}, \\ \boxed{X_{1,1} \oplus \pi_1, \{L_{1,i,X_{1,1}[i]}\}_i} \end{array} \right), \\ Z_{1,t} : X_{1,t+1} \oplus \pi_{t+1}, \{L_{t+1,i,X_{1,t+1}[i]}\}_i \text{ in the clear, } \widetilde{\text{rct}}_t, \widetilde{\text{wct}}_t, \quad \text{if } t \leq T_{\max}. \end{array} \right.$$

This change does not alter the distribution, i.e.,  $G_1 \equiv G_2$ , because  $\pi_t$  is a one-time pad that hides  $X_{0,t}$  or  $X_{1,t}$  and either set of  $\{L_{t,i,b}\}_{i,b}$  chosen by  $X_{0,t}$  and  $X_{1,t}$  is simply random.

- $G_3$ . In this hybrid, we undo the simulation of  $\widetilde{\text{rct}}_t$  and  $\widetilde{\text{wct}}_t$ , i.e., the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t, \quad \mathfrak{s}, \quad \mathring{k}_{\{t\} \cup \mathfrak{s}},$$

$$\{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \mathfrak{s}}, \quad \widehat{\text{AugCPU}}_t,$$

where

$$(\pi_t, \{L_{t,i,b}\}_{i,b}, \boxed{r_t^{\text{LOT}}}) \stackrel{\mathfrak{s}}{\leftarrow} \text{uniformly random};$$

$$\left\{ \begin{array}{l} \widehat{M}, \text{ LGRAM labels} \leftarrow X_{1,1} \oplus \pi_1, \{L_{1,i,X_{1,1}[i]}\}_i, \dots, \quad \text{if } t = 1; \\ \widehat{\text{AugCPU}}_{t-1} \stackrel{\mathfrak{s}}{\leftarrow} \text{GC.Garble} \left( \begin{array}{c} 1^{|\text{AugCPU}|}, Z_{0,t-1}, \\ X_{0,t-1} \oplus \pi_{t-1}, \{L_{t-1,i,X_{0,t-1}[i]}\}_i \end{array} \right), \\ Z_{0,t-1} : X_{1,t} \oplus \pi_1, \{L_{t,i,X_{1,t}[i]}\}_i \text{ in the clear, } \widetilde{\text{rct}}_{t-1}, \widetilde{\text{wct}}_{t-1}, \quad \text{if } t-1 \in \mathfrak{s}; \\ \widehat{\text{AugCPU}}_t \stackrel{\mathfrak{s}}{\leftarrow} \text{GC.Garble} \left( \begin{array}{c} 1^{|\text{AugCPU}|}, Z_{1,t}, \\ X_{1,1} \oplus \pi_1, \{L_{1,i,X_{1,1}[i]}\}_i \end{array} \right), \\ Z_{1,t} : \boxed{\pi_{t+1}, \{L_{t+1,i,b}\}_{i,b}} \text{ in the clear, } \boxed{\text{rct}_t, \text{wct}_t}, \quad \text{if } t \leq T_{\max}. \end{array} \right.$$

$G_2 \approx G_3$  by the read/write security of LOT.

- $G_4$ . In this hybrid, we undo the simulation of  $\widehat{\text{AugCPU}}_t$ , i.e., the hardwired information is

$$\{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t, \quad \mathfrak{s}, \quad \mathring{k}_{\{t\} \cup \mathfrak{s}},$$

$$\{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \mathfrak{s}}, \quad \widehat{\text{AugCPU}}_t,$$

where

$$\begin{aligned}
& (\pi_t, \{L_{t,i,b}\}_{i,b}, r_t^{\text{LOT}}, r_t^{\text{GC}}) \stackrel{\$}{\leftarrow} \text{uniformly random}; \\
& \left\{ \begin{array}{ll}
\widehat{M}, \text{ LGRAM labels} \leftarrow X_{1,1} \oplus \pi_1, \{L_{1,i,X_{1,1}[i]}\}_i, \dots, & \text{if } t = 1; \\
\widehat{\text{AugCPU}}_{t-1} \stackrel{\$}{\leftarrow} \text{GC.Garble} \left( \begin{array}{l} 1^{|\text{AugCPU}|}, Z_{0,t-1}, \\ X_{0,t-1} \oplus \pi_{t-1}, \{L_{t-1,i,X_{0,t-1}[i]}\}_i \end{array} \right), \\
Z_{0,t-1} : X_{1,t} \oplus \pi_1, \{L_{t,i,X_{1,t}[i]}\}_i \text{ in the clear, } \widetilde{\text{rct}}_{t-1}, \widetilde{\text{wct}}_{t-1}, & \text{if } t-1 \in \mathcal{S}; \\
\text{AugCPU}_t \leftarrow \text{AugCPU} \left[ \begin{array}{l} \{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \\ \pi_{t+1}, \{L_{t+1,i,b}\}_{i,b}, r_t^{\text{LOT}} \end{array} \right], \\
\widehat{\text{AugCPU}}_t \leftarrow \text{GC.Garble}(\text{AugCPU}_t, \pi_t, \{L_{t,i,b}\}_{i,b}; r_t^{\text{GC}}), & \text{if } t \leq T_{\max}.
\end{array} \right.
\end{aligned}$$

Since  $\text{AugCPU}_t(X_{1,t}) = Z_{1,t}$ , it follows from the security of GC that  $G_3 \approx G_4$ .

- $G_5$ . In this hybrid, we change the randomness for step  $t$  to the PPRF evaluation, i.e., the hardwired information is

$$\begin{aligned}
& \{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t, \quad \mathcal{S}, \quad \mathring{k}_{\{t\} \cup \mathcal{S}}, \\
& \{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \mathcal{S}}, \quad \widehat{\text{AugCPU}}_t,
\end{aligned}$$

where

$$\begin{aligned}
& (\pi_t, \{L_{t,i,b}\}_{i,b}, r_t^{\text{LOT}}, r_t^{\text{GC}}) \stackrel{\$}{\leftarrow} \text{PPRF.Eval}(k, t); \\
& \left\{ \begin{array}{ll}
\widehat{M}, \text{ LGRAM labels} \leftarrow X_{1,1} \oplus \pi_1, \{L_{1,i,X_{1,1}[i]}\}_i, \dots, & \text{if } t = 1; \\
\widehat{\text{AugCPU}}_{t-1} \stackrel{\$}{\leftarrow} \text{GC.Garble} \left( \begin{array}{l} 1^{|\text{AugCPU}|}, Z_{0,t-1}, \\ X_{0,t-1} \oplus \pi_{t-1}, \{L_{t-1,i,X_{0,t-1}[i]}\}_i \end{array} \right), \\
Z_{0,t-1} : X_{1,t} \oplus \pi_1, \{L_{t,i,X_{1,t}[i]}\}_i \text{ in the clear, } \widetilde{\text{rct}}_{t-1}, \widetilde{\text{wct}}_{t-1}, & \text{if } t-1 \in \mathcal{S}; \\
\text{AugCPU}_t \leftarrow \text{AugCPU} \left[ \begin{array}{l} \{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \\ \pi_{t+1}, \{L_{t+1,i,b}\}_{i,b}, r_t^{\text{LOT}} \end{array} \right], \\
\widehat{\text{AugCPU}}_t \leftarrow \text{GC.Garble}(\text{AugCPU}_t, \pi_t, \{L_{t,i,b}\}_{i,b}; r_t^{\text{GC}}), & \text{if } t \leq T_{\max}.
\end{array} \right.
\end{aligned}$$

$G_4 \approx G_5$  by the security of PPRF.

- $G_6$ . In this hybrid, we no longer puncture the PPRF key at  $t$  and no longer hardwire the randomness nor the garbled step for  $t$ , i.e., the hardwired information is

$$\begin{aligned}
& \{\text{hk}_\tau, h_\tau, |D_\tau|\}_{\tau \in [\mathcal{T}]}, \text{hk}_{\text{work}}, \quad t-1, \quad \mathcal{S} \setminus \{t-1\}, \quad \mathring{k}_{\{t-1\} \cup (\mathcal{S} \setminus \{t-1\})}, \\
& \{\pi_t, \{L_{t,i,b}\}_{i,b}, \widehat{\text{AugCPU}}_t\}_{t \in \{t-1\} \cup (\mathcal{S} \setminus \{t-1\})},
\end{aligned}$$

and  $\widehat{M}$  and the LGRAM labels given to the adversary contain  $X_{1,1} \oplus \pi_1, \{L_{1,i,X_{1,1}[i]}\}_i$  if  $t = 1$ . By the security of  $i\mathcal{O}$  for polynomial-sized domains,  $G_5 \approx G_6$ .

By inspection,  $G_6$  is just  $H_{t-1, \mathcal{S} \setminus \{t-1\}}$ . Therefore,  $H_{t, \mathcal{S}} \equiv G_0 \approx G_6 \equiv H_{t-1, \mathcal{S} \setminus \{t-1\}}$ .  $\square$

## 5 Transformations of LGRAM

In this section, we describe the transformations of LGRAM to upgrade its security and functionality. We adapt known transformations [CH16] to our syntax of bounded LGRAM to go from fixed-memory security, to fixed-address section (Section 5.1), and to full security (Section 5.2). We employ the powers-of-two transformation [AL18] to construct an unbounded LGRAM from a bounded one (Section 5.3).

### 5.1 Fixed-Memory to Fixed-Address

Given an LGRAM with fixed-memory security, we construct one with fixed-address security. Recall that in such a scheme, we want to hide the write sequence, consisting of the contents written to the working tape. The addresses where each read and write occurs are not protected, though. We employ a transformation due to [CH16] upgrading fixed-memory security to fixed-address security, by encrypting the data before they are written to the working tape. In more details, given a RAM  $M$ , we construct another machine  $M'$ , which runs two copies of  $M$  in parallel and encrypts the two logical working tapes of  $M$  using two puncturable PRF keys. Normally, only one copy is used and the other mimics its memory access pattern. The idle copy is used in the security proof.

**Ingredients of Construction 2.** Let

- LGRAM' = (Compress', Garble', Eval') be an LGRAM with fixed-memory security, and
- PPRF = (PPRF.Puncture, PPRF.Eval) a puncturable PRF.

**Construction 2** (LGRAM with fixed-address security). Our scheme works as follows:

- Compress( $1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, \tau, D_\tau$ ) is the same as Compress'.
- Garble( $T_{\text{max}}, M, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}$ ) prepares  $M'$  from  $M$  as shown in Figure 7, samples  $\beta \xleftarrow{\$} \{0, 1\}$  and two random PPRF keys  $k_0, k_1$ , runs

$$(\widehat{M}', \{L'_{i,b}\}_{i,b}) \xleftarrow{\$} \text{Garble}'(T_{\text{max}}, M', \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]})$$

and splits  $\{L'_{i,b}\}_{i,b}$ , according to the part of input to  $M'$  they select, into

$$\{L_b^{\text{B}}\}_b, \{L_{i,b}^{\text{W}_0}\}_{i,b}, \{L_{i,b}^{\text{W}_1}\}_{i,b}, \{L_{i,b}^{\text{K}_0}\}_{i,b}, \{L_{i,b}^{\text{K}_1}\}_{i,b}, \{L_{i,b}^{\text{T}'}\}_{i,b}, \{L_{i,b}^{\text{wdata}'}\}_{i,b}.$$

The algorithm sets and outputs

$$\begin{aligned} \widehat{M} &= (\widehat{M}', \{\text{labels for } \text{B} = \beta, \text{K}_0 = k_0, \text{K}_1 = k_1, \text{T}' = 0, \text{wdata}' = \perp\}), \\ L_{i,b} &= L_{i,b}^{\text{W}_0} \| L_{i,b}^{\text{W}_1} \quad \text{for } i \in [\ell_{\text{in}}], b \in \{0, 1\}. \end{aligned}$$

- Eval $^{D_1, \dots, D_\tau}(T_{\text{max}}, M, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}, \widehat{M}, \{L_i\}_i)$  combines the labels for  $\widehat{M}'$  in  $\widehat{M}$  with  $\{L_i\}_i$  to recover  $\{L'_i\}_i$ , and runs and outputs

$$(\text{Eval}')^{D_1, \dots, D_\tau}(T_{\text{max}}, M', \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}, \widehat{M}', \{L'_i\}_i).$$

### Machine $M'$ Transformed from $M$

**Lengths.**  $\ell'_{\text{in}} = \text{poly}(\lambda, |M|, \log T_{\text{max}}), \quad \ell'_{\text{st}} = \text{poly}(\lambda, |M|, \log T_{\text{max}}),$   
 $\ell'_{\text{addr}} = \ell_{\text{addr}}, \quad \ell'_{\text{cell}} = \ell_{\text{cell}},$   
 $\ell'_{\text{ADDR}} = \ell_{\text{ADDR}}, \quad \ell'_{\text{CELL}} = 2\ell_{\text{CELL}} + O(\log T_{\text{max}}).$

**Input ( $w'$ ).**  $B,$  bit indicating the active copy;  
 $W_0, W_1,$  short inputs to the two copies of  $M$ ;  
 $K_0, K_1,$  PPRF keys for the two copies of  $M$ ;  
 $T',$  progress of security proof;  
 $wdata',$  hardwired content of one cell of  $M'$ .

**State ( $st'$ ).**  $t - 1,$  current time step minus one (initially 0);  
 $\tilde{st}_{0,t-1}, \tilde{st}_{1,t-1},$  optional old states of the two copies of  $M$ ;  
 $\tau_{t-1}, i_{t-1},$  location of the last-read cell.

