

## MIT Open Access Articles

*Indistinguishability Obfuscation for RAM  
Programs and Succinct Randomized Encodings*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

**Citation:** Bitansky, Nir, Canetti, Ran, Garg, Sanjam, Holmgren, Justin, Jain, Abhishek et al. 2018. "Indistinguishability Obfuscation for RAM Programs and Succinct Randomized Encodings." 47 (3).

**As Published:** 10.1137/15m1050963

**Publisher:** Society for Industrial & Applied Mathematics (SIAM)

**Persistent URL:** <https://hdl.handle.net/1721.1/137817>

**Version:** Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

**Terms of Use:** Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



## INDISTINGUISHABILITY OBFUSCATION FOR RAM PROGRAMS AND SUCCINCT RANDOMIZED ENCODINGS\*

NIR BITANSKY<sup>†</sup>, RAN CANETTI<sup>‡</sup>, SANJAM GARG<sup>§</sup>, JUSTIN HOLMGREN<sup>¶</sup>,  
ABHISHEK JAIN<sup>||</sup>, HUIJIA LIN<sup>#</sup>, RAFAEL PASS<sup>††</sup>, SIDHARTH TELANG<sup>††</sup> AND  
VINOD VAIKUNTANATHAN<sup>¶</sup>

**Abstract.** We show how to construct indistinguishability obfuscation (**iO**) for RAM programs with bounded space, assuming **iO** for circuits and one-way functions, both with subexponential security. That is, given a RAM program whose computation requires space  $s(n)$  in the worst case for inputs of length at most  $n$ , we generate an obfuscated RAM program that, for inputs of size at most  $n$ , runs in roughly the same time as the original program, using space roughly  $s(n)$ . The obfuscation process is quasi-linear in the description length of the input program and  $s(n)$ . At the heart of our construction are succinct randomized encodings for RAM programs. We present two very different constructions of such encodings, each with its own unique properties. Beyond their use as a tool in obfuscation for RAM programs, we show that succinct randomized encodings are interesting objects in their own right. We demonstrate the power of succinct randomized encodings in applications such as publicly verifiable delegation, functional encryption for RAMs, and key-dependent security amplification.

**Key words.** cryptography, randomized encodings, obfuscation, bootstrapping

**AMS subject classifications.** 68P25, 68Q99

**DOI.** 10.1137/15M1050963

---

\*Received by the editors December 2, 2015; accepted for publication (in revised form) March 12, 2018; published electronically June 26, 2018. This paper is based on results from two works, [BGL<sup>+</sup>15] and [CHJV15], preliminary versions of which appeared in the proceedings of ACM STOC 2015. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.  
<http://www.siam.org/journals/sicomp/47-3/M105096.html>

**Funding:** The work of the first author was supported by the Alon Young Faculty Fellowship and by Len Blavatnik and the Blavatnik Family foundation. The work of the second author was supported by the Check Point Institute for Information Security, ISF grant 1523/14, and NSF Frontier grants CNS1413920 and 1218461. The work of the fourth and ninth authors was supported by NSF grants CNS-1350619 and CNS-1414119, a DARPA Safeware grant, a Alfred P. Sloan Research Fellowship, a Microsoft Faculty Fellowship, the NEC Corporation, and a Steven and Renee Finn Career Development Chair from MIT. The work of the sixth author was supported in part by NSF grants CNS-1528178 and CNS-1514526. The work of the seventh and eighth authors was supported in part by an Alfred P. Sloan Fellowship, a Microsoft New Faculty Fellowship, NSF award CNS-1217821, NSF CAREER award CCF-0746990, NSF award CCF-1214844, AFOSR YIP award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2-0211.

<sup>†</sup>School of Computer Science, Tel Aviv University, Tel Aviv, 69978 Israel (nirbitan@tau.ac.il). Member of the Check Point Institute of Information Security. Part of this research was done while at IBM.

<sup>‡</sup>Department of Computer Science, Boston University, Boston, MA 02215, and School of Computer Science, Tel-Aviv University, Tel-Aviv, 69978 Israel (canetti@bu.edu).

<sup>§</sup>Department of Electrical Engineering and Computer Science, University of California at Berkeley, Berkeley, CA 94720 (sanjam@berkeley.edu).

<sup>¶</sup>Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA 02142 (holmgren@csail.mit.edu, vinodv@csail.mit.edu).

<sup>||</sup>Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218 (abhishek@cs.jhu.edu).

<sup>#</sup>Department of Computer Science, University of California at Santa Barbara, Santa Barbara, CA 93106 (rachel.lin@cs.ucsb.edu).

<sup>††</sup>Department of Computer Science, Cornell University, Ithaca, NY 14853 (rafael@cs.cornell.edu, sidtelang@cs.cornell.edu).

**1. Introduction.** The ability to cryptographically obfuscate general programs holds great prospects for securing the future digital world [BGI<sup>+</sup>01]. However, general-purpose obfuscation mechanisms prior to this work were inherently inefficient. One of the main sources of inefficiency is the fact that existing solutions consider different models of computation from those used to write modern computer programs. Indeed, the first candidate indistinguishability obfuscator (**iO**) of Garg et al. [GGH<sup>+</sup>13] and most general-purpose obfuscators that followed shortly after were designed for Boolean circuits. Using a circuit obfuscator for different models of computation such as *Turing machines* or *random access machines* (RAMs) requires first translating the given program into a circuit. This comes with clear loss in efficiency aspects: the succinct description of the program (as Turing machine or RAM) is not preserved and now becomes proportional to its running time on the worst-case input. Similarly, executing the program on any specific input now requires the same time and space resources as the worst-case input, while the *input-specific* resources required by the original program could potentially be much smaller.

Assuming the existence of succinct noninteractive arguments and circuit obfuscators that satisfy a notion of obfuscation called differing-input obfuscation, Boyle, Chung, and Pass [BCP14] and Ananth et al. [AJS17] show how to construct Turing machine obfuscators—namely, obfuscators that take a Turing machine as input and produce an obfuscated Turing machine whose runtime and space requirements per input, as well as their size, are proportional to those of the input machine (up to a polynomial factor in the security parameter). However, differing-input obfuscation is a significantly stronger security requirement than **iO**; in fact some strong variants of differing-input obfuscation are believed to be impossible [GGHW13].

Furthermore, while preferable to circuits in terms of succinct representation and input-specific usage of resources, Turing machines still do not achieve the full gains of RAM programs, which for some tasks may be much faster due to their ability to randomly access memory. To apply any of the above obfuscators to a RAM program, one has to first translate the program to a circuit or a Turing machine, which loses these advantages.

### 1.1. This work: **iO** for RAMs through succinct randomized encodings.

In this work, we show how to directly obfuscate *RAM programs with bounded input and space*. That is, our obfuscator  $\mathcal{O}$  takes as input a RAM program  $\Pi$ , input length  $n$ , and memory bound  $s$ . It outputs another RAM program  $\mathcal{O}(\Pi)$  such that  $\mathcal{O}(\Pi)$  has the same functionality as that of  $\Pi$  on all inputs  $x$  of length up to  $n$  and such that the execution of  $\Pi$  on  $x$  uses space at most  $s$ . The running time of  $\mathcal{O}(\Pi)$  per each specific input is the same as that of  $\Pi$ , and its space usage is proportional to  $s$  (all up to polynomial factors in the security parameter and input bound  $n$ ). The running time of the obfuscator, and in particular the size of the obfuscated  $\mathcal{O}(\Pi)$ , is quasi-linear in that the description  $\Pi$  and the space bound  $s$ .

We show security of the construction assuming one-way functions and **iO** that are subexponentially secure. These assumptions are somewhat better understood and “more straightforward” than differing-input obfuscation. In particular, it is easier to demonstrate that an obfuscator is not subexponentially secure **iO** (simply demonstrate a pair of programs along with subexponential-time distinguisher) than it is to demonstrate that it is not a differing-input obfuscator.

The central piece of our construction and conceptual heart of this paper is the notion of a *succinct randomized encoding*. We then extend these schemes to full-fledged obfuscators in a relatively straightforward way. We thus start by reviewing

randomized encodings and its history next.

**Garbling and randomized encodings.** Garbling [Yao86, Rog91] and randomized encoding [IK00] schemes are a pillar of cryptographic protocol design, with numerous applications such as secure two-party and multiparty computation, verifiable delegation schemes, and one-time programs, to name a few (see [Rog91, App11b] for a survey).

Such schemes partition the process of evaluating a function  $f$  on input  $x$  into two stages: a randomized *encoding* (or, *garbling*) stage, which results in an encoded version  $\widehat{f(x)}$  of the computation, and a *decoding stage*, where the output  $f(x)$  can be recovered from the encoding  $\widehat{f(x)}$ . The key property shared by all such schemes is that the encoding  $\widehat{f(x)}$  should hide all information regarding the computation except for the output  $f(x)$ . This *privacy* requirement is naturally captured by an efficient simulator  $\text{Sim}(f(x))$ , which, given only the output  $f(x)$ , produces a simulated encoding (perfectly, statistically, or computationally) indistinguishable from  $\widehat{f(x)}$ .

Whereas privacy alone could be trivially achieved by having the encoding be  $f(x)$ , garbling and randomized encoding schemes feature additional properties that (in conjunction with privacy) turn them into a powerful tool. These additional properties may vary. The seminal work of Yao on *garbled circuits* [Yao86] leverages the ability to encode  $(f, x)$  with low input-locality (each output bit of the encoding depends on at most one bit of  $x$ ). Another property of Yao's solution often used in applications is that encoding the input can be done separately from encoding the function and independently of the function's complexity (this is sometimes known as separability or decomposability).<sup>1</sup> The work of Ishai, Kushilevitz, and Applebaum on *randomized encodings* [IK00, IK02, AIK04, AIK06] asks more generally, how much simpler can the encoding procedure be compared to direct computation of  $f(x)$ ?

Capturing what it means to “simplify the computation of  $f(x)$ ” may take different forms according to the complexity measure of interest. The work of Ishai, Kushilevitz, and Applebaum focuses on computing the randomized encoding  $\widehat{f(x)}$  with lower parallel-time complexity than required for computing the original function  $f$ . This line of work has been extremely fruitful both in terms of constructions and applications (again, see [App11b] for a survey).

**Succinct randomized encodings.** In this work, we address another natural complexity measure: *the time required to compute  $\widehat{f(x)}$* . Specifically, we require that, given the description of  $f$  and the input  $x$ , we can compute the encoding  $\widehat{f(x)}$  in time  $\widehat{T}$  that is significantly smaller than the time  $T$  required to compute  $f(x)$ . Decoding time, in contrast, would be as large as  $T$ , with some tolerable overhead. For this goal to be achievable,  $f$  has to be given in some succinct representation that is smaller than  $T$ , and cannot be given by, say, a size- $T$  circuit. With the goal of obtaining obfuscation for RAM machines, we shall focus on the case that  $f$  is represented by a RAM machine  $\Pi$ . As a step toward this, we shall also consider the simpler case of Turing machines, which will already capture much of the challenge.

We will also be interested in the complexity of the decoding process. Ideally, this decoding should not be much more complicated than the original, nonencoded computation. In particular, decoding the result of a given RAM computation should

<sup>1</sup>Accordingly, the term “garbling” is usually identified with schemes that have such separate input encoding, whereas “randomized encodings” refer to schemes where the input is encoded together with the function.

amount to executing this computation (in the RAM model).

**The existence of succinct randomized encodings.** Under commonly believed complexity-theoretic assumptions, perfectly private randomized encodings for all of  $\mathbf{P}$  are unlikely to be computable in fixed polynomial time.<sup>2</sup> In contrast, restricting our attention to privacy against computationally bounded adversaries, no lower bounds or barriers are known. Accordingly, we will restrict our attention to computational privacy—a simulated encoding  $\text{Sim}(f(x))$  is computationally indistinguishable from a real encoding  $\widehat{f(x)}$ .

We note that  $\mathbf{iO}$  for RAMs (respectively, Turing machines) would directly imply corresponding succinct randomized encodings.

**Our contributions.** We first show that subexponentially secure succinct randomized encodings for RAMs imply  $\mathbf{iO}$  for RAMs with similar complexity parameters.

**INFORMAL THEOREM 1.1** ( $\mathbf{iO}$  for RAMs from succinct randomized encodings). *Assume there exist one-way functions, (nonsuccinct)  $\mathbf{iO}$  for  $\mathbf{P/poly}$ , and succinct randomized encodings for RAM programs, all subexponentially secure. Then there exists  $\mathbf{iO}$  for RAM programs with inputs bounded by  $n$ . The time to obfuscate a program  $\Pi$  and the size of the obfuscated program are proportional to the time of encoding  $(\Pi, x)$  for inputs  $x$  of length  $n$ . The time and space requirements of the obfuscated program, per input, are proportional to those of the decoding process.*

The theorem is in a sense the succinct analog of previous bootstrapping theorems [App14, CLTV15] that show how (nonsuccinct) randomized encodings and pseudo-random functions in  $\mathbf{NC}^1$ , together with obfuscation for  $\mathbf{NC}^1$  circuits, imply obfuscation for  $\mathbf{P/poly}$ . Here, through succinct randomized encodings, we reduce  $\mathbf{iO}$  for arbitrarily long computations to  $\mathbf{iO}$  for circuits of fixed polynomial size. Indeed, the proof of this theorem follows closely ideas from there. We note that this theorem is rather general; in particular it is used by [KLW15, CH16, CCC<sup>+</sup>16] to obtain fully succinct  $\mathbf{iO}$  for Turing and RAM machines.

Our core technical contribution consists of two succinct randomized encoding schemes for RAM programs. Our schemes rely on (nonsuccinct)  $\mathbf{iO}$  for *circuits* and incur encoding overhead that is polynomial in  $s(n)$ , the worst-case space requirement of the input machine, but is essentially independent of its time complexity. Decoding amounts to a RAM computation with the same complexity as the original computation, up to polynomial factors in the security parameter.

**INFORMAL THEOREM 1.2** (succinct randomized encodings from circuit  $\mathbf{iO}$ ). *Assume the existence of  $\mathbf{iO}$  for  $\mathbf{P/poly}$  and one-way functions. Then there exists a succinct randomized encoding for all polynomial-time RAM  $\Pi$  with bounded space complexity  $s(n)$ . The time to encode is quasi-linear in the size of  $\Pi$ , input and output lengths  $n, m$ , and the space bound  $s(n)$ . The time to decode  $\widehat{\Pi(x)}$  is polynomial in the size of  $\Pi$ , and quasi-linear in  $s(n)$  and the time  $T$  for evaluating  $\Pi(x)$ .*

<sup>2</sup>Specifically, it can be shown that, for a language  $\mathcal{L}$ , recognized by a given  $T(n)$ -time machine  $\Pi$ , succinct randomized encodings with perfect-privacy computable in time  $t(n) \ll T(n)$  would imply that  $\mathcal{L}$  has 2-message interactive proofs with an  $O(t(n))$ -time verifier, which already suggests that  $t(n)$  should at least depend on the space (or depth) of the computation. Furthermore, under commonly believed derandomization assumptions (used to show that  $\mathbf{AM} \subseteq \mathbf{NP}$  [Kv99, MV99]), the above would imply that  $\mathcal{L}$  can be nondeterministically decided in time  $\text{poly}(t(n))$  for some fixed polynomial  $\text{poly}$ . Thus, any speedup in encoding would imply related speedup by nondeterminism, whereas significant speedup is believed to be unlikely.

We give two alternative proofs of Theorem 1.2, by way of two quite different randomized encoding schemes. The first scheme starts with a relatively simple succinct randomized encoding scheme for Turing machines, based on Yao’s garbled circuit method. The idea is to avoid direct garbling of the machine, and instead obfuscate the circuit that generates the garbled machine. The evaluator then runs this circuit, obtains the garbled machine, and then runs it to obtain the desired output value. Here we use the strong locality properties of Yao’s garbling scheme to make sure that the obfuscated garbling circuit is small. We then show how to extend it to RAMs based on any (nonsuccinct) RAM garbling scheme [LO13, GHL<sup>+</sup>14, GLOS15]. An additional feature of this construction is the following. Assuming puncturable pseudo-random functions in  $\text{NC}^1$  (known based on various hardness assumptions, such as the hardness of the learning with errors problem [BLMR13]), and restricting our attention to any class of computations with a-priori-bounded running time  $t(n)$ , it is enough to assume existence of nonsuccinct  $\mathbf{iO}$  for circuits in  $\text{NC}^1$  with input size  $O(\log(t(n)))$ . This appears to be a qualitatively weaker assumption than  $\mathbf{iO}$  for arbitrary circuits. In particular, it is a  $\text{poly}(t(n))$ -time falsifiable assumption in the terminology of [Nao03].

The second scheme takes a more direct approach. Rather than obfuscating the circuit that garbles the machine, we directly obfuscate the circuit that implements the basic computational step of the machine itself. That is, we obfuscate the circuit that reads the contents of the appropriate cell on the machine’s tape, writes a new value to the cell, moves the head to a new tape location, and updates the machine’s state according to its transition function. To make sure that the computation remains meaningful and that the evaluator does not learn the intermediate values that the machine writes on its tape, we provide a mechanism for the obfuscated machine to encrypt and authenticate the contents of the tape. Also here we first describe the scheme for the simpler case of Turing machines. We then extend the scheme to the case of RAM machines; here we also provide a mechanism for hiding the pattern of memory accesses of the machine. This construction approach is more “low level” and more technically involved than the first one. Still its directness provides some additional power; in particular it was extended in subsequent work to obtain fully succinct randomized encoding of Turing machines, RAM machines, and RAM machines with persistent memory [KLW15, CH16, CCC<sup>+</sup>16, CCHR15, ACC<sup>+</sup>15] (see further details in section 4.4).

We note that both constructions also satisfy the more general promise of a garbling scheme where encodings of the input and of the program can be done separately. More detailed overviews of the two constructions can be found within.

**Dependence on the output length.** As stated above, the size of our basic randomized encodings (or  $\mathbf{iO}$ ) grows with the output of the underlying computation. Such dependence can be easily shown to be inherent as long as we require simulation-based security (using a standard incompressibility argument). Nevertheless, we show that this dependence can be removed if we settle for a weaker indistinguishability-based guarantee saying that randomized encodings of two computations leading to the same output are indistinguishable.

**1.2. Additional applications of succinct randomized encodings.** We demonstrate the power of succinct randomized encodings in several additional applications, some new, and some analogous to previous applications of randomized encodings, but with new succinctness properties.

**Succinct functional encryption and reusable garbling.** Recent advancement in the study of obfuscation has brought with it a corresponding advancement in functional encryption (FE). Today, (indistinguishability-based) functional encryption for all circuits can be constructed from  $\mathbf{iO}$  [GGH<sup>+</sup>13, Wat14], or even from concrete (and efficiently falsifiable) assumptions on composite order multilinear graded-encodings [GGHZ16]. For models of computation with succinct representations, we may hope to have *succinct FE*, where a secret key  $\text{sk}_\Pi$ , allowing us to decrypt  $\Pi(x)$  from an encryption of  $x$ , can be computed faster than the running time of  $\Pi$ . However, here the state of the art was similar to succinct randomized encodings, or succinct  $\mathbf{iO}$ , requiring essentially the same strong (nonfalsifiable) assumptions.

One can replace  $\mathbf{iO}$  for circuits in the above FE constructions, with the succinct  $\mathbf{iO}$  from Theorem 1.1, and obtain FE where computing  $\text{sk}_\Pi$  is comparable to (succinctly) obfuscating  $\Pi$ . This, however, will require the same subexponential hardness of  $\mathbf{iO}$  for circuits. Based on existing nonsuccinct functional encryption schemes, we show that succinct FE can be constructed without relying on subexponentially hard primitives.

As observed in previous work [GHRW14, CIJ<sup>+</sup>13, GKP<sup>+</sup>13], FE (even indistinguishability based) directly implies an enhanced version of randomized encodings known as *reusable garbling*. Here reusability means that an encoding consists of two parts: The first part,  $\hat{\Pi}$ , is independent of any specific input and depends only on the machine  $\Pi$ .  $\hat{\Pi}$  can then be “reused” together with the second part,  $\hat{x}$ , encoding any input  $x$ . We get succinct reusable garbling for space-bounded computations: encoding  $\Pi$  depends on the space, but is done once; subsequent input encodings depend only on the input size  $n$  and not on space.

**Publicly verifiable delegation and succinct NIZKs.** Succinct randomized encodings directly imply a one-round delegation scheme for polynomial-time computations with bounded space complexity. A main feature of the scheme is *public-verifiability*, meaning that given the verifier’s message  $\sigma$  anyone can verify the proof  $\pi$  from the prover, without requiring any secret verification state. Previous schemes based on standard assumptions could only achieve this goal in an amortized sense (allowing an expensive preprocessing step) via reusable garbled circuits or RAMs [GKP<sup>+</sup>13, GHRW14]. Other schemes relied on strong knowledge assumptions [BCC<sup>+</sup>17, BCCT12, DFH12, BCCT13] or proven secure only in the random oracle model [Mic00].<sup>3</sup> Another prominent feature of the scheme is that it guarantees input privacy for the verifier. While this can generically be guaranteed with fully homomorphic encryption, the generic solution requires the prover to convert the computation into a circuit, which could incur quadratic blowup; in our solution, the complexity of the prover corresponds to decoding complexity, which could be made quasi-linear. See further discussion below.

The delegation scheme is based only on randomized encodings (and one-way functions) and thus, as explained above, can be based only on polynomial assumptions. Assuming also  $\mathbf{iO}$ , we can make the verifier’s message reusable; namely, the verifier can publish his message  $\sigma$  once and for all, and then get noninteractive proofs for multiple computations. Our transformation for reusing the verifier’s message is, in fact, a generic one that can be applied to any delegation scheme, including privately verifiable schemes (e.g., [KRR14]). For privately verifiable schemes, the transformation has an additional advantage: it removes what is known as the *verifier rejection problem*; specifically, in the transformed scheme, soundness holds even against provers

<sup>3</sup>Notably, in the setting of private-verification, Kalai, Raz, and Rothblum give a solution based on the subexponential learning with errors assumption [KRR14].

with a verification oracle.

The delegation schemes mentioned above are restricted to  $\mathbf{P}$ . Plugging our succinct  $\mathbf{iO}$  into the perfect noninteractive zero-knowledge (NIZK) arguments of Sahai and Waters [SW14] directly yields a construction of perfect succinct NIZK for bounded space  $\mathbf{NP}$  from  $\mathbf{iO}$  for  $\mathbf{P/poly}$  and one-way functions that are both sub-exponentially secure. The NIZK has a succinct common reference string whose size is independent of the time required to verify the  $\mathbf{NP}$  statement to be proven, and depends only on the space and on the size of the input and witness (verification time depends only on the length of the statement, as in [SW14]).

**Other applications.** We reexamine previous applications of (nonsuccinct) randomized encodings and note the resulting succinctness features.

One application is to multiparty computation [IK00, IK02], where we can reduce the communication overhead from depending on the circuit size required to compute a multiparty function  $f(x_1, \dots, x_m)$  to depending on the space required to compute  $f$ , which can be much smaller. (In this solution, the parties simply compute  $\hat{f}$  instead of  $f$  itself.) By now, beyond its conceptual simplicity, this solution does not give any advantage over recent one-round solutions based on multikey fully homomorphic encryption [Gen09, AJL<sup>+</sup>12, LTV12, GGHR14, MW16, DHRW16].

Another application is the amplification of *key-dependent message* (KDM) security. In KDM encryption schemes, semantic security needs to hold, even when the adversary obtains encryptions of functions of the secret key taken from a certain class  $\mathcal{F}$ . Applebaum [App11a] shows that any scheme that is KDM secure with respect to some class of functions  $\mathcal{F}$  can be made resilient to a bigger class  $\mathcal{F}' \supseteq \mathcal{F}$  if functions in  $\mathcal{F}'$  can be randomly encoded in  $\mathcal{F}$ . Our succinct randomized encodings will essentially imply that KDM security for circuits of any fixed polynomial size  $s(\cdot)$  (such as the scheme of [BHHI10]) can be amplified to KDM security for functions that can be computed by programs with space  $S \ll s(n)$ , but could potentially have larger running time.

**Obfuscation with quasi-polynomial blowup and a new bootstrapping theorem.** As an application of our technical approach (rather than a direct application of our succinct randomized encodings), we show a new bootstrapping theorem for  $\mathbf{iO}$  for circuits. This transformation shows that obfuscation for an arbitrary circuit  $C$  can be reduced to obfuscation of  $O(|C|)$  circuits whose size depends only on the security parameter and input length (the circuits themselves are also simple and are dominated by a constant number of applications of a pseudorandom function). In particular, the size of the resulting obfuscated  $C$  is  $\tilde{O}(|C|)$ .

**1.3. Subsequent work.** In a beautiful subsequent work, Koppula, Lewko, and Waters [KLW15] construct fully succinct randomized encodings for Turing machines from  $\mathbf{iO}$ . They overcome the “space barrier” by introducing a clever “selective enforcement mechanism.” This was extended by Canetti and Holmgren [CH16] and independently by Chen et al. [CCC<sup>+</sup>16] to produce a fully succinct randomized encoding for RAM machines. Canetti et al. [CCHR15] and independently Ananth et al. [ACC<sup>+</sup>15] subsequently extended this scheme to provide security even when RAM machines to be garbled are chosen adaptively after an initial database is encoded. Ananth, Jain, and Sahai [AJS15] gave a direct construction of succinct  $\mathbf{iO}$  for Turing machines that avoids invoking Theorem 1.1 and achieves a constant multiplicative size overhead.



**1.4. Road map.** Most of the remainder of the paper is dedicated to our constructions of succinct randomized encodings, covered in sections 3 and 4. More detailed technical overviews of the approaches can be found in the beginning of the respective sections. Section 5 covers the different applications of succinct randomized encodings. The required background and preliminaries are given in section 2.

**2. Preliminaries.** Let  $\mathbb{N}$  denote the set of positive integers, and let  $[n]$  denote the set  $\{1, 2, \dots, n\}$ . We write PPT as shorthand for probabilistic polynomial time Turing machines. The term *negligible* is used for denoting functions that are (asymptotically) smaller than any inverse polynomial. More precisely, a function  $\nu : \mathbb{N} \rightarrow \mathbb{R}$  is called negligible if, for every constant  $c > 0$  and all sufficiently large  $n$ , it holds that  $\nu(n) < n^{-c}$ .