**Steps.** Each step of  $M'$  executes one step in each copy of  $M$ .

1. Translate  $rdata'$  for  $M'$  into  $rdata$  for  $M$ :

**if**  $t = 1$  or  $\tau_{t-1} \in [T]$ :

**parse**  $rdata'_{t-1}$  into  $\widetilde{rdata}_{0,t-1} \| 0^{\ell'_{\text{CELL}} - \ell_{\text{CELL}}} = \widetilde{rdata}_{1,t-1} \| 0^{\ell'_{\text{CELL}} - \ell_{\text{CELL}}}$

**if**  $t \neq 1$  and  $\tau_{t-1} = \text{work}$ :

**compute**  $\widetilde{rdata}_{0,t-1}, \widetilde{rdata}_{1,t-1}$

$$\left\{ \begin{array}{ll} \text{from } rdata'_{t-1} = t' \| (\widetilde{rdata}_{0,t-1} \oplus \text{PPRF.Eval}(k_0, t')) \| (\widetilde{rdata}_{1,t-1} \oplus \text{PPRF.Eval}(k_1, t')), & \text{if } t' \neq 0; \\ \text{as } 0^{\ell_{\text{CELL}}}, 0^{\ell_{\text{CELL}}}, & \text{otherwise;} \end{array} \right.$$

2. Execute one step in the active copy:

$(\text{done}_t, \tilde{st}_{B,t}, \tau_t, i_t, \widetilde{wdata}_{B,t}, \text{out}_t) \leftarrow \text{CPU}(\ell_1, \dots, \ell_T, w_B, \tilde{st}_{B,t-1}, \widetilde{rdata}_{B,t-1})$

3. Optionally execute one step in the inactive copy:

$$(\tilde{st}_{1-B,t}, \widetilde{wdata}_{1-B,t}) \leftarrow \begin{cases} (0^{\ell_{\text{st}}}, 0^{\ell_{\text{CELL}}}), & \text{if } t > T'; \\ \text{from CPU} \left( \ell_1, \dots, \ell_T, w_{1-B}, \tilde{st}_{1-B,t-1}, \widetilde{rdata}_{1-B,t-1} \right), & \text{if } t \leq T'. \end{cases}$$

4. Translate  $wdata$  for  $M$  into  $wdata'$  for  $M'$ :

$$wdata'_t \leftarrow \begin{cases} \perp, & \text{if } \tau_t \in [T]; \\ wdata', & \text{if } t = \text{work} \text{ and } t = T' \text{ and } wdata' \neq \perp; \\ t \| (\widetilde{wdata}_{0,t} \oplus \text{PPRF.Eval}(k_0, t)) \| (\widetilde{wdata}_{1,t} \oplus \text{PPRF.Eval}(k_1, t)), & \text{otherwise;} \end{cases}$$

5. Output and finish:

**output**  $(\text{done}_t, \text{newst}', \tau_t, i_t, wdata'_t, \text{out}_t)$

where  $\text{newst}'$  contains incremented  $t$  and  $\tilde{st}_{0,t}, \tilde{st}_{1,t}, \tau_t, i_t$

**Figure 7.** The machine  $M'$  in Construction 2.



**Correctness and Efficiency.** By construction, during evaluation, the labels selected for  $M'$  correspond to

$$B = \beta, \quad W_0 = w, W_1 = w, \quad K_0 = k_0, K_1 = k_1, \quad T' = 0, \text{wdata}' = \perp.$$

The execution of  $M'$  has the following important invariants:

- $\tilde{\text{st}}_{\beta,t} = \text{st}_t$ , where  $\text{st}_t$  is from  $M^{\dots}(w)$ .
- $D'_{\text{work},t}[i]$  for all  $t, i$  is either all-zero if never touched (never read nor written to) or

$$t' \parallel (D_{0,\text{work},t}[i] \oplus \text{PPRF.Eval}(k_0, t')) \parallel (D_{1,\text{work},t}[i] \oplus \text{PPRF.Eval}(k_1, t')),$$

where  $t' \neq 0$  is the last time when cell  $i$  on the working tape of  $\underline{M}$  was touched,  $D_{\beta,\text{work},t}$  is  $D_{\text{work},t}$  of  $M^{\dots}(w)$ , and  $D_{1-\beta,\text{work},t}$  is all-zero. In all cases,  $\text{rdata}_{\beta,t} = \text{rdata}_t$ , where  $\text{rdata}_t$  is from  $M^{\dots}(w)$ .

Correctness follows from that of LGRAM' and those invariants, together with the other logics in  $M'$ . For efficiency, clearly  $|M'| = \text{poly}(\lambda, |M|, \log T_{\max})$ , and  $M'$  has the same running time as  $M$  when supplied with the input specified above.

**Theorem 16 (¶).** *Suppose in Construction 2, LGRAM' is fixed-memory secure (Definition 3 or 4) and PPRF is secure (Definition 17), then the constructed scheme is fixed-address secure (Definition 3 or 4) and inherits the (un-)boundedness of LGRAM'.*

*Proof* (Theorem 16).  $\text{Exp}_{\text{LGRAM}}^b$  of the constructed scheme corresponds to  $\text{Exp}_{\text{LGRAM}}^b$  of LGRAM' with machine  $M'$  and input

$$B = \beta \stackrel{\$}{\leftarrow} \{0, 1\}, \quad W_0 = w_b, K_0 = k_0, \quad W_1 = w_b, K_1 = k_1, \quad T' = 0, \text{wdata}' = \perp,$$

where  $k_0, k_1$  are random PPRF keys. We will consider hybrids that effectively changes the input to  $M'$  by altering  $\tilde{M}$  and  $\{L_i\}_i$  given to the adversary, so we describe them by specifying the changed input.

- $H_0^b$ . This is just  $\text{Exp}_{\text{LGRAM}}^b$ , i.e.,

$$B = \beta, \quad W_{\beta} = w_b, \quad K_{\beta} = k_{\beta}, \quad W_{1-\beta} = w_b, \quad K_{1-\beta} = k_{1-\beta}, \quad T' = 0, \quad \text{wdata}' = \perp.$$

- $H_1^b$ . In this hybrid, we change  $W_{1-B}$  to  $w_{1-b}$ , i.e.,

$$B = \beta, \quad W_{\beta} = w_b, \quad K_{\beta} = k_{\beta}, \quad W_{1-\beta} = \boxed{w_{1-b}}, \quad K_{1-\beta} = k_{1-\beta}, \quad T' = 0, \quad \text{wdata}' = \perp.$$

The only place  $M'$  uses  $W_{1-B}$  is the branch  $t > T'$  in Step 3 of  $M'$ . But  $T' = 0$  and that branch is never taken. Therefore, for each  $b \in \{0, 1\}$ , the executions of  $M'$  in  $H_0^b$  and  $H_1^b$  satisfy the condition of fixed-memory security of LGRAM', and  $H_0^b \approx H_1^b$ .

- $H_{2,t}^b$  for  $t = 0, \dots, T$ , where  $T$  is the (common) running time of  $M^{\dots}(w_0)$  and  $M^{\dots}(w_1)$ . In this hybrid, we increase  $T'$  to  $t$ , i.e.,

$$B = \beta, \quad W_{\beta} = w_b, \quad K_{\beta} = k_{\beta}, \quad W_{1-\beta} = w_{1-b}, \quad K_{1-\beta} = k_{1-\beta}, \quad T' = \boxed{t}, \quad \text{wdata}' = \perp.$$

Clearly,  $H_1^b \equiv H_{2,0}^b$  for each  $b \in \{0, 1\}$ .

Claim 17 (¶).  $H_{2,t-1}^b \equiv H_{2,t}^b$  for all  $t \in [T]$  and  $b \in \{0, 1\}$ .

- $H_3^b$ . In this hybrid, we set  $T'$  to  $T$ , i.e.,

$$B = \beta, \quad W_\beta = w_b, \quad K_\beta = k_\beta, \quad W_{1-\beta} = w_{1-b}, \quad K_{1-\beta} = k_{1-\beta}, \quad T' = \boxed{T}, \quad \text{wdata}' = \perp.$$

By definition,  $H_{2,T}^b \equiv H_3^b$  for each  $b \in \{0, 1\}$ .

- $H_4^b$ . In this hybrid, we rename  $(\beta \oplus b)$  to  $\gamma$ . It can be described as

$$B = \boxed{\gamma \oplus b}, \quad W_{\gamma} = w_{\boxed{0}}, \quad K_{\gamma} = k_{\boxed{0}}, \quad W_{\boxed{1-\gamma}} = w_{\boxed{1}}, \quad K_{\boxed{1-\gamma}} = k_{\boxed{1}}, \quad T' = T, \quad \text{wdata}' = \perp.$$

Since the change is syntactical,  $H_3^b \equiv H_4^b$  for each  $b \in \{0, 1\}$ .

In  $H_4^0$  and  $H_4^1$ , the machine  $M'$  fully executes both  $M^{\dots}(w_0)$  on the copy  $\gamma$  and  $M^{\dots}(w_1)$  on the copy  $(1 - \gamma)$ . Their only difference is from which execution  $M'$  takes  $(\text{done}_t, \tau_t, i_t, \text{out}_t)$ . The two executions of  $M$  satisfy the constraint of fixed-address security, so their  $(\text{done}_t, \tau_t, i_t, \text{out}_t)$  sequences are identical. Therefore,  $H_4^0 \approx H_4^1$  by the fixed-memory security of LGRAM'.

We conclude that  $\text{Exp}_{\text{LGRAM}}^0 \equiv H_4^0 \approx H_4^1 \equiv \text{Exp}_{\text{LGRAM}}^1$  for the constructed scheme.

It is clear that the time of Eval is instance-specific if so is that of Eval'. Note that the number of hybrids is polynomial in (a polynomial upper bound of) the instance running time  $T$ , not  $T_{\max}$ . Therefore, the proof shows that LGRAM inherits the (un-)boundedness of the security of LGRAM'.  $\square$

*Proof (Claim 17).* We further alter the input to  $M'$  to show  $H_{2,t-1}^b \approx H_{2,t}^b$ .

- $G_0$ . This is just  $H_{2,t-1}^b$ , described by

$$B = \beta, \quad W_\beta = w_b, \quad K_\beta = k_\beta, \quad W_{1-\beta} = w_{1-b}, \quad K_{1-\beta} = k_{1-\beta}, \quad T' = t - 1, \quad \text{wdata}' = \perp.$$

- $G_1$ . In this hybrid, we puncture  $k_{1-\beta}$  at  $t$ , hardwire  $\text{wdata}'_t$ , and increment  $T'$  to activate the hardwiring at the right time step. The hybrid is described by

$$B = \beta, \quad W_\beta = w_b, \quad K_\beta = k_\beta, \quad W_{1-\beta} = w_{1-b}, \quad K_{1-\beta} = \boxed{\overset{\circ}{k}_{1-\beta, \{t\}}}, \quad T' = \boxed{t},$$

$$\text{wdata}' = \begin{cases} \perp, & \text{if } i_t \in [T]; \\ t \parallel (\text{wdata}_{b,t} \oplus \text{PRF}(k_0, t)) \parallel \text{PRF}(k_1, t), & \text{if } i_t = \text{work and } \beta = 0; \\ t \parallel \text{PRF}(k_0, t) \parallel (\text{wdata}_{b,t} \oplus \text{PRF}(k_1, t)), & \text{if } i_t = \text{work and } \beta = 1; \end{cases}$$

where  $\text{wdata}_{b,t}$  is from  $M^{\dots}(w_b)$  and  $\widetilde{\text{wdata}}_{\beta,t} = \text{wdata}_{b,t}$  in the execution of  $M'$ .  $G_0 \approx G_1$  by the fixed-memory security of LGRAM'.

- $G_2$ . In this hybrid, we make the portion of  $\text{wdata}'$  for the copy  $(1 - \beta)$  random, i.e.,

$$B = \beta, \quad W_\beta = w_b, \quad K_\beta = k_\beta, \quad W_{1-\beta} = w_{1-b}, \quad K_{1-\beta} = \overset{\circ}{k}_{1-\beta, \{t\}}, \quad T' = t,$$

$$\text{wdata}' = \begin{cases} \perp, & \text{if } i_t \in [T]; \\ t \parallel (\text{wdata}_{b,t} \oplus \text{PRF}(k_0, t)) \parallel \boxed{\text{random}}, & \text{if } i_t = \text{work and } \beta = 0; \\ t \parallel \boxed{\text{random}} \parallel (\text{wdata}_{b,t} \oplus \text{PRF}(k_1, t)), & \text{if } i_t = \text{work and } \beta = 1. \end{cases}$$

Note that the altered portion is  $\text{PPRF.Eval}(k_{1-\beta}, t)$  in  $G_1$ . By the security of PPRF, it follows that  $G_1 \approx G_2$ .

- $G_3$ . In this hybrid, we encode  $wdata_{1-b,t}$  in the portion of  $wdata'$  for the copy  $(1 - \beta)$ . It is described by

$$B = \beta, \quad W_\beta = w_b, \quad K_\beta = k_\beta, \quad W_{1-\beta} = w_{1-b}, \quad K_{1-\beta} = \overset{\circ}{k}_{1-\beta,\{t\}}, \quad T' = t,$$

$$wdata' = \begin{cases} \perp, & \text{if } i_t \in [\mathcal{T}]; \\ t \parallel (wdata_{b,t} \oplus \text{PRF}(k_0, t)) \parallel (wdata_{1-b,t} \oplus \text{PRF}(k_1, t)), & \text{if } i_t = \text{work and } \beta = 0; \\ t \parallel (wdata_{1-b,t} \oplus \text{PRF}(k_0, t)) \parallel (wdata_{b,t} \oplus \text{PRF}(k_1, t)), & \text{if } i_t = \text{work and } \beta = 1. \end{cases}$$

$G_2 \approx G_3$  by the security of PPRF.

- $G_4$ . In this hybrid, we undo puncturing and hardwiring, i.e.,

$$B = \beta, \quad W_\beta = w_b, \quad K_\beta = k_\beta, \quad W_{1-\beta} = w_{1-b}, \quad K_{1-\beta} = k_{1-\beta}, \quad T' = t,$$

$$wdata' = \perp.$$

$G_3 \approx G_4$  by the fixed-memory security of LGRAM', analogously to  $G_0 \approx G_1$ .

By inspection,  $G_4$  is just  $H_{2,t}^b$ . Therefore,  $H_{2,t-1}^b \equiv G_0 \approx G_4 \equiv H_{2,t}^b$ .  $\square$

## 5.2 Fixed-Address to Full Security

Given an LGRAM with fixed-address security, we construct one with full security with a transformation due to [CH16] using ORAM.

There are two technical differences of our LGRAM from existing notions of garbled RAM that are relevant to this transformation. First, the input tapes are compressed, and their digests cannot be an ORAM'ed version of their contents, because that would make the digests as long as the tape contents and our digests must be reusable while the usual ORAM is not. Second, the garbling procedure is only allowed to run in time  $\text{poly}(\lambda, |M|, \log T_{\max})$ , which is not sufficient to initialize an ORAM for the working tape.

The first issue is resolved by copying the input tapes to the working tape before the actual computation is performed, and reading from the copy on the working tape whenever the actual computation wants to read from an input tape. During the copying stage, the access pattern to the tapes is simply sequential reading and writing, independent of the computation. The second issue is resolved by delaying ORAM initialization to evaluation time. Recall that our notion of ORAM (Definition 19) starts the physical execution with an empty physical tape, implicitly requiring on-demand initialization. This is matched by our notion of LGRAM (Definition 2), where the initial working tape is all-zero.

Similar to the transformation from fixed-memory security to fixed-address security (Section 5.1), we run two copies of the underlying machine in parallel to facilitate the security proof.

**Ingredients of Construction 3.** Let

- LGRAM' = (Compress', Garble', Eval') be an LGRAM with fixed-address security,
- MakeORAM an ORAM with localized randomness with (PartRnd, SimORAM), and
- PPRF = (PPRF.Puncture, PPRF.Eval) a puncturable PRF.

### Machine $M'$ Transformed from $M$

<b>Lengths.</b>	$\ell'_{\text{in}} = \text{poly}(\lambda,  M , \log T_{\text{max}}),$	$\ell'_{\text{st}} = \text{poly}(\lambda,  M , \log T_{\text{max}}),$
	$\ell'_{\text{addr}} = \ell_{\text{addr}},$	$\ell'_{\text{cell}} = \ell_{\text{cell}},$
	$\ell'_{\text{ADDR}} = \lceil \log_2(2S'_{\text{max}}) \rceil,$	$\ell'_{\text{CELL}}$ determined by MakeORAM.
<b>Input (<math>w'</math>).</b>	<p><math>B, W_0, W_1, K_0, K_1, T',</math> see Figure 7;</p> <p><math>R',</math> hardwired ORAM randomness if <math>K_B</math> is punctured;</p> <p><math>K',</math> PPRF key for ORAM simulation;</p> <p><math>\text{addrS}',</math> <math>T_0</math> hardwired addresses of <math>M'</math>.</p>	
<b>State (<math>st'</math>).</b>	<p><math>e,</math> flag indicating whether an error is detected;</p> <p><math>t - 1,</math> current time period minus one (initially 0);</p> <p><math>t_0 - 1,</math> current time step minus one (initially 0) modulo the period <math>(2T_0 + 4);</math></p> <p><math>\widetilde{\text{st}}_{0,t-1}, \widetilde{\text{rdata}}_{0,t-1}, \text{ost}_{0,t-1,t_0-1}, \widetilde{\text{st}}_{1,t-1}, \widetilde{\text{rdata}}_{1,t-1}, \text{ost}_{1,t-1,t_0-1},</math> for the two copies of <math>M</math> and their ORAM;</p> <p><math>\text{done}'_t, \text{out}'_t,</math> output of <math>M</math> from the copy <math>B</math> of <math>M</math>.</p>	
<b>Steps.</b>	<p>Each step of the two copies of <math>M</math> is usually simulated by <math>(2T_0 + 4)</math> steps of <math>M'</math>.</p> <ol style="list-style-type: none"> <li>1. Decide what to do in this period: <ul style="list-style-type: none"> <li><b>if</b> <math>\widetilde{\ell}_{&lt;\tau} = \sum_{\pi &lt; \tau}  D_\pi  &lt; t \leq \sum_{\pi \leq \tau}  D_\pi </math> for some <math>\tau \in [\mathcal{T}]</math>: <ul style="list-style-type: none"> <li><b>copy</b> <math>D_\tau[t - \widetilde{\ell}_{&lt;\tau}]</math> to ORAM for both copies of <math>M</math></li> </ul> </li> <li><b>else:</b> <b>execute</b> or <b>simulate</b> one step of <math>M</math></li> </ul> </li> <li>2a. Copy one cell from input tape to ORAM: <ul style="list-style-type: none"> <li>1 step to <b>read</b> the cell</li> <li><math>(T_0 + T_0)</math> steps to ORAM-<b>write</b> for the copy 0 and 1</li> <li>1 step to <b>report</b> potential error, 2 step unused</li> </ul> </li> <li>2b. Execute or simulate one step of <math>M</math>: <ul style="list-style-type: none"> <li><math>2(1 + T_0)</math> step for one step of <math>M</math> then ORAM-R/W for the copy 0 and/or 1</li> <li>the copy <math>(1 - B)</math> is <math>\begin{cases} \text{simulated using } K', &amp; \text{if } t - \widetilde{\ell}_{&lt;\mathcal{T}+1} &gt; T'; \\ \text{hardwired in } \text{addrS}', &amp; \text{if } t - \widetilde{\ell}_{&lt;\mathcal{T}+1} = T'; \end{cases}</math></li> <li>1 step to <b>report</b> potential error, 1 step to <b>output</b> <math>\text{done}'_t, \text{out}'_t</math></li> </ul> </li> <li>3. Error handling and other details: <ul style="list-style-type: none"> <li>ORAM randomness for non-simulated steps is from <math>K_0, K_1, R'</math></li> <li>addresses of ORAM for the copy 0 (resp. 1) are mapped to those of <math>M'</math> by <math>i \mapsto 2i - 1</math> (resp. <math>i \mapsto 2i</math>; i.e., interleaved)</li> <li><b>if</b> an error is detected by ORAM <ul style="list-style-type: none"> <li><b>idle</b> and <b>report</b> it at the end of this period</li> <li>use <math>\ell'_{\text{in}}</math> extra steps to <b>output</b> <math>W_B</math> and <b>halt</b></li> </ul> </li> </ul> </li> </ol>	

**Figure 8.** The machine  $M'$  in Construction 3.

**Construction 3.** Our scheme works as follows:

- $\text{Compress}(1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, \tau, D_\tau)$  runs

$$\text{digest}'_\tau \stackrel{\S}{\leftarrow} \text{Compress}'(1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, \tau, D_\tau)$$

and outputs  $\text{digest}_\tau = (|D_\tau|, \text{digest}'_\tau)$ .

- $\text{Garble}(T_{\max}, M, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]})$  runs

$$(S'_{\max}, 1^{\ell_{\text{CELL}}}, 1^{\ell_r}, 1^{\ell_{\text{ost}}}, \{\text{ORW}_{t_0}\}_{t_0 \in [T_0]}) \leftarrow \text{MakeORAM} \left( 1^{\ell_{\text{CELL}}}, 1^{\ell_{\text{ADDR}}}, \sum_{\tau \in [\mathcal{T}]} |D_\tau| + T_{\max} \right),$$

prepares  $M'$  as shown in Figure 8. It sets

$$T'_{\max} = \max \left\{ 2S'_{\max}, \ell_{\text{in}} + (2T_0 + 4) \left( T_{\max} + \sum_{\tau \in [\mathcal{T}]} |D_\tau| \right) \right\},$$

runs

$$(\widehat{M}', \{L'_{i,b}\}_{i,b}) \stackrel{\S}{\leftarrow} \text{Garble}'(T'_{\max}, M', \{\text{digest}'_\tau\}_{\tau \in [\mathcal{T}]})$$

splits the labels according to the part of input to  $M'$  they select into

$$\{L_b^B\}_b, \{L_{i,b}^{W_0}\}_{i,b}, \{L_{i,b}^{W_1}\}_{i,b}, \{L_{i,b}^{K_0}\}_{i,b}, \{L_{i,b}^{K_1}\}_{i,b}, \{L_{i,b}^{T'}\}_{i,b}, \{L_{i,b}^{R'}\}_{i,b}, \{L_{i,b}^{K'}\}_{i,b}, \{L_{i,b}^{\text{addr}S'}\}_{i,b},$$

samples  $\beta \stackrel{\S}{\leftarrow} \{0, 1\}$  and three random PPRF keys  $k_0, k_1, k'$ , and sets and outputs

$$\begin{aligned} \widehat{M} &= (\widehat{M}', \{\text{labels for } B = \beta, K_0 = k_0, K_1 = k_1, T' = 0, R' = \perp, K' = k', \text{addr}S' = \perp\}), \\ L_{i,b} &= L_{i,b}^{W_0} \| L_{i,b}^{W_1} \quad \text{for } i \in [\ell_{\text{in}}], b \in \{0, 1\}. \end{aligned}$$

- $\text{Eval}^{D_1, \dots, D_{\mathcal{T}}}(T_{\max}, M, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}, \widehat{M}, \{L_i\}_i)$  combines the labels for  $\widehat{M}'$  in  $\widehat{M}$  with  $\{L_i\}_i$  to recover  $\{L'_i\}_i$ , and runs

$$(\text{Eval}')^{D_1, \dots, D_{\mathcal{T}}}(T_{\max}, M', \{\text{digest}'_\tau\}_{\tau \in [\mathcal{T}]}, \widehat{M}', \{L'_i\}_i).$$

If the evaluation does not report any ORAM error, the algorithm removes the empty, placeholder output (all steps in the periods of copying the input tapes, and  $(2T_0 + 3)$  steps in the periods of executing  $M$ ) from  $\text{Eval}'$ , and use the stripped version as *the* output. Otherwise, the algorithm will obtain  $w$  after the error is reported, and it evaluates  $M^{D_1, \dots, D_{\mathcal{T}}}(w)$  in the clear and outputs it.

**Correctness and Efficiency.** Correctness and efficiency follow from those of the underlying ingredients as well as the error reporting mechanism of ORAM used by  $M'$ .

**Theorem 18** (¶). *Suppose in Construction 3, LGRAM' is fixed-address secure (Definition 3 or 4), MakeORAM has localized randomness with PartRnd, SimORAM (Definition 20), and PPRF is secure (Definition 17), then the constructed scheme is (fully) secure (Definition 3 or 4) and inherits the (un-)boundedness of LGRAM'.*

*Proof* (Theorem 18). The proof is analogous to that of Theorem 16. We only demonstrate the key differences:

- When the randomness (e.g., after puncturing a PPRF key) is altered,  $M'$  could halt early due to error arising from the unfortunate randomness. We must argue that such case only happens with negligible probability. ORAM guarantees  $2^{-\lambda}$  error (overflow) probability over the entire execution if the randomness is truly random, and by PRF security, the error probability is negligible when the randomness is pseudorandom (or a mix of pseudorandomness and true randomness, as needed in certain hybrids).
- The proof corresponding to Claim 17 is slightly more involved. Let  $R_t$  be the set produced by PartRnd, which is the index set of the randomness precisely affecting the access pattern for logical step  $t$ . The core idea is to puncture  $k_{1-B}$  at  $R_t$  and argue it can be simulated using SimORAM due to the localized randomness property. However,  $\text{PPRF.Eval}(k_{1-B}, i)$  for  $i \in R_t$  is written into the working tape of  $M'$  at various time steps and locations, which are difficult to track. This is where the fixed-address property helps. It can be used to “decouple” or “couple” what is written during previous steps and where is read for logical step  $t$ . Following [CH19; Section 6.3], the transition sequence is as follows:
  1. (Fixed-address security.) Puncture  $k_{1-B}$  at  $R_t$ , puncture  $k'$  at  $t$ , hardwire  $\{\text{PPRF.Eval}(k_{1-B}, i)\}_{i \in R_t}$  into  $R'$ , and hardwire  $\text{SimORAM}(\dots, \text{PPRF.Eval}(k', t))$  into  $\text{addrS}'$ .
  2. (PPRF security.) Change  $\text{PPRF.Eval}(k', t)$  in  $\text{addrS}'$  into true randomness.
  3. (Fixed-address security.) Change  $\{\text{PPRF.Eval}(k_{1-B}, i)\}_{i \in R_t}$  in  $R'$  into the same true randomness used by  $\text{addrS}'$ . This step “couples” the previously written content and the later physical addresses.
  4. (PPRF security.) Change the common true randomness in both  $R$  and  $\text{addrS}'$  into  $\{\text{PPRF.Eval}(k_{1-B}, i)\}_{i \in R_t}$ .
  5. (Fixed-address security.) Undo puncturing and hardwiring. □

### 5.3 Bounded to Unbounded

So far we only obtain bounded LGRAM from Constructions 1, 2, and 3. While the actual evaluation time of the resultant scheme is instance-dependent (scaling with  $T$  instead of  $T_{\max}$ ) and it works even if the machine access memory at large addresses (both are syntactical restrictions for bounded LGRAM), its security proof requires  $T_{\max}$  to be polynomially bounded. To get rid of this restriction, we apply the transformation due to [AL18]. The idea is to generate a series of LGRAM garblings with  $T'_{\max, e} = 2^e$  for  $e \in [\log T_{\max}]$ , each encrypted under a different key. The  $e^{\text{th}}$  garbling simulates the execution, but if the time exceeds  $2^e$ , it outputs the secret key that decrypts the next garbling, and the evaluation can be retried. This avoids the exponential security loss because we can apply the LGRAM security to the garblings with  $T'_{\max} < 2T$  and argue that the other, larger garblings are hidden by encryption, removing the need to apply the LGRAM security on those large garblings.