**2.1. Models of computation.** In this work we will consider different models of computation. Below we define formally different classes of algorithms; we will start by defining classes of deterministic algorithms of fixed polynomial size, and then move to define classes of randomized algorithms and classes of algorithms of arbitrary polynomial size.

**Classes of deterministic algorithms of fixed polynomial size.**

**POLYNOMIAL-TIME CIRCUITS.** For every polynomial  $D$ , the class  $\text{CIR}[D] = \{\mathcal{C}_\lambda\}$  includes all deterministic circuits of size at most  $D(\lambda)$ .

**NC<sup>1</sup> CIRCUITS.** For every constant  $c$  and polynomial  $D$ , the class  $\text{NC}_c^1[D] = \{\mathcal{C}_\lambda\}$  of polynomial-sized circuits of depth  $c \log \lambda$  includes all deterministic circuits of size  $D(\lambda)$  and depth at most  $c \log \lambda$ .

**EXPONENTIAL-TIME TURING MACHINES.** We consider a canonical representation of Turing machines  $M = (M', n, m, S, T)$ , where the bit length of  $n, m, S, T$  is  $\lambda$  and  $n, m \leq S \leq T$ ;  $M$  takes input  $x$  of length  $n$ , runs  $M'(x)$  using  $S$  space for at most  $T$  steps, and finally outputs the first  $m$  bits of the output of  $M'$ . (If  $M'(x)$  does not halt in time  $T$  or if  $M'$  requires more than  $S$  space, then  $M$  outputs  $\perp$ .) In other words, given the description  $M$  of a Turing machine in this representation, one can efficiently read off its bound parameters denoted by  $(M.n, M.m, M.S, M.T)$ .

Now we define the class of exponential-time Turing machines. For every polynomial  $D$ , the class  $\text{TM}[D] = \{\mathcal{M}_\lambda\}$  includes all deterministic Turing machines  $\Pi_M$  containing the canonical representation of a Turing machine  $M$  of size  $D(\lambda)$ ;  $\Pi_M(x, t)$  takes input  $x$  and  $t$  of length  $M.n$  and  $\lambda$ , respectively, runs  $M(x)$  for  $t$  steps, and finally outputs what  $M$  returns.

*Remark.* Note that machine  $\Pi_M(x, t)$  on any input terminates in  $t < 2^\lambda$ , and hence its output is well-defined. Furthermore, we say that two Turing machines  $M_1$  and  $M_2$  have the same functionality if and only if they produce identical outputs and run for the same number of steps for every input  $x$ . This property is utilized when defining and constructing indistinguishability obfuscation for Turing machines, as in previous work [BCP14].

**EXPONENTIAL-TIME RAMS.** We consider a canonical representation of RAM machines  $R = (R', n, m, S, T)$  identical to the canonical representation of Turing machines above.

For every polynomial  $D$ , the class  $\text{RAM}[D] = \{\mathcal{R}_\lambda\}$  of polynomial-sized RAM machines include all deterministic RAM machines  $\Pi_R$ , defined as  $\Pi_M$  above for Turing machines, except that the Turing machine  $M$  is replaced with a RAM machine  $R$ .

**Classes of randomized algorithms.** The above-defined classes contain only deterministic algorithms. We define analogously these classes for their corresponding randomized algorithms. Letting  $\mathcal{X}[D]$  be any class defined above, we denote by  $\text{r}\mathcal{X}[D]$  the corresponding class of randomized algorithms. For example,  $\text{rCIR}[D]$  denotes all randomized circuits of size  $D(\lambda)$ , and  $\text{rTM}[D]$  denotes all randomized Turing machines of size  $D(\lambda)$ .

**Classes of (arbitrary) polynomial-sized algorithms.** The above-defined classes consist of algorithms of a fixed polynomial  $D$  description size. We define corresponding classes of arbitrary polynomial size. Letting  $\mathcal{X}[D]$  be any class defined above, we simply denote by  $\mathcal{X} = \cup_{\text{poly } D} \mathcal{X}[D]$  the corresponding class of algorithms of arbitrary polynomial size. For instance, CIR and rCIR denotes all deterministic and randomized polynomial-sized circuits, and TM denotes all polynomial-sized Turing machines.

In the rest of the paper, when we write a family of algorithms  $\{AL_\lambda\} \in \mathcal{X}$ , we mean  $\{AL_\lambda\} \in \mathcal{X}[D]$  for some polynomial  $D$ . This means the size of the family of algorithms is bounded by some polynomial. Below, for convenience of notation, when  $\mathcal{X}$  is a class of algorithms of arbitrary polynomial size, we write  $AL \in \mathcal{X}_\lambda$  as shorthand for  $\{AL_\lambda\} \in \{\mathcal{X}_\lambda\}$ .

**Classes of well-formed algorithms.** In the rest of the preliminary, we define various cryptographic primitives. In order to avoid repeating the definitions for different classes of machines, we provide definitions for general classes of algorithms  $\{\mathcal{AL}_\lambda\}$  that can be instantiated with specific classes defined above. In particular, we will work with classes of algorithms that are *well formed*, satisfying the following properties:

1. For every  $AL \in \mathcal{AL}_\lambda$ , and input  $x$ ,  $AL$  on input  $x$  terminates in  $2^\lambda$  steps. Note that this also implies that  $AL$  has bounded input and output lengths.
2. The size of every ensemble of algorithms  $\{AL_\lambda\} \in \{\mathcal{AL}_\lambda\}$  is bounded by some polynomial  $D$  in  $\lambda$ .
3. Given the description of an algorithm  $AL \in \mathcal{AL}_\lambda$ , one can efficiently read off the bound parameters  $AL.n, AL.m, AL.S, AL.T$ .

All of the above-defined algorithm classes are well formed. Below, we denote by  $T_{AL}(x)$  the running time of  $AL$  on input  $x$ , and by  $T_{AL}$  the worst-case running time of  $AL$ . Note that well-formed algorithm classes are not necessarily efficient; for instance, the class of polynomial-sized Turing machines TM contain Turing machines that run for exponential time. In order to define cryptographic primitives for only polynomial-time algorithms, we will use the notation  $\text{ALG}^T = \{\mathcal{AL}_\lambda^T\}$  to denote the class of algorithms in  $\text{ALG} = \{\mathcal{AL}_\lambda\}$  that run in time  $T(\lambda)$  (in particular, these with  $AL_\lambda.T < T(\lambda)$ ).

In the rest of the paper, all algorithm classes are well formed.

**2.2. Garbling schemes.** In this section, we define garbling schemes, following in most part the definitions in [BHR12b]. As explained in the introduction, the main difference between garbling schemes and randomized encodings is that garbling schemes allow the input to be encoded separately from the program. This slight strengthening is not needed in our applications, which are described in section 5.

**DEFINITION 2.1** (garbling scheme). *A garbling scheme for a class of (well-formed) deterministic algorithms  $\{\mathcal{AL}_\lambda\}_{\lambda \in \mathbb{N}}$  is a tuple  $\mathcal{GS} = (\text{Garb}, \text{Encode}, \text{Eval})$  satisfying the following properties:*

**SYNTAX:** *For every  $\lambda \in \mathbb{N}$ ,  $AL \in \mathcal{AL}_\lambda$ , and input  $x$ ,*

- **Garb** is probabilistic and on input  $(1^\lambda, AL)$  outputs a pair  $(\widehat{AL}, \mathbf{key})$ ;<sup>4</sup>
- **Encode** is deterministic and on input  $(\mathbf{key}, x)$  outputs  $\hat{x}$ ;
- **Eval** is deterministic and on input  $(\widehat{AL}, \hat{x})$  produced by **Garb**, **Encode** outputs  $y$ .

**CORRECTNESS:** For every polynomial  $T$  and every family of algorithms  $\{AL_\lambda\} \in \{\mathcal{AL}_\lambda^T\}$  and sequence of inputs  $\{x_\lambda\}$ , there exists a negligible function  $\mu$ , such that, for every  $\lambda \in \mathbb{N}$ ,

$$\Pr[\text{Eval}(\widehat{AL}, \hat{x}) \neq AL_\lambda(x_\lambda)] \leq \mu(\lambda)$$

in the probability space defined by sampling

$$\begin{aligned} (\widehat{AL}, \mathbf{key}) &\stackrel{\$}{\leftarrow} \text{Garb}(1^\lambda, AL_\lambda), \\ \hat{x} &\stackrel{\$}{\leftarrow} \text{Encode}(\mathbf{key}, x_\lambda). \end{aligned}$$

**DEFINITION 2.2** (security of a garbling scheme). We say that a garbling scheme  $\mathcal{GS}$  for a class of deterministic algorithms  $\{\mathcal{AL}_\lambda\}_{\lambda \in \mathbb{N}}$  is secure if the following holds.

There is a uniform machine **Sim** such that for every nonuniform PPT distinguisher  $\mathcal{D}$ , polynomial  $T'$ , sequence of algorithms  $\{AL_\lambda\} \in \{\mathcal{AL}_\lambda^{T'}\}$ , and sequence  $\{x_\lambda\}$  of inputs where  $x_\lambda \in \{0, 1\}^{(AL_\lambda)^{T'}.n}$ , there exists a negligible function  $\mu$ , such that, for every  $\lambda \in \mathbb{N}$  the following holds:

$$|\Pr[\mathcal{D}(\widehat{AL}, \hat{x}) = 1] - \Pr[\mathcal{D}(\widetilde{AL}, \tilde{x}) = 1]| \leq \mu(\lambda)$$

in the probability space defined by sampling

$$\begin{aligned} (\widehat{AL}, \mathbf{key}) &\stackrel{\$}{\leftarrow} \text{Garb}(1^\lambda, AL_\lambda), \\ \hat{x} &\stackrel{\$}{\leftarrow} \text{Encode}(\mathbf{key}, x), \\ (\widetilde{AL}, \tilde{x}) &\stackrel{\$}{\leftarrow} \text{Sim}(1^\lambda, |x|, |AL_\lambda|, (n, m, S, T), T_{AL_\lambda}(x), AL_\lambda(x)), \end{aligned}$$

where  $n, m, S$ , and  $T$  are, respectively, shorthand for  $AL.n, AL.m, AL.S$ , and  $AL.T$ , and **Sim** runs in time  $\text{poly}(\lambda, T)$ . The function  $\mu$  is called the distinguishing gap.

Furthermore, we say that  $\mathcal{GS}$  is  $\delta$ -indistinguishable if the above security condition holds with a distinguishing gap  $\mu$  bounded by  $\delta$ . Especially,  $\mathcal{GS}$  is subexponentially indistinguishable if  $\mu(\lambda)$  is bounded by  $2^{-\lambda^\epsilon}$  for a constant  $\epsilon$ .

We note that the subexponentially indistinguishability defined above is weaker than usual subexponential hardness assumptions in that the distinguishing gap only need to be small for PPT distinguisher, rather than subexponential time distinguisher.

We remark that in the above definition simulator **Sim** receives many inputs, meaning that a garbled pair  $(\widehat{AL}, \hat{x})$  reveals nothing but the following: the output  $AL(x)$ , instance running time  $T_{AL}(x)$ , input length  $|x|$ , and machine size  $|AL|$ , together with various parameters  $(n, m, S, T)$  of  $AL$ . We note that the leakage of the instance running time is necessary in order to achieve instance-based efficiency (see efficiency guarantees below). The leakage of  $|AL|$  can be avoided by padding machines if an upper bound on their size is known. The leakage of parameters  $(n, m, S, T)$  can be avoided by setting them to  $2^\lambda$ ; see Remark 2.14 for more details. In particular, when the algorithms are circuits, inputs to the simulation algorithm can be simplified to  $(1^\lambda, |x|, |C|, AL(x))$ , since all bound parameters  $n, m, S, T$  can be set to  $2^\lambda$ .

<sup>4</sup>Note that as the algorithm class is well formed, **Garb** implicitly has all bound parameters of  $AL$ .

**Efficiency guarantees.** We proceed to describe the efficiency requirements for garbling schemes. When considering only circuit classes, all algorithms **Garb**, **Encode**, **Eval** should be polynomial-time machines; that is, the complexity of **Garb**, **Eval** scales with the size of the circuit  $|C|$ , and that of **Encode** with the input length  $|x|$ . However, when considering general algorithm classes, since the description size  $|AL|$  could be much smaller than the running time  $AL.T$ , or even other parameters  $AL.S$ ,  $AL.n$ ,  $AL.m$ , there could be different variants of efficiency guarantees, depending on what parameters the complexity of the algorithms depends on. Below we define different variants.

**DEFINITION 2.3** (different levels of efficiency of garbling schemes). *We say that a garbling scheme  $\mathcal{GS}$  for a class of deterministic algorithms  $\{\mathcal{AL}_\lambda\}_{\lambda \in \mathbb{N}}$  has succinctness or I/O- / space- / time-dependent complexity if the following holds.*

**OPTIMAL EFFICIENCY:** *There exist universal polynomials  $p_{\text{Garb}}$ ,  $p_{\text{Encode}}$ , and  $p_{\text{Eval}}$  such that for every  $\lambda \in \mathbb{N}$ ,  $AL \in \mathcal{AL}_\lambda$ , and input  $x \in \{0, 1\}^{AL.n}$*

- $(\widehat{AL}, \text{key}) \stackrel{s}{\leftarrow} \text{Garb}(1^\lambda, AL)$  runs in time  $p_{\text{Garb}}(\lambda, |AL|, AL.m)$ ;<sup>5</sup>
- $\hat{x} = \text{Encode}(\text{key}, x)$  runs in time  $p_{\text{Encode}}(\lambda, |x|, AL.m)$ ; and
- $y = \text{Eval}(\widehat{AL}, \hat{x})$  runs in time  $p_{\text{Eval}}(\lambda, |AL|, |x|, AL.m) \times T_{AL}(x)$ , with overwhelming probability over the random coins of **Garb**. We note that **Eval** has instance-based efficiency.

**I/O-DEPENDENT COMPLEXITY:** *A relaxed version of optimal efficiency, in which the polynomials  $p_{\text{Garb}}$ ,  $p_{\text{Encode}}$ , and  $p_{\text{Eval}}$  are additionally functions of  $AL.n$  and  $AL.m$ .*

**SPACE-DEPENDENT COMPLEXITY:** *A relaxed version of optimal efficiency, in which the polynomials  $p_{\text{Garb}}$ ,  $p_{\text{Encode}}$ , and  $p_{\text{Eval}}$  are additionally functions of  $AL.S$ .*

**LINEAR-TIME-DEPENDENT COMPLEXITY:** *A relaxed version of optimal efficiency, in which the polynomials  $p_{\text{Garb}}$  and  $p_{\text{Encode}}$  additionally depend (quasi-)linearly on  $AL.T$ , and the running time of **Eval** is bounded by  $AL.T \cdot p_{\text{Eval}}(\lambda, |AL|, |x|)$ .*

*Furthermore, we say that the garbling scheme  $\mathcal{GS}$  has succinct input encodings if the encoding algorithm  $\text{Encode}(\text{key}, x)$  runs in time  $p_{\text{Encode}}(1^\lambda, |x|)$ .*

We note that the above four succinctness requirements are ordered from the strongest one (optimal) to the weakest one (time-dependent). We say that a garbling scheme is *succinct* if its complexity depends only polylogarithmically on the time bound. Thus a scheme with space-dependent complexity is succinct for a class of algorithms whose space usage is bounded by a fixed polynomial.

**On output dependence.** Note that in the optimal efficiency defined above, the complexity of the algorithms depends on the length of their respective inputs and the bound on their output lengths  $AL.m$ . We argue that this is necessary as long as we require simulation-based security. This follows from a standard incompressibility argument. Indeed, assume the existence of a pseudorandom generator  $G$ , and consider the encoding of  $G$  and a random input seed  $s$ . We claim that the size of the garbled function  $\widehat{G}$  or encoded input  $\hat{s}$  must be as large as the output  $|G(s)|$ . Otherwise, the efficient simulator can “compress” random strings, as it cannot distinguish the actual output  $G(s)$  from a uniformly random string.

The dependence on the output size can be eliminated if we settle for a weaker notion of indistinguishability security, meaning that the garbling of two (equal-length)

<sup>5</sup>Note that the running time of **Garb**, and similarly other algorithms that takes  $AL$  as an input, implicitly depends logarithmically on the time bound of  $AL$ , as its description contains the time bound  $AL.T$ .

program-input pairs  $(\widehat{AL}_0, \widehat{x}_0), (\widehat{AL}_1, \widehat{x}_1)$  are computationally indistinguishable, provided that  $AL_0(x_0)$  and  $AL_1(x_1)$  output the same result after a similar number of steps. In section 5.2, we show how this can be achieved assuming **iO**.

**Static vs. adaptive security.** Throughout this work, we consider statically secure garbling schemes; that is, the privacy guarantees hold only when the entire computation  $(AL, x)$  to be garbled is chosen statically. In the literature, stronger privacy guarantees have been considered [BHR12a, BHK13], allowing the input  $x$  to be chosen maliciously and adaptively depending on the garbled  $\widehat{AL}$ .

We leave open the question of constructing succinct adaptively secure garbling schemes.

**Garbling schemes for specific algorithm classes.** Next we instantiate the above general definition of a garbling scheme with concrete classes.

**DEFINITION 2.4** (garbling scheme for polynomial-sized circuits). *A triplet of algorithms  $\mathcal{GS}_{\text{CIR}} = (\text{Garb}_{\text{CIR}}, \text{Encode}_{\text{CIR}}, \text{Eval}_{\text{CIR}})$  is a garbling scheme (with linear-time-dependent complexity) for polynomial-sized circuits if it is a garbling scheme for class CIR (with linear-time-dependent complexity).*

We note that in the case of circuits, succinctness means the complexity scales polynomially in  $|C|$ , whereas linear-time-dependency means the complexity scales linearly with  $|C|$ .

**DEFINITION 2.5** (garbling schemes for polynomial-time Turing machines). *A triplet  $\mathcal{GS}_{\text{TM}} = (\text{Garb}_{\text{TM}}, \text{Encode}_{\text{TM}}, \text{Eval}_{\text{TM}})$  of algorithms is a garbling scheme with optimal efficiency or I/O- / space- / linear-time-dependent complexity (and succinct input encodings) for Turing machines if it is a garbling scheme for class TM, with the same level of efficiency.*

Different efficiency requirements impose qualitatively different restrictions. In this work, we will construct a garbling scheme for Turing machines with space-dependent complexity assuming indistinguishability obfuscation for circuits. The construction of a garbling scheme from **iO** for Turing machines, sketched in the introduction, has I/O-dependent complexity. On the other hand, we show that a scheme with I/O-independent complexity is impossible; in particular, the complexity of the scheme must scale with the bound on the output length.

**DEFINITION 2.6** (garbling schemes for polynomial-time RAM machines). *A triplet  $\mathcal{GS}_{\text{RAM}} = (\text{Garb}_{\text{RAM}}, \text{Encode}_{\text{RAM}}, \text{Eval}_{\text{RAM}})$  of algorithms is a garbling scheme for polynomial-time RAM machines with optimal efficiency or I/O- / space- / linear-time-dependent-complexity (and succinct input encodings) if it is a garbling scheme for class RAM, with the same level of efficiency.*

Recently, the works [LO13, GHL<sup>+</sup>14, GLOS15] provided construction of a garbling scheme for RAM machines with linear-time-dependent complexity and succinct input encodings, assuming only one-way functions.

**Garbled circuits with independent input encoding.** In this work, we will make use of a garbling scheme for circuits with a special structural property. In Definition 2.4, the key **key** for garbling inputs is generated depending on the circuit (by  $\text{Garb}(1^\lambda, C)$ ); the special property of a circuit garbling scheme is that the key can be generated depending only on the length of the input  $1^{|x|}$  and the security parameter, which implies that the garbled inputs  $\widehat{x}$  can also be generated depending only on the plain input  $x$  and the security parameter  $\lambda$ , independently of the circuit—we call this

independent input encoding.

DEFINITION 2.7 (garbling scheme for circuits with independent input encoding). A garbling scheme  $\mathcal{GS} = (\text{Garb}, \text{Encode}, \text{Eval})$  for a deterministic circuit class  $\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$  has independent input encoding if the following holds: For every  $\lambda \in \mathbb{N}$ , and every  $C \in \mathcal{C}_\lambda$

- the algorithm  $\text{Garb}$  on input  $(1^\lambda, C)$  invokes first  $\mathbf{key} \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda, 1^{|x|})$  and then  $\tilde{C} \stackrel{\$}{\leftarrow} \text{Gb}(\mathbf{key}, C)$ , where  $\text{Gen}$  and  $\text{Gb}$  are all PPT algorithms;
- the security condition holds with respect to a simulator  $\text{Sim}$  that on input  $(1^\lambda, 1^{|x|}, 1^{|C|}, T_C(x), C(x))$  outputs  $(\tilde{x}, \tilde{C})$  that are obtained by sampling

$$\begin{aligned} (\tilde{x}, \mathbf{st}) &\stackrel{\$}{\leftarrow} \text{Sim.Gen}(1^\lambda, |x|), \\ \tilde{C} &\stackrel{\$}{\leftarrow} \text{Sim.Gb}((1^\lambda, |x|, |C|, C(x), \mathbf{st}), \end{aligned}$$

where  $\text{Sim.Gen}$  and  $\text{Sim.Gb}$  run in times  $\text{poly}(\lambda, |x|)$  and  $\text{poly}(\lambda, |C|)$ , respectively.

It is easy to check that many known circuit garbling schemes, in particular the construction by Yao [Yao86], have independent input encoding.

PROPOSITION 2.8. Assume the existence of one-way functions that are hard to invert in  $\Gamma$  time. Then there exists a garbling scheme  $\mathcal{GS}_{\text{CIR}}$  for polynomial-sized circuits with independent input encoding that is  $\Gamma^{-\varepsilon}$ -indistinguishable for some constant  $\varepsilon \in (0, 1)$ .

**2.3. Indistinguishability obfuscation.** We recall the definition of indistinguishability obfuscation, adapting it to arbitrary classes of polynomial-time algorithms. As before, we first define the syntax, correctness and security of  $\mathbf{iO}$  and then discuss different efficiency guarantees.

DEFINITION 2.9 (indistinguishability obfuscator ( $i\mathcal{O}$ )). A uniform machine  $i\mathcal{O}$  is an indistinguishability obfuscator for a class of deterministic algorithms  $\{\mathcal{AL}_\lambda\}_{\lambda \in \mathbb{N}}$  if the following conditions are satisfied.

CORRECTNESS: For all security parameters  $\lambda \in \mathbb{N}$ , for all  $AL \in \mathcal{AL}_\lambda$ , and for all input  $x$ , we have that

$$\Pr[AL' \leftarrow i\mathcal{O}(1^\lambda, AL) : AL'(x) = AL(x)] = 1.$$

SECURITY: For every polynomial  $T$ , every nonuniform PPT sampleable distribution ensemble  $\mathcal{D} = \{\mathcal{D}_\lambda\}$  over  $\{\mathcal{AL}_\lambda^T \times \mathcal{AL}_\lambda^T \times \{0, 1\}^{\text{poly}(\lambda)}\}$ , and adversary  $\mathcal{A}$ , the following hold:

If when sampling  $(AL_1, AL_2, z) \stackrel{\$}{\leftarrow} \mathcal{D}(1^\lambda)$ , it holds with probability 1 that

$$\begin{aligned} \forall x, AL_1(x) &= AL_2(x), \\ T_{AL_1}(x) &= T_{AL_2}(x), \end{aligned}$$

and

$$(|AL_1|, (AL_1).n, (AL_1).m, (AL_1).S) = (|AL_2|, (AL_2).n, (AL_2).m, (AL_2).S),$$

then

$$\left| \Pr[\mathcal{A}(AL'_1, z) = 1] - \Pr[\mathcal{A}(AL'_2, z) = 1] \right| \leq \mu(\lambda)$$

in the probability space defined by sampling

$$\begin{aligned} (AL_1, AL_2, z) &\stackrel{\$}{\leftarrow} \mathcal{D}(1^\lambda), \\ AL'_1 &\stackrel{\$}{\leftarrow} i\mathcal{O}(1^\lambda, AL_1), \\ AL'_2 &\stackrel{\$}{\leftarrow} i\mathcal{O}(1^\lambda, AL_2), \end{aligned}$$

where  $\mu$  is called the distinguishing gap for  $\mathcal{D}$  and  $\mathcal{A}$ .

Furthermore, we say that  $i\mathcal{O}$  is  $\delta$ -indistinguishable if the above security condition holds with a distinguishing gap  $\mu$  bounded by  $\delta$ . We say that  $i\mathcal{O}$  is subexponentially indistinguishable if  $\mu(\lambda)$  is bounded by  $2^{-\lambda^\epsilon}$  for a constant  $\epsilon$ .

Note that in the security guarantee above, the distribution  $\mathcal{D}$  samples algorithms  $AL_1$  and  $AL_2$  that have the same functionality and matching bound parameters. This means that an obfuscated machine “reveals” the functionality (as desired) and these bound parameters. We remark that the leakage of the latter is without loss of generality: In the case of circuits, all bound parameters are set to  $2^\lambda$ . In the case of other algorithm classes, say Turing and RAM machines, if an  $i\mathbf{O}$  scheme must ensure that  $AL.S$  or  $AL.T$  are not revealed, one can simply pad them to  $2^\lambda$ ; then the above security definition automatically ensures privacy of that parameter. See Remark 2.14 for more details.

**DEFINITION 2.10** (different levels of efficiency of IO). *We say that an indistinguishability obfuscator  $i\mathcal{O}$  of a class of algorithms  $\{\mathcal{AL}_\lambda\}$  has optimal efficiency if there is a universal polynomial  $p$  such that for every  $\lambda \in \mathbb{N}$  and every  $AL \in \mathcal{AL}_\lambda$ ,  $i\mathcal{O}(1^\lambda, AL)$  runs in time  $p(\lambda, |AL|)$ .*

*Additionally, we say  $i\mathcal{O}$  has input- / space- / linear-time- dependent complexity if  $i\mathcal{O}(1^\lambda, AL)$  runs in time  $\text{poly}(\lambda, |AL|, AL.n) / \text{poly}(\lambda, |AL|, AL.S) / \text{poly}(\lambda, |AL|) AL.T$ .*

We note that, unlike the case of garbling schemes, the optimal efficiency of an  $i\mathbf{O}$  scheme does not need to depend on the length of the output. Loosely speaking, this stems from the fact that indistinguishability-based security does not require “programming” outputs, which is the case in simulation-based security for garbling.