<b>Machine <math>M'</math> Transformed from <math>M</math> and <math>T_{\max}</math></b>	
<b>Lengths.</b>	$\ell'_{\text{in}} = \ell_{\text{in}} + \text{poly}(\lambda),$ $\ell'_{\text{st}} = \ell_{\text{st}} + O(\log T_{\max}) + \text{poly}(\lambda,  M ),$ $\ell'_{\text{addr}} = \ell_{\text{addr}},$ $\ell'_{\text{cell}} = \ell_{\text{cell}},$ $\ell'_{\text{ADDR}} = \text{poly}(\lambda,  M , \log T_{\max}),$ $\ell'_{\text{CELL}} = \text{poly}(\lambda,  M , \log T_{\max}).$
<b>Input (<math>w'</math>).</b>	$W,$ short input to $M$ ; $E,$ attempted time limit of $M$ ; $K,$ secret key of SKE.
<b>State (<math>\text{st}'</math>).</b>	$t - 1,$ current time period minus one (initially 0); $t_0 - 1,$ current time step minus one (initially 0) modulo the period; $\text{st},$ state of $M$ .
<b>Steps.</b>	Each step of $M$ is simulated by one period of $M'$ .
	1. When $t \leq 2^E$ : 1 step for one step of $M$ $\text{poly}(\lambda,  M , \log T_{\max})$ steps for the memory operation of $M$ (using a deterministic balanced binary tree of capacity $2^E$ to implement a virtual sparse memory for $M$ ) <b>halt</b> if $M$ halts 2. When $2^E < t < 2^E + \text{poly}(\lambda)$ : <b>report</b> <i>runningtimeexceeded</i> <b>output</b> $K$

**Figure 9.** The machine  $M'$  in Construction 4.

**Ingredients of Construction 4.** Let

- LGRAM' = (Compress', Garble', Eval') be a bounded LGRAM, and
- SKE = (SKE.Gen, SKE.Enc, SKE.Dec) a secret-key encryption scheme.

**Construction 4.** Our scheme works as follows:

- $\text{Compress}(1^{\ell_{\text{addr}}}, 1^{\ell_{\text{cell}}}, \tau, D_\tau)$  is the same as  $\text{Compress}'$ .
- $\text{Garble}(T_{\max}, M, \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]})$  prepares  $M'$  from  $M$  and  $T_{\max}$  as shown in Figure 9. For  $e \in [\lceil \log_2 T_{\max} \rceil]$ , it samples SKE key  $k_e$  and runs

$$(\widehat{M}'_e, \{L'_{e,i,b}\}_{i,b}) \stackrel{\$}{\leftarrow} \text{Garble}'(T'_{\max,e}, M', \{\text{digest}_\tau\}_{\tau \in [\mathcal{T}]}) \quad \text{for } e \in [\lceil \log_2 T_{\max} \rceil],$$

where  $T'_{\max,e} = 2^e + \text{poly}(\lambda, |M|, \log T_{\max})$ . The algorithm encrypts  $\widehat{M}'_e$  and  $L'_{e,i,b}$  by

$$\widetilde{M}'_e \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_e, \widehat{M}'_e), \quad \widetilde{L}'_{e,i,b} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_e, L'_{e,i,b}),$$

and splits the encrypted labels into  $\{\widetilde{L}_{e,i,b}^W\}_{e,i,b}$ ,  $\{\widetilde{L}_{e,i,b}^E\}_{e,i,b}$ ,  $\{\widetilde{L}_{e,i,b}^K\}_{e,i,b}$  by the portion of input to  $M'$  they correspond to. Defining  $k_{\lceil \log_2 T_{\max} \rceil + 1} = \perp$ , it sets and outputs

$$\widehat{M} = \left( k_1, \{\widetilde{M}'_e, \widetilde{L}_{e,i,b}^E \text{'s and } \widetilde{L}_{e,i,b}^K \text{'s for } E = e, K = k_{e+1}\}_{e \in [\lceil \log_2 T_{\max} \rceil]} \right),$$

$$L_{i,b} = \widetilde{L}_{1,i,b}^W \parallel \cdots \parallel \widetilde{L}_{\lceil \log_2 T_{\max} \rceil, i, b}^W.$$

- $\text{Eval}^{D_1, \dots, D_T}(T_{\max}, M, \{\text{digest}_\tau\}_{\tau \in [T]}, \widehat{M}, \{L_i\}_i)$  parses  $\widehat{M}$  and  $L_i$ 's as defined by Garble. For  $e = 1, \dots, \lceil \log_2 T_{\max} \rceil$ , it retrieves  $k_e$ , either from  $\widehat{M}$  if  $e = 1$ , or from the previous call to  $\text{Eval}'$  if  $e > 1$ , and runs

$$\begin{aligned} & (\text{Eval}')^{D_1, \dots, D_T}(T_{\max}, M, \{\text{digest}_\tau\}_{\tau \in [T]}, \text{SKE.Dec}(k_e \widetilde{M}'_e), \{\text{SKE.Dec}(k_e, \widetilde{L}_{e,i})\}_i) \\ & \rightarrow \begin{cases} \text{outs}, & \text{if running time is not exceeded;} \\ (\text{prefix of outs}, k_{e+1}), & \text{if running time is exceeded.} \end{cases} \end{aligned}$$

The algorithm removes the empty, placeholder output from  $\text{outs}$ , use the stripped version as the output, and halt, if the running time is not exceeded. Otherwise, it attempts the next  $e$ .

**Correctness and Efficiency.** Those are clear by inspection.

**Theorem 19** (♣). *Suppose in Construction 4, LGRAM' is a (fixed-memory, fixed-address, fully) secure bounded LGRAM (Definition 4) and SKE is secure (Definition 18), then the constructed scheme is an unbounded LGRAM with the same level of security as LGRAM' (Definition 3).*

*Proof* (Theorem 19). Since  $M'$  uses a deterministic balanced binary tree to implement the virtual sparse memory for  $M$ , it preserves the constraints of fixed-memory, fixed-address, or full security. Let  $T$  be the instance running time and  $\bar{e} = \lceil \log_2 T \rceil$ . We consider the following hybrids:

- $H_{-1}^b$ . This is just  $\text{Exp}_{\text{LGRAM}}^b$  for the constructed scheme, where

$$\begin{aligned} \widehat{M} & : k_1, \{\widetilde{M}'_e, \widetilde{L}_{e,i,b}^E \text{'s and } \widetilde{L}_{e,i,b}^K \text{'s for } E = e, K = k_{e+1}\}_{e \in [\lceil \log_2 T_{\max} \rceil]}, \\ L_i & : \{\widetilde{L}_{e,i,b}^W \text{'s for } W = w_b\}_{e \in [\lceil \log_2 T_{\max} \rceil]}. \end{aligned}$$

- $H_e^b$  for  $e = 0, \dots, \lceil \log_2 T_{\max} \rceil$ . In this hybrid, we remove  $k_{\bar{e}+1}$  as well as all the unused garblings up to time bound  $2^e$  as follows.

$$\widehat{M} : k_1, \begin{cases} \widetilde{M}'_e, \widetilde{L}_{e,i,b}^E \text{'s and } \widetilde{L}_{e,i,b}^K \text{'s for } E = e, K = k_{e+1}, & \text{if } e < \bar{e}; \\ \widetilde{M}'_{\bar{e}}, \widetilde{L}_{\bar{e},i,b}^E \text{'s and } \widetilde{L}_{\bar{e},i,b}^K \text{'s for } E = \bar{e}, K = \perp, & \text{if } e = \bar{e}; \\ \text{random}, & \text{if } e > \bar{e} \text{ and } e \leq e; \\ \widetilde{M}'_e, \widetilde{L}_{e,i,b}^E \text{'s and } \widetilde{L}_{e,i,b}^K \text{'s for } E = e, K = k_{e+1}, & \text{if } e > \bar{e} \text{ and } e > e; \end{cases}$$

$$L_i : \begin{cases} \widetilde{L}_{e,i,b}^W \text{'s for } W = w_b, & \text{if } e \leq \bar{e}; \\ \text{random}, & \text{if } e > \bar{e} \text{ and } e \leq e; \\ \widetilde{L}_{e,i,b}^W \text{'s for } W = w_b, & \text{if } e > \bar{e} \text{ and } e > e. \end{cases}$$

To see indistinguishability:



- $H_{-1}^b \approx H_0^b$ . Compared to  $H_{-1}^b$ , in  $H_0^b$ , the garbling  $\widehat{M}'_e$  has  $K$  in its input changed from  $k_{\bar{e}+1}$  to  $\perp$ . Since  $2^{\bar{e}} \geq T$ , the execution of  $M$  in that garbling does not output  $K$  and this change satisfies the constraint of LGRAM security. Moreover,  $T'_{\max, \bar{e}} = 2^{\bar{e}} \leq 2T$  is polynomially bounded. Therefore,  $H_{-1}^b \approx H_0^b$  by the bounded security of LGRAM'.
- $H_{\epsilon-1}^b \approx H_\epsilon^b$ . The two hybrids are different only when  $\epsilon > \bar{e}$ , in which case the ciphertexts under  $k_\epsilon$  changes into random strings and  $k_\epsilon$  is not used in the two hybrids. Therefore,  $H_{\epsilon-1}^b \approx H_\epsilon^b$  by the ciphertext pseudorandomness of SKE.
- $H_{\lfloor \log_2 T_{\max} \rfloor}^0 \approx H_{\lfloor \log_2 T_{\max} \rfloor}^1$ . The difference is that the  $\bar{e}$  non-erased garblings have  $W$  in their inputs changed between  $w_0$  and  $w_1$ . By how  $M'$  works, this change satisfies the constraint of LGRAM security. Moreover, all of them have  $T'_{\max, e} \leq 2T$  polynomially bounded. Therefore,  $H_{\lfloor \log_2 T_{\max} \rfloor}^0 \approx H_{\lfloor \log_2 T_{\max} \rfloor}^1$  follows from a standard hybrid argument by the bounded security of LGRAM'.

We conclude that  $\text{Exp}_{\text{LGRAM}}^0 \equiv H_{-1}^0 \approx H_{-1}^1 \equiv \text{Exp}_{\text{LGRAM}}^1$ .  $\square$

## 6 PHFE for RAM

### 6.1 Bounded Private Input

In this section, we build a PHFE for RAM with bounded private input that is  $f$ - and  $x$ -succinct and has linear-time KeyGen and Enc and whose Dec runs in time  $O(T + |f| + |x|)$ , ignoring polynomial factors in the security parameter. (See Definitions 5, 7, and 9.)

Recall that in PHFE for RAM with bounded private input, the functionality (resp. key, ciphertext) is associated with some RAM  $M$  and (up to exponential) time bound  $T_{\max}$  (resp.  $f$  of arbitrary length,  $x$  of arbitrary length and  $y$  of some fixed polynomial length) and decryption yields  $M^{f,x}(y)$  if the execution halts in time at most  $T_{\max}$ .

**Ingredients of Construction 5.** Let

- $\text{ckt} = (\text{ckt.Setup}, \text{ckt.KeyGen}, \text{ckt.Enc}, \text{ckt.Dec})$  be an FE scheme for circuits,
- $\text{LGRAM} = (\text{LGRAM.Compress}, \text{LGRAM.Garble}, \text{LGRAM.Eval})$  a 2-tape LGRAM scheme,
- $\text{PPRF} = (\text{PPRF.Puncture}, \text{PPRF.Eval})$  a puncturable PRF, and
- $\text{SKE} = (\text{SKE.Gen}, \text{SKE.Enc}, \text{SKE.Dec})$  a secret-key encryption scheme.

**Construction 5** (PHFE for RAM with bounded private input). It works as follows:

- $\text{Setup}(M, T_{\max})$  runs

$$(\text{ckt.mpk}, \text{ckt.msk}) \stackrel{\$}{\leftarrow} \text{ckt.Setup}(1^{\dots}, 1^{\dots}),$$

and outputs  $(\text{mpk}, \text{msk}) = ((T_{\max}, M, \text{ckt.mpk}), \text{ckt.msk})$ , where the input/circuit sizes given to  $\text{ckt.Setup}$  are appropriately chosen (see below).

- $\text{KeyGen}(\text{msk}, f)$  samples random strings  $s_{\text{PPRF}}, \text{id}_{\text{sk}}', \text{lgram}'_{\text{sk}}$ , runs

$$\text{digest}_f \stackrel{\$}{\leftarrow} \text{LGRAM.Compress}(1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, 1, f),$$

$$\text{ckt.sk} \stackrel{\$}{\leftarrow} \text{ckt.KeyGen}(\text{ckt.msk}, \text{GenLGRAM}[\text{digest}_f, s_{\text{PPRF}}, \text{id}_{\text{sk}}', \text{lgram}'_{\text{sk}}]),$$

and outputs  $sk = (\text{digest}_f, s_{\text{PPRF}}, \text{idx}'_{sk}, \text{lgram}'_{sk}, \text{ckt.sk})$ . where GenLGRAM is in Figure 10.

- $\text{Enc}(\text{mpk}, x, y)$  samples random string  $k_{\text{PPRF}}$ , runs

$$\begin{aligned} \text{digest}_x &\stackrel{\$}{\leftarrow} \text{LGRAM.Compress}(1^{\ell_{\text{cell}}}, 1^{\ell_{\text{addr}}}, 2, x), & y', \text{idx}_{\text{ct}}, \text{idx}_{\text{H}}, k'_{\text{idx}}, k'_{\text{lgram}}, \text{lgram}_{\text{ct}} \\ \text{ckt.ct} &\stackrel{\$}{\leftarrow} \text{ckt.Enc}(\text{ckt.mpk}, \perp, (\text{digest}_x, y, k_{\text{PPRF}}, \text{normal}, \underbrace{\perp, \perp, \perp, \perp, \perp, \perp}), \end{aligned}$$

and outputs  $\text{ct} = (\text{digest}_x, \text{ckt.ct})$ .

- $\text{Dec}^{f,x,sk,ct}(\text{mpk})$  parses  $\text{mpk}, sk, \text{ct}$  as defined earlier, runs

$$(\widehat{M}, \{L_i\}_i) \stackrel{\$}{\leftarrow} \text{ckt.Dec}^{\text{GenLGRAM}[\text{digest}_f, s_{\text{PPRF}}, \text{idx}'_{sk}, \text{lgram}'_{sk}], \perp, \text{ckt.sk}, \text{ckt.ct}}(\text{ckt.mpk}),$$

and outputs  $\text{LGRAM.Eval}^{f,x}(T_{\text{max}}, M, \text{digest}_f, \text{digest}_x, \widehat{M}, \{L_i\}_i)$ .

**Correctness and Efficiency.** To ensure correctness, it suffices to set the input/circuit sizes of the underlying FE for circuits to be  $\text{poly}(\lambda, M, \log T_{\text{max}})$  — the normal branch of GenLGRAM is just LGRAM.Eval, which runs in time  $\text{poly}(\lambda, M, \log T_{\text{max}})$ . This (order of) value is also sufficient for the security proof to go through. By the efficiency guarantees of its ingredients, Construction 5 is  $f$ - and  $x$ -succinct, has linear-time KeyGen and Enc, and its Dec runs in time  $(T + |f| + |x|) \text{poly}(\lambda, M, T_{\text{max}})$ .

**Theorem 20 (¶).** *Suppose in Construction 5, ckt, LGRAM, PPRF are secure (Definitions 7, 3, and 17) and SKE has pseudorandom ciphertexts (Definition 18), then the constructed scheme is secure (Definition 7).*

We prove security by hybridizing over the keys. We denote by

$$\text{digest}_{f_q, s_{\text{PPRF},q}}, \text{idx}'_{sk,q}, \text{lgram}'_{sk,q}$$

the content hardwired into GenLGRAM for the  $q^{\text{th}}$  secret key  $sk_q$ . In the hybrids,

$$y', \text{idx}_{\text{ct}}, \text{idx}_{\text{H}}, \text{lgram}_{\text{ct}}, k'_{\text{idx}}, k'_{\text{lgram}}$$

in the challenge ciphertext  $\text{ct}$  are used, and its decryption behavior is controlled by mode and those values. The strategies of handling pre- and post-challenge keys are different. At a high level, they work as follows:

- $\text{idx}'_{sk,q}$  is an encryption of  $q$  so that each key “knows” its ordinal number.
- $\text{idx}_{\text{ct}}$  is the ordinal number of the challenge ciphertext.
- $\text{idx}_{\text{H}}$  indicates the progress of the proof and increases from 0 to  $Q$ :
  - initially,  $q > \text{idx}_{\text{H}}$ , decrypting  $\text{ct}$  by  $sk_q$  yields an LGRAM garbling for  $y_0$ ;
  - when transitioning,  $q = \text{idx}_{\text{H}}$ , the behavior depends on mode and the relative timing of  $\text{ct}$  and  $sk_q$ ,
    - \* when mode = hardwire, decryption takes the hardwired LGRAM garbling from either  $\text{lgram}'_{sk,q}$  (in  $sk_q$ ) or  $\text{lgram}_{\text{ct}}$  (in  $\text{ct}$ ), whichever is generated later the security game (decided by comparing  $\text{idx}_{\text{ct}}$  and  $q$ );

$$\text{GenLGRAM} \left[ \begin{array}{l} \text{digest}_f, s_{\text{PPRF}}, \\ \text{idx}'_{\text{sk}}, \text{lgram}'_{\text{sk}} \end{array} \right] \left( \begin{array}{l} \text{digest}_x, y, k_{\text{PPRF}}, \text{mode}, \\ y', \text{idx}_{\text{ct}}, \text{idx}_{\text{H}}, k'_{\text{idx}}, k'_{\text{lgram}}, \text{lgram}_{\text{ct}} \end{array} \right)$$

<b>Hardwired.</b>	$\text{digest}_f,$ $s_{\text{PPRF}},$ $\text{idx}'_{\text{sk}},$ $\text{lgram}'_{\text{sk}},$	LGRAM digest of $f$ (tape 1); PPRF input for LGRAM randomness; encrypted ordinal ( $q^{\text{th}}$ query) of this key; encrypted hardwired LGRAM garbling.
<b>Input.</b>	$\text{digest}_x, y,$ $k_{\text{PPRF}},$ $\text{mode},$ $y',$ $\text{idx}_{\text{ct}},$  $\text{idx}_{\text{H}},$  $k'_{\text{idx}}, k'_{\text{lgram}},$ $\text{lgram}_{\text{ct}},$	LGRAM digest of $x$ (tape 2), short input; PPRF key for LGRAM randomness; mode of ciphertext and security proof; alternative short input; ordinal of the challenge ciphertext (number of pre-challenge keys); progress of security proof (number of keys using $y'$ ); SKE keys to decrypt $\text{idx}'_{\text{sk}}, \text{lgram}'_{\text{sk}};$ hardwired LGRAM garbling.
<b>Output.</b>	$\text{lgram},$	LGRAM garbling, computed as follows.

1. If mode = normal (garbling is for  $y$ ):

$$(\widehat{M}, \{L_{i,b}\}_{i,b}) \leftarrow \text{LGRAM.Garble} \left( \begin{array}{l} T_{\text{max}}, M, \text{digest}_f, \text{digest}_x; \\ \text{PPRF.Eval}(k_{\text{PPRF}}, s_{\text{PPRF}}) \end{array} \right)$$

$$\text{output lgram} = (\widehat{M}, \{L_{i,y[i]}\}_i)$$

2. Otherwise, this key is decrypting the challenge ciphertext:

$$\text{idx}_{\text{sk}} \leftarrow \text{SKE.Dec}(k'_{\text{idx}}, \text{idx}'_{\text{sk}})$$

3a. If mode = hybrid or  $\text{idx}_{\text{sk}} \neq \text{idx}_{\text{H}}$  (garbling is for  $y$  or  $y'$ ):

$$(\widehat{M}, \{L_{i,b}\}_{i,b}) \leftarrow \text{LGRAM.Garble} \left( \begin{array}{l} T_{\text{max}}, M, \text{digest}_f, \text{digest}_x; \\ \text{PPRF.Eval}(k_{\text{PPRF}}, s_{\text{PPRF}}) \end{array} \right)$$

$$\text{output lgram} = \begin{cases} (\widehat{M}, \{L_{i,y[i]}\}_i), & \text{if } \text{idx}_{\text{H}} < \text{idx}_{\text{sk}}; \\ (\widehat{M}, \{L_{i,y'[i]}\}_i), & \text{if } \text{idx}_{\text{H}} \geq \text{idx}_{\text{sk}}. \end{cases}$$

3b. If mode = hardwire and  $\text{idx}_{\text{sk}} = \text{idx}_{\text{H}}$  (garbling is hardwired):

$$\text{output lgram} = \begin{cases} \text{lgram}_{\text{ct}}, & \text{if } \text{idx}_{\text{ct}} \geq \text{idx}_{\text{sk}}; \\ \text{SKE.Dec}(k'_{\text{lgram}}, \text{lgram}'_{\text{sk}}), & \text{if } \text{idx}_{\text{ct}} < \text{idx}_{\text{sk}}. \end{cases}$$

**Figure 10.** The circuit GenLGRAM in Construction 5.

- \* when mode = hybrid, the behavior is the same as when  $q < \text{idx}_H$ ;
- eventually,  $q < \text{idx}_H$ , decrypting ct by  $\text{sk}_q$  yields an LGRAM garbling for  $y_1$ .

*Proof* (Theorem 20). Let  $Q_1, Q_2$  be the number of secret key queries in **Query I** and **Query II**, respectively. Writing  $k_{\text{idx}}$  for a random key of SKE and  $k$  for a random (non-punctured) key of PPRF, we list our hybrids.

- $H_0$ . This is just  $\text{Exp}_{\text{PHFE}}^0$ , described as

$$\begin{array}{lll} \text{sk}_q : & s_{\text{PPRF},q} \xleftarrow{\$} \text{random}, & \text{idx}'_{\text{sk},q} \xleftarrow{\$} \text{random}, & \text{lgram}'_{\text{sk},q} \xleftarrow{\$} \text{random}; \\ \text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{normal}, \\ & y' \leftarrow \perp, & \text{idx}_{\text{ct}} \leftarrow \perp, & \text{idx}_H \leftarrow \perp, \\ & k'_{\text{idx}} \leftarrow \perp, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp. \end{array}$$

- $H_1$ . In this hybrid, we sample  $s_{\text{PPRF},q}$ 's without replacement, i.e.,

$$\begin{array}{lll} \text{sk}_q : & s_{\text{PPRF},q} \xleftarrow{\$} \text{distinct}, & \text{idx}'_{\text{sk},q} \xleftarrow{\$} \text{random}, & \text{lgram}'_{\text{sk},q} \xleftarrow{\$} \text{random}; \\ \text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{normal}, \\ & y' \leftarrow \perp, & \text{idx}_{\text{ct}} \leftarrow \perp, & \text{idx}_H \leftarrow \perp, \\ & k'_{\text{idx}} \leftarrow \perp, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp. \end{array}$$

$H_0$  and  $H_1$  are statistically indistinguishable.

- $H_2$ . In this hybrid, we change  $\text{idx}'_{\text{sk},q}$  into an encryption of  $q$  (padded to some fixed  $\text{poly}(\lambda)$  bits), i.e.,

$$\begin{array}{lll} \text{sk}_q : & s_{\text{PPRF},q} \xleftarrow{\$} \text{distinct}, & \text{idx}'_{\text{sk},q} \xleftarrow{\$} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \xleftarrow{\$} \text{random}; \\ \text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{normal}, \\ & y' \leftarrow \perp, & \text{idx}_{\text{ct}} \leftarrow \perp, & \text{idx}_H \leftarrow \perp, \\ & k'_{\text{idx}} \leftarrow \perp, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp. \end{array}$$

$H_1 \approx H_2$  follows from the ciphertext pseudorandomness of SKE, as the key  $k_{\text{idx}}$  is not used anywhere else.

- $H_3$ . In this hybrid, we prepare for hybridizing over the keys.  $H_3$  is described as

$$\begin{array}{lll} \text{sk}_q : & s_{\text{PPRF},q} \xleftarrow{\$} \text{distinct}, & \text{idx}'_{\text{sk},q} \xleftarrow{\$} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \xleftarrow{\$} \text{random}; \\ \text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{hybrid}, \\ & y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_H \leftarrow 0, \\ & k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp. \end{array}$$

It is readily verified that the decryption outcome (as perceived by ckt) does not change. Therefore,  $H_2 \approx H_3$  follows from the security of ckt.

- $H_{4,q}$  for  $q = 0, \dots, Q_1 + Q_2$ . In this hybrid,  $\text{idx}_H$  is incremented to  $q$ , i.e.,

$$\text{sk}_q : \quad s_{\text{PPRF},q} \xleftarrow{\$} \text{distinct}, \quad \text{idx}'_{\text{sk},q} \xleftarrow{\$} \text{SKE.Enc}(k_{\text{idx}}, q), \quad \text{lgram}'_{\text{sk},q} \xleftarrow{\$} \text{random};$$

$$\begin{array}{lll}
\text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{hybrid}, \\
& y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_{\text{H}} \leftarrow \boxed{q}, \\
& k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp.
\end{array}$$

$H_{4,0}$  is just  $H_3$ . The proofs of indistinguishability between  $H_{4,q-1}$  and  $H_{4,q}$  depend on whether  $q \leq Q_1$  or  $q > Q_1$ .

*Claim 21* (¶).  $H_{4,q-1} \approx H_{4,q}$  for  $q = 1, \dots, Q_1$ .

*Claim 22* (¶).  $H_{4,q-1} \approx H_{4,q}$  for  $q = Q_1 + 1, \dots, Q_1 + Q_2$ .

- $H_5$ . In this hybrid, we finish the hybrid argument over the keys.  $H_5$  is described as

$$\begin{array}{lll}
\text{sk}_q : & s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{hybrid}, \\
& y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_{\text{H}} \leftarrow \boxed{Q_1 + Q_2}, \\
& k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp.
\end{array}$$

$H_5$  is just  $H_{4,Q_1+Q_2}$ .

- $H_6$ . In this hybrid, the challenge ciphertext becomes a normal one for  $y_1$ , i.e.,

$$\begin{array}{lll}
\text{sk}_q : & s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & y \leftarrow \boxed{y_1}, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \boxed{\text{normal}}, \\
& y' \leftarrow \boxed{\perp}, & \text{idx}_{\text{ct}} \leftarrow \boxed{\perp}, & \text{idx}_{\text{H}} \leftarrow \boxed{\perp}, \\
& k'_{\text{idx}} \leftarrow \boxed{\perp}, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp.
\end{array}$$

$H_5 \approx H_6$  follows from the security of ckt, analogously to  $H_2 \approx H_3$ .