**$i\mathbf{O}$  for specific algorithm classes.** We recall the definition of  $i\mathbf{O}$  for polynomial-sized circuits,  $\text{NC}^1$  [BGI<sup>+</sup>01], and give definitions of  $i\mathbf{O}$  for polynomial-time Turing machines [BCP14] and RAM machines with different efficiency guarantees.

**DEFINITION 2.11** (indistinguishability obfuscator for polynomial-sized circuits and for  $\text{NC}^1$ ). *A uniform PPT machine  $i\mathcal{O}_{\text{CIR}}(\cdot, \cdot)$  is an indistinguishability obfuscator for polynomial-sized circuits if it is an indistinguishability obfuscator for CIR with optimal efficiency.*

*A uniform PPT machine  $i\mathcal{O}_{\text{NC}^1}(\cdot, \cdot, \cdot)$  is an indistinguishability obfuscator for  $\text{NC}^1$  circuits if for all constants  $c \in \mathbb{N}$ ,  $i\mathcal{O}_{\text{NC}^1}(c, \cdot, \cdot)$  is an indistinguishability obfuscator for  $\text{NC}_c$  with optimal efficiency.*

**DEFINITION 2.12** ( $i\mathbf{O}$  for Turing machines). *A uniform machine  $i\mathcal{O}_{\text{TM}}(\cdot, \cdot)$  is an indistinguishability obfuscator for polynomial-time Turing machines, with optimal efficiency or input- / space-dependent complexity, if it is an indistinguishability obfuscator for the class TM with the same efficiency.*

Recently, the works [BCP14, AJS17] gave constructions of  $i\mathbf{O}$  for Turing machines<sup>6</sup> with input-dependent complexity assuming fully homomorphic encryption

<sup>6</sup>Their works actually realize the stronger notion of differing-input, or extractability, obfuscation for Turing machines.

(FHE), differing-input obfuscation for circuits, and P-certificates [CLP13]; furthermore, the dependency on input length can be removed—leading to a scheme with optimal efficiency—if assuming succinct noninteractive argument of knowledge (*SNARK*) instead of P-certificates.

**DEFINITION 2.13 (iO for RAM machines).** *A uniform machine  $i\mathcal{O}_{\text{TM}}(\cdot, \cdot)$  is an indistinguishability obfuscator for polynomial-time RAM machines, with optimal efficiency or linear-time-dependent complexity, if it is an indistinguishability obfuscator for the class RAM with the same efficiency.*

*Remark 2.14 (explicit vs. implicit bound parameters).* In the above definitions of garbling schemes and **iO** for general algorithms, we considered a canonical representation of algorithms  $AL$  that gives information of various bound parameters of the algorithm, specifically the size  $|AL|$ , bound on input and output lengths  $AL.n$ ,  $AL.m$ , space complexity  $AL.S$ , and time complexity  $AL.T$ . This representation allows us to define, in a *unified* way, different garbling and **iO** schemes that depend on different subsets of parameters. The following are some examples of such schemes:

- The garbling and **iO** schemes for TM that we construct in section 3 and subsection 5.2 (from **iO** and subexp **iO** for circuits, respectively) have complexity  $\text{poly}(|AL|, AL.S, \log(AL.T))$ . (In particular, the size of the garbled Turing machine and obfuscated Turing machine is of this order.)
- The garbling scheme for TM (from **iO** for TM) sketched in the introduction has complexity  $\text{poly}(|AL|, AL.n, AL.m, \log(AL.T))$ .
- The garbling scheme for RAM from one-way functions by [LO13, GHL<sup>+</sup>14, GLOS15] has complexity that scales polynomially in  $(|AL|, AL.n, AL.m)$  and quasi-linearly in  $AL.T$ . This construction leads to an **iO** for RAM (from subexp **iO** for circuits) of the same complexity as in 5.2.

By using the canonical representation, our general definition allows the garbling or **iO** scheme to depend on any subset of parameters flexibly. Naturally, if a scheme depends on a subset of parameters, the resulting garbled or obfuscated machines may “leak” these parameters (in the three examples above, the size of the garbled or obfuscated machines leaks the parameters they depend on); thus, the security definitions must reflect this “leakage” correspondingly. The general security definitions, Definitions 2.2 and 2.9, capture this by allowing leakage of all parameters  $|AL|, AL.n, AL.m, AL.S, AL.T$ . However, this seems to “overshoot,” since if a specific scheme does not depend on a particular parameter (e.g.,  $AL.S$ ), then this parameter should be kept private. This can be easily achieved by simply considering an algorithm representation that always sets that parameter to  $2^\lambda$  (e.g.  $AL.S = 2^\lambda$ ).

**2.4. Puncturable pseudorandom functions.** We recall the definition of puncturable pseudorandom functions (PRF) from [SW14]. Since in this work we only use puncturing at one point, the definition below is restricted to puncturing only at one point instead of at polynomially many points.

**DEFINITION 2.15 (puncturable PRFs).** *A puncturable family of PRFs is given by a triple of uniform PPT machines  $(\text{PRF.Gen}, \text{PRF.Punc}, F)$  and a pair of computable functions  $n(\cdot)$  and  $m(\cdot)$ , satisfying the following conditions:*

- CORRECTNESS.** *For all outputs  $K$  of  $\text{PRF.Gen}(1^\lambda)$ , all points  $i \in \{0, 1\}^{n(\lambda)}$ , and  $K(-i) = \text{PRF.Punc}(K, i)$ , we have that  $F(K(-i), x) = F(K, x)$  for all  $x \neq i$ .*
- PSEUDORANDOM AT PUNCTURED POINT.** *For every PPT adversary  $(\mathcal{A}_1, \mathcal{A}_2)$ , there is a negligible function  $\mu$ , such that in an experiment where  $\mathcal{A}_1(1^\lambda)$  outputs a point  $i \in \{0, 1\}^{n(\lambda)}$  and a state  $\sigma$ ,  $K \stackrel{\$}{\leftarrow} \text{PRF.Gen}(1^\lambda)$  and  $K(i) =$*



PRF.Punc( $K, i$ ), the following holds:

$$|\Pr[\mathcal{A}_2(\sigma, K(i), i, F(K, i)) = 1] - \Pr[\mathcal{A}_2(\sigma, K(i), i, U_{m(\lambda)}) = 1]| \leq \mu(\lambda),$$

where  $\mu$  is called the distinguishing gap for  $(\mathcal{A}_1, \mathcal{A}_2)$ .

Furthermore, we say that the puncturable PRF is  $\delta$ -indistinguishable if the above pseudorandom property holds with a distinguishing gap  $\mu$  bounded by  $\delta$ . In particular, the puncturable PRF is subexponentially indistinguishable if  $\mu(\lambda)$  is bounded by  $2^{-\lambda^\epsilon}$  for a constant  $\epsilon$ .

As observed by [BW13, BGI14, KPTZ13], the GGM construction of PRFs [GGM86] from pseudorandom generators (PRGs) yields puncturable PRFs for any (efficiently computable)  $n(\cdot)$  and  $m(\cdot)$ . Furthermore, it is easy to see that if the PRG underlying the GGM construction is subexponentially hard (and this can in turn be built from subexponentially hard one-way functions (OWFs)), then the resulting puncturable PRF is subexponentially pseudorandom.

**2.5. Injective noninteractive commitment.** An injective noninteractive bit commitment scheme is a pair of polynomials  $n(\cdot)$  and  $m(\cdot)$  and a polynomial-sized ensemble of injective functions  $\text{Commit}_\lambda : \{0, 1\} \times \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)}$  such that for all polynomial time adversaries  $\mathcal{A}$ ,

$$\Pr \left[ \mathcal{A}(1^\lambda, \text{Commit}_\lambda(b; r)) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ r \leftarrow \{0, 1\}^{n(\lambda)} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

The scheme is said to be subexponentially secure if the  $\text{negl}(\lambda)$  above can in fact be replaced by  $2^{-\lambda^\epsilon}$  for some constant  $\epsilon > 0$ .

We can construct a (subexponentially secure) injective noninteractive commitment scheme given a (subexponentially secure) injective one-way function

$$f : \{0, 1\}^{n'(\lambda)} \rightarrow \{0, 1\}^{m'(\lambda)},$$

and we give the construction here without proof. Let  $h$  be a hard-core bit of  $f$  (note that the standard construction of a one-way function with a hard-core bit preserves injectivity). Then define  $n(\lambda) = n'(\lambda)$ ,  $m(\lambda) = m'(\lambda) + 1$ , and

$$\text{Commit}_\lambda(b; r) = f(r) \| b \oplus h(r).$$

**3. Succinct garbling via garbled circuits and garbled RAMs.** In this section (based on [BGL<sup>+</sup>15]) we first show a succinct garbling scheme for Turing machines based on  $\mathbf{iO}$  for circuits and Yao's garbled circuit method. We then show how to extend it to RAMs based on circuit  $\mathbf{iO}$  and any (nonsuccinct) RAM garbling scheme.

**3.1. Overview.** We first provide an overview of our garbling scheme for the class of Turing machines  $\text{TM}$  with space-dependent complexity. In particular, we show the following theorem (formally proven in section 3.2).

**THEOREM 3.1.** *Assuming the existence of  $\mathbf{iO}$  for circuits and one-way functions, there exists a garbling scheme for  $\text{TM}$  with space-dependent complexity.*

Toward this goal, we proceed in two steps: In the first step, we construct a *non-succinct* garbling scheme for  $\text{TM}$ , which satisfies the correctness and security requirements of Definitions 2.1 and 2.2, except that the garbling and evaluation algorithms

can run in time polynomial in both the time and space complexity,  $M.T$  and  $M.S$ , of the Turing machine being garbled,  $M$  (as well as the simulation algorithm); the produced garbled Turing machine is of size in the same order. In contrast, the time to compute the encoded input  $\hat{x}$  depends on the length of the input  $x$ , but not on the typically larger running time of  $M$ . This feature of “independent input encoding” is crucial for our construction. Another crucial property of the nonsuccinct garbling scheme is that the garbled machine consists of many “small garbled pieces” that can be generated separately. In the second step, we use  $\mathbf{iO}$  to “compress” the size of the garbled program by providing an obfuscated program that takes an index as input and generates the “garbled piece” corresponding to that index. As a result, the final garbled program (namely, the obfuscated program) is small (i.e., depends on the space complexity  $M.S$ ). It is only at evaluation time that the underlying nonsuccinct garbled program is unraveled, by running the obfuscated program on every index, and decoded.

**Nonsuccinct garbling scheme.** We outline the *nonsuccinct garbling scheme* for Turing machines, based on any one-way function. A “trivial” approach toward such garbling is to simply transform any polynomial-time Turing machine into a circuit and then garble the circuit. While our construction in essence relies on this principle, it will in fact invoke garbling for “small” fixed-size circuits. Concretely, we rely on the existence of a circuit garbling scheme satisfying two additional properties. First, we require that the shared string  $\mathbf{key}$ , and thus also the input encoding, be generated independently of the circuit being garbled (e.g.,  $\mathbf{key}$  is sampled at random and given to both the input encoding and circuit garbling procedures). Second, we require that encoded inputs can be simulated, given only the input size, whereas the garbled program is simulated using the result  $M(x)$  of the computation (and the randomness used to simulate the encoded input). We refer to such schemes as *garbling schemes with independent input encoding* (Definition 2.7) and note that Yao’s basic scheme [Yao86] satisfies the two properties.

Our nonsuccinct garbling scheme now proceeds as follows. Let  $M$  be a Turing machine with bounded space complexity  $M.S$ , running-time  $M.T$ , and inputs of length  $M.n$ . We construct a “chain” of  $M.T$  garbled circuits that evaluate  $M$  step by step. More precisely, we first generate keys  $\mathbf{key}_1, \dots, \mathbf{key}_{M.T}$  for the  $M.T$  garbled circuits. The  $j$ th garbled circuit (which is computed using key  $\mathbf{key}_j$ ) takes as input some state of  $M$  and computes the next state (i.e., the state after one computation step); if the next state is a final state, it returns the output generated by  $M$ ; otherwise it outputs an *encoding* of this new state using the next key  $\mathbf{key}_{j+1}$ . (Note that after  $M.T$  steps we are guaranteed to get to a final state, and thus this process is well-defined.)

To encode the input, we simply encode the initial configuration of  $M$ , including the input  $x$ , using  $\mathbf{key}_1$ . To evaluate the garbled program, we sequentially evaluate each garbled circuit using the encodings generated in the previous one as inputs to the next one, and so on until the output is generated.

**Security of the nonsuccinct scheme.** To show that this construction is a secure (nonsuccinct) garbling scheme we need to exhibit a simulator that, given just the output  $y = M(x)$  of the program  $M$  on input  $x$  and *the number of steps  $t^*$  taken by  $M(x)$* , can simulate the encoded input and program. (The reason we provide the simulation with the number of steps  $t^*$  is that we desire a garbling scheme with a “per-instance efficiency”—that is, the evaluation time is polynomial in the actual running time  $t^*$  and not just the worst-case running time. To achieve such “per-instance efficiency” requires leaking the running time, which is why the simulator

gets access to it.) Toward this goal, we start by simulating the  $t^*$ th garbled circuit with the output being set to  $y$ ; this simulation generates a garbled program  $\tilde{M}_{t^*}$  and an encoded input  $\text{conf}_{t^*-1}$  (which are the simulated input keys).

We then iteratively, in descending order, simulate the  $j$ th garbled circuits  $\tilde{M}_j$  with the output being set to  $\text{conf}_{j+1}$  generated in the previously simulated garbled circuit. We finally simulate the remaining  $j > t^*$  garbled circuits  $\tilde{M}_j$  with the output being set to some arbitrary output in the range of the circuit (e.g., the output  $y$ ). The simulated encoded input is then  $\text{conf}_1$  and the simulated garbled program is  $(\tilde{M}_1, \dots, \tilde{M}_{M.T})$ .<sup>7</sup>

To prove indistinguishability of the simulated garbling and the real garbling, we consider a sequence of hybrid experiments  $H_0, \dots, H_{M.T}$ , where in  $H_j$  the first  $j$  garbled circuits are simulated, and the remaining  $M.T - j$  garbled circuits are honestly generated. To “stitch together” the simulated circuits with the honestly generated ones, the  $j$ th garbled circuit is simulated using as output an honest encoding  $\widehat{\text{conf}}_j$  of the actual configuration  $\text{conf}_j$  of the Turing machine  $M$  after  $j$  steps.

It follows from the security of the garbling scheme that hybrids  $H_j$  and  $H_{j+1}$  are indistinguishable and thus also  $H_0$  (i.e., the real experiment) and  $H_{M.T}$ .

Let us finally note a useful property of the above-mentioned simulation. Due to the fact that we rely on a garbling scheme with *independent input encoding*, each garbled circuit can in fact be *independently simulated*—recall that the independent input encoding property guarantees that encoded inputs can be simulated without knowledge of the circuit to be computed, and thus all simulated encoded inputs  $\text{conf}_1, \dots, \text{conf}_{M.T}$  can be generated in an initial step. Next, the garbled circuits can be simulated in any order.

**Succinct garbling scheme.** We now show how to make this garbling scheme succinct. The idea is simple: instead of providing the actual garbled circuits in the clear, we provide an obfuscation of the *randomized* program that generates these garbled circuits. More precisely, we provide an  $\mathbf{iO}$  of a program  $\mathbf{M}^{s,s'}(\cdot)$ , where  $s$  and  $s'$  are seeds for a PRF  $F: \mathbf{M}^{s,s'}(j)$ , given a “time step”  $j \in [M.T]$ , which generates the  $j$ th garbled circuit in the nonsuccinct garbling of  $M$  using pseudorandom coins generated by the PRF with seed  $s$  and  $s'$ . Specifically, it uses  $F(s, j)$  and  $F(s, j + 1)$  as randomness to generate  $\text{key}_j$  and  $\text{key}_{j+1}$  (recall that the functionality of the  $j$ th circuit depends on  $\text{key}_{j+1}$ ) and uses  $F(s', j)$  as randomness for garbling the  $j$ th circuit.

The succinct garbled program is then the obfuscated program  $\Lambda \stackrel{s}{\leftarrow} \mathbf{iO}(\mathbf{M}^{s,s'})$ , and the encoding  $\hat{x}$  of  $x$  remains the same as before, except that now it is generated using pseudorandom coins  $F(s, 1)$ . Given such a garbled pair  $\Lambda$  and  $\hat{x}$ , one can compute the output by gradually generating the nonsuccinct garbled program, *one garbled circuit at a time*, by computing  $\Lambda$  on every time step  $j$  and evaluating the produced garbled circuit with  $\hat{x}$  until the output is produced. (This way, the evaluation still has “per-instance efficiency.”)

**Security of the succinct scheme.** Given that the new succinct garbled program  $\Lambda$  produces “pieces” of the nonsuccinct garbled program, the natural idea for simulating the succinct garbled program is to obfuscate a program that produces “pieces” of the simulated nonsuccinct garbled program. The above-mentioned “independent simulation” property of the nonsuccinct garbling (following from independent input encoding) enables us to fulfill this idea.

<sup>7</sup>This “layered” simulation strategy resembles that of Applebaum, Ishai, and Kushilevitz in the context of arithmetic garbling [AIK11].

More precisely, given an output  $y$  and the running time  $t^*$  of  $M(x)$ , the simulator outputs the obfuscation  $\widetilde{\Lambda}$  of a program  $\widetilde{M}^{y,t^*,s,s'}$  that, given input  $\underline{j}$ , outputs a simulated  $j$ th garbled circuit, using randomness  $F(s, j+1)$  to generate  $\text{conf}_{j+1}$  as the output, and  $F(s, j)$  and  $F(s', j)$  as the extra randomness needed to simulate the input  $\text{conf}_j$  and the garbled  $M_j$ .<sup>8</sup> The encoded input  $\tilde{x}$  is simulated as in the nonsuccinct garbling scheme, but using pseudorandom coins  $F(s, 1)$ .

It is not hard to see that this simulation works if the obfuscation is virtually black-box secure, as (nonsuccinct) garbling security guarantees that the entire truth tables of the two programs  $M^{s,s'}$  and  $\widetilde{M}^{y,t^*,s,s'}$  are indistinguishable given an encoding of  $x$ , when the hardwired PRF keys  $s, s'$  are chosen at random. Our goal, however, is to show that **iO** suffices. Toward this goal, we consider a sequence of hybrid experiments  $H'_0, \dots, H'_{M.T}$  with a corresponding sequence of obfuscated programs  $\widetilde{M}_0^{s,s'}, \dots, \widetilde{M}_{M.T}^{s,s'}$  that “morph” gradually from the real  $M$  to the fully simulated  $\widetilde{M}$ . Specifically, the program  $\widetilde{M}_j^{s,s'}$  obfuscated in  $H'_j$  produces a nonsuccinct *hybrid* garbled program as in hybrid  $H_j$  in the proof of the nonsuccinct garbling scheme, except that pseudorandom coins are used instead of truly random coins. That is, for the first  $j$  inputs,  $\widetilde{M}_j$  produces simulated garbled circuits, and for the rest of the inputs, it produces honestly generated garbled circuits, having hardwired the true configuration  $\text{conf}_{j+1}$ .

To prove indistinguishability of any two consecutive hybrids  $H'_j$  and  $H'_{j+1}$ , we rely on the punctured program technique of Sahai and Waters [SW14] to replace pseudorandom coins  $F(s, j+1), F(s', j+1)$  for generating the  $j+1$ st simulated garbled circuit with truly random coins, and then rely on the indistinguishability of the simulation of the  $j+1$ st garbled circuit. A bit more concretely, at each step we puncture the seeds  $s, s'$  only on the (three) points corresponding to the  $j+1$ st step, and hardwire instead the corresponding outputs generated by  $\widetilde{M}_j^{s,s'}$ ; next, relying on the puncturing guarantee, we can sample these outputs using true independent randomness. At this point, we can already replace the real hardwired garbling with a simulated one. Finally, we go back to generating the hardwired value pseudorandomly as part of the circuit’s logic, now identical to  $\widetilde{M}_{j+1}^{s,s'}$ , and “unpuncture” the seeds  $s, s'$ . We note that each such step requires hardwiring a new (real) intermediate configuration  $\text{conf}_{j+1}$  (used to simulate the  $j+1$ st garbling), but now the previous hardwired configuration  $\text{conf}_j$  can be “forgotten” and blowup is avoided.

**iO for a simple class of circuits is enough.** The obfuscated circuits in the construction are of a special kind—their input size is  $O(\log M.T)$ . Canetti et al. [CLTV15] show that **iO** for  $\text{NC}^1$  can be bootstrapped to obtain **iO** for all circuits, assuming puncturable PRFs in  $\text{NC}^1$  [BLMR13], and incurring a security loss that is exponential in the size of the input. Accordingly, for polynomial  $M.T$ , it suffices to assume (polynomially secure) **iO** for classes in  $\text{NC}^1$  with logarithmic-size inputs.

**A garbling scheme for any bounded space computation.** We observe that the above approach for constructing a succinct garbling scheme for bounded space Turing machines applies generally to any *bounded space computation* (e.g., bounded space RAM). In particular, we show the following theorem (formally proven in section 3.3).

**THEOREM 3.2.** *Assume the existence of **iO** for circuits and one-way functions. There exists a garbling scheme for any abstract model of sequential computation, such*

<sup>8</sup>Recall that simulating a garbled circuit requires both the output and the randomness for simulating the input encoding.

as TM and RAM, with space-dependent complexity.

The succinct garbling scheme described above for Turing machines does not apply uniquely to Turing machines, but rather to any model of computation that can be divided into sequential steps using one memory, for instance, RAMs. Given an underlying circuit garbling scheme  $\mathcal{GS} = (\text{Garb}, \text{Encode}, \text{Eval})$  with independent input encoding, to construct a garbling scheme  $\mathcal{GS}^A$  for  $\{\mathcal{AL}_\lambda\}$ , proceed with the following two steps.

- Step 1. Construct a nonsuccinct garbling scheme.* Observe that the computation of a machine  $AL$  of  $AL.T$  steps can be divided into  $AL.T$  1-step “blocks” that transforms the current configuration to the next; therefore, to garble  $AL$ , it suffices to produce a sequence of “garbled blocks,” one for each 1-step block. The actual programs being garbled is an “augmented block,” whose execution consists of a 1-step block followed by the encoding algorithm of  $\mathcal{GS}$  that encodes the output configuration for the next garbled block (when an output is produced, it is output directly without encoding). The final garbling then consists of a sequence of  $T$  garbled blocks.
- Step 2. Compress the size using  $\mathbf{iO}$ .* As before, we then use  $\mathbf{iO}$  to “compress” the size of the nonsuccinct garbling constructed in the first step by giving the obfuscation of the algorithm that on input  $t$  runs  $\text{Garb}$  to garble the  $t$ th augmented block, producing the  $t$ th garbled block. The obfuscated program is the succinct garbled program.

**Improved construction and analysis.** We also note that in the case of RAM, our construction can be made more efficient, relying on previous garbled RAM solutions [LO13, GHL<sup>+</sup>14, GLOS15]. In the solution described above, we can in fact replace the underlying circuit garbling scheme with any garbling scheme, as long as it admits independent input encoding. For instance, in the case where the program  $AL$  is a RAM, we may use previous *garbled RAM* solutions [LO13, GHL<sup>+</sup>14, GLOS15]. The benefit is that this allows optimizing the efficiency of our scheme. Indeed, in the solution described above, each step of the machine is translated to a garbled circuit of size  $O(AL.S)$  (up to polynomial factors in the security parameter), which means that the complexity of encoding is  $\text{poly}_{\mathbf{iO}}(AL.S)$ , where  $\text{poly}_{\mathbf{iO}}(\cdot)$  is the overhead due to obfuscation, and the complexity of decoding for a  $T$ -time computation  $AL(x)$  is at least  $T \cdot \text{poly}_{\mathbf{iO}}(AL.S)$ , which may be significantly larger than the original computation.

In contrast, known garbled RAM solutions provide a more efficient way of garbling RAMs than converting them into circuits, taking into consideration the RAM structure, and guaranteeing that encoding and decoding require essentially the same time and space as the original RAM computation. Aiming to leverage this efficiency in our solution, instead of partitioning a RAM computation into  $AL.T$  steps, each implemented by a circuit of size  $AL.S$ , we can partition it into  $AL.T/AL.S$  pieces, where each piece is an  $AL.S$ -step RAM. The encoding and decoding time for each piece are essentially linear in its running time  $O(AL.S)$  (whereas a circuit implementing any such piece might be of size  $\Omega(AL.S^2)$ ).

This modification on its own may still be insufficient; indeed, obfuscating the circuit that produces the garbled RAM may incur nonlinear overhead  $\text{poly}_{\mathbf{iO}}(\cdot)$ , so that eventually decoding may take time  $\text{poly}_{\mathbf{iO}}(AL.S) \cdot AL.T/AL.S$ , which may be again as large as  $AL.T \cdot AL.S$ . To circumvent this blowup, and as a result of independent interest, we show how to bootstrap any  $\mathbf{iO}$  for circuits to one that has quasi-linear blowup. Overall, in the new solution, for a  $T$ -time computation  $AL(x)$ , encoding

takes time  $\tilde{O}(AL.S)$  and decoding  $\widehat{AL}(x)$  takes time  $O(T + AL.S)$ .

**3.2. Succinct garbling for Turing machines via garbled circuits.** In this section we present our succinct garbling scheme for Turing machines. As explained in the overview, we first construct a nonsuccinct garbling scheme based on Yao’s garbled circuit construction. Next, we show how to rely on this nonsuccinct scheme and **iO** to construct a succinct garbling scheme.

**3.2.1. A nonsuccinct garbling scheme.** We now describe formally our non-succinct garbling scheme  $\mathcal{GS}_{ns} = (\text{Garb}_{ns}, \text{Encode}_{ns}, \text{Eval}_{ns})$ . We rely on a garbling scheme for polynomial-sized circuits with independent input encoding.