- $H_7$ . In this hybrid, we revert  $\text{idx}'_{\text{sk},q}$  to random, i.e.,

$$\begin{array}{lll}
\text{sk}_q : & s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \boxed{\text{random}}, & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & y \leftarrow y_1, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{normal}, \\
& y' \leftarrow \perp, & \text{idx}_{\text{ct}} \leftarrow \perp, & \text{idx}_{\text{H}} \leftarrow \perp, \\
& k'_{\text{idx}} \leftarrow \perp, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp.
\end{array}$$

$H_6 \approx H_7$  by the ciphertext pseudorandomness of SKE, analogously to  $H_1 \approx H_2$ .

- $H_8$ . In this hybrid,  $s_{\text{PPRF},q}$ 's are sampled as specified by KeyGen, i.e.,

$$\begin{array}{lll}
\text{sk}_q : & s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \boxed{\text{random}}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}, & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & y \leftarrow y_1, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{normal}, \\
& y' \leftarrow \perp, & \text{idx}_{\text{ct}} \leftarrow \perp, & \text{idx}_{\text{H}} \leftarrow \perp, \\
& k'_{\text{idx}} \leftarrow \perp, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp.
\end{array}$$

$H_8$  is statistically indistinguishable from  $H_7$ .

By inspection,  $H_8$  is just  $\text{Exp}_{\text{PHFE}}^1$ , and therefore,  $\text{Exp}_{\text{PHFE}}^0 \equiv H_0 \approx H_8 \equiv \text{Exp}_{\text{PHFE}}^1$ .  $\square$

*Proof (Claim 21).* To show indistinguishability between  $H_{4,q-1}$  and  $H_{4,q}$  when  $q \leq Q_1$ , we temporarily hardwire the LGRAM garbling yielded by decryption  $\text{ct}$  using  $\text{sk}_q$  into  $\text{lgram}_{\text{ct}}$ , which is generated when both  $f_q$  and  $x, y_0, y_1$  are known.

Let  $\hat{k}_{\{q\}}$  be  $k$  punctured at the singleton set  $\{q\}$ . We specify the hybrids.

- $G_0$ . This is just  $H_{4,q-1}$ , described as

$$\begin{array}{lll} \text{sk}_q : & s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\ \text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{hybrid}, \\ & y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_H \leftarrow q-1, \\ & k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp. \end{array}$$

- $G_1$ . In this hybrid, we puncture  $k_{\text{PPRF}}$ , hardwire the decryption result of  $\text{ct}$  by  $\text{sk}_q$  into  $\text{ct}$  at  $\text{lgram}_{\text{ct}}$ , indicating hardwiring using  $\text{mode}$ , and increment  $\text{idx}_H$ , i.e.,

$$\begin{array}{lll} \text{sk}_q : & s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\ \text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow \hat{k}_{\{q\}}, & \text{mode} \leftarrow \text{hardwire}, \\ & y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_H \leftarrow q, \\ & k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} \leftarrow \perp, & \\ \text{lgram}_{\text{ct}} \leftarrow & \left( \widehat{M}, \{L_{i,y_0[i]}\}_i \right) \text{ from LGRAM.Garble} \left( \begin{array}{l} T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \text{PPRF.Eval}(k, s_{\text{PPRF},q}) \end{array} \right). \end{array}$$

$G_0 \approx G_1$  follows from the security of  $\text{ckt}$ .

- $G_2$ . In this hybrid, we change the LGRAM randomness from  $\text{PPRF.Eval}(k, q)$  to true randomness, i.e.,

$$\begin{array}{lll} \text{sk}_q : & s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\ \text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow \hat{k}_{\{q\}}, & \text{mode} \leftarrow \text{hardwire}, \\ & y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_H \leftarrow q, \\ & k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} \leftarrow \perp, & \\ \text{lgram}_{\text{ct}} \leftarrow & \left( \widehat{M}, \{L_{i,y_0[i]}\}_i \right) \text{ from LGRAM.Garble} \left( \begin{array}{l} T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \text{random} \end{array} \right). \end{array}$$

$G_1 \approx G_2$  follows from the security of PPRF.

- $G_3$ . In this hybrid, we change the hardwired LGRAM garbling into one for  $y_1$ , i.e.,

$$\begin{array}{lll} \text{sk}_q : & s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\ \text{ct} : & y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow \hat{k}_{\{q\}}, & \text{mode} \leftarrow \text{hardwire}, \\ & y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_H \leftarrow q, \end{array}$$

$$\begin{aligned}
k'_{\text{idX}} &\leftarrow k_{\text{idX}}, & k'_{\text{lgram}} &\leftarrow \perp, \\
\text{lgram}_{\text{ct}} &\leftarrow (\widehat{M}, \{L_{i, \boxed{y_1}}[i]\}_i) \text{ from } \text{LGRAM.Garble} \left( \begin{array}{l} T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \text{random} \end{array} \right).
\end{aligned}$$

$G_2 \approx G_3$  follows from the security of LGRAM.

- $G_4$ . In this hybrid, the hardcoded LGRAM garbling is reverted to be generated with pseudorandomness  $\text{PPRF.Eval}(k, q)$ , i.e.,

$$\begin{aligned}
\text{sk}_q : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idX}'_{\text{sk},q} &\stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idX}}, q), & \text{lgram}'_{\text{sk},q} &\stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} &\leftarrow \dot{k}_{\{q\}}, & \text{mode} &\leftarrow \text{hardware}, \\
& \quad y' \leftarrow y_1, & \text{idX}_{\text{ct}} &\leftarrow Q_1, & \text{idX}_{\text{H}} &\leftarrow q, \\
& \quad k'_{\text{idX}} \leftarrow k_{\text{idX}}, & k'_{\text{lgram}} &\leftarrow \perp, \\
\text{lgram}_{\text{ct}} &\leftarrow (\widehat{M}, \{L_{i, y_1}[i]\}_i) \text{ from } \text{LGRAM.Garble} \left( \begin{array}{l} T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \boxed{\text{PPRF.Eval}(k, s_{\text{PPRF},q})} \end{array} \right).
\end{aligned}$$

$G_3 \approx G_4$  follows from the security of PPRF.

- $G_5$ . In this hybrid,  $k_{\text{PPRF}}$  is no longer punctured and hardwiring is undone, i.e.,

$$\begin{aligned}
\text{sk}_q : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idX}'_{\text{sk},q} &\stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idX}}, q), & \text{lgram}'_{\text{sk},q} &\stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} &\leftarrow \boxed{k}, & \text{mode} &\leftarrow \text{hybrid}, \\
& \quad y' \leftarrow y_1, & \text{idX}_{\text{ct}} &\leftarrow Q_1, & \text{idX}_{\text{H}} &\leftarrow q, \\
& \quad k'_{\text{idX}} \leftarrow k_{\text{idX}}, & k'_{\text{lgram}} &\leftarrow \perp, & \text{lgram}_{\text{ct}} &\leftarrow \boxed{\perp}.
\end{aligned}$$

$G_4 \approx G_5$  follows from the security of ckt.

By inspection,  $G_5$  is exactly  $H_{4,q}$ . Therefore,  $H_{4,q-1} \equiv G_0 \approx G_5 \equiv H_{4,q}$ .  $\square$

*Proof (Claim 22).* To show indistinguishability between  $H_{4,q-1}$  and  $H_{4,q}$  when  $\boxed{q > Q_1}$ , we temporarily hardwire the LGRAM garbling yielded by decryption ct using  $\text{sk}_q$  into  $\boxed{\text{lgram}'_{\text{sk},q}}$ , generated when both  $x, y_0, y_1$  and  $f_q$  are known. Let  $\dot{k}_{\{q\}}$  be  $k$  punctured at  $\{q\}$  and  $k_{\text{lgram}}$  a random secret key of SKE. We specify the hybrids.

- $G_0$ . This is just  $H_{4,q-1}$ , described as

$$\begin{aligned}
\text{sk}_{q \neq q} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idX}'_{\text{sk},q} &\stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idX}}, q), & \text{lgram}'_{\text{sk},q} &\stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} &\leftarrow k, & \text{mode} &\leftarrow \text{hybrid}, \\
& \quad y' \leftarrow y_1, & \text{idX}_{\text{ct}} &\leftarrow Q_1, & \text{idX}_{\text{H}} &\leftarrow q-1, \\
& \quad k'_{\text{idX}} \leftarrow k_{\text{idX}}, & k'_{\text{lgram}} &\leftarrow \perp, & \text{lgram}_{\text{ct}} &\leftarrow \perp; \\
\text{sk}_{q > Q_1} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idX}'_{\text{sk},q} &\stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idX}}, q), & \text{lgram}'_{\text{sk},q} &\stackrel{\$}{\leftarrow} \text{random}.
\end{aligned}$$

- $G_1$ . In this hybrid, we encrypt the LGRAM garbling into  $\text{lgram}'_{\text{sk},q}$ , i.e.,

$$\text{sk}_{q \neq q} : \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, \quad \text{idX}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idX}}, q), \quad \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random};$$

$$\begin{aligned}
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow k, & \text{mode} \leftarrow \text{hybrid}, \\
& \quad y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_{\text{H}} \leftarrow q - 1, \\
& \quad k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} \leftarrow \perp, & \text{lgram}_{\text{ct}} \leftarrow \perp; \\
\text{sk}_{q > Q_1} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), \\
& \quad \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc} \left( k_{\text{lgram}}, \text{LGRAM.Garble} \left( \begin{array}{l} (\widehat{M}, \{L_{i,y_0[i]}\}_i) \text{ from} \\ T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \text{PPRF.Eval}(k, s_{\text{PPRF},q}) \end{array} \right) \right).
\end{aligned}$$

$G_0 \approx G_1$  by the ciphertext pseudorandomness of SKE.

- $G_2$ . In this hybrid, we puncture  $k$  at  $\{q\}$ , activate the hardwiring using  $k'_{\text{lgram}}$ , mode, and increment  $\text{idx}_{\text{H}}$ , i.e.,

$$\begin{aligned}
\text{sk}_{q \neq q} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow \overset{\circ}{k}_{\{q\}}, & \text{mode} \leftarrow \text{hardwire}, \\
& \quad y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_{\text{H}} \leftarrow q, \\
& \quad k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} \leftarrow \overset{\circ}{k}_{\text{lgram}}, & \text{lgram}_{\text{ct}} \leftarrow \perp; \\
\text{sk}_{q > Q_1} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), \\
& \quad \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc} \left( k_{\text{lgram}}, \text{LGRAM.Garble} \left( \begin{array}{l} (\widehat{M}, \{L_{i,y_0[i]}\}_i) \text{ from} \\ T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \text{PPRF.Eval}(k, s_{\text{PPRF},q}) \end{array} \right) \right).
\end{aligned}$$

$G_1 \approx G_2$  follows from the security of ckt.

- $G_3$ . In this hybrid, we generate the hardwired garbling with true randomness instead of  $\text{PPRF.Eval}(k, q)$ , i.e.,

$$\begin{aligned}
\text{sk}_{q \neq q} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} \leftarrow \overset{\circ}{k}_{\{q\}}, & \text{mode} \leftarrow \text{hardwire}, \\
& \quad y' \leftarrow y_1, & \text{idx}_{\text{ct}} \leftarrow Q_1, & \text{idx}_{\text{H}} \leftarrow q, \\
& \quad k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} \leftarrow k_{\text{lgram}}, & \text{lgram}_{\text{ct}} \leftarrow \perp; \\
\text{sk}_{q > Q_1} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), \\
& \quad \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc} \left( k_{\text{lgram}}, \text{LGRAM.Garble} \left( \begin{array}{l} (\widehat{M}, \{L_{i,y_0[i]}\}_i) \text{ from} \\ T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \text{random} \end{array} \right) \right).
\end{aligned}$$

$G_2 \approx G_3$  follows from the security of PPRF.

- $G_4$ . In this hybrid, the hardwired garbling becomes one for  $y_1$ , i.e.,

$$\text{sk}_{q \neq q} : \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, \quad \text{idx}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), \quad \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{random};$$



$$\begin{aligned}
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} & \leftarrow \hat{k}_{\{q\}}, & \text{mode} & \leftarrow \text{hardware}, \\
& \quad y' \leftarrow y_1, & \text{idx}_{\text{ct}} & \leftarrow Q_1, & \text{idx}_{\text{H}} & \leftarrow q, \\
& \quad k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} & \leftarrow k_{\text{lgram}}, & \text{lgram}_{\text{ct}} & \leftarrow \perp; \\
\text{sk}_{q>Q_1} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} & \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), \\
& \quad \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc} \left( k_{\text{lgram}}, \begin{array}{l} (\widehat{M}, \{L_{i,y_1}[i]\}_i) \text{ from} \\ \text{LGRAM.Garble} \left( \begin{array}{l} T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \text{random} \end{array} \right) \end{array} \right).
\end{aligned}$$

$G_3 \approx G_4$  follows from the security of LGRAM.