Let  $\mathcal{GS}_{\text{CIR}} = (\text{Garb}_{\text{CIR}}, \text{Encode}_{\text{CIR}}, \text{Eval}_{\text{CIR}})$  be a garbling scheme for polynomial-sized circuits, and  $\text{Sim}_{\text{CIR}}$  the simulation algorithm. We require  $\mathcal{GS}_{\text{CIR}}$  to have independent input encoding; that is,

$$\begin{aligned} \text{Garb}_{\text{CIR}} &= (\text{Gen}_{\text{CIR}}, \text{Gb}_{\text{CIR}}), \\ \text{Sim}_{\text{CIR}} &= (\text{Sim.Gen}_{\text{CIR}}, \text{Sim.Gb}_{\text{CIR}}), \end{aligned}$$

as described in Definition 2.7.

The execution of a Turing machine  $M$  consists of a sequence of steps in which each step  $t$  depends on the description of the machine  $M$  and its current configuration  $\text{conf}_t$ , and produces the next configuration  $\text{conf}_{t+1}$ . In the Turing machine model, each step takes constant time, independent of the size of the Turing machine and its configuration. However, each step can be implemented using a circuit  $\text{Next}^{D,S}$  that on input  $(M, \text{conf}_t)$  with  $|M| \leq D, |\text{conf}_t| \leq S$ , outputs the next configuration  $\text{conf}_{t+1}$ —we call this circuit the “universal next-step circuit.” The size of the circuit is a fixed polynomial  $p_{\text{Next}}$  in the size of the machine and the configuration, that is,  $p_{\text{Next}}(D, S)$ . The whole execution of  $M(x)$  can be carried out by performing at most  $M.T$  evaluations of  $\text{Next}^{D,S}(M, \cdot)$ , producing a chain of configurations denoted by

$$\text{CONFIG}(M, x) = (T^*, \text{conf}_1, \dots, \text{conf}_{M.T}, \text{conf}_{M.T+1}),$$

where  $T^* = T_M(x)$  and  $\text{conf}_1, \dots, \text{conf}_{T^*-1}, \text{conf}_{T^*}$  are the sequence of configurations of  $M(x)$  until it halts ( $\text{conf}_t$  is the configuration before the  $t$ th step starts).  $\text{conf}_{T^*}, \dots, \text{conf}_{M.T+1}$  are set for simplicity to the output  $y = M(x)$ .

We note that the initial configuration  $\text{conf}_1$  can be derived efficiently from  $x$ ;  $\text{conf}_{T^*}$  is called the final configuration, which can be efficiently recognized and from which an output  $y$  can be extracted efficiently.

For every  $\lambda$  and  $M \in \text{TM}_\lambda$ , our scheme proceeds as follows.

**THE GARBLING ALGORITHM  $\text{Garb}_{ns}(1^\lambda, M)$ :** Let  $S = M.S, T = M.T$ , and  $D = |M|$ .

Sample  $2T$  sufficiently long random strings  $\alpha_1, \dots, \alpha_t$  and  $\beta_1, \dots, \beta_t$ ; produce a chain of  $T$  garbled circuits using  $\text{Garb}_{\text{CIR}}$  by running the following program for every  $t \in [T]$ .

*Program  $\mathbf{P}^{\lambda,S,M}(t; (\alpha_t, \alpha_{t+1}, \beta_t))$ :*

1. *Generate the key  $\text{key}_{t+1}$  for the next garbled circuit:*  
 If  $t < T$ , compute the key for the  $t + 1$ st garbled circuit  $\text{key}_{t+1} = \text{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha_{t+1})$  using randomness  $\alpha_{t+1}$ . (Note that  $\text{key}_t$  is generated for inputs of length  $S$ .)
2. *Prepare the step-circuit  $\mathbf{C}_t$ :*  
 $\mathbf{C}_t$  on an  $S$ -bit input  $\text{conf}_t$  (i) computes  $\text{conf}_{t+1} = \text{Next}^{D,S}(M, \text{conf}_t)$ ;  
 (ii) if  $\text{conf}_{t+1}$  is a final configuration, simply outputs the output  $y$  con-

tained in it;<sup>9</sup> (iii) otherwise, translate  $\text{conf}_{t+1}$  to the garbled inputs of the  $t+1$ st garbled circuit by computing  $\widehat{\text{conf}}_{t+1} = \text{Encode}_{\text{CIR}}(\text{key}_{t+1}, \text{conf}_{t+1})$ .

3. *Garble the step-circuit  $\mathbf{C}_t$ :*

Compute the key using randomness  $\alpha_t$ ,  $\text{key}_t = \text{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha_t)$ , and garble  $\mathbf{C}_t$  using randomness  $\beta_t$ ,  $\widehat{\mathbf{C}}_t = \text{Gb}_{\text{CIR}}(\text{key}_t, \mathbf{C}_t; \beta_t)$ .

4. *Output  $\widehat{\mathbf{C}}_t$ .*

Generate **key** as follows: Compute the key for the first garbled circuit using randomness  $\alpha_1$ ,  $\text{key}_1 = \text{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha_1)$ ; set  $\mathbf{key} = \text{key}_1 \| 1^S$ .

Finally, output  $\widehat{M} = (\widehat{\mathbf{C}}_1, \dots, \widehat{\mathbf{C}}_T)$ , **key**.

THE ENCODING ALGORITHM  $\text{Encode}_{ns}(\mathbf{key}, x)$ : Let  $\text{conf}_1 \in \{0, 1\}^S$  be the initial configuration of  $M$  with input  $x$ ; compute  $\widehat{x} = \widehat{\text{conf}}_1 = \text{Encode}_{\text{CIR}}(\text{key}_1, \text{conf}_1)$ .

THE EVALUATION ALGORITHM  $\text{Eval}_{ns}(\widehat{M}, \widehat{x})$ : Evaluate the chain of garbled circuits  $\widehat{M} = (\widehat{\mathbf{C}}_1, \dots, \widehat{\mathbf{C}}_T)$  in sequence in  $T$  iterations: At iteration  $t$ , compute  $z = \text{Eval}_{\text{CIR}}(\widehat{\mathbf{C}}_t, \widehat{\text{conf}}_t)$ ; if  $z$  is the garbled input  $\widehat{\text{conf}}_{t+1}$  for the next garbled circuit  $\widehat{\mathbf{C}}_{t+1}$ , proceed to the next iteration; otherwise, terminate and output  $y = z$ .

Next, we proceed to show that  $\mathcal{GS}_{ns}$  is a nonsuccinct garbling scheme for TM.

**Efficiency.** We summarize the complexity of different algorithms of the nonsuccinct scheme. It is easy to see that for any Turing machine  $M$  with  $D = |M|$ ,  $S = M \cdot S$ , and  $T = M \cdot T$ , the garbling algorithm  $\text{Garb}_{ns}$  runs in time  $\text{poly}(\lambda, D, S) \times T$  and produces a garbling machine whose size is of the same order. Thus the garbling scheme is nonsuccinct. On the other hand, the encoding and evaluation algorithms  $\text{Encode}_{ns}$  and  $\text{Eval}_{ns}$  are all deterministic polynomial-time algorithms. Finally, the simulation runs in time  $\text{poly}(\lambda, D, S) \times T$ , as does the garbling algorithm.

**Correctness.** We show that for every polynomial  $T'$ , every sequence of algorithms  $\{M = M_\lambda\} \in \{\text{TM}_\lambda^{T'}\}$ , and sequence of inputs  $\{x = x_\lambda\}$ , where  $x_\lambda \in \{0, 1\}^{M \cdot n}$ , there exists a negligible function  $\mu$ , such that

$$\Pr[(\mathbf{key}, \widehat{M}) \stackrel{\$}{\leftarrow} \text{Garb}_{ns}(1^\lambda, M), \widehat{x} = \text{Encode}_{ns}(\mathbf{key}, x) : \text{Eval}_{ns}(\widehat{M}, \widehat{x}) \neq M(x)] \leq \mu(\lambda).$$

Let  $\text{CONFIG}(M, x) = (T^*, \text{conf}_1, \dots, \text{conf}_T, \text{conf}_{T+1})$  be the sequence of configurations generated in the computation of  $M(x)$ , where  $T \leq T'(\lambda)$ . It follows from the correctness of the circuit garbling scheme  $\text{Garb}_{\text{CIR}}$  that with overwhelming probability (over the randomness of  $\text{Garb}_{ns}$ ), the following is true: (1) for every  $t < T^*$ , the garbled circuit  $\widehat{\mathbf{C}}_t$ , if given the garbled input  $\widehat{\text{conf}}_t$  corresponding to  $\text{conf}_t$ , computes the correct garbled inputs  $\widehat{\text{conf}}_{t+1}$  corresponding to  $\text{conf}_{t+1}$ , and (2) for  $t = T^*$ , the garbled circuit  $\widehat{\mathbf{C}}_{T^*}$ , if given the garbled input  $\widehat{\text{conf}}_{T^*-1}$  corresponding to  $\text{conf}_{T^*-1}$ , produces the correct output  $y$ . (Note that the evaluation procedure terminates after  $T^*$  iterations, and circuits  $\widehat{\mathbf{C}}_t$  for  $t > T^*$  are never evaluated.) Then since the garbled input  $\widehat{x}$  equals the garbled initial configuration  $\widehat{\text{conf}}_1$ , by conditions (1) and (2), the evaluation procedure produces the correct output with overwhelming probability.

**Security.** Fix any polynomial  $T'$ , any sequence of algorithms  $\{M = M_\lambda\} \in \{\text{TM}_\lambda^{T'}\}$ , and any sequence of inputs  $\{x = x_\lambda\}$ , where  $x_\lambda \in \{0, 1\}^{M \cdot n}$ . To show the security of  $\mathcal{GS}_{ns}$ , we construct a simulation algorithm  $\text{Sim}_{ns}$  and show that the following two ensembles are indistinguishable: For convenience of notation, we suppress

<sup>9</sup>Pad  $y$  with 0 if it is not long enough.

the appearance of  $M.n$  and  $M.m$  as input to  $\text{Sim}$ :

$$(3.1) \quad \{\text{real}_{ns}(1^\lambda, M, x)\} = \left\{ (\hat{M}, \mathbf{key}) \stackrel{\$}{\leftarrow} \text{Garb}_{ns}(1^\lambda, M), \hat{x} = \text{Encode}_{ns}(\mathbf{key}, x) : (\hat{M}, \hat{x}) \right\}_\lambda,$$

$$(3.2) \quad \{\text{simu}_{ns}(1^\lambda, M, x)\} = \left\{ (\tilde{M}, \tilde{x}) \stackrel{\$}{\leftarrow} \text{Sim}_{ns}(1^\lambda, 1^{|x|}, 1^{|M|}, S, T, T_M(x), M(x)) : (\tilde{M}, \tilde{x}) \right\}_\lambda.$$

Below we describe the simulation algorithm. Observe that the garbled machine  $\hat{M}$  consists of  $T$  garbled circuits  $(\hat{\mathbf{C}}_1, \dots, \hat{\mathbf{C}}_T)$ , and the garbled input  $\hat{x}$  is simply the garbled input of the initial configuration  $\text{conf}_0$  (corresponding to  $x$ ) for the first garbled circuit  $\hat{\mathbf{C}}_1$ . Naturally, to simulate them, the algorithm  $\text{Sim}_{ns}$  needs to utilize the simulation algorithm  $\text{Sim}_{\text{CIR}} = (\text{Sim.Gen}_{\text{CIR}}, \text{Sim.Gb}_{\text{CIR}})$  of the circuit garbling scheme, which requires knowing the output of each garbled circuit. In a real evaluation with  $\hat{M}, \hat{x}$ , the output of the  $(T^*)$ th garbled circuit is  $y = M(x)$ , the output of the garbled circuits  $t < T^*$  is the garbled input  $\widehat{\text{conf}}_{t+1}$  for next garbled circuit  $t + 1$ , and the garbled circuits  $t > T^*$  are not evaluated, but for which  $y$  is a valid output. Thus, in the simulation, garbled circuits  $t = T^*, \dots, T$  can be simulated using output  $y$ , whereas garbled circuits  $t = 1, \dots, T^* - 1$  will be simulated using the *simulated garbled inputs* for circuit  $t + 1$ . More precisely, the simulation algorithm is as follows.

THE SIMULATION ALGORITHM  $\text{Sim}_{ns}(1^\lambda, 1^{|x|}, 1^{|M|}, S, T, T^* = T_M(x), y = M(x))$ :

Sample  $2T$  sufficiently long random strings  $\alpha_1, \dots, \alpha_T, \beta_1, \dots, \beta_T$ . Simulate the chain of garbled circuits by running the following program for every  $t \in [T]$ .

Program  $\mathbf{Q}^{\lambda, S, |M|, T^*, y}(t; (\alpha_t, \alpha_{t+1}, \beta_t))$ :

1. Prepare the output  $out_t$  for the  $t$ th simulated circuit  $\tilde{\mathbf{C}}_t$ :  
 If  $t \geq T^*$ ,  $out_t = y$ . Otherwise, if  $t < T^*$ , set the output as the garbled input for the next garbled circuits, that is,  $out_t = \widehat{\text{conf}}_{t+1}$  computed from  $(\widehat{\text{conf}}_{t+1}, \mathbf{st}_{t+1}) = \text{Sim.Gen}_{\text{CIR}}(1^\lambda, S; \alpha_{t+1})$  using randomness  $\alpha_{t+1}$ .
2. Simulate the  $t$ th step-circuit  $\tilde{\mathbf{C}}_t$ :  
 Given the output  $out_t$ , simulate the  $t$ th garbled circuit  $\tilde{\mathbf{C}}_t$  by computing first  $(\widehat{\text{conf}}_t, \mathbf{st}_t) = \text{Sim.Gen}_{\text{CIR}}(1^\lambda, S; \alpha_t)$  and then  $\tilde{\mathbf{C}}_t = \text{Sim.Gb}_{\text{CIR}}(1^\lambda, S, q, out_t, \mathbf{st}_t; \beta_t)$ , using randomness  $\alpha_t, \beta_t$ , where  $q = q(\lambda, S)$  is the size of the circuit  $\mathbf{C}_t$ .
3. Output  $\tilde{\mathbf{C}}_t$ .

Simulate the garbled input  $\tilde{x}$  by computing again  $(\widehat{\text{conf}}_1, \mathbf{st}_1) = \text{Sim.Gen}_{\text{CIR}}(1^\lambda, S; \alpha_1)$  using randomness  $\alpha_1$  and setting  $\tilde{x} = \widehat{\text{conf}}_1$ .

Finally, output  $(\tilde{M} = (\tilde{\mathbf{C}}_1, \dots, \tilde{\mathbf{C}}_T), \tilde{x})$ .

Toward showing the indistinguishability between honestly generated garbling  $(\hat{M}, \hat{x})$  and the simulation  $(\tilde{M}, \tilde{x})$ , we will consider a sequence of hybrids  $\text{hyb}_{ns}^0, \dots, \text{hyb}_{ns}^T$ , where  $\text{hyb}_{ns}^0$  samples  $(\hat{M}, \hat{x})$  honestly, while  $\text{hyb}_{ns}^T$  generates the simulated garbling  $(\tilde{M}, \tilde{x})$ . In every intermediate hybrid  $\text{hyb}_{ns}^\gamma$ , a hybrid simulator  $\text{HSim}_{ns}^\gamma$  is invoked, producing a pair  $(\tilde{M}_\gamma, \tilde{x}_\gamma)$ . At a high level, the  $\gamma$ th hybrid simulator on input  $(1^\lambda, M, x)$  simulates the first  $\gamma - 1$  garbled circuits using the program  $\mathbf{Q}$ , generates the last  $T - \gamma$  garbled circuits honestly using the program  $\mathbf{P}$ , and simulates the  $\gamma$ th garbled circuits using the program  $\mathbf{R}$  described below, which “stitches” together the first  $\gamma - 1$  simulated circuits with the last  $T - \gamma$  honest circuits into a chain that evaluates to the correct output. More precisely, we will denote by



COMBINE $[(P_1, S_1), \cdot, (P_\ell, S_\ell)]$  a merged circuit that on input  $x$  in the domain  $X$  computes  $P_j(x)$  if  $x \in S_j$ , where  $S_1, \dots, S_\ell$  is a partition of the domain  $X$ .

THE HYBRID SIMULATION ALGORITHM  $\text{HSim}_{ns}^\gamma(1^\lambda, M, x)$  FOR  $\gamma = 0, \dots, T$ :

Compute  $T^* = T_M(x)$  and  $y = M(x)$ , and the intermediate configuration  $\text{conf}_{\gamma+1}$  as defined by  $\text{CONFIG}(M, x)$ .

Sample  $2T$  sufficiently long random strings  $\{\alpha_t, \beta_t\}_{t \in [T]}$ . Simulate the chain of garbled circuits by running the following program for every  $t \in [T]$ , which combines programs  $\mathbf{P}$ ,  $\mathbf{Q}$ , and  $\mathbf{R}$  as below.

Program  $\mathbf{M}^\gamma = \text{COMBINE}[(\mathbf{Q}, [\gamma-1]), (\mathbf{R}, \{\gamma\}), (\mathbf{P}, [\gamma+1, T])](t; (\alpha_t, \alpha_{t+1}, \beta_t))$ :

- If  $t \leq \gamma - 1$ , compute  $\tilde{\mathbf{C}}_t = \mathbf{Q}^{\lambda, S, |M|, T^* \cdot y}(t; (\alpha_t, \alpha_{t+1}, \beta_t))$ ; output  $\tilde{\mathbf{C}}_t$ .
- If  $t \geq \gamma + 1$ , compute  $\tilde{\mathbf{C}}_t = \mathbf{P}^{\lambda, S, M}(t; (\alpha_t, \alpha_{t+1}, \beta_t))$ ; output  $\tilde{\mathbf{C}}_t$ .
- If  $t = \gamma$ , compute  $\tilde{\mathbf{C}}_t = \mathbf{R}^{\lambda, S, \text{conf}_{\gamma+1}}(\gamma; (\alpha_\gamma, \alpha_{\gamma+1}, \beta_\gamma))$  using program  $\mathbf{R}$  defined below; output  $\tilde{\mathbf{C}}_t$ .

Program  $\mathbf{R}^{\lambda, S, \text{conf}_{\gamma+1}}(\gamma; (\alpha_\gamma, \alpha_{\gamma+1}, \beta_\gamma))$ :

1. Prepare the output  $out_\gamma$  of the simulated  $\gamma$ th circuit  $\tilde{\mathbf{C}}_t$ :

Set the output  $out_\gamma$  to  $y$  if  $\text{conf}_{\gamma+1}$  is a final configuration. Otherwise, the output should be the garbled input corresponding to  $\text{conf}_{\gamma+1}$  for the next garbled circuit; since the  $\gamma + 1$ st circuit is generated honestly, we compute  $out_\gamma = \widehat{\text{conf}}_{\gamma+1}$  by first computing  $\text{key}_{\gamma+1} = \text{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha_{\gamma+1})$ , and then encoding  $\widehat{\text{conf}}_{\gamma+1} = \text{Encode}_{\text{CIR}}(\text{key}_{\gamma+1}, \text{conf}_{\gamma+1})$ .

(Note that the difference between program  $\mathbf{Q}$  and  $\mathbf{R}$  is that the former prepares the output  $out_\gamma$  using simulated garbled input  $\widehat{\text{conf}}_{t+1}$ , whereas the latter uses honestly generated garbled input  $\widehat{\text{conf}}_{\gamma+1}$ .)

2. Simulate the  $\gamma$ th circuit  $\tilde{\mathbf{C}}_t$ :

Given the output  $out_\gamma$ , simulate the  $\gamma$ th garbled circuit  $\tilde{\mathbf{C}}_\gamma$  by computing  $(\widehat{\text{conf}}_\gamma, \text{st}_\gamma) = \text{Sim.Gen}_{\text{CIR}}(1^\lambda, S; \alpha_\gamma)$  and  $\tilde{\mathbf{C}}_t = \text{Sim.Gb}_{\text{CIR}}(1^\lambda, S, q, out_\gamma, \text{st}_\gamma; \beta_\gamma)$ , where  $q = q(\lambda, S)$  is the size of the circuit  $\mathbf{C}_t$ .

3. Output  $\tilde{\mathbf{C}}_t$ .

If  $\gamma > 0$ , simulate the garbled input  $\tilde{x}_\gamma$  as  $\text{Sim}_{ns}$  does. Otherwise, if  $\gamma = 0$ , generate the garbled input  $\tilde{x}_0$  honestly as in  $\text{Garb}_{ns}$  and  $\text{Encode}_{ns}$ . Finally, output  $(\tilde{M}_\gamma = (\tilde{\mathbf{C}}_1, \dots, \tilde{\mathbf{C}}_\gamma, \tilde{\mathbf{C}}_{\gamma+1} \tilde{\mathbf{C}}_T), \tilde{x}_\gamma)$ .

We use notation  $\text{hyb}_{ns}^\gamma(1^\lambda, M, x)$  to denote the output distribution of the hybrid simulator  $\text{HSim}_{ns}^\gamma$ . By construction, in  $\text{HSim}_{ns}^\gamma$ , when  $\gamma = 0$ ,  $\mathbf{M}^0 = \mathbf{P}$  and the garbled input  $\tilde{x}_0$  is generated honestly, that is,  $\{\text{hyb}_{ns}^0(1^\lambda, M, x)\} = \{\text{real}_{ns}(1^\lambda, M, x)\}$  (where  $\text{real}_{ns}$  is the distribution of honestly generated garbling; see (3.1)); furthermore, when  $\gamma = T$ ,  $\mathbf{M}^T = \mathbf{Q}$  and the garbled input  $\tilde{x}_\gamma$  is simulated, that is,  $\{\text{hyb}_{ns}^\gamma(1^\lambda, M, x)\} = \{\text{simu}_{ns}(1^\lambda, M, x)\}$  (where  $\text{simu}_{ns}$  is the distribution of simulated garbling; see (3.2)). Thus to show the indistinguishability between  $\{\text{real}_{ns}(1^\lambda, M, x)\}$  and  $\{\text{simu}_{ns}(1^\lambda, M, x)\}$ , it suffices to show the following claim.

CLAIM 3.3. For every  $\gamma \in \mathbb{N}$ , the following holds

$$\{\text{hyb}_{ns}^{\gamma-1}(1^\lambda, M, x)\}_\lambda \approx \{\text{hyb}_{ns}^\gamma(1^\lambda, M, x)\}_\lambda.$$

*Proof.* Fix a  $\gamma \in \mathbb{N}$ , a sufficiently large  $\lambda \in \mathbb{N}$ , an  $M = M_\lambda$ , and a  $x = x_\lambda$ . The only difference between the garbling  $(\tilde{M}_{\gamma-1}, \tilde{x}_{\gamma-1})$  sampled by  $\text{hyb}_{ns}^{\gamma-1}(1^\lambda, M, x)$  and

the garbling  $(\tilde{M}_\gamma, \tilde{x}_\gamma)$  sampled by  $\text{hyb}_{ns}^\gamma(1^\lambda, M, x)$  is the following: Let  $\text{conf}_\gamma$  be the intermediate configuration at the beginning of step  $\gamma$ .

- In  $\text{hyb}_{ns}^{\gamma-1}$ , the  $\gamma$ th garbled circuit  $\tilde{\mathbf{C}}_\gamma$  is generated honestly using program  $\mathbf{P}$ . The circuit  $\mathbf{C}_\gamma$  (as described in algorithm  $\text{Garb}_{ns}$ ) is the composition of the circuit  $\text{Next}^{\lambda, S}(M, \cdot)$  and the encoding algorithm  $\text{Encode}_{\text{CIR}}(\text{key}_{\gamma+1}, \cdot)$ , where  $\text{key}_{\gamma+1} = \text{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha_{\gamma+1})$  is generated honestly. Furthermore, the first  $\gamma-1$  garbled circuits are simulated using  $\mathbf{R}$  and  $\mathbf{Q}$ . The simulation of the first  $\gamma-1$  circuits as well as the generation of the garbled input  $\tilde{x}_\gamma$  depend potentially on the garbled input  $\widehat{\text{conf}}_\gamma$  corresponding to  $\text{conf}_\gamma$  for  $\tilde{\mathbf{C}}_\gamma$  (when  $\text{conf}_\gamma$  is not a final configuration; see step 1 in  $\mathbf{R}$ ). In other words, the output of  $\text{hyb}_{ns}^{\gamma-1}$  can be generated by the following alternative sampling algorithm:
  - Generate garbled circuits  $\gamma+1, \dots, T$  honestly using program  $\mathbf{P}$ ; prepare the  $\gamma$ th circuit  $\mathbf{C}_\gamma$  using  $\text{key}_{\gamma+1}$ .
  - Receive externally honest garbling  $(\tilde{\mathbf{C}}_\gamma, \widehat{\text{conf}}_\gamma)$  of  $(\mathbf{C}_\gamma, \text{conf}_\gamma)$ .
  - Simulate the first  $\gamma-1$  circuits using  $\mathbf{R}$  and  $\mathbf{Q}$ , with  $\widehat{\text{conf}}_\gamma$  hardwired in  $\mathbf{R}$ .
- In  $\text{hyb}_{ns}^\gamma$ , the  $\gamma$ th garbled circuit  $\tilde{\mathbf{C}}_\gamma$  is simulated using program  $\mathbf{R}$ ; the output  $\text{out}_\gamma$  used for simulation is set to either  $y$  (if  $\text{conf}_{\gamma+1}$  is a final configuration) or the honestly generated garbled input  $\widehat{\text{conf}}_{\gamma+1}$ . In other words,  $\text{out}_\gamma = \mathbf{C}_\gamma(\text{conf}_\gamma)$ , where  $\mathbf{C}_\gamma$  is prepared in the same way as above. Furthermore, the previous  $\gamma-1$  garbled circuits are also simulated using program  $\mathbf{Q}$ . Their simulation as well as the generation of the garbled input  $\tilde{x}_{\gamma+1}$  depend potentially on the corresponding simulated garbled input  $\text{conf}_\gamma$  of  $\tilde{\mathbf{C}}_\gamma$ . In other words, the output of  $\text{hyb}_{ns}^\gamma$  can be generated by the same alternative sampling algorithm above, except that the second step is modified as follows:
  - Receive externally simulated garbling  $(\tilde{\mathbf{C}}_\gamma, \text{conf}_\gamma)$  generated using output  $\mathbf{C}_\gamma(\text{conf}_\gamma)$ .

Then it follows from the security of the circuit garbling scheme  $\mathcal{GS}_{\text{CIR}}$  that the distributions of  $(\tilde{\mathbf{C}}_\gamma, \widehat{\text{conf}}_\gamma)$  and  $(\tilde{\mathbf{C}}_\gamma, \text{conf}_\gamma)$  received externally by the alternative sampling algorithm above are computationally indistinguishable, and thus the distributions of outputs of  $\text{hyb}_{ns}^{\gamma-1}$  and  $\text{hyb}_{ns}^\gamma$ , which can be efficiently constructed from them, are also indistinguishable.  $\square$

Finally, by the above claim, it follows from a hybrid argument over  $\gamma$  that  $\{\text{real}_{ns}(1^\lambda, M, x)\}$  and  $\{\text{simu}_{ns}(1^\lambda, M, x)\}$  are indistinguishable. Hence,  $\mathcal{GS}_{ns}$  is a secure garbling scheme for TM.

### 3.2.2. A garbling scheme for TM with space-dependent complexity.

In this section, we construct a garbling scheme  $\mathcal{GS} = (\text{Garb}, \text{Encode}, \text{Eval})$  for TM with space-dependent complexity. Our construction relies on the following building blocks:

- A garbling scheme for polynomial-sized circuits, with independent input encoding:  $\mathcal{GS}_{\text{CIR}} = (\text{Garb}_{\text{CIR}}, \text{Encode}_{\text{CIR}}, \text{Eval}_{\text{CIR}})$ , where  $\text{Garb}_{\text{CIR}} = (\text{Gen}_{\text{CIR}}, \text{Gb}_{\text{CIR}})$  and whose simulation algorithm is  $\text{Sim}_{\text{CIR}} = (\text{Sim.Gen}_{\text{CIR}}, \text{Sim.Gb}_{\text{CIR}})$ .
- An indistinguishability obfuscator  $i\mathcal{O}_{\text{CIR}}(\cdot, \cdot)$  for polynomial-sized circuits.
- A puncturable PRF  $(\text{PRF.Gen}, \text{PRF.Punc}, \text{F})$  with input length  $n(\lambda)$  and output length  $m(\lambda)$ , where  $n(\lambda)$  can be set to any superlogarithmic function  $n(\lambda) = \omega(\log \lambda)$ , and  $m$  is a sufficiently large polynomial in  $\lambda$ .

**Circuit**  $\mathbb{P} = \mathbb{P}^{\lambda, S, M, K_\alpha, K_\beta}$ : On input  $t \in [T]$ , does:  
 Generates pseudorandom strings  $\alpha_t = F(K_\alpha, t)$ ,  $\alpha_{t+1} = F(K_\alpha, t+1)$  and  $\beta_t = F(K_\beta, t)$ ;  
 Compute  $\widehat{\mathbf{C}}_t = \mathbf{P}^{\lambda, S, M}(t; (\alpha_t, \alpha_{t+1}, \beta_t))$  and output  $\widehat{\mathbf{C}}_t$ .

**Circuit**  $\mathbb{Q} = \mathbb{Q}^{\lambda, S, |M|, T^*, y, K_\alpha, K_\beta}$ : On input  $t \in [T]$ , does:  
 Generate pseudorandom strings  $\alpha_t = F(K_\alpha, t)$ ,  $\alpha_{t+1} = F(K_\alpha, t+1)$  and  $\beta_t = F(K_\beta, t)$ ;  
 Compute  $\widetilde{\mathbf{C}}_t = \mathbf{Q}^{\lambda, S, |M|, T^*, y}(t; (\alpha_t, \alpha_{t+1}, \beta_t))$  and output  $\widetilde{\mathbf{C}}_t$ .

The circuits in Figures 3.1, 3.2, and 3.3 are padded to their maximum size.

FIG. 3.1. Circuits used in the construction and simulation of  $\mathcal{GS}$ .

For every  $\lambda$  and  $M \in \text{TM}_\lambda$ , the garbling scheme  $\mathcal{GS}$  proceeds as follows.

THE GARBLING ALGORITHM  $\text{Garb}(1^\lambda, M)$ :

1. Sample the PRF keys  $K_\alpha \xleftarrow{\$} \text{PRF.Gen}(1^\lambda)$  and  $K_\beta \xleftarrow{\$} \text{PRF.Gen}(1^\lambda)$ .
2. Obfuscate the circuit  $\mathbb{P}$ :  
 Obfuscate the circuit  $\mathbb{P}(t) = \mathbb{P}^{\lambda, S, M, K_\alpha, K_\beta}(t)$  as described in Figure 3.1, which is essentially a wrapper program that evaluates  $\mathbf{P}$  on  $t$  using pseudorandom coins generated using  $K_\alpha$  and  $K_\beta$  as described above. Obtain  $\overline{\mathbb{P}} \xleftarrow{\$} i\mathcal{O}(1^\lambda, \mathbb{P})$ .
3. Generate the key for garbling input:  
 Compute  $\mathbf{key}$  in the same way as the garbling scheme  $\text{Garb}_{ns}$  does, but using pseudorandom coins generated using  $K_\alpha$ . That is, compute the key for the first garbled circuit using randomness  $\alpha_1 = F(K_\alpha, 1)$ ,  $\mathbf{key}_1 = \text{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha_1)$ ; set  $\mathbf{key} = \mathbf{key}_1 \parallel 1^S$ .
4. Finally, output  $(\overline{\mathbb{P}}, \mathbf{key})$ .

THE ENCODING ALGORITHM  $\text{Encode}(\mathbf{key}, x)$ : Compute  $\hat{x} = \text{Encode}_{ns}(\mathbf{key}, x)$ .

THE EVALUATION ALGORITHM  $\text{Eval}(\overline{\mathbb{P}}, \hat{x})$ : Generate and evaluate the garbled circuits in the nonsuccinct garbling  $\widehat{M}$  one by one; terminate as soon as an output is produced. More precisely, evaluation proceeds in  $T$  iterations as follows: At the beginning of iteration  $t \in [T]$ , the previous  $t-1$  garbled circuits have been generated and evaluated, producing garbled input  $\widehat{\text{conf}}_t$  ( $\widehat{\text{conf}}_1 = \hat{x}$ ). Then compute  $\widehat{\mathbf{C}}_t = \overline{\mathbb{P}}(t)$ ; evaluate  $z = \text{Eval}_{\text{CIR}}(\widehat{\mathbf{C}}_t, \widehat{\text{conf}}_t)$ ; if  $z$  is a valid output, terminate and output  $y = z$ ; otherwise, proceed to the next iteration  $t+1$  with  $\widehat{\text{conf}}_{t+1} = z$ .

Next, we proceed to show that  $\mathcal{GS}$  is a garbling scheme for TM with space-dependent complexity.

**Correctness.** Fix any machine  $M \in \text{TM}$  and input  $x$ . Recall that the garbling algorithm  $\text{Garb}$  generates a pair  $(\overline{\mathbb{P}}, \mathbf{key})$ ; the latter is used later by the encoding algorithm  $\text{Encode}$  to obtain garbled input  $\hat{x}$ , while the former is used later by the evaluation algorithm  $\text{Eval}$  to create the nonsuccinct garbling  $\widehat{M} = \{\widehat{\mathbf{C}}_t = \overline{\mathbb{P}}(t)\}_{t \in [T]}$ ; the nonsuccinct garbling  $\widehat{M}$  is then evaluated with  $\hat{x}$  using algorithm  $\text{Eval}_{ns}$ . The distribution of the garbled input and the nonsuccinct garbling recovered by  $\text{Eval}$  is as follows:

$$\mathcal{D}_1 = \left\{ (\overline{\mathbb{P}}, \mathbf{key}) \xleftarrow{\$} \text{Garb}(1^\lambda, M) : \left( \hat{x} = \text{Encode}(\mathbf{key}, x), \widehat{M} = \left\{ \widehat{\mathbf{C}}_t = \overline{\mathbb{P}}(t) \right\}_{t \in [T]} \right) \right\}.$$

It follows from the construction of  $\mathbf{Garb}$ ,  $\mathbf{Encode}$  and the correctness of the indistinguishability obfuscator that the above distribution  $\mathcal{D}_1$  is identical to the distribution  $\mathcal{D}_2$  of a garbled pair  $(\hat{M}', \hat{x}')$  generated by the algorithms  $\mathbf{Garb}_{ns}$ ,  $\mathbf{Encode}_{ns}$  of the nonsuccinct scheme, using pseudorandom coins, formalized below:

$$\mathcal{D}_2 = \left\{ K_\alpha, K_\beta \stackrel{\$}{\leftarrow} \text{PRF.Gen}(1^\lambda) \quad \forall t \in [T], \alpha_t = F(K_\alpha, t), \beta_t = F(K_\beta, t) : \right. \\ \left. \left( \hat{x}' = \mathbf{Encode}_{ns}(\mathbf{key}' = \mathbf{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha_1), x), \hat{M}' = \left\{ \hat{\mathbf{C}}_t = \mathbf{P}(t; \alpha_t, \alpha_{t+1}, \beta_t) \right\}_{t \in [T]} \right) \right\}.$$

By the pseudorandomness of PRF, distribution  $\mathcal{D}_2$  is computationally indistinguishable from the garbled pair generated by  $\mathbf{Garb}_{ns}$ ,  $\mathbf{Encode}_{ns}$  using truly random coins:

$$\mathcal{D}_3 = \left\{ (\hat{M}'', \mathbf{key}'') \stackrel{\$}{\leftarrow} \mathbf{Garb}_{ns}(1^\lambda, M) : \left( \hat{x}'' = \mathbf{Encode}_{ns}(\mathbf{key}'', x), \hat{M}'' \right) \right\}.$$

The correctness of the nonsuccinct garbling scheme  $\mathcal{GS}_{ns}$  guarantees that, with overwhelming probability, evaluating  $\hat{M}''$  with  $\hat{x}''$  produces the correct output  $y = M(x)$ ; furthermore, the correct output  $y$  is produced after evaluating only the first  $T^* = T_M(x)$  garbled circuits. Thus, it follows from the indistinguishability between  $\mathcal{D}_1$  and  $\mathcal{D}_3$  that when evaluating a garbled pair  $(\hat{M}, \hat{x})$  sampled from  $\mathcal{D}_1$ , the correct output  $y$  is also produced after evaluating the first  $T^*$  garbled circuits. Given that  $\mathcal{D}_1$  is exactly the distribution of the nonsuccinct garbled pairs generated in  $\mathbf{Eval}$ , we have that correctness holds.

**Efficiency.** We show that the garbling scheme  $\mathcal{GS}$  has space-dependent complexity.

- The garbling algorithm  $\mathbf{Garb}(1^\lambda, M)$  runs in time  $\text{poly}(\lambda, |M|, S)$ . This is because  $\mathbf{Garb}$  produces an obfuscation of the program  $\mathbb{P}$  (a derandomized version of  $P$ ) which garbles circuits  $\mathbf{C}_t$  using pseudorandom coins for every input  $t \in [T]$ . Since the program  $\mathbf{C}_t$  has size  $q = \text{poly}(\lambda, |M|, S)$  as analyzed in the nonsuccinct garbling scheme, so do  $\mathbf{P}$  and  $\mathbb{P}$  (note that the input range  $T$  of these two programs is contained as part of the description of  $M$ , and hence  $|M| > \log T$ ). Therefore,  $\mathbf{Garb}$  takes time  $\text{poly}(\lambda, |M|, S)$  to produce the obfuscation of  $\mathbb{P}$ . Additionally, notice that  $\mathbf{Garb}$  generates the  $\mathbf{key}$  as the algorithm  $\mathbf{Garb}_{ns}$  does, which in turn runs  $\mathbf{Garb}_{\text{CIR}}(1^\lambda, 1^S)$  and takes time  $\text{poly}(\lambda, S)$ . Overall,  $\mathbf{Garb}$  runs in time  $\text{poly}(\lambda, |M|, S)$  as claimed.
- $\mathbf{Encode}$  run in the same time as the  $\mathbf{Encode}_{ns}$  algorithm, which is  $\text{poly}(\lambda, |M|, S)$ .
- Evaluation algorithm  $\mathbf{Eval}$  on input  $(\overline{\mathbb{P}}, \hat{x})$  produced by  $(\overline{\mathbb{P}}, \mathbf{key}) \stackrel{\$}{\leftarrow} \mathbf{Garb}(1^\lambda, 1^S)$  and  $\hat{x} = \mathbf{Encode}(\mathbf{key}, x)$  runs in time  $\text{poly}(\lambda, |M|, S) \times T^*$ ,  $T^* = T_M(x)$ , with overwhelming probability.

It follows from the analysis of correctness of  $\mathcal{GS}$  that, with overwhelming probability over the coins of  $\mathbf{Garb}$ , the nonsuccinct garbling  $\hat{M}$  defined by  $\overline{\mathbb{P}}$  satisfies that when evaluated with  $\hat{x}$ , the correct output is produced after  $T^*$  iterations. Since  $\mathbf{Eval}$  does not compute the entire nonsuccinct garbling  $\hat{M}$  in one shot, but rather generates and evaluates the garbled circuits in  $\hat{M}$  one by one, it thus terminates after producing and evaluating  $T^*$  garbled circuits. Since the generation and evaluation of each garbled circuit takes  $\text{poly}(\lambda, |M|, S)$  time, overall  $\mathbf{Eval}$  runs in time  $T_M(x) \times \text{poly}(\lambda, |M|, S)$ , as claimed.

**Security.** Fix any polynomial  $T'$ , any sequence of algorithms  $\{M = M_\lambda\} \in \{\text{TM}_\lambda^{T'}\}$ , and any sequence of inputs  $\{x = x_\lambda\}$ , where  $x_\lambda \in \{0, 1\}^{M.n}$ . Toward show-

ing the security of  $\mathcal{GS}$ , we construct a simulator  $\text{Sim}$  satisfying that the following two ensembles are indistinguishable in  $\lambda$ :

$$(3.3) \quad \{\text{real}(1^\lambda, M, x)\} = \left\{ (\overline{\mathbb{P}}, \mathbf{key}) \stackrel{\$}{\leftarrow} \text{Garb}(1^\lambda, M), \hat{x} = \text{Encode}(\mathbf{key}, x) : (\overline{\mathbb{P}}, \hat{x}) \right\}_\lambda,$$

$$(3.4) \quad \{\text{simu}(1^\lambda, M, x)\} = \left\{ (\overline{\mathbb{Q}}, \tilde{x}) \stackrel{\$}{\leftarrow} \text{Sim}(1^\lambda, |x|, |M|, S, T, T_M(x), M(x)) : (\overline{\mathbb{Q}}, \tilde{x}) \right\}_\lambda.$$

As discussed in the overview, the simulation will obfuscate the program  $\mathbf{Q}$  used for simulating the nonsuccinct garbled machine  $\tilde{M} = (\tilde{\mathbf{C}}_1, \dots, \tilde{\mathbf{C}}_T)$ , shown more precisely as follows.

THE SIMULATION ALGORITHM  $\text{Sim}(1^\lambda, |x|, |M|, S, T, T^* = T_M(x), y = M(x))$ :

1. *Sample PRF keys:*  $K_\alpha \stackrel{\$}{\leftarrow} \text{PRF.Gen}(1^\lambda)$  and  $K_\beta \stackrel{\$}{\leftarrow} \text{PRF.Gen}(1^\lambda)$ .
2. *Obfuscate the circuit  $\mathbb{Q}$ :*  
Obfuscate the circuit  $\mathbb{Q}(t) = \mathbb{Q}^{\lambda, S, |M|, T^*, y, K_\alpha, K_\beta}(t)$  as described in Figure 3.1, which is essentially a wrapper program that evaluates  $\mathbf{Q}$  on  $t$ , using pseudorandom coins  $\{\alpha_t, \beta_t\}$  generated by evaluating  $F$  on keys  $K_\alpha$  and  $K_\beta$  and inputs  $t \in [T]$ . Obtain  $\overline{\mathbb{Q}} \stackrel{\$}{\leftarrow} i\mathcal{O}(1^\lambda, \mathbb{Q})$ .
3. *Simulate the garbled input:*  
Simulate the garbled input  $\tilde{x}$  in the same way that simulator  $\text{Sim}_{ns}$  does, but using pseudorandom coins. That is, compute  $(\widetilde{\text{conf}}_1, \mathbf{st}_1) = \text{Sim.Gen}_{\text{CIR}}(1^\lambda, S; \alpha_1)$ , where  $\alpha_1 = F(K_\alpha, 1)$ ; set  $\tilde{x} = \widetilde{\text{conf}}_1$ .
4. *Finally, output  $(\overline{\mathbb{Q}}, \tilde{x})$ .*

The simulator  $\text{Sim}(1^\lambda, |x|, |M|, S, T, T^*, y = M(x))$  runs in time  $\text{poly}(\lambda, |M|, S)$ . This follows because the simulator simulates the garbled Turing machine by obfuscating the program  $\mathbb{Q}$ . As the program  $\mathbb{Q}$  simply runs  $\mathbf{Q}$  using pseudorandom coins, its size is  $\text{poly}(\lambda, |M|, S)$ ; thus obfuscation takes time in the same order. On the other hand,  $\text{Sim}$  simulates the garbled input  $\tilde{x}$  as the simulator  $\text{Sim}_{ns}$  does, which simply invokes  $\text{Sim}_{\text{CIR}}(1^\lambda, S)$  of the circuit garbling scheme, which takes time  $\text{poly}(\lambda, S)$ . Therefore, overall the simulation takes time  $\text{poly}(\lambda, |M|, S)$  as claimed.

To show the indistinguishability between honestly generated garbling  $(\overline{\mathbb{P}}, \hat{x}) \stackrel{\$}{\leftarrow} \text{real}(1^\lambda, M, x)$  and the simulation  $(\overline{\mathbb{Q}}, \tilde{x}) \stackrel{\$}{\leftarrow} \text{simu}(1^\lambda, M, x)$  (see (3.3) and (3.4) for formal definitions of  $\text{real}$  and  $\text{simu}$ ), we will consider a sequence of hybrids  $\text{hyb}^0, \dots, \text{hyb}^T$ , where the output distribution of  $\text{hyb}^0$  is identical to  $\text{real}$ , while that of  $\text{hyb}^T$  is identical to  $\text{simu}$ . In every intermediate hybrid  $\text{hyb}^\gamma$ , a hybrid simulator  $\text{HSim}^\gamma$  is invoked, producing a pair  $(\overline{\mathbb{M}}^\gamma, \tilde{x}^\gamma)$ , where  $\overline{\mathbb{M}}^\gamma$  is the obfuscation of (the derandomized wrapper of) a merged program  $\mathbf{M}^\gamma$  that produces a hybrid chain of garbled circuits as in the security proof of the nonsuccinct garbling scheme, where the first  $\gamma$  garbled circuits are simulated and the rest are generated honestly. This hybrid simulation algorithm is given more precisely as follows.

THE HYBRID SIMULATION ALGORITHM  $\text{HSim}^\gamma(1^\lambda, M, x)$  FOR  $\gamma = 0, \dots, T$ :

Compute  $T^* = T_M(x)$  and  $y = M(x)$  and the intermediate configuration  $\text{conf}_{\gamma+1}$  as defined by  $\text{CONFIG}(M, x)$ .

1. *Sample the PRF keys*  $K_\alpha \stackrel{\$}{\leftarrow} \text{PRF.Gen}(1^\lambda)$  and  $K_\beta \stackrel{\$}{\leftarrow} \text{PRF.Gen}(1^\lambda)$ .
2. *Obfuscate the circuit  $\mathbb{M}^\gamma$ :*  
Obfuscate the circuit  $\mathbb{M}^\gamma(t) = (\mathbb{M}^\gamma)^{\lambda, S, M, T^*, y, \text{conf}_{\gamma+1}, K_\alpha, K_\beta}(t)$  as described in Figure 3.2, which is essentially a wrapper program that evaluates the combined program  
 $\mathbf{M}^\gamma = \text{COMBINE}[(\mathbf{Q}, [\gamma - 1]), (\mathbf{R}, \{\gamma\}), (\mathbf{P}, [\gamma + 1, T])](t; (\alpha_t, \alpha_{t+1}, \beta_t))$ ,

using pseudorandom coins  $\{\alpha_t, \beta_t\}$  generated using  $K_\alpha$  and  $K_\beta$ . Obtain  $\overline{\mathbb{M}}^\gamma \stackrel{s}{\leftarrow} i\mathcal{O}(1^\lambda, \mathbb{M}^\gamma)$ .

3. *Simulate the garbled input:*

If  $\gamma > 0$ , simulate the garbled input  $\tilde{x}^\gamma$  in the same way as in Sim. Otherwise, if  $\gamma = 0$ , generate  $\tilde{x}^0$  honestly, using Garb and Encode.

4. *Finally, output  $(\overline{\mathbb{M}}^\gamma, \tilde{x}^\gamma)$ .*

**Circuit**  $\mathbb{M}^\gamma = (\mathbb{M}^\gamma)^{\lambda, S, M, T^*, y, \text{conf}_{\gamma+1}, K_\alpha, K_\beta}$ : On input  $t \in [T]$ , does:  
 Generate pseudorandom strings  $\alpha_t = F(K_\alpha, t)$ ,  $\alpha_{t+1} = F(K_\alpha, t + 1)$  and  $\beta_t = F(K_\beta, t)$ ;  
 Compute  $\tilde{\mathbf{C}}_t = \mathbf{M}^\gamma(t; (\alpha_t, \alpha_{t+1}, \beta_t))$  and output  $\tilde{\mathbf{C}}_t$ , where  $\mathbf{M}^\gamma$  is

$$(\mathbf{M}^\gamma)^{\lambda, S, M, T^*, y, \text{conf}_{\gamma+1}}$$

$$= \text{COMBINE}[(\mathbf{Q}, [\gamma - 1]), (\mathbf{R}, \{\gamma\}), (\mathbf{P}, [\gamma + 1, T])](t; (\alpha_t, \alpha_{t+1}, \beta_t)).$$

The circuits in Figures 3.1, 3.2, and 3.3 are padded to their maximum size.

FIG. 3.2. Circuits used in the security analysis of  $\mathcal{GS}$ .

We overload the notation  $\text{hyb}^\gamma(1^\lambda, M, x)$  as the output distribution of the  $\gamma$ th hybrid. By construction, when  $\gamma = 0$ ,  $\mathbf{M}^0 = \mathbf{P}$  and the garbled input  $\tilde{x}^0$  is generated honestly; thus,  $\{\text{hyb}^0(1^\lambda, M, x)\} = \{\text{real}(1^\lambda, M, x)\}$ ; furthermore, when  $\gamma = T$ ,  $\mathbf{M}^T = \mathbf{Q}$  and the garbled input  $\tilde{x}^T$  is simulated; thus  $\{\text{hyb}^T(1^\lambda, M, x)\} = \{\text{simu}(1^\lambda, M, x)\}$ . Therefore, to show the security of  $\mathcal{GS}$ , it boils down to proving the following claim.

CLAIM 3.4. *For every  $\gamma \geq 0$ , the following holds:*

$$\{\text{hyb}^\gamma(1^\lambda, M, x)\}_\lambda \approx \{\text{hyb}^{\gamma+1}(1^\lambda, M, x)\}_\lambda.$$

*Proof.* Fix a  $\gamma \in \mathbb{N}$ , a sufficiently large  $\lambda \in \mathbb{N}$ , an  $M = M_\lambda$ , and a  $x = x_\lambda$ . Note that the only difference between  $(\overline{\mathbb{M}}^\gamma, \tilde{x}^\gamma) \stackrel{s}{\leftarrow} \text{hyb}^\gamma$  and  $(\overline{\mathbb{M}}^{\gamma+1}, \tilde{x}^{\gamma+1}) \stackrel{s}{\leftarrow} \text{hyb}^{\gamma+1}$  is the following:

- For every  $\gamma$ , the underlying obfuscated programs  $\mathbb{M}^\gamma, \mathbb{M}^{\gamma+1}$  differ in their implementation by at most two inputs, namely,  $\gamma, \gamma + 1$ .
- When  $\gamma = 0$ , the garbled input  $\tilde{x}^0$  is generated honestly in  $\text{hyb}^0$ , whereas  $\tilde{x}^1$  is simulated in  $\text{hyb}^1$ .

To show the indistinguishability of the two hybrids, we consider a sequence of sub-hybrids from  $\text{H}_0^\gamma = \text{hyb}^\gamma$  to  $\text{H}_7^\gamma = \text{hyb}^{\gamma+1}$ . Below we describe these hybrids  $\text{H}_0^\gamma, \dots, \text{H}_7^\gamma$  and argue that the output distributions of any two subsequent hybrids are indistinguishable. We denote by  $(\overline{\mathbb{M}}_i^\gamma, \tilde{x}_i^\gamma)$  the garbled pair produced in hybrid  $\text{H}_i^\gamma$  for  $i = 0, \dots, 7$ . For convenience, below we suppress the superscript  $\gamma$  and simply use the notation  $\text{H}_i = \text{H}_i^\gamma$ ,  $\overline{\mathbb{M}}_i = \overline{\mathbb{M}}_i^\gamma$ ,  $\mathbb{M}_i = \mathbb{M}_i^\gamma$ , and  $\tilde{x}_i = \tilde{x}_i^\gamma$ .

HYBRID  $\text{H}_1$ : Generate a garbled pair  $(\overline{\mathbb{M}}_1, \tilde{x}_1)$  by running a simulation procedure that proceeds identically to  $\text{HSim}^\gamma$ , except with the following modifications:

- In the first step, puncture the two PRF keys  $K_\alpha, K_\beta$  at input  $\gamma + 1$  and obtain  $K_\alpha(\gamma + 1) = \text{PRF.Punc}(K_\alpha, \gamma + 1)$  and  $K_\beta(\gamma + 1) = \text{PRF.Punc}(K_\beta, \gamma + 1)$ . Furthermore, compute  $\alpha_{\gamma+1} = F(K_\alpha, \gamma + 1)$  and  $\beta_{\gamma+1} = F(K_\beta, \gamma + 1)$ .
- In the second step, obfuscate a circuit  $\mathbb{M}_1$  slightly modified from  $\mathbb{M}^\gamma$ : Instead of having the full PRF keys  $K_\alpha, K_\beta$  hardwired in,  $\mathbb{M}_1$  has the

punctured keys  $K_\alpha(\gamma + 1), K_\beta(\gamma + 1)$  and the PRF values  $\alpha_{\gamma+1}, \beta_{\gamma+1}$  hardwired in;  $\mathbb{M}_1$  proceeds identically to  $\mathbb{M}_1$ , except that it uses the punctured PRF keys to generate pseudorandom coins corresponding to input  $t \neq \gamma + 1$  and directly uses  $\alpha_{\gamma+1}, \beta_{\gamma+1}$  as the coins for input  $t = \gamma + 1$ . See Figure 3.3 for a description of  $\mathbb{M}_1 = \mathbb{M}_1^\gamma$ .