- $G_5$ . In this hybrid, the randomness for generating the hardwired LGRAM garbling is reverted to be pseudorandom, i.e.,

$$\begin{aligned}
\text{sk}_{q \neq q} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} & \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} & \stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} & \leftarrow \hat{k}_{\{q\}}, & \text{mode} & \leftarrow \text{hardware}, \\
& \quad y' \leftarrow y_1, & \text{idx}_{\text{ct}} & \leftarrow Q_1, & \text{idx}_{\text{H}} & \leftarrow q, \\
& \quad k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} & \leftarrow k_{\text{lgram}}, & \text{lgram}_{\text{ct}} & \leftarrow \perp; \\
\text{sk}_{q>Q_1} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} & \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), \\
& \quad \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc} \left( k_{\text{lgram}}, \begin{array}{l} (\widehat{M}, \{L_{i,y_1}[i]\}_i) \text{ from} \\ \text{LGRAM.Garble} \left( \begin{array}{l} T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \text{PPRF.Eval}(k, q) \end{array} \right) \end{array} \right).
\end{aligned}$$

$G_4 \approx G_5$  follows from the security of PPRF.

- $G_6$ . In this hybrid, the hardwiring is deactivated by reverting most of the changes made in the transition from  $G_1$  to  $G_2$ , i.e.,

$$\begin{aligned}
\text{sk}_{q \neq q} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} & \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} & \stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} & \leftarrow k, & \text{mode} & \leftarrow \text{hybrid}, \\
& \quad y' \leftarrow y_1, & \text{idx}_{\text{ct}} & \leftarrow Q_1, & \text{idx}_{\text{H}} & \leftarrow q, \\
& \quad k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} & \leftarrow \perp, & \text{lgram}_{\text{ct}} & \leftarrow \perp; \\
\text{sk}_{q>Q_1} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} & \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), \\
& \quad \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc} \left( k_{\text{lgram}}, \begin{array}{l} (\widehat{M}, \{L_{i,y_1}[i]\}_i) \text{ from} \\ \text{LGRAM.Garble} \left( \begin{array}{l} T_{\text{max}}, M, \text{digest}_{f_q}, \text{digest}_x; \\ \text{PPRF.Eval}(k, q) \end{array} \right) \end{array} \right).
\end{aligned}$$

$G_5 \approx G_6$  follows from the security of ckt.

- $G_7$ . In this hybrid, we revert  $\text{lgram}'_{\text{sk},q}$  to random, i.e.,

$$\begin{aligned}
\text{sk}_{q \neq q} : & \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, & \text{idx}'_{\text{sk},q} & \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{idx}}, q), & \text{lgram}'_{\text{sk},q} & \stackrel{\$}{\leftarrow} \text{random}; \\
\text{ct} : & \quad y \leftarrow y_0, & k_{\text{PPRF}} & \leftarrow k, & \text{mode} & \leftarrow \text{hybrid}, \\
& \quad y' \leftarrow y_1, & \text{idx}_{\text{ct}} & \leftarrow Q_1, & \text{idx}_{\text{H}} & \leftarrow q, \\
& \quad k'_{\text{idx}} \leftarrow k_{\text{idx}}, & k'_{\text{lgram}} & \leftarrow \perp, & \text{lgram}_{\text{ct}} & \leftarrow \perp;
\end{aligned}$$

$$\text{sk}_{q>Q_1} : \quad s_{\text{PPRF},q} \stackrel{\$}{\leftarrow} \text{distinct}, \quad \text{id}_{\text{sk},q} \stackrel{\$}{\leftarrow} \text{SKE.Enc}(k_{\text{id}_x}, q), \quad \text{lgram}'_{\text{sk},q} \stackrel{\$}{\leftarrow} \boxed{\perp}.$$

$G_6 \approx G_7$  by the ciphertext pseudorandomness of SKE.

By inspection,  $G_7$  is exactly  $H_{4,q}$ . Therefore,  $H_{4,q-1} \equiv G_0 \approx G_7 \equiv H_{4,q}$ .  $\square$

## 6.2 Full-Fledged PHFE for RAM

In this section, we build a full-fledged PHFE for RAM that is  $f$ -succinct and has rate-2 ciphertext and linear-time KeyGen and Enc and whose Dec runs in time  $O(T + |f| + |x| + |y|)$ , ignoring polynomial factors in the security parameter. (See Definitions 5, 7, and 10.)

Recall that in full-fledged PHFE for RAM, the functionality (resp. key, ciphertext) is associated with some RAM  $M$  and (up to exponential) time bound  $T_{\max}$  (resp.  $f$ ,  $(x, y)$ , each of arbitrary length) and decryption yields  $M^{f,x||y}()$  if the execution halts in time at most  $T_{\max}$ .

**Ingredients of Construction 6.** Let

- PHFE' = (Setup', KeyGen', Enc', Dec') be a PHFE scheme for RAM with bounded private input that is  $f'$ - and  $x'$ -succinct and has linear-time KeyGen' and Enc' and whose Dec' runs in time  $(T' + |f'| + |x'|) \text{poly}(\lambda, M', \log T'_{\max})$ , and
- PRF a pseudorandom function.

**Construction 6** (full-fledged PHFE for RAM). Our scheme works as follows:

- Setup( $M, T_{\max}$ ) runs and outputs

$$(\text{mpk}, \text{msk}) = (\text{mpk}', \text{msk}') \stackrel{\$}{\leftarrow} \text{Setup}'(M', T'_{\max}),$$

where  $M'$  is shown in Figure 11 and  $T'_{\max} = T_{\max} + \text{poly}(\lambda, \log T_{\max})$ .

- KeyGen( $\text{msk}, f$ ) pads  $f$  by setting  $\text{one cell on an input tape of } M'$

$$f' \leftarrow (f[1] || 0^{\ell_{\text{cell}}}) || \cdots || \overbrace{(f[|f|] || 0^{\ell_{\text{cell}}})}$$

and runs and outputs

$$\text{sk} = \text{sk}' \stackrel{\$}{\leftarrow} \text{KeyGen}'(\text{msk}', f').$$

- Enc( $\text{mpk}, x, y$ ) samples  $\beta \stackrel{\$}{\leftarrow} \{0, 1\}$ , a PRF key  $k_\beta$ , and a string  $w_{1-\beta}$  of the same length as  $y$ . It pads  $x$  and appends to it an interleaved version of  $w_0, w_1$ , where  $w_\beta$  encrypts  $y$ , by setting

$$\begin{aligned} w_\beta[i] &\leftarrow y[i] \oplus \text{PRF}(k_\beta, i) && \text{for } i \in [|y|], \\ w &\leftarrow \underbrace{(w_0[1] || w_1[1])}_{\text{one cell on an input tape of } M'} || \cdots || (w_0[|y|] || w_1[|y|]), \\ x' &\leftarrow \underbrace{(x[0] || 0^{\ell_{\text{cell}}})}_{\text{one cell on an input tape of } M'} || \cdots || (x[|x|] || 0^{\ell_{\text{cell}}}) || w, && y' \leftarrow (|x|, \beta, k_\beta). \end{aligned}$$

The algorithm runs  $\text{ct}' \stackrel{\$}{\leftarrow} \text{Enc}'(\text{mpk}', x', y')$  and outputs  $\text{ct} = (\text{ct}', w)$ .

- Dec $^{f,x,\text{sk},\text{ct}}$ ( $\text{mpk}$ ) prepares oracles for  $f', x'$  from  $f, x, \text{ct}$  as specified in KeyGen, Enc. It runs and outputs

$$(\text{Dec}')^{f',x',\text{sk}',\text{ct}'}(\text{mpk}').$$

### 2-Tape RAM $M'$

- Lengths.**  $\ell'_{\text{st}} = \ell_{\text{st}} + \text{poly}(\lambda, |M|, \log T_{\text{max}})$ ,  
the state of  $M$ , the location of the last-read cell;  
 $\ell'_{\text{in}} = \text{poly}(\lambda, |M|, \log T_{\text{max}}) + 1 + \text{poly}(\lambda, |M|, \log T_{\text{max}})$ ,  
to encode  $|x|$ , choice bit  $\beta$ , and PRF key  $k_\beta$ ;  
 $\ell'_{\text{addr}} = \ell_{\text{addr}} + 1$ ,  $\ell'_{\text{cell}} = 2\ell_{\text{cell}}$ ,  
twice as many cells with each cell twice as large  
on an input tape (to encode  $x, y_0, y_1$  in  $x'$ );  
 $\ell'_{\text{ADDR}} = \ell_{\text{ADDR}}$ ,  $\ell'_{\text{CELL}} = \ell_{\text{CELL}}$ ,  
exactly the working tape of  $M$ .
- Input.**  $w' = (|x|, \beta, k_\beta)$ , length of  $x$ , choice bit, PRF key for  $y$  or  $y_\beta$ ;  
 $D'_1 = f'$ ,  $\ell'_1 = \ell_1$ , padded version of  $f$ ;  
 $D'_2 = x'$ ,  $\ell'_2 = |x| + |y|$ ,  
padded version of  $x$  followed by  
interleaved encryption of  $y$  or  $y_0, y_1$ .
- Steps.** Each step of  $M'$  replicates one step of  $M$   
by translating its memory access as follows.
1. If  $M$  reads  $D_1[i] = f[i]$ :  
**read**  $D'_1[i] = f'[i] = f[i] || 0^{\ell_{\text{cell}}}$  in this step  
**provide**  $f[i]$  in the next step (similar below)
  - 2a. If  $M$  reads  $D_2[i] = x[i]$  for  $i \in [|x|]$ :  
**read**  $D'_2[i] = x'[i] = x[i] || 0^{\ell_{\text{cell}}}$   
**provide**  $x[i]$
  - 2b. If  $M$  reads  $D_2[|x| + i] = y[i]$  for  $i \in [|y|]$ :  
**read**  $D'_2[|x| + i] = w[i] = w_0[i] || w_1[i]$   
**provide**  $w_\beta[i] \oplus \text{PRF}(k_\beta, i)$
  3. If  $M$  reads  $D_{\text{work}}[i]$ :  
**read and provide**  $D'_{\text{work}}[i]$

**Figure 11.** The machine  $M'$  in Construction 6.

**Correctness and Efficiency.** Correctness is immediate by inspection. The construction scheme incurs an additional storage twice as long as  $y$  in ct, no additional machine time, and constant-factor additional decryption time. Since  $|M'| = \text{poly}(\lambda, |M|, \log T_{\max})$  and  $T'_{\max} = T_{\max} + \text{poly}(\lambda, M, \log T_{\max})$ , with parameters for PHFE' denoted with primes,

$$\begin{aligned}
|\text{sk}| &= |\text{sk}'| = \text{poly}(\lambda, |M'|, \log T'_{\max}) \\
&= \text{poly}(\lambda, |M|, \log T_{\max}), \\
|\text{ct}| &= 2|y| + |\text{ct}'| = 2|y| + \text{poly}(\lambda, |M'|, \log T'_{\max}) \\
&= 2|y| + \text{poly}(\lambda, |M|, \log T_{\max}), \\
T_{\text{KeyGen}} &= T'_{\text{KeyGen}} = |f'| \text{poly}(\lambda, |M'|, \log T'_{\max}) \\
&= |f| \text{poly}(\lambda, |M|, \log T_{\max}), \\
T_{\text{Enc}} &= T'_{\text{Enc}} = |x'| \text{poly}(\lambda, |M'|, \log T'_{\max}) \\
&= (|x| + |y|) \text{poly}(\lambda, |M|, \log T_{\max}), \\
T_{\text{Dec}} &= T'_{\text{Dec}} = (T' + |f'| + |x'|) \text{poly}(\lambda, |M'|, \log T'_{\max}) \\
&= (T + |f| + |x| + |y|) \text{poly}(\lambda, |M|, \log T_{\max}).
\end{aligned}$$

**Theorem 23** (¶). *Suppose in Construction 6, PHFE', PRF are secure (Definitions 7 and 17), then the constructed scheme is secure (Definition 7).*

*Proof* (Theorem 23). Let  $\beta \stackrel{\$}{\leftarrow} \{0, 1\}$  be a random bit and  $k_\beta, k_{1-\beta}$  two random PRF keys. We describe how the challenge ciphertext is generated in each hybrid.

- $H_0^b$ . This is just  $\text{Exp}_{\text{PHFE}}^b$ , where

$$\begin{aligned}
w_\beta[i] &\leftarrow y_b[i] \oplus \text{PRF}(k_\beta, i), & w_{1-\beta}[i] &\stackrel{\$}{\leftarrow} \text{random}, \\
x' &\leftarrow (x[1] \| 0^{\ell_{\text{cell}}}) \| \dots \| (x[|x|] \| 0^{\ell_{\text{cell}}}) \| (w_0[1] \| w_1[i]) \| \dots \| (w_0[|y|] \| w_1[|y|]), \\
\text{ct}' &\stackrel{\$}{\leftarrow} \text{Enc}'(\text{mpk}, x', (|x|, \beta, k_\beta)).
\end{aligned}$$

- $H_1^b$ . In this hybrid, we make  $w_{1-\beta}$  an encryption of  $y_{1-b}$  under  $k_{1-\beta}$ , i.e.,

$$\begin{aligned}
w_\beta[i] &\leftarrow y_b[i] \oplus \text{PRF}(k_\beta, i), & w_{1-\beta}[i] &\leftarrow y_{1-b}[i] \oplus \text{PRF}(k_{1-\beta}, i), \\
x' &\leftarrow (x[1] \| 0^{\ell_{\text{cell}}}) \| \dots \| (x[|x|] \| 0^{\ell_{\text{cell}}}) \| (w_0[1] \| w_1[i]) \| \dots \| (w_0[|y|] \| w_1[|y|]), \\
\text{ct}' &\stackrel{\$}{\leftarrow} \text{Enc}'(\text{mpk}, x', (|x|, \beta, k_\beta)).
\end{aligned}$$

$H_0^b \approx H_1^b$  for each  $b \in \{0, 1\}$  by the security of PRF.

- $H_2^b$ . In this hybrid, we rename  $(\beta \oplus b)$  to  $\gamma$ , making it

$$\begin{aligned}
w_{\gamma}[i] &\leftarrow y_{\gamma}[i] \oplus \text{PRF}(k_{\gamma}, i), & w_{1-\gamma}[i] &\leftarrow y_{1-\gamma}[i] \oplus \text{PRF}(k_{1-\gamma}, i), \\
x' &\leftarrow (x[1] \| 0^{\ell_{\text{cell}}}) \| \dots \| (x[|x|] \| 0^{\ell_{\text{cell}}}) \| (w_0[1] \| w_1[i]) \| \dots \| (w_0[|y|] \| w_1[|y|]), \\
\text{ct}' &\stackrel{\$}{\leftarrow} \text{Enc}'(\text{mpk}, x', (|x|, \gamma \oplus b, k_{\gamma \oplus b})),
\end{aligned}$$

where  $\gamma \stackrel{\$}{\leftarrow} \{0, 1\}$ . This change is conceptual, hence  $H_1^b \equiv H_2^b$  for each  $b \in \{0, 1\}$ .

$H_2^0 \approx H_2^1$  follows from the security of PHFE'. Therefore,  $\text{Exp}_{\text{PHFE}}^0 \equiv H_0^0 \approx H_0^1 \equiv \text{Exp}_{\text{PHFE}}^1$ .  $\square$

**Acknowledgement.** The authors were supported by NSF grants CNS-1528178, CNS-1929901, CNS-1936825 (CAREER), CNS-2026774, a Hellman Fellowship, a JP Morgan AI Research Award, the Defense Advanced Research Projects Agency (DARPA) and Army Research Office (ARO) under Contract No. W911NF-15-C-0236, and a subcontract No. 2017-002 through Galois. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

## References

- [ACC<sup>+</sup>16] Prabhanjan Ananth, Yu-Chi Chen, Kai-Min Chung, Huijia Lin, and Wei-Kai Lin. Delegating RAM computations with adaptive soundness and privacy. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 3–30. Springer, Heidelberg, October / November 2016.
- [AFS19] Prabhanjan Ananth, Xiong Fan, and Elaine Shi. Towards attribute-based encryption for RAMs from LWE: Sub-linear decryption, and more. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 112–141. Springer, Heidelberg, December 2019.
- [AIKW13] Benny Applebaum, Yuval Ishai, Eyal Kushilevitz, and Brent Waters. Encoding functions with constant online rate or how to compress garbled circuits keys. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 166–184. Springer, Heidelberg, August 2013.
- [AJL<sup>+</sup>19] Prabhanjan Ananth, Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. Indistinguishability obfuscation without multilinear maps: New paradigms via low degree weak pseudorandomness and security amplification. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 284–332. Springer, Heidelberg, August 2019.
- [AJS17] Prabhanjan Ananth, Abhishek Jain, and Amit Sahai. Indistinguishability obfuscation for turing machines: Constant overhead and amortization. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 252–279. Springer, Heidelberg, August 2017.
- [AL18] Prabhanjan Ananth and Alex Lombardi. Succinct garbling schemes from functional encryption through a local simulation paradigm. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 455–472. Springer, Heidelberg, November 2018.
- [ALdP11] Nuttapong Attrapadung, Benoît Libert, and Elie de Panafieu. Expressive key-policy attribute-based encryption with constant-size ciphertexts. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011*, volume 6571 of *LNCS*, pages 90–108. Springer, Heidelberg, March 2011.
- [AM18] Shweta Agrawal and Monosij Maitra. FE and iO for turing machines from minimal assumptions. In Amos Beimel and Stefan Dziembowski, editors,

*TCC 2018, Part II*, volume 11240 of *LNCS*, pages 473–512. Springer, Heidelberg, November 2018.

- [AS16] Prabhanjan Vijendra Ananth and Amit Sahai. Functional encryption for turing machines. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part I*, volume 9562 of *LNCS*, pages 125–153. Springer, Heidelberg, January 2016.
- [AT20] Nuttapon Attrapadung and Junichi Tomida. Unbounded dynamic predicate compositions in ABE from standard assumptions. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 405–436. Springer, Heidelberg, December 2020.
- [Att16] Nuttapon Attrapadung. Dual system encryption framework in prime-order groups via computational pair encodings. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 591–623. Springer, Heidelberg, December 2016.
- [Ben89] Charles Henry Bennett. Time/space trade-offs for reversible computation. *SIAM Journal on Computing*, 18(4):766–776, 1989.
- [BGG<sup>+</sup>14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Heidelberg, May 2014.
- [BGL<sup>+</sup>15] Nir Bitansky, Sanjam Garg, Huijia Lin, Rafael Pass, and Sidharth Telang. Succinct randomized encodings and their applications. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 439–448. ACM Press, June 2015.
- [BHMW21] Elette Boyle, Justin Holmgren, Fermi Ma, and Mor Weiss. On the security of doubly efficient PIR. Cryptology ePrint Archive, Report 2021/1113, 2021. <https://eprint.iacr.org/2021/1113>.
- [BHR12] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 784–796. ACM Press, October 2012.
- [BIPW17] Elette Boyle, Yuval Ishai, Rafael Pass, and Mary Wootters. Can we access a database both locally and privately? In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part II*, volume 10678 of *LNCS*, pages 662–693. Springer, Heidelberg, November 2017.
- [BMR90] Donald Beaver, Silvio Micali, and Phillip Rogaway. The round complexity of secure protocols (extended abstract). In *22nd ACM STOC*, pages 503–513. ACM Press, May 1990.
- [BS15] Zvika Brakerski and Gil Segev. Function-private functional encryption in the private-key setting. In Yevgeniy Dodis and Jesper Buus Nielsen, editors,

*TCC 2015, Part II*, volume 9015 of *LNCS*, pages 306–324. Springer, Heidelberg, March 2015.

- [BSW11] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, Heidelberg, March 2011.
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013.
- [CCC<sup>+</sup>16] Yu-Chi Chen, Sherman S. M. Chow, Kai-Min Chung, Russell W. F. Lai, Wei-Kai Lin, and Hong-Sheng Zhou. Cryptography for parallel RAM from indistinguishability obfuscation. In Madhu Sudan, editor, *ITCS 2016*, pages 179–190. ACM, January 2016.
- [CCHR16] Ran Canetti, Yilei Chen, Justin Holmgren, and Mariana Raykova. Adaptive succinct garbled RAM or: How to delegate your database. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 61–90. Springer, Heidelberg, October / November 2016.
- [CDG<sup>+</sup>17] Chongwon Cho, Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Antigoni Polychroniadou. Laconic oblivious transfer and its applications. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 33–65. Springer, Heidelberg, August 2017.
- [CH16] Ran Canetti and Justin Holmgren. Fully succinct garbled RAM. In Madhu Sudan, editor, *ITCS 2016*, pages 169–178. ACM, January 2016.
- [CH19] Ran Canetti and Justin Holmgren. Succinct garbled RAM. Cryptology ePrint Archive, Report 2015/388, version 20190517:011532, 2019. <https://eprint.iacr.org/archive/2015/388/20190517:011532>.
- [CHJV15] Ran Canetti, Justin Holmgren, Abhishek Jain, and Vinod Vaikuntanathan. Succinct garbling and indistinguishability obfuscation for RAM programs. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 429–437. ACM Press, June 2015.
- [CHR17] Ran Canetti, Justin Holmgren, and Silas Richelson. Towards doubly efficient private information retrieval. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part II*, volume 10678 of *LNCS*, pages 694–726. Springer, Heidelberg, November 2017.
- [DIJ<sup>+</sup>13] Angelo De Caro, Vincenzo Iovino, Abhishek Jain, Adam O’Neill, Omer Paneth, and Giuseppe Persiano. On the achievability of simulation-based security for functional encryption. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 519–535. Springer, Heidelberg, August 2013.
- [DTT10] Anindya De, Luca Trevisan, and Madhur Tulsiani. Time space tradeoffs for attacks against one-way functions and PRGs. In Tal Rabin, editor,

*CRYPTO 2010*, volume 6223 of *LNCS*, pages 649–665. Springer, Heidelberg, August 2010.

- [GGH<sup>+</sup>13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.
- [GKP<sup>+</sup>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, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 536–553. Springer, Heidelberg, August 2013.
- [GOS18] Sanjam Garg, Rafail Ostrovsky, and Akshayaram Srinivasan. Adaptive garbled RAM from laconic oblivious transfer. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 515–544. Springer, Heidelberg, August 2018.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 89–98. ACM Press, October / November 2006. Available as Cryptology ePrint Archive Report 2006/309.
- [GS16] Sanjam Garg and Akshayaram Srinivasan. Single-key to multi-key functional encryption with polynomial loss. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 419–442. Springer, Heidelberg, October / November 2016.
- [GS18a] Sanjam Garg and Akshayaram Srinivasan. Adaptively secure garbling with near optimal online complexity. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 535–565. Springer, Heidelberg, April / May 2018.
- [GS18b] Sanjam Garg and Akshayaram Srinivasan. A simple construction of iO for turing machines. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 425–454. Springer, Heidelberg, November 2018.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 545–554. ACM Press, June 2013.
- [GVW15] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Predicate encryption for circuits from LWE. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 503–523. Springer, Heidelberg, August 2015.
- [GWZ22] Jiaxin Guan, Daniel Wichs, and Mark Zhandry. Incompressible cryptography. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 700–730. Springer, Heidelberg, May / June 2022.



- [HJO<sup>+</sup>16] Brett Hemenway, Zahra Jafargholi, Rafail Ostrovsky, Alessandra Scafuro, and Daniel Wichs. Adaptively secure garbled circuits from one-way functions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 149–178. Springer, Heidelberg, August 2016.
- [JLMS19] Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. How to leverage hardness of constant-degree expanding polynomials over  $\mathbb{R}$  to build  $iO$ . In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 251–281. Springer, Heidelberg, May 2019.
- [JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 60–73, New York, NY, USA, 2021. Association for Computing Machinery.
- [JLS22] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from LPN over  $\mathbb{F}_p$ , DLIN, and PRGs in  $NC^0$ . In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 670–699. Springer, Heidelberg, May / June 2022.
- [KLW15] Venkata Koppula, Allison Bishop Lewko, and Brent Waters. Indistinguishability obfuscation for turing machines with unbounded memory. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 419–428. ACM Press, June 2015.
- [KNT18] Fuyuki Kitagawa, Ryo Nishimaki, and Keisuke Tanaka. Simple and generic constructions of succinct functional encryption. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part II*, volume 10770 of *LNCS*, pages 187–217. Springer, Heidelberg, March 2018.
- [KNTY19] Fuyuki Kitagawa, Ryo Nishimaki, Keisuke Tanaka, and Takashi Yamakawa. Adaptively secure and succinct functional encryption: Improving security and efficiency, simultaneously. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 521–551. Springer, Heidelberg, August 2019.
- [LL20] Huijia Lin and Ji Luo. Succinct and adaptively secure ABE for ABP from  $k$ -lin. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 437–466. Springer, Heidelberg, December 2020.
- [LLL22] Hanjun Li, Huijia Lin, and Ji Luo. ABE for circuits with constant-size secret keys and adaptive security. Cryptology ePrint Archive, Report 2022/659, 2022. <https://eprint.iacr.org/2022/659>.
- [LM16] Baiyu Li and Daniele Micciancio. Compactness vs collusion resistance in functional encryption. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 443–468. Springer, Heidelberg, October / November 2016.
- [LT17] Huijia Lin and Stefano Tessaro. Indistinguishability obfuscation from trilinear maps and block-wise local PRGs. In Jonathan Katz and Hovav

- Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 630–660. Springer, Heidelberg, August 2017.
- [Luo22] Ji Luo. Ad hoc (decentralized) broadcast, trace, and revoke. Cryptology ePrint Archive, Report 2022/925, 2022. <https://eprint.iacr.org/2022/925>.
- [LZ17] Qipeng Liu and Mark Zhandry. Decomposable obfuscation: A framework for building applications of obfuscation from polynomial hardness. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 138–169. Springer, Heidelberg, November 2017.
- [NY15] Moni Naor and Eylon Yogev. Bloom filters in adversarial environments. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 565–584. Springer, Heidelberg, August 2015.
- [O’N10] Adam O’Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010. <https://eprint.iacr.org/2010/556>.
- [QWW18] Willy Quach, Hoeteck Wee, and Daniel Wichs. Laconic function evaluation and applications. In Mikkel Thorup, editor, *59th FOCS*, pages 859–870. IEEE Computer Society Press, October 2018.
- [SW05] Amit Sahai and Brent R. Waters. Fuzzy identity-based encryption. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Heidelberg, May 2005.
- [Tak14] Katsuyuki Takashima. Expressive attribute-based encryption with constant-size ciphertexts from the decisional linear assumption. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 298–317. Springer, Heidelberg, September 2014.
- [YAHK14] Shota Yamada, Nuttapon Attrapadung, Goichiro Hanaoka, and Noboru Kunihiro. A framework and compact constructions for non-monotonic attribute-based encryption. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 275–292. Springer, Heidelberg, March 2014.
- [Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd FOCS*, pages 160–164. IEEE Computer Society Press, November 1982.
- [ZGT<sup>+</sup>16] Kai Zhang, Junqing Gong, Shaohua Tang, Jie Chen, Xiangxue Li, Haifeng Qian, and Zhenfu Cao. Practical and efficient attribute-based encryption with constant-size ciphertexts in outsourced verifiable computation. In Xiaofeng Chen, XiaoFeng Wang, and Xinyi Huang, editors, *ASIACCS 16*, pages 269–279. ACM Press, May / June 2016.