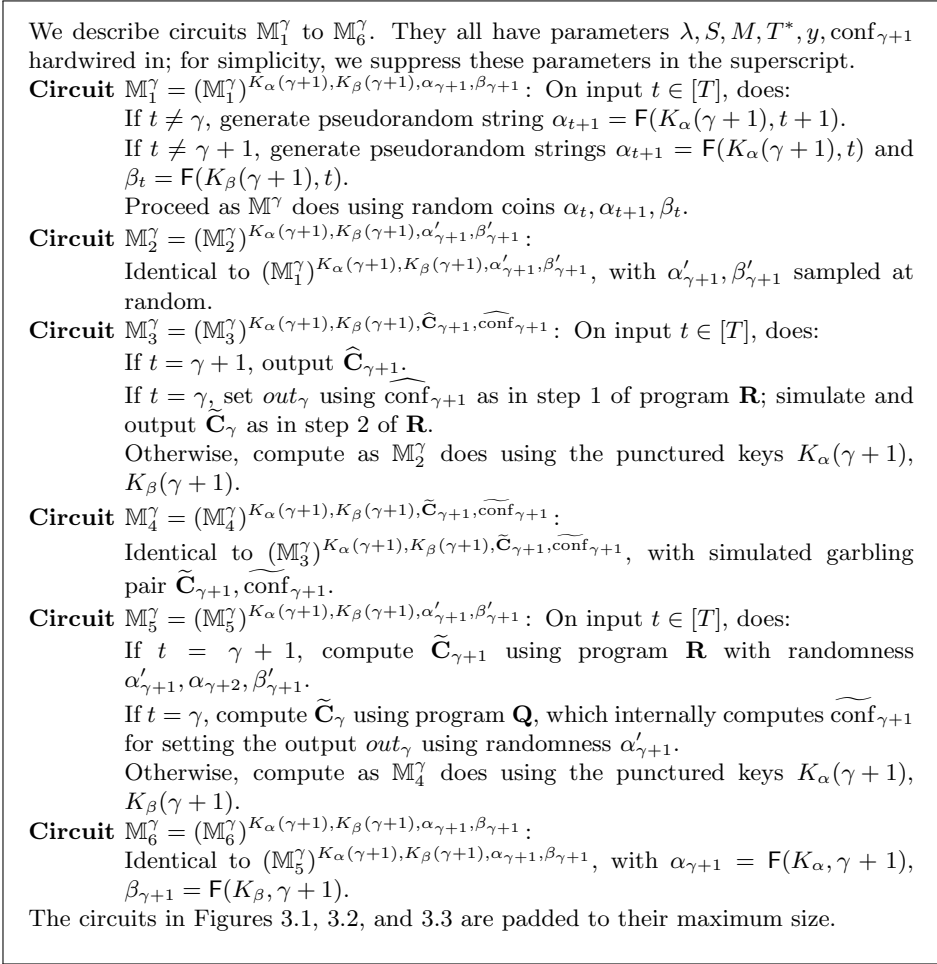


FIG. 3.3. Circuits used in the security analysis of  $\mathcal{GS}$ , continued.

By construction,  $\mathbf{H}_1$  only differs from  $\text{hyb}^\gamma$  in the underlying program that is obfuscated; this program  $\mathbb{M}_1$  has the same functionality as  $\mathbb{M}^\gamma$ . Thus it follows from the security of indistinguishability obfuscator  $i\mathcal{O}$  that the obfuscated programs  $\overline{\mathbb{M}}^\gamma$  and  $\overline{\mathbb{M}}_1$  are indistinguishable. (Furthermore, the garbled inputs  $\tilde{x}^\gamma$  and  $\tilde{x}_1$  in these two hybrids are generated in the same way.) Thus, we have that the output  $(\overline{\mathbb{M}}_1, \tilde{x}_1)$  of  $\mathbf{H}_1$  is indistinguishable from the output  $(\overline{\mathbb{M}}^\gamma, \tilde{x}^\gamma)$  of  $\text{hyb}^\gamma$ . That is,

$$\{\text{hyb}^\gamma(1^\lambda, M, x)\}_\lambda \approx \{\mathbf{H}_0(1^\lambda, M, x)\}_\lambda.$$

**HYBRID  $\mathbf{H}_2$ :** Generate a garbled pair  $(\overline{\mathbb{M}}_2, \tilde{x}_2)$  by running the same simulation procedure as in  $\mathbf{H}_1$ , except with the following modifications: Instead of using

pseudorandom coins  $\alpha_{\gamma+1}$  and  $\beta_{\gamma+1}$ , hybrid  $H_2$  samples two sufficiently long truly random string  $\alpha'_{\gamma+1}, \beta'_{\gamma+1} \xleftarrow{\$} \{0, 1\}^{\text{poly}(\lambda)}$  and replaces  $\alpha_{\gamma+1}, \beta_{\gamma+1}$  with these truly random strings. More specifically,  $H_2$  obfuscates a program  $M_2$  that is identical to  $M_1$ , but with  $(K_\alpha(\gamma + 1), K_\beta(\gamma + 1), \alpha'_{\gamma+1}, \beta'_{\gamma+1})$  hardwired in; furthermore, if  $\gamma = 0$ ,  $\alpha'_1$  (as opposed to  $\alpha_1$ ) is used to generate the garbled input  $\tilde{x}_2$ . Since only the punctured keys  $K_\alpha(\gamma + 1), K_\beta(\gamma + 1)$  are used in the whole simulation procedure, it follows from the pseudorandomness of the punctured PRF that the output  $(\overline{M}_2, \tilde{x}_2)$  of  $H_2$  is indistinguishable from  $(\overline{M}_1, \tilde{x}_1)$  of  $\text{hyb}_1$ . That is,

$$\{H_1(1^\lambda, M, x)\}_\lambda \approx \{H_2(1^\lambda, M, x)\}_\lambda.$$

**HYBRID  $H_3$ :** Generate a garbled pair  $(\overline{M}_3, \tilde{x}_3)$  by running the same simulation procedure as in  $H_2$  with the following modifications:

- Observe that in program  $M_2$ ,  $\alpha'_{\gamma+1}, \beta'_{\gamma+1}$  are used in the evaluation of at most two inputs,  $\gamma$  and  $\gamma + 1$ :

For input  $\gamma + 1$ , program  $P$  is invoked with input  $\gamma + 1$  and randomness  $\alpha'_{\gamma+1}, \alpha_{\gamma+2}, \beta'_{\gamma+1}$ , in which a circuit  $C_{\gamma+1}$  is prepared depending on  $\alpha_{\gamma+2}$ , and then obfuscated by computing

$$\text{key}_{\gamma+1} = \text{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha'_{\gamma+1}), \quad \widehat{C}_{\gamma+1} = \text{Gb}_{\text{CIR}}(\text{key}_{\gamma+1}, C_{\gamma+1}; \beta'_{\gamma+1}).$$

If  $\gamma > 0$ , for input  $\gamma$ , program  $R$  is invoked with input  $\gamma$  and randomness  $\alpha_\gamma, \alpha'_{\gamma+1}, \beta_\gamma$ , in which a garbled circuit  $\widehat{C}_\gamma$  is simulated; the output  $\text{out}_\gamma$  used for the simulation depends potentially on an honest garbling of  $\text{conf}_{\gamma+1}$ , that is,

$$\widehat{\text{conf}}_{\gamma+1} = \text{Encode}_{\text{CIR}}(\text{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha'_{\gamma+1}), \text{conf}_{\gamma+1}).$$

Using  $\text{out}_\gamma, \widehat{C}_\gamma$  is simulating using randomness  $\alpha_\gamma, \beta_\gamma$ .

*First modification:* Hybrid  $H_3$  receives externally the above pair  $(\widehat{C}_{\gamma+1}, \widehat{\text{conf}}_{\gamma+1})$ . Instead of obfuscating  $M_2$  (which computes  $\widehat{C}_{\gamma+1}, \widehat{\text{conf}}_{\gamma+1}$  internally),  $H_3$  obfuscates  $M_3$  that has  $\widehat{C}_{\gamma+1}, \widehat{\text{conf}}_{\gamma+1}$  directly hardwired in (as well as  $K_\alpha(\gamma + 1), K_\beta(\gamma + 1)$ ).  $M_3$  on input  $\gamma + 1$  directly outputs  $\widehat{\text{conf}}_{\gamma+1}$ ; on input  $\gamma$ , it uses  $\widehat{\text{conf}}_{\gamma+1}$  to compute  $\widehat{C}_\gamma$ ; on all other inputs, it proceeds identically as  $M_2$ . (See Figure 3.1 for a description of  $M_3$ .) It is easy to see that when the correct values  $\widehat{C}_{\gamma+1}, \widehat{\text{conf}}_{\gamma+1}$  are hardwired, the program  $M_3$  has the same functionality as  $M_2$ .

- In  $H_2$ , if  $\gamma = 0$ ,  $\alpha'_1$  is used for garbling the input,

$$\text{key}_1 = \text{Gen}_{\text{CIR}}(1^\lambda, 1^S; \alpha'_1), \quad \widehat{\text{conf}}_1 = \text{Encode}_{\text{CIR}}(\text{key}_1, \text{conf}_1),$$

where  $\text{conf}_1$  is the initial state corresponding to  $x$ .

*Second modification:* Instead, if  $\gamma = 0$ , hybrid  $H_3$  receives  $\widehat{\text{conf}}_1$  externally and directly outputs it as the garbled inputs  $\hat{x}_3 = \widehat{\text{conf}}_1$ .

When  $H_3$  receives the correct values of  $(\widehat{\text{conf}}_{\gamma+1}, \widehat{C}_{\gamma+1})$  externally, it follows from the security of  $i\mathcal{O}$  that the output distribution of  $H_3$  is indistinguishable from that of  $H_2$ . That is,

$$\{H_2(1^\lambda, M, x)\}_\lambda \approx \{H_3(1^\lambda, M, x)\}_\lambda.$$



HYBRID  $H_4$ : Generate a garbled pair  $(\widetilde{\mathbb{M}}_4, \tilde{x}_4)$  by running the same procedure as in  $H_3$ , except that  $H_4$  receives externally a simulated pair  $(\widetilde{\text{conf}}_{\gamma+1}, \widetilde{\mathbf{C}}_{\gamma+1})$  produced as follows:

$$(3.5) \quad (\widetilde{\text{conf}}_{\gamma+1}, \mathbf{st}_{\gamma+1}) = \text{Sim.Gen}_{\text{CIR}}(1^\lambda, S; \alpha'_{\gamma+1}),$$

$$(3.6) \quad \widetilde{\mathbf{C}}_{\gamma+1} = \text{Sim.Gb}_{\text{CIR}}(1^\lambda, S, 1^q, \text{out}_{\gamma+1}, \mathbf{st}_{\gamma+1}; \beta'_{\gamma+1}),$$

where  $\text{out}_{\gamma+1}$  is set to be the output of circuit  $\mathbf{C}_{\gamma+1}$  on input  $\text{conf}_{\gamma+1}$ . Thus, it follows from the security of the circuit garbling scheme  $\mathcal{GS}_{\text{CIR}}$  that the simulated pair  $(\widetilde{\text{conf}}_{\gamma+1}, \widetilde{\mathbf{C}}_{\gamma+1})$  that hybrid  $H_4$  receives externally is indistinguishable from the honest pair  $(\widetilde{\text{conf}}_{\gamma+1}, \widehat{\mathbf{C}}_{\gamma+1})$  that  $H_3$  receives externally. Since these two hybrids only differ in which pair they receive externally, it follows that

$$\{H_3(1^\lambda, M, x)\}_\lambda \approx \{H_4(1^\lambda, M, x)\}_\lambda.$$

HYBRID  $H_5$ : Generate a garbled pair  $(\widetilde{\mathbb{M}}_5, \tilde{x}_5)$  by running the same procedure as in  $H_4$ , except that instead of receiving  $(\widetilde{\text{conf}}_{\gamma+1}, \widetilde{\mathbf{C}}_{\gamma+1})$  externally,  $H_5$  computes them internally using truly random coins  $\alpha'_{\gamma+1}, \beta'_{\gamma+1}$ . More precisely, it does the following:

- It obfuscates a program  $\mathbb{M}_5$  that has  $K_\alpha(\gamma+1), K_\beta(\gamma+1), \alpha'_{\gamma+1}, \beta'_{\gamma+1}$  hardwired in:  
On input  $\gamma+1$ , it computes  $\widetilde{\mathbf{C}}_{\gamma+1}$  using the program  $\mathbf{R}$  with randomness  $\alpha'_{\gamma+1}, \alpha_{\gamma+2}, \beta'_{\gamma+1}$  (which computes  $\widetilde{\mathbf{C}}_{\gamma+1}$  as described in (3.5) and (3.6)).  
On input  $\gamma$ , it computes  $\widetilde{\mathbf{C}}_\gamma$  using the program  $\mathbf{Q}$  with randomness  $\alpha_\gamma, \alpha'_{\gamma+2}, \beta_\gamma$  (which computes internally  $\widetilde{\text{conf}}_{\gamma+1}$  as described in (3.5)).  
On other inputs  $t \neq \gamma, \gamma+1$ , it computes as  $\mathbb{M}_4$  does.
- If  $\gamma = 0$ ,  $\alpha'_1$  is used for computing  $\widetilde{\text{conf}}_1$  as described in (3.5), and then outputs  $\tilde{x}_4 = \widetilde{\text{conf}}_1$ .

It follows from the fact that  $\mathbb{M}_5$  computes  $(\widetilde{\text{conf}}_{\gamma+1}, \widetilde{\mathbf{C}}_{\gamma+1})$  correctly internally that it has the same functionality as  $\mathbb{M}_4$ ; thus, the obfuscations of these two programs are indistinguishable. Combined with the fact that the distribution of the garbled inputs  $\tilde{x}_4$  is identical to  $\tilde{x}_3$ , we have that

$$\{H_4(1^\lambda, M, x)\}_\lambda \approx \{H_5(1^\lambda, M, x)\}_\lambda.$$

HYBRID  $H_6$ : Generate a garbled pair  $(\widetilde{\mathbb{M}}_6, \tilde{x}_6)$  by running the same procedure as in  $H_5$ , except that instead of using truly random coins  $\alpha'_{\gamma+1}, \beta'_{\gamma+1}$ , use pseudo-random coins  $\alpha_{\gamma+1} = F(K_\alpha, \gamma+1)$  and  $\beta_{\gamma+1} = F(K_\beta, \gamma+1)$ . In particular,  $H_6$  obfuscates a program  $\mathbb{M}_6$  that is identical to  $\mathbb{M}_5$  except that  $K_\alpha(\gamma+1), K_\beta(\gamma+1), \alpha_{\gamma+1}, \beta_{\gamma+1}$  are hardwired in, and if  $\gamma = 0$ ,  $\alpha_1$  is used to generate the garbled input  $\tilde{x}_6$ . It follows from the pseudorandomness of the punctured PRF that

$$\{H_6(1^\lambda, M, x)\}_\lambda \approx \{H_5(1^\lambda, M, x)\}_\lambda.$$

HYBRID  $H_7$ : Generate a garbled pair  $(\widetilde{\mathbb{M}}_7, \tilde{x}_7)$  by running hybrid simulator  $\text{HSim}^{\gamma+1}$ . Note that the only difference between  $\text{HSim}^{\gamma+1}$  and the simulation procedure in  $H_6$  is that instead of obfuscating  $\mathbb{M}_6$  that has tuple  $(K_\alpha(\gamma+1), K_\beta(\gamma+1), \alpha_{\gamma+1}, \beta_{\gamma+1})$  hardwired in,  $\text{HSim}^{\gamma+1}$  obfuscates  $\mathbb{M}^{\gamma+1}$  that has the full PRF keys  $K_\alpha, K_\beta$  hardwired in and evaluates  $\alpha_{\gamma+1}, \beta_{\gamma+1}$  internally.

Since  $\mathbb{M}^{\gamma+1}$  and  $\mathbb{M}_6^\gamma$  have the same functionality, it follows from the security of  $i\mathcal{O}$  that

$$\{\mathbb{H}_6(1^\lambda, M, x)\}_\lambda \approx \{\mathbb{H}_5(1^\lambda, M, x)\}_\lambda.$$

Finally, by a hybrid argument, we conclude the claim.  $\square$

Given the above claim, by a hybrid argument over  $\gamma$ , we have that  $\{\text{real}(1^\lambda, M, x)\}$  and  $\{\text{simu}(1^\lambda, M, x)\}$  are indistinguishable. Hence,  $\mathcal{GS}$  is a secure garbling scheme for TM.

**3.3. Succinct garbling for RAMs via garbled RAMs.** Note that our construction of succinct garbling in the previous section uses the underlying circuit garbling scheme  $\mathcal{GS}$  in a black-box way. In fact, the scheme does not even require the underlying garbling scheme to be for circuits—*any garbling scheme for any class of algorithms that is “complete,” in particular that can be used to implement the augmented blocks, suffices.* Below we show that by plugging in the one-time garbled RAM of [LO13, GHL<sup>+</sup>14] and modifying the construction of Theorem 3.2 slightly, we can improve the efficiency of our garbling scheme when the algorithm class is RAM. More precisely, we show the following theorem.

**THEOREM 3.5.** *Assume the existence of  $i\mathcal{O}$  for circuits and one-way functions. There exists a garbling scheme  $\mathcal{GS}^{\text{RAM}}$  for RAM with linear-space-dependent complexity. Furthermore, for any RAM  $R$  and input  $x$ , evaluation of a garbled pair  $(\hat{R}, \hat{x})$  produced by  $\mathcal{GS}^{\text{RAM}}$  takes time  $\text{poly}(\lambda, |R|) \times (T_R(x) + S)$ .*

*Proof.* Our garbling scheme  $\mathcal{GS}$  supports garbling RAM computations, given by a RAM program  $R$  and an input  $x$ . The RAM program  $R$  on input  $x$  has access to a memory (initially empty) that  $R$  can read from and write to at various locations throughout its execution. If a RAM program  $R$  has space complexity  $S$ , its memory size is bounded by  $S$ .

Below, we start with a brief review of schemes in [LO13, GHL<sup>+</sup>14].

**The RAM garbling schemes of [LO13, GHL<sup>+</sup>14].** These schemes support garbling RAM computations of a more general form, given by a RAM program  $R$ , an input  $x$ , and additional initial data  $D$  contained in the memory; denote this evaluation by  $R^D(x)$ . This formulation is able to capture evaluation of multiple RAM programs with persistent memory, denoted by  $(R_1(x_1), \dots, R_k(x_k))^D$ , where the memory  $D$  is updated by each RAM evaluation and the updates persist to the next evaluation. Though we do not show that our RAM garbling scheme  $\mathcal{GS}$  supports persistent memory, as we will see later, the fact that schemes in [LO13, GHL<sup>+</sup>14] can garble initial data  $D$  in linear time in its size  $|D|$  is the key toward showing the linear-space-dependent complexity of  $\mathcal{GS}$ .

Below, we recall the syntax and efficiency of schemes in [LO13, GHL<sup>+</sup>14] and note several special properties that will be useful for our construction later. Let  $R$  be a RAM machine with parameters  $n, m, S, T$ .

**ALGORITHMS:** The schemes contain the following algorithms.

- $k^d \xleftarrow{\$} \text{Gen}^d(1^\lambda)$  samples a data garbling key  $k^d$  in  $\text{poly}(\lambda)$  steps.
- $\hat{D} \xleftarrow{\$} \text{GData}^D(k^d)$  takes  $D$  as initial data and a data garbling key  $k^d \in \{0, 1\}^\lambda$  as input and outputs garbled data  $\hat{D}$  in linear time in the data size  $\text{poly}(\lambda)|D|$ . The RAM program  $\text{GData}$  has size  $\text{poly}(\lambda)$ .
- $(\hat{R}, k^i) \xleftarrow{\$} \text{GProg}(k^d, R)$  takes the description of a RAM machine  $R$  and a data garbling key  $k^d$  and outputs a garbled program  $\hat{R}$  and an input garbling key  $k^i$  in linear time in the time bound  $\text{poly}(\lambda, n, |R|)T$ .

*Independent input/data encoding:* Note that the data garbling algorithm  $\text{GData}$  does not depend on the program or the input. Moreover, the program garbling algorithm  $\text{GProg}$  can be decomposed into two subroutines,  $k^i \stackrel{\$}{\leftarrow} \text{Gen}^i(1^\lambda, 1^n)$  and  $\widehat{R} \stackrel{\$}{\leftarrow} \text{GProg}'(k^d, k^i, R)$ , which runs, respectively, in time  $\text{poly}(\lambda, n)$  and  $\text{poly}(\lambda, n, |R|)T$ . (The garbled RAM schemes of [LO13, GHL<sup>+</sup>14] are based on Yao's garbled circuits and inherit the property that the input garbling key depends only on the input length and security parameter.) Thus, the input garbling algorithm  $\text{GInput}$  below also does not depend on the program.

- $\hat{x} \stackrel{\$}{\leftarrow} \text{GInput}(k^i, x)$  takes input  $x \in \{0, 1\}^n$  and the input garbling key  $k^i$ , and outputs a garbled input  $\hat{x}$  in time  $\text{poly}(\lambda, n)$ .
- $y = \text{GEval}^{\widehat{D}}(\widehat{R}, \hat{x})$  takes a garbled program  $\widehat{R}$  and a garbled input  $\hat{x}$ , with access to the garbled data  $\widehat{D}$ , and computes the output  $y = R^D(x)$  in linear time in the time bound  $\text{poly}(\lambda, n, |R|)T$ .<sup>10</sup>

**SIMULATION:**  $(\widetilde{R}, \widetilde{x}, \widetilde{D}) \stackrel{\$}{\leftarrow} \text{GSim}(1^\lambda, (|x|, |D|, |R|, n, m, S, T), R(x))$  simulates the garbled program, input, and data.<sup>11</sup>

*Independent garbled input/data simulation:* The simulation algorithm can be decomposed into the following two PPT procedures:

$$\begin{aligned} (\widetilde{x}, \widetilde{D}, \mathbf{st}) &\stackrel{\$}{\leftarrow} \text{GSim.Gen}(1^\lambda, |x|, |D|), \\ \widetilde{R} &\stackrel{\$}{\leftarrow} \text{GSim.Gb}(1^\lambda, (|x|, |D|, |R|, n, m, S, T), R(x), \mathbf{st}). \end{aligned}$$

Finally, we note that the schemes of [LO13, GHL<sup>+</sup>14] assumes the existence of an identity-based encryption (IBE), which is implied by  $\mathbf{iO}$  and one-way functions.

**The basic RAM garbling scheme  $\mathcal{GS}'$ .** From the nonsuccinct RAM garbling schemes of [LO13, GHL<sup>+</sup>14] described above, we derive a basic RAM garbling scheme  $\mathcal{GS}' = (\text{Garb}', \text{Encode}', \text{Eval}')$  with similar syntax to a circuit garbling scheme with independent input encoding. Then it becomes easier to see why we can use  $\mathcal{GS}'$  to replace a circuit garbling scheme in the construction of Theorem 3.2. Furthermore, the efficiency guarantee of  $\mathcal{GS}'$  makes the construction more efficient, leading to linear-space-dependent complexity.

**ALGORITHMS:** The scheme  $\mathcal{GS}'$  consists of the following algorithms:

- The garbling algorithm  $(\mathbf{key}, \widehat{R}) \stackrel{\$}{\leftarrow} \text{Garb}'(1^\lambda, R)$  consists of two subroutines,  $\mathbf{key} \stackrel{\$}{\leftarrow} \text{Gen}'(1^\lambda)$  and  $\widehat{R} \stackrel{\$}{\leftarrow} \text{Gb}'(\mathbf{key}, R)$ , defined below:

$$\begin{aligned} \mathbf{key} \stackrel{\$}{\leftarrow} \text{Gen}'(1^\lambda, 1^n) : \mathbf{key} &= (k^d, k^i), \quad k^d \stackrel{\$}{\leftarrow} \text{Gen}^d(1^\lambda), \quad k^i \stackrel{\$}{\leftarrow} \text{Gen}^i(1^\lambda, 1^n), \\ \widehat{R} \stackrel{\$}{\leftarrow} \text{Gb}'(\mathbf{key}, R) : \widehat{R} &\stackrel{\$}{\leftarrow} \text{GProg}'(k^d, k^i, R). \end{aligned}$$

The runtime of  $\text{Garb}'$  is bounded by  $\text{poly}(\lambda, n, |R|)T$ .

- The input/data encoding algorithm  $\text{Encode}'$  runs the  $\text{GInput}$  and  $\text{GData}$  algorithms:

$$(\hat{x}, \widehat{D}) \stackrel{\$}{\leftarrow} \text{Encode}'^D(\mathbf{key} = (k^d, k^i), x) : \hat{x} \stackrel{\$}{\leftarrow} \text{GInput}(k^i, x), \quad \widehat{D} \stackrel{\$}{\leftarrow} \text{GData}^D(k^d).$$

<sup>10</sup>Note that the evaluation runs in linear time in the time bound  $T$ , as opposed to the instance running time  $T_R(x)$ . The weaker efficiency guarantee suffices.

<sup>11</sup>Note that the simulation procedure does not receive the instance running time  $T_R(x)$ . This matches the fact that evaluation procedure does not have instance-based efficiency.

The runtime of  $\text{Encode}'$  is bounded by  $\text{poly}(\lambda, n)|D|$ , and the size of  $\text{Encode}'$  is bounded by  $\text{poly}(\lambda, n)$ .

- The evaluation algorithm  $\text{Eval}'$  runs  $\text{GEval}$  and outputs  $y = R^D(x)$ :

$$y = \text{Eval}'^{\widehat{D}}(\widehat{R}, \widehat{x}) : y = \text{GEval}^{\widehat{D}}(\widehat{R}, \widehat{x}).$$

The runtime of  $\text{Eval}'$  is bounded by  $\text{poly}(\lambda, n, |R|)T$ .

**SIMULATION:** The simulation procedure  $\text{Sim}'$  runs  $\text{Sim.Gen}' := \text{GSim.Gen}$  and  $\text{Sim.Gb}' := \text{GSim.Gb}$ :

$$\begin{aligned} (\tilde{R}, (\tilde{x}, \tilde{D})) &\stackrel{\$}{\leftarrow} \text{Sim}'(1^\lambda, (|x|, |D|, |R|, n, m, S, T), R(x)) : \\ &((\tilde{x}, \tilde{D}), \mathbf{st}) \stackrel{\$}{\leftarrow} \text{Sim.Gen}'(1^\lambda, |x|, |D|), \\ &\tilde{R} \stackrel{\$}{\leftarrow} \text{Sim.Gb}'(1^\lambda, (|x|, |D|, |R|, n, m, S, T), R(x), \mathbf{st}). \end{aligned}$$

Let  $\mathcal{GS}' = (\text{Garb}' = (\text{Gen}', \text{Gb}'), \text{Encode}', \text{Eval}')$  be a basic garbling scheme as described above, with simulation procedure  $\text{Sim}' = (\text{Sim.Gen}', \text{Sim.Gb}')$ . We now construct a garbling scheme  $\mathcal{GS}$  for bounded space RAM with improved efficiency. In particular, it has (1) *linear*-space-dependent complexity and (2) produces garbled RAM with  $\text{poly}(\lambda, |R|)$  overhead (that is, evaluation of  $\widehat{R}, \widehat{x}$  takes  $\text{poly}(\lambda, |R|)T_R(x)$  steps). In comparison, the previous general construction has *polynomial*-space-dependent complexity and  $\text{poly}(\lambda, |R|, S)$  overhead. Toward this goal, we plug  $\mathcal{GS}'$  and  $\text{Sim}'$  into our general construction and make the following modifications.

**MODIFICATION TO STEP 1:** As before, the first step is constructing a nonsuccinct garbling scheme by dividing a RAM computation into small blocks and garbling all of them using  $\mathcal{GS}'$ .

The only, and key, difference is, instead of dividing a  $T$ -step RAM computation into  $T$  1-step “blocks,” divide it into  $\lceil T/S \rceil$   $S$ -step “blocks.” Each block, say the  $t$ th block  $W_t$ , is a RAM algorithm that (i) takes as input a state  $\mathbf{st}_t$  and has access to a memory with content  $D_t$ , (ii) runs  $R$  for  $S$  steps, and (iii) outputs a state  $\mathbf{st}_{t+1}$  for the next block and updates the memory content to  $D_{t+1}$ . (When  $t = 1$ , the memory is initialized with the input  $D_1 = x$ , and  $\mathbf{st}_1$  is set to the initial state.) That is,  $\mathbf{st}_{t+1} = W_t^{D_t}(\mathbf{st}_t)$ . The states have size at most  $|R|$ ,<sup>12</sup> and the memory has size at most  $|S|$ .

As before, each block  $W_t$  is augmented with the encoding algorithm  $\text{Encode}'$  that garbles the state  $\mathbf{st}_{t+1}$  and data  $D_{t+1}$  produced by  $W_t$ , implementing a **key** <sub>$t+1$</sub>  used for garbling the next augmented data, that is, the  $t$ th augmented block  $B^{D_t}(t, \mathbf{key}_{t+1}, \mathbf{st}_t)$  does the following:

1. Run  $W_t^{D_t}(\mathbf{st}_t)$  to obtain output  $\mathbf{st}_{t+1}$  and updated data  $D_{t+1}$ .
2. Run  $\text{Encode}'^{D_{t+1}}(\mathbf{key}_{t+1}, \mathbf{st}_{t+1})$  and obtain garbled state  $\widehat{\mathbf{st}}_{t+1}$  and data  $\widehat{D}_{t+1}$ .
3. Output  $\widehat{\mathbf{st}}_{t+1}, \widehat{D}_{t+1}$ . In addition, if  $\mathbf{st}_{t+1}$  is a final state, output  $y$ .

Then the program  $B^*(t, \mathbf{key}_{t+1}, \star)$  for each augmented block is garbled using  $\text{Garb}'$ , producing garbled blocks.

*Efficiency.* We now analyze various efficiency parameters.

<sup>12</sup>The states are part of the inputs and outputs of the circuit  $C_{\text{CPU}}$  implementing the next step function of  $R$ . The description size of  $R$  is at least  $|C_{\text{CPU}}|$ . Therefore, the sizes of states are bounded by  $|R|$ .

- Each augmented block basically runs  $W_t$  followed by  $\text{Encode}'$ . Since  $\text{key}_{t+1}$  has size  $\text{poly}(\lambda)$ , we have

$$\Psi = |B| = \text{poly}(\lambda, |R|), \quad T_B = \text{poly}(\lambda, |R|)S.$$

The latter follows since  $\text{Encode}'$  runs in linear time in the data size and polynomially in the size of the states (which are inputs to  $\text{Encode}'$ ).

- By the efficiency of  $\text{Gb}'$ , each garbled block has size

$$\Phi = \text{poly}(\lambda, |R|, |B|)T_B = \text{poly}(\lambda, |R|)S.$$

- Overall, there are  $\lceil T/S \rceil$  blocks, resulting in a nonsuccinct garbled RAM  $\widehat{R}$  of size

$$|\widehat{R}| = \lceil T/S \rceil \times \Phi = \text{poly}(\lambda, |R|)T.$$

- We note that for any input  $x$  that takes  $T^*$  steps to evaluate, the output  $R(x)$  is produced after evaluating  $\lceil T^*/S \rceil$  garbled blocks, taking  $\text{poly}(\lambda, |R|)(T^* + S)$  steps.

**MODIFICATION TO STEP 2:** As before, the second step is using obfuscation to “compress” the size of the nonsuccinct garbling scheme constructed in step 1. However, using any obfuscator to obfuscate the program that generates each of  $\lceil T/S \rceil$  garbled blocks leads to an obfuscated program of size at least  $\text{poly}(\lambda, \Phi) = \text{poly}(\lambda, |R|, S)$ . In this case, the complexity of the new garbling scheme is not linear in  $S$ , and the overhead of the produced garbled RAM is at least  $\text{poly}(\lambda, |R|, S)$ .

*Better efficiency.* To avoid the polynomial overhead due to obfuscation, we instead use an  $i\mathcal{O}$  for circuits with quasi-linear complexity  $|C| \text{poly}(\lambda, n)$ , where  $|C|$  is the size of the circuit obfuscated and  $n$  is the length of the input. As shown in section 5.4, such a scheme can be constructed generically from any  $i\mathcal{O}$  (for circuits), puncturable PRF, and randomized encoding that is local (as defined in section 5.4 and satisfied, for instance, by Yao’s garbled circuits), all with  $2^{-(n+\omega(\log \lambda))}$  security.

*Efficiency.* Since the obfuscated programs  $\mathbb{P}_i$ ,  $\mathbb{Q}_i$ , and  $\mathbb{R}_i$  take input a time index  $t$  of length  $O(\log T)$  and output a garbled block computed in time  $\text{poly}(\lambda, |R|)S$  (roughly the same as  $\Phi$ ), the size of the new garbled RAM (and the complexity for generating it) is therefore

$$\text{size of garbled RAM} = \text{poly}(\lambda, |R|)S \times \text{poly}(\lambda, \log T) = \text{poly}(\lambda, |R|) \times S,$$

which is linear in the space complexity of  $R$ .

Moreover, evaluation of an input  $x$  of instance complexity  $T^*$  requires generating and evaluating  $\lceil T^*/S \rceil$  garbled blocks, which takes time

$$\text{run time of garbled RAM} = \lceil T^*/S \rceil \times \text{poly}(\lambda, |R|) \times S = \text{poly}(\lambda, |R|) \times (S + T^*).$$

This concludes the proof of Theorem 3.5.  $\square$

**3.3.1. RAM garbling scheme with complexity linear in the program size.** The RAM garbling scheme of Theorem 3.5 produces garbled RAM of size  $\text{poly}(\lambda, |R|) \times S$  and running time  $\text{poly}(\lambda, |R|)T^*$  (for an input of instance complexity  $T^*$ ), both depending polynomially in the description size of the underlying RAM  $|R|$ . We show that the complexity can be improved to depending linearly on  $|R|$ , that is, the garbled RAM has size  $\text{poly}(\lambda) \times (|R| + S)$  and running time  $\text{poly}(\lambda) \times (|R| + S + T^*)$ .

To achieve this, we need to rely on a basic RAM garbling scheme that satisfies the properties—*independent input encoding and linear complexity*—described above, and the following strengthening: The complexity of the garbling algorithm  $\text{Gb}'$  depends linearly on  $|R|$ , that is,  $\text{poly}(\lambda)(|R|+T)$  (as opposed to  $\text{poly}(\lambda, |R|)T$ ). To obtain such a basic RAM garbling scheme, we observe that there is a universal RAM  $M$ , such that any RAM computation  $R(x)$  can be transformed into computing  $M^R(x, |R|)$ , where the description of  $R$  is provided as a part of the initial memory. The universal machine  $M$  has constant size and  $M^R(x, |R|)$  takes at most  $cT_R(x)$  steps for some constant  $c$  (since each step of  $R$  depends on at most a constant number of bits of the description of  $R$ ). Then applying the construction of [LO13, GHL<sup>+</sup>14] to  $M$  with a persistent database  $R$  yields a garbled RAM of size  $\text{poly}(\lambda)(T + |R|)$  (where  $\text{poly}(\lambda)T$  corresponds to the size of garbling of  $M$  and  $\text{poly}(\lambda)|R|$  corresponds to the garbling of the persistent database  $R$ ).

Now, instantiate the construction of Theorem 3.5 with such a basic RAM garbling scheme, with an additional modification: In step 1, instead of dividing a  $T$ -step RAM computation into  $\lceil T/S \rceil$   $S$ -step blocks, divide it into  $\lceil T/(S + |R|) \rceil$   $(S + |R|)$ -step blocks; the rest of the construction follows identically. We now argue that this construction indeed has complexity linear in  $|R|$ . Each augmented block has the same size as before  $\text{poly}(\lambda) + |R|$ , but a longer running time of  $\text{poly}(\lambda)(S + |R|)$ . By the complexity of the (new) basic RAM garbling scheme, each of the garbled blocks has size  $\text{poly}(\lambda)(S + |R|)$ . Therefore, when obfuscating using an  $i\mathcal{O}$  with quasi-linear complexity, the program that produces the garbled blocks, it leads to a new garbled RAM of size  $\text{poly}(\lambda)(S + |R|)$ . The evaluation of such a garbled RAM with an input  $x$  of instance complexity  $T^*$  takes time  $\lceil T^*/(S + |R|) \rceil \times \text{poly}(\lambda)(S + |R|) = \text{poly}(\lambda)(S + |R| + T^*)$ . Since the construction and analysis is essentially the same as in Theorem 3.5, we omit the details here.

**4. Succinct garbling via reusable obfuscation.** In this section (based on [CHJV15]) we provide alternative proofs for Theorems 3.1 and 3.2. As mentioned in the introduction, the approach here is more direct: We obfuscate the circuit that implements a single computational step of the machine itself. More specifically, this is the circuit that reads the contents of the appropriate cell on the machine's tape, writes a new value to the cell, moves the head to a new tape location, and updates the machine's state according to its transition function. To make sure that the computation remains meaningful and that the evaluator does not learn the intermediate values that the machine writes on its tape, we provide a mechanism for the obfuscated machine to encrypt and authenticate the contents of the tape.

We proceed in three steps. First, in section 4.1, we present a special, **iO**-friendly authenticated encryption scheme, ACE. This scheme, which is a central piece in our solution, will be used to encrypt and authenticate data written to the tape.

Second, in section 4.2 we present and analyze the scheme for garbling Turing machines using **iO** and an ACE scheme.

Finally, in section 4.3 we present and analyze our scheme for garbling RAM machines. The main additional component here is an “**iO**-friendly oblivious RAM” mechanism.

**4.1. Special authenticated encryption (ACE).** In the following sections we give an alternate method for garbling Turing machines and RAM machines from **iO**. These methods rely heavily on a special authenticated encryption scheme, which we define here and call asymmetrically constrained encryption (ACE). An ACE scheme has the following special properties:

1. Ciphertexts are unique: for each message  $m$  and key  $K$ , there is at most a

single string that decrypts to  $m$  under key  $K$ .

2. Both encryption and decryption capabilities under the key  $K$  are *puncturable*. That is, let  $S$  be a subset of the message space which is decidable by a bounded-size circuit  $C$  (i.e.,  $S = \{m : C(m) = 1\}$ ). One can generate a “punctured” encryption key  $EK_C$  which is only able to encrypt messages  $m$  for which  $m \notin S$ .

Similarly, one can generate a constrained *decryption* key  $DK_C$  which, given an encryption of  $m$ , outputs  $m$  only if  $m \notin S$  (and otherwise outputs  $\perp$ ).

We also require that  $DK_C$  be indistinguishable from  $DK$ , even given certain “allowed” auxiliary information. In particular, they remain indistinguishable even given an *encryption* key which is constrained on a superset of  $S$ .

These properties are central in the analysis of our garbling schemes.

#### 4.1.1. Definition.

**DEFINITION 4.1.** *An asymmetrically constrained encryption (ACE) scheme is a 5-tuple of PPT algorithms (Setup, GenEK, GenDK, Enc, Dec) satisfying syntax, correctness, security of constrained decryption, and selective indistinguishability of ciphertexts as described below.*

**Syntax.** The algorithms (Setup, GenEK, GenDK, Enc, Dec) have the following syntax.

- *Setup.*  $\text{Setup}(1^\lambda, 1^n, 1^s)$  is a randomized algorithm that takes as input the security parameter  $\lambda$ , the message length  $n$ , and a “circuit succinctness” parameter  $s$ , all in unary. Setup then outputs a secret key  $SK$ . We think of secret keys as consisting of two parts: an encryption key  $EK$  and a decryption key  $DK$ .

Let  $\mathcal{M} = \{0, 1\}^n$  denote the message space.

- *(Constrained) key generation.* Let  $S \subset \mathcal{M}$  be any set whose membership is decidable by a circuit  $C_S$ . We say that  $S$  is *admissible* if  $|C_S| \leq s$ . Intuitively, the set size parameter  $s$  denotes the upper bound on the size of circuit description of sets to which encryption and decryption keys can be constrained.
  - $\text{GenEK}(SK, C_S)$  takes as input the secret key  $SK$  of the scheme and the description of circuit  $C_S$  for an admissible set  $S$ . It outputs an encryption key  $EK\{S\}$ . We write  $EK$  to denote  $EK\{\emptyset\}$ .
  - $\text{GenDK}(SK, C_S)$  also takes as input the secret key  $SK$  of the scheme and the description of circuit  $C_S$  for an admissible set  $S$ . It outputs a decryption key  $DK\{S\}$ . We write  $DK$  to denote  $DK\{\emptyset\}$ .

Unless mentioned otherwise, we will only consider admissible sets  $S \subset \mathcal{M}$ .

- *Encryption.*  $\text{Enc}(EK', m)$  is a deterministic algorithm that takes as input an encryption key  $EK'$  (that may be constrained) and a message  $m \in \mathcal{M}$  and outputs a ciphertext  $c$  or the reject symbol  $\perp$ .
- *Decryption.*  $\text{Dec}(DK', c)$  is a deterministic algorithm that takes as input a decryption key  $DK'$  (that may be constrained) and a ciphertext  $c$  and outputs a message  $m \in \mathcal{M}$  or the reject symbol  $\perp$ .

**Correctness.** An ACE scheme is correct if the following properties hold.

1. *Correctness of decryption.* For all  $n$ , all  $m \in \mathcal{M}$ , and all sets  $S, S' \subset \mathcal{M}$  such that  $m \notin S \cup S'$ ,

$$\Pr \left[ \text{Dec}(DK, \text{Enc}(EK, m)) = m \mid \begin{array}{l} SK \leftarrow \text{Setup}(1^\lambda), \\ EK \leftarrow \text{GenEK}(SK, C_{S'}), \\ DK \leftarrow \text{GenDK}(SK, C_S) \end{array} \right] = 1.$$

Informally, this says that  $\text{Dec} \circ \text{Enc}$  is the identity on messages which are in neither of the punctured sets.

2. *Equivalence of constrained encryption.* Let  $SK \leftarrow \text{Setup}(1^\lambda)$ . For any message  $m \in \mathcal{M}$  and any sets  $S, S' \subset \mathcal{M}$  with  $m$  not in the symmetric difference  $S \Delta S'$  (i.e., we require that  $m$  is in both  $S$  and  $S'$  or  $m$  is in neither  $S$  nor  $S'$ ),

$$\Pr \left[ \text{Enc}(EK, m) = \text{Enc}(EK', m) \mid \begin{array}{l} SK \leftarrow \text{Setup}(1^\lambda), \\ EK \leftarrow \text{GenEK}(SK, C_S), \\ EK' \leftarrow \text{GenEK}(SK, C_{S'}) \end{array} \right] = 1.$$

3. *Unique ciphertexts.* With high probability over  $SK \leftarrow \text{Setup}(1^\lambda)$ , it holds for any  $c$  and  $c'$  that if  $\text{Dec}(DK, c) = \text{Dec}(DK, c') \neq \perp$ , then  $c = c'$ .
4. *Safety of constrained decryption.* For all strings  $c$  and all  $S \subset \mathcal{M}$ ,

$$\Pr [\text{Dec}(DK, c) \in S \mid SK \leftarrow \text{Setup}(1^\lambda), DK \leftarrow \text{GenDK}(SK, C_S)] = 0.$$

This says that a punctured key  $DK\{S\}$  will never decrypt a string  $c$  to a message in  $S$ .

5. *Equivalence of constrained decryption.* If  $\text{Dec}(DK\{S\}, c) = m \neq \perp$  and  $m \notin S'$ , then  $\text{Dec}(DK\{S'\}, c) = m$ .

**Security of constrained decryption.** Intuitively, this property says that for any two sets  $S_0, S_1$ , no adversary can distinguish between the constrained key  $DK\{S_0\}$  and  $DK\{S_1\}$ , even given additional auxiliary information in the form of a constrained encryption key  $EK'$  and ciphertexts  $c_1, \dots, c_t$ . To rule out trivial attacks,  $EK'$  is constrained at least on  $S_0 \Delta S_1$ . Similarly, each  $c_i$  is an encryption of a message  $m \notin S_0 \Delta S_1$ .

Formally, we describe security of constrained decryption as a multistage game between an adversary  $\mathcal{A}$  and a challenger.

- *Setup.*  $\mathcal{A}$  chooses sets  $S_0, S_1, U$  such that  $S_0 \Delta S_1 \subseteq U \subseteq \mathcal{M}$  and sends its circuit descriptions  $(C_{S_0}, C_{S_1}, C_U)$  to the challenger.  $\mathcal{A}$  also sends arbitrary polynomially many messages  $m_1, \dots, m_t$  such that  $m_i \notin S_0 \Delta S_1$ .

The challenger chooses a bit  $b \in \{0, 1\}$  and computes the following:

1.  $SK \leftarrow \text{Setup}(1^\lambda)$ .
2.  $DK\{S_b\} \leftarrow \text{GenDK}(SK, C_{S_b})$ .
3.  $EK \leftarrow \text{GenEK}(SK, \emptyset)$ .
4.  $c_i \leftarrow \text{Enc}(EK, m_i)$  for every  $i \in [t]$ .
5.  $EK\{U\} \leftarrow \text{GenEK}(SK, C_U)$ .

Finally, it sends the tuple  $(EK\{U\}, DK\{S_b\}, \{c_i\})$  to  $\mathcal{A}$ .

- *Guess.*  $\mathcal{A}$  outputs a bit  $b' \in \{0, 1\}$ .

The advantage of  $\mathcal{A}$  in this game (on security parameter  $\lambda$ ) is defined as  $\text{adv}_{\mathcal{A}} = |\Pr[b' = b] - \frac{1}{2}|$ . We require that for all PPT  $\mathcal{A}$ ,  $\text{adv}_{\mathcal{A}}(\lambda)$  is negligible in  $\lambda$ .



**Selective indistinguishability of ciphertexts.** Intuitively, this property says that no adversary can distinguish encryptions of  $m_0$  from encryptions of  $m_1$ , even given certain auxiliary information. The auxiliary information corresponds to constrained encryption and decryption keys  $EK', DK'$ , as well as some ciphertexts  $c_1, \dots, c_t$ . In order to rule out trivial attacks,  $EK'$  and  $DK'$  should both be punctured on at least  $\{m_0, m_1\}$ , and none of  $c_1, \dots, c_t$  should be an encryption of  $m_0$  or  $m_1$ . Let both  $\mathcal{F}_1$  and  $\mathcal{F}_2$  be subexponentially secure.

Formally, we require that for all sets  $S, U \subset \mathcal{M}$ , for all  $m_0^*, m_1^* \in S \cap U$ , and for all  $m_1, \dots, m_t \in \mathcal{M} \setminus \{m_0^*, m_1^*\}$ , the distribution

$$EK\{S\}, DK\{U\}, c_0^*, c_1^*, c_1, \dots, c_t$$

is computationally indistinguishable from

$$EK\{S\}, DK\{U\}, c_1^*, c_0^*, c_1, \dots, c_t$$

in the probability space defined by sampling  $SK \leftarrow \text{Setup}(1^\lambda)$ ,  $EK \leftarrow \text{GenEK}(SK, \emptyset)$ ,  $EK\{S\} \leftarrow \text{GenEK}(SK, C_S)$ ,  $DK\{U\} \leftarrow \text{GenDK}(SK, C_U)$ ,  $c_b^* \leftarrow \text{Enc}(EK, m_b^*)$ , and  $c_i \leftarrow \text{Enc}(EK, m_i)$ .

**4.1.2. Construction.** We now present a construction of an asymmetrically constrainable encryption scheme  $ACE = (\text{Setup}, \text{GenEK}, \text{GenDK}, \text{Enc}, \text{Dec})$ . Our scheme is based on **iO** and one-way functions that are subexponentially secure and is based on the “hidden triggers” mechanism in the deniable encryption scheme of [SW14].

$\text{Setup}(1^\lambda, 1^n, 1^s)$ : Let  $\mathcal{F}_1 = (\text{PRF.Gen}_1, \text{PRF.Punc}_1, F_1)$  be a subexponentially secure puncturable PRF (PPRF) family with domain  $\{0, 1\}^n$  and codomain  $\{0, 1\}^{3n+\lambda}$  such that  $\mathcal{F}_1$  is (statistically) injective with failure probability  $2^{-n-\lambda}$ , as defined and constructed in [SW14, Theorem 2]. Let  $\mathcal{F}_2 = (\text{PRF.Gen}_2, \text{PRF.Punc}_2, F_2)$  be a (noninjective) subexponentially secure PPRF family with domain  $\{0, 1\}^{3n+\lambda}$  and codomain  $\{0, 1\}^n$ .

Let  $\lambda'$  denote a security parameter chosen as a function of  $\lambda$  and  $n$  to ensure that any nonuniform  $\text{poly}(\lambda)$ -size adversary  $\mathcal{A}$  has  $2^{-n} \cdot \text{negl}(\lambda)$  advantage in the PRF games associated to  $\mathcal{F}_1$  and  $\mathcal{F}_2$ , as well as in the **iO** game associated with  $i\mathcal{O}(1^{\lambda'}, \cdot)$ . Since we assumed that  $\mathcal{F}_1, \mathcal{F}_2$ , and  $i\mathcal{O}$  were subexponentially secure, this is achievable (indeed with *subexponential*  $\text{negl}$ ) with some  $\lambda' = \text{poly}(\lambda)$ .

$\text{Setup}$  samples keys  $K_1 \leftarrow \text{PRF.Gen}_1(1^{\lambda'})$  and  $K_2 \leftarrow \text{PRF.Gen}_2(1^{\lambda'})$ , and outputs  $(K_1, K_2, p)$ , where  $p = \text{poly}(\lambda', n, s)$  is a fixed polynomial “padding” chosen to be larger than  $\mathcal{G}_{\text{enc}}$  and  $\mathcal{G}_{\text{dec}}$  in each hybrid distribution of the security proof (see Remarks 4.4 and 4.6). We will write  $F_i$  to denote the function  $F_i(K_i, \cdot)$ .

$\text{GenEK}((K_1, K_2, p), C_S)$ : The encryption key generation algorithm takes as input keys  $K_1, K_2$  and the circuit description  $C_S$  of an admissible set  $S$ . It prepares a circuit representation of  $\mathcal{G}_{\text{enc}}$  (Figure 4.1), padded to be of size  $p$ . Next, it computes the encryption key  $EK\{S\} \leftarrow i\mathcal{O}(\mathcal{G}_{\text{enc}})$  and outputs the result.

$\text{GenDK}((K_1, K_2, p), C_S)$ : The decryption key generation algorithm takes as input keys  $K_1, K_2$  and the circuit description  $C_S$  of an admissible set  $S$ . It prepares a circuit representation of  $\mathcal{G}_{\text{dec}}$  (Figure 4.2), padded to be of size  $p$ . It then computes the decryption key  $DK\{S\} \leftarrow i\mathcal{O}(1^{\lambda'}, \mathcal{G}_{\text{dec}})$  and outputs the result.

$\text{Enc}(EK', m)$ : The encryption algorithm simply runs the encryption key program  $EK'$  on message  $m$  to compute the ciphertext  $c \leftarrow EK'(m)$ .

**Input:** Message  $m \in \{0, 1\}^n$ .  
**Constants:**  $K_1, K_2$ , circuit  $C_S$ .  
 1. If  $m \in S$ , output  $\perp$ .  
 2. Otherwise, output  $\alpha\|\beta$ , where  $\alpha = F_1(m)$  and  $\beta = F_2(\alpha) \oplus m$ .

FIG. 4.1. (Constrained) encryption  $\mathcal{G}_{\text{enc}}$ .

**Input:** Ciphertext  $c \in \{0, 1\}^{3n+\lambda}$ .  
**Constants:**  $K_1, K_2$ , circuit  $C_S$ .  
 1. Parse  $c$  as  $\alpha\|\beta$  with  $\beta \in \{0, 1\}^n$ , and compute  $m = F_2(\alpha) \oplus \beta$ .  
 2. If  $m \in S$ , or if  $\alpha \neq F_1(m)$ , then output  $\perp$ .  
 3. Otherwise, output  $m$ .

FIG. 4.2. (Constrained) decryption  $\mathcal{G}_{\text{dec}}$ .

$\text{Dec}(DK', c)$ : The decryption algorithm simply runs the decryption key program  $DK'$  on the input ciphertext  $c$  and returns the output  $DK'(c)$ .

*Remark 4.2.* Subexponential security of **iO** and one-way functions was required only for  $2^{-n} \cdot \text{negl}(\lambda)$  security against  $\text{poly}(\lambda)$ -time adversaries with relatively small  $\lambda'$  ( $= \text{poly}(\lambda)$ ). If we only had polynomially secure **iO** and one-way functions, we would need  $\lambda' = \lambda \cdot 2^{\epsilon n}$  for any  $\epsilon > 0$ , which in general could be extremely large.

However, in our applications to garbling polynomial-time Turing machine and RAM computations on inputs of length  $\text{poly}(\lambda)$ , we use ACE with  $n = O(\log \lambda)$ . Thus in this case we can achieve the requisite security with  $\lambda' = \text{poly}(\lambda)$ , assuming only polynomially secure **iO** and one-way functions.

This completes the description of our construction of ACE.

#### 4.1.3. Proof of correctness and security.

**Correctness.** We first argue correctness:

1. *Correctness of decryption.* This follows directly from the definitions of  $\mathcal{G}_{\text{enc}}$  and  $\mathcal{G}_{\text{dec}}$  given in Figures 4.1 and 4.2, along with the perfect correctness of  $i\mathcal{O}$ .
2. *Equivalence of constrained encryption.* This follows directly from the definition of  $\mathcal{G}_{\text{enc}}$  and the perfect correctness of  $i\mathcal{O}$ .
3. *Uniqueness of encryptions.* A ciphertext  $c = \alpha\|\beta$  decrypts to  $m \neq \perp$  only if  $\alpha = F_1(m)$  and  $\beta = F_2(\alpha)$ . Since  $F_1$  is injective with probability at least  $1 - 2^{-\lambda}$ , there can be only one ciphertext which decrypts to  $m$ .
4. *Safety of constrained decryption.* This follows directly from the definition of  $\mathcal{G}_{\text{dec}}$  and the perfect correctness of  $i\mathcal{O}$ .
5. *Equivalence of constrained decryption.* This follows directly from the definition of  $\mathcal{G}_{\text{dec}}$  and the perfect correctness of  $i\mathcal{O}$ .

**Security of constrained decryption.** We now prove that ACE satisfies security of constrained decryption.

**LEMMA 4.3.** *The proposed scheme ACE satisfies security of constrained decryption.*

*Proof.* Let  $S_0, S_1, U$  be arbitrary subsets of  $\{0, 1\}^n$  such that  $S_0 \Delta S_1 \subseteq U$ , and let  $C_{S_0}, C_{S_1}, C_U$  be their circuit descriptions. Let  $m_1, \dots, m_t$  be arbitrary messages such

that every  $m_i \in M \setminus (S_0 \Delta S_1)$ . We argue that no PPT distinguisher can distinguish between

$$(EK\{U\}, DK\{S_0\}, \{c_i\})$$

and

$$(EK\{U\}, DK\{S_1\}, \{c_i\})$$

in the probability space defined by sampling the following:

$$\begin{aligned} SK &\leftarrow \text{Setup}(1^\lambda), \\ EK &\leftarrow \text{GenEK}(SK, \emptyset), \\ EK\{U\} &\leftarrow \text{GenEK}(SK, C_U). \\ \text{For every } i \in [t] & \\ c_i &:= \text{Enc}(EK, m_i). \\ \text{For } b \in \{0, 1\} & \\ DK\{S_b\} &:= \text{GenDK}(SK, C_{S_b}). \end{aligned}$$

Without loss of generality, we will suppose that  $S_0 \subseteq S_1$ . The general case follows because then

$$\begin{aligned} (EK\{U\}, DK\{S_0\}, \{c_i\}) &\approx_c (EK\{U\}, DK\{S_0 \cap S_1\}, \{c_i\}) \\ &\approx_c (EK\{U\}, DK\{S_1\}, \{c_i\}). \end{aligned}$$

We now prove the lemma by exhibiting a sequence of  $|S_1 \setminus S_0| + 1$  hybrid experiments  $H_0, \dots, H_{|S_1 \setminus S_0|+1}$ , each of which is  $2^{-n} \cdot \text{negl}(\lambda)$ -indistinguishable from the previous, such that

$$H_0 \approx_c (EK\{U\}, DK\{S_0\}, \{c_i\})$$

and

$$H_{|S_1 \setminus S_0|} \approx_c (EK\{U\}, DK\{S_1\}, \{c_i\}).$$

We now proceed to give details. Let  $u_i$  denote the lexicographically  $i$ th element of  $S_1 \setminus S_0$ . Throughout the experiments, we will refer to the encryption key and decryption key programs given to the distinguisher as  $EK'$  and  $DK'$ , respectively. Similarly, (unless stated otherwise) we will refer to the unobfuscated algorithms underlying  $EK'$  and  $DK'$  as  $\mathcal{G}'_{\text{enc}}$  and  $\mathcal{G}'_{\text{dec}}$ , respectively.

**HYBRID  $H_i$  (FOR  $1 \leq i \leq |S_1 \setminus S_0|$ ):** In the  $i$ th hybrid, the decryption key program  $\mathcal{G}'_{\text{dec}}$  first checks whether  $m \in S_1$  and  $m \leq u_i$ . If this is the case, then it simply outputs  $\perp$ . Otherwise, it behaves in the same manner as  $DK\{S_0\}$ . The underlying unobfuscated program  $\mathcal{G}'_{\text{dec}}$  is described in Figure 4.3.

Therefore, in hybrid  $H_0$ ,  $\mathcal{G}'_{\text{dec}}$  has the same functionality as  $DK\{S_0\}$ , and in  $H_{|S_1 \setminus S_0|}$ ,  $\mathcal{G}'_{\text{dec}}$  has the same functionality as  $DK\{S_1\}$ .

**Input:** Ciphertext  $c \in \{0, 1\}^{4n+\lambda}$ .

**Constants:** PPRF keys  $K_1, K_2$ , circuits  $C_{S_0}, C_{S_1}$ .

1. Parse  $c$  as  $\alpha\|\beta$  with  $\beta \in \{0, 1\}^n$ , and compute  $m = F_2(\alpha) \oplus \beta$ .
2. If  $m \leq u_i$  and  $m \in S_1$ , or  $m > u_i$  and  $m \in S_0$ , or  $\alpha \neq F_1(m)$ , output  $\perp$ .
3. Otherwise, output  $m$ .

FIG. 4.3. (Constrained) decryption  $\mathcal{G}'_{\text{dec}}$  in hybrid  $i$ .

We now construct a series of intermediate hybrid experiments  $H_{i,0}, \dots, H_{i,7}$ , where  $H_{i,0}$  is the same as  $H_i$  and  $H_{i,7}$  is the same as  $H_{i+1}$ . For every  $j$ , we will prove that  $H_{i,j}$  and  $H_{i,j+1}$  are  $2^{-n} \cdot \text{negl}(\lambda)$ -indistinguishable, which establishes that  $H_i$  and  $H_{i+1}$  are also  $2^{-n} \cdot \text{negl}(\lambda)$ -indistinguishable.

HYBRID  $H_{i,0}$ : This is the same as experiment  $H_i$ .

HYBRID  $H_{i,1}$ : This is the same as experiment  $H_{i,0}$  except that we modify  $\mathcal{G}'_{\text{dec}}$  as in Figure 4.4. If the decrypted message  $m$  is  $u_{i+1}$ , then instead of checking whether  $\alpha \neq F_1(m)$  in line 5,  $\mathcal{G}'_{\text{dec}}$  now checks whether  $\text{Commit}(0; \alpha) \neq \text{Commit}(0; F_1(u_{i+1}))$ , where  $\text{Commit}$  is an injective noninteractive commitment.

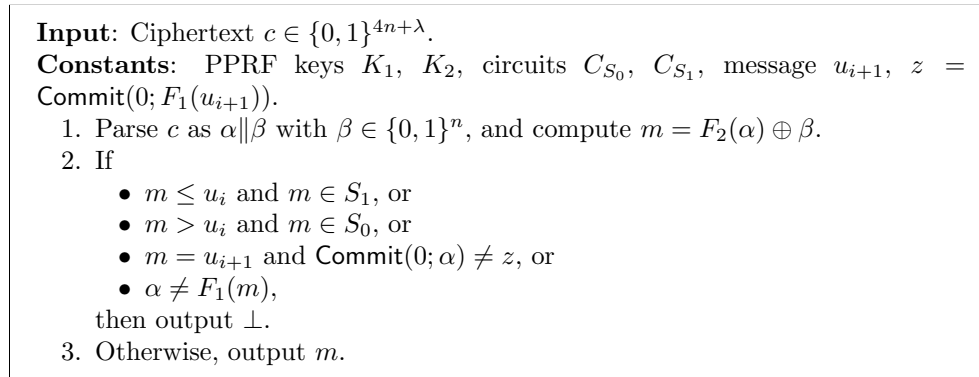


FIG. 4.4.  $\mathcal{G}'_{\text{dec}}$  in hybrid  $H_{i,1}$ .

HYBRID  $H_{i,2}$ : This is the same as experiment  $H_{i,1}$  except that we modify  $\mathcal{G}'_{\text{dec}}$  such that the PRF key  $K_1$  in  $\mathcal{G}'_{\text{dec}}$  is punctured at  $u_{i+1}$ ; i.e.,  $K_1$  is replaced with  $K_1\{u_{i+1}\} \leftarrow \text{Puncture}_{\text{PRF}}(K_1, u_{i+1})$ .

HYBRID  $H_{i,3}$ : This is the same as experiment  $H_{i,2}$  except that the PRF key  $K_1$  hardwired in  $\mathcal{G}'_{\text{enc}}$  is replaced with the same punctured key  $K_1\{u_{i+1}\} \leftarrow \text{Puncture}_{\text{PRF}}(K_1, u_{i+1})$  as is used in  $\mathcal{G}'_{\text{dec}}$ .

HYBRID  $H_{i,4}$ : This is the same as experiment  $H_{i,3}$  except that the hardwired value  $z$  in  $\mathcal{G}'_{\text{dec}}$  is now computed as  $\text{Commit}(0; r)$ , where  $r$  is a randomly chosen string in  $\{0, 1\}^{3n+\lambda}$ .

HYBRID  $H_{i,5}$ : This is the same as experiment  $H_{i,4}$  except that the hardwired value  $z$  in  $\mathcal{G}'_{\text{dec}}$  is now set to  $\text{Commit}(1; r)$ , where  $r$  is a randomly chosen string in  $\{0, 1\}^{3n+\lambda}$ .

HYBRID  $H_{i,6}$ : This is the same as experiment  $H_{i,5}$  except that we now modify  $\mathcal{G}'_{\text{dec}}$  such that it outputs  $\perp$  when the decrypted message  $m$  is  $u_{i+1}$ . An equivalent description of  $\mathcal{G}'_{\text{dec}}$  is that in line 3 it now checks whether  $m \leq u_{i+1}$  instead of  $m \leq u_i$ .

HYBRID  $H_{i,7}$ : This is the same as experiment  $H_{i,6}$  except that the PRF key corresponding to  $F_1$  is unpunctured in both  $\mathcal{G}'_{\text{enc}}$  and  $\mathcal{G}'_{\text{dec}}$ . That is, we replace  $K_1\{u_{i+1}\}$  with  $K_1$  in both  $\mathcal{G}'_{\text{enc}}$  and  $\mathcal{G}'_{\text{dec}}$ . Note that experiment  $H_{i,7}$  is the same as experiment  $H_{i+1}$ .

This completes the description of the hybrid experiments.

*Remark 4.4.* There is a fixed polynomial  $\text{poly}$  such that in each hybrid experiment  $H_{i,j}$ ,  $|\mathcal{G}_{\text{enc}}|$  and  $|\mathcal{G}_{\text{dec}}|$  are bounded by  $\text{poly}(\lambda', n, s)$ .

We now argue their  $2^{-n} \cdot \text{negl}(\lambda)$ -indistinguishability.

*Indistinguishability of  $H_{i,0}$  and  $H_{i,1}$ .* Since  $\text{Commit}$  is injective, the following two checks are equivalent:  $\alpha \neq F_1(m)$  and  $\text{Commit}(0; \alpha) \neq \text{Commit}(0; F_1(m))$ . Thus, the algorithms  $\mathcal{G}'_{\text{dec}}$  in  $H_{i,0}$  and  $H_{i,1}$  are functionally equivalent. Therefore, the indistinguishability of  $H_{i,0}$  and  $H_{i,1}$  follows from the security of  $i\mathcal{O}$ .

*Indistinguishability of  $H_{i,1}$  and  $H_{i,2}$ .* Let  $\mathcal{G}'_{\text{dec}}$  (resp.,  $\mathcal{G}''_{\text{dec}}$ ) denote the unobfuscated algorithms underlying the decryption key program  $DK'$  in experiments  $H_{i,1}$  (resp.,  $H_{i,2}$ ). We will argue that  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  are functionally equivalent. The indistinguishability of  $H_{i,0}$  and  $H_{i,1}$  then follows from the security of the indistinguishability obfuscator  $i\mathcal{O}$ .

Let  $c_{i+1} = \alpha_{i+1} \parallel \beta_{i+1}$  denote the *unique* ciphertext such that  $\text{Dec}(DK, c_{i+1}) = u_{i+1}$  (where  $DK$  denotes the unconstrained decryption key program). First note that on any input  $c \neq c_{i+1}$ , both  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  have identical behavior, except that  $\mathcal{G}'_{\text{dec}}$  uses the PRF key  $K_1$ , while  $\mathcal{G}''_{\text{dec}}$  uses the punctured PRF key  $K_1\{u_{i+1}\}$ . Since the punctured PRF scheme preserves functionality under puncturing, we have that  $\mathcal{G}'_{\text{dec}}(c) = \mathcal{G}''_{\text{dec}}(c)$ . Now, on input  $c_{i+1}$ , after decrypting to obtain  $u_{i+1}$ ,  $\mathcal{G}'_{\text{dec}}$  computes  $\text{Commit}(0; F_1(u_{i+1}))$  and then checks whether  $\text{Commit}(0; \alpha_{i+1}) \neq \text{Commit}(0; F_1(u_{i+1}))$ , whereas  $\mathcal{G}''_{\text{dec}}$  simply checks whether  $\text{Commit}(0; \alpha_i) \neq z$ . But since the value  $z$  hardwired in  $\mathcal{G}''_{\text{dec}}$  is equal to  $\text{Commit}(0; F_1(u_{i+1}))$ , we have that  $\mathcal{G}'_{\text{dec}}(c_i) = \mathcal{G}''_{\text{dec}}(c_i)$ .

Thus we have that  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  are functionally equivalent.

*Indistinguishability of  $H_{i,2}$  and  $H_{i,3}$ .* Let  $\mathcal{G}'_{\text{enc}}$  denote the circuit such that in  $H_{i,2}$ , the encryption key  $EK'$  is  $i\mathcal{O}(1^{\lambda'}, \mathcal{G}'_{\text{enc}})$ .  $H_{i,3}$  is obtained from  $H_{i,2}$  by replacing the hard-coded PPRF key  $K_1$  with the punctured key  $K_1\{u_{i+1}\}$ . By the security of  $i\mathcal{O}$ , it suffices to show that this is a functionality-preserving change. Since the punctured PRF preserves functionality at unpunctured points, it suffices to show that  $\mathcal{G}'_{\text{enc}}$  does not depend on the value  $F_1(u_{i+1})$  (this notion of dependence is well-defined because  $\mathcal{G}'_{\text{enc}}$  uses  $F_1$  as a black box). This is true simply because  $\mathcal{G}'_{\text{enc}}$  evaluates  $F_1$  on  $u_{i+1}$  only when its input  $m = u_{i+1}$ . But when this happens,  $\mathcal{G}'_{\text{enc}}$  outputs  $\perp$  unconditionally because  $u_{i+1} \in S_1 \subseteq U$ .

*Indistinguishability of  $H_{i,3}$  and  $H_{i,4}$ .* From the security of the punctured PRF, it follows immediately that  $H_{i,3}$  and  $H_{i,4}$  are computationally indistinguishable.

*Indistinguishability of  $H_{i,4}$  and  $H_{i,5}$ .*  $H_{i,4}$  and  $H_{i,5}$  are computationally indistinguishable because of the hiding properties of  $\text{Commit}$ .

*Indistinguishability of  $H_{i,5}$  and  $H_{i,6}$ .* Let  $\mathcal{G}'_{\text{dec}}$  (resp.,  $\mathcal{G}''_{\text{dec}}$ ) denote the unobfuscated algorithms underlying the decryption key program  $DK'$  in experiments  $H_{i,5}$  and  $H_{i,6}$ . We will argue that with all but negligible probability,  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  are functionally equivalent. The indistinguishability of  $H_{i,5}$  and  $H_{i,6}$  then follows from the security of the indistinguishability obfuscator  $i\mathcal{O}$ .

Let  $c_i$  denote the *unique* ciphertext corresponding to the message  $u_i$ . We note that with overwhelming probability, the random string  $z$  (hardwired in both  $H_{i,5}$  and  $H_{i,6}$ ) is not in the image of the PRG. Thus, except with negligible probability, there does not exist an  $\alpha_i$  such that  $\text{PRG}(\alpha_i) = z$ . This implies that except with negligible probability,  $\mathcal{G}'_{\text{dec}}(c_i) = \perp$ . Since  $\mathcal{G}''_{\text{dec}}$  also outputs  $\perp$  on input  $c_i$  and  $\mathcal{G}'_{\text{dec}}, \mathcal{G}''_{\text{dec}}$  behave identically on all other input ciphertexts, we have that  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  are functionally equivalent.

*Indistinguishability of  $H_{i,6}$  and  $H_{i,7}$ .* This follows in the same manner as the indistinguishability of experiments  $H_{i,2}$  and  $H_{i,3}$ . We omit the details.  $\square$

**Selective indistinguishability of ciphertexts.** We now prove that *ACE* satisfies indistinguishability of ciphertexts.

LEMMA 4.5. *The proposed scheme ACE satisfies selective indistinguishability of ciphertexts.*

*Proof.* The proof of the lemma proceeds in a sequence of hybrid distributions where we make indistinguishable changes to  $EK\{U\}$ ,  $DK\{S\}$ , and the challenge ciphertexts  $(c_0^*, c_1^*)$ . The “extra” ciphertexts  $c_1, \dots, c_t$  remain unchanged throughout the hybrids.

HYBRID  $H_0$ : This is the real-world distribution. For completeness (and to ease the presentation of the subsequent hybrid distributions), we describe the sampling process here. Let  $S, U \subset \mathcal{M} = \{0, 1\}^n$  be the sets chosen by the adversary, and let  $C_S, C_U$  be their corresponding circuit descriptions. Let  $m_0^*, m_1^*$  be the challenge messages in  $S \cap U$ , and let  $(m_1, \dots, m_t)$  be the extra messages in  $\{0, 1\}^n$ . Then the following hold:

1. Sample the PRF keys  $K_1 \leftarrow \text{Gen}(1^\lambda)$ ,  $K_2 \leftarrow \text{Gen}(1^\lambda)$ ,  $b \leftarrow \{0, 1\}$ .
2. For  $b \in \{0, 1\}$ , compute  $\alpha_b^* \leftarrow F_1(m_b^*)$ ,  $\gamma_b^* \leftarrow F_2(\alpha_b^*)$ , and  $\beta_b^* = \gamma_b^* \oplus m_b^*$ . Let  $c_b^* = \alpha_b^* \parallel \beta_b^*$ .
3. For every  $j \in [t]$ , compute  $\alpha_j \leftarrow F_1(m_j)$ ,  $\gamma_j \leftarrow F_2(\alpha_j)$ , and  $\beta_j = \gamma_j \oplus m_j$ . Let  $c_j = \alpha_j \parallel \beta_j$ .
4. Compute  $EK\{U\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}'_{\text{enc}})$ , where  $\mathcal{G}'_{\text{enc}}$  is described in Figure 4.5, but padded to size  $p$ .
5. Compute  $DK\{S\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}'_{\text{dec}})$ , where  $\mathcal{G}'_{\text{dec}}$  is described in Figure 4.6, but padded to size  $p$ .
6. Output the following tuple:

$$(EK\{S\}, DK\{U\}, c_b^*, c_{1-b}^*, c_1, \dots, c_t).$$

**Input:** Message  $m \in \{0, 1\}^n$ .  
**Constants:**  $K_1, K_2$ , circuit  $C_U$ .

1. If  $m \in U$ , output  $\perp$ .
2. Otherwise, output  $\alpha \parallel \beta$  where  $\alpha = F_1(m)$  and  $\beta = F_2(\alpha) \oplus m$ .


FIG. 4.5.  $\mathcal{G}'_{\text{enc}}$  in hybrid  $H_0$ .

**Input:** Ciphertext  $c \in \{0, 1\}^{3n+\lambda}$ .  
**Constants:**  $K_1, K_2$ , circuit  $C_S$ .

1. Parse  $c$  as  $\alpha \parallel \beta$  with  $\beta \in \{0, 1\}^n$ , and compute  $m = F_2(\alpha) \oplus \beta$ .
2. If  $m \in S$ , or if  $\alpha \neq F_1(m)$ , then output  $\perp$ .
3. Otherwise, output  $m$ .

FIG. 4.6.  $\mathcal{G}'_{\text{dec}}$  in hybrid  $H_0$ .

HYBRID  $H_1$ : Modify  $\mathcal{G}'_{\text{enc}}$ : the hardwired PRF key  $K_1$  is replaced with a punctured key  $K_1\{m_0^*, m_1^*\} \leftarrow \text{Puncture}_{\text{PRF}}(K_1, \{m_0^*, m_1^*\})$ .

FIG. 4.7.  $\mathcal{G}'_{\text{dec}}$  in hybrid 3.


HYBRID H<sub>2</sub>: Modify  $\mathcal{G}'_{\text{dec}}$ : the hardwired PRF key  $K_1$  is replaced with a punctured key  $K_1\{m_0^*, m_1^*\} \leftarrow \text{Puncture}_{\text{PRF}}(K_1, \{m_0^*, m_1^*\})$ .

HYBRID H<sub>3</sub>: Modify  $\mathcal{G}'_{\text{dec}}$ : Perform the following check in the beginning. If input ciphertext  $c = c_b^*$  for  $b \in \{0, 1\}$ , then output  $\perp$ . The modified  $\mathcal{G}'_{\text{dec}}$  is described in Figure 4.7.

HYBRID H<sub>4</sub>: Modify challenge ciphertexts  $c_b^* = \alpha_b^* \|\beta_b^*$ : Generate each  $\alpha_b^*$  as a truly random string.

HYBRID H<sub>5</sub>: Modify  $\mathcal{G}'_{\text{enc}}$ : the hardwired PRF key  $K_2$  is replaced with a punctured key  $K_2\{\alpha_0^*, \alpha_1^*\} \leftarrow \text{Puncture}_{\text{PRF}}(K_2, \{\alpha_0^*, \alpha_1^*\})$ .

HYBRID H<sub>6</sub>: Modify  $\mathcal{G}'_{\text{dec}}$ : we change the check performed in line 1 of Figure 4.7. For any input ciphertext  $c = \alpha\|\beta$ , if  $\alpha = \alpha_b^*$  for  $b \in \{0, 1\}$ , then output  $\perp$ . Note that  $\mathcal{G}'_{\text{dec}}$  only has  $\alpha_b^*$  hardwired, as opposed to  $c_b^*$ . The modified  $\mathcal{G}'_{\text{dec}}$  is described in Figure 4.8.

FIG. 4.8.  $\mathcal{G}'_{\text{dec}}$  in hybrid 6.

HYBRID H<sub>7</sub>: Modify  $\mathcal{G}'_{\text{dec}}$ : the hardwired PRF key  $K_2$  is replaced with the same punctured key  $K_2\{\alpha_0^*, \alpha_1^*\} \leftarrow \text{Puncture}_{\text{PRF}}(K_2, \{\alpha_0^*, \alpha_1^*\})$  as was used in  $\mathcal{G}'_{\text{enc}}$ .

HYBRID H<sub>8</sub>: Modify challenge ciphertexts  $c_b^* = \alpha_b^* \|\beta_b^*$ : For  $b \in \{0, 1\}$ , generate  $\beta_b^*$  as a truly random string.

This completes the description of the hybrid experiments.

*Remark 4.6.* There is a fixed polynomial  $\text{poly}$  such that in each hybrid experiment  $H_0, \dots, H_8$ ,  $|\mathcal{G}_{\text{enc}}|$  and  $|\mathcal{G}_{\text{dec}}|$  are bounded by  $\text{poly}(\lambda', n, s)$ .

We will now first prove indistinguishability of experiments  $H_i$  and  $H_{i+1}$  for every  $i$ . We will then observe that no adversary can guess bit  $b$  in the final hybrid  $H_8$  with probability better than  $\frac{1}{2}$ . This suffices to prove the claim.

*Indistinguishability of  $H_0$  and  $H_1$ .* Let  $\mathcal{G}'_{\text{enc}}$  denote the circuit such that in  $H_0$ ,

$EK\{U\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}'_{\text{enc}})$ .  $H_1$  is obtained from  $H_0$  by replacing the PPRF key  $K_1$  in  $\mathcal{G}'_{\text{enc}}$  by the punctured key  $K_1\{m_0^*, m_1^*\}$ . By the security of  $i\mathcal{O}$ , it suffices to show that this change is functionality-preserving. Since the punctured PRF scheme preserves functionality at unpunctured points, it suffices to show that  $\mathcal{G}'_{\text{enc}}$ 's functionality does not depend on  $F_1(m_0^*)$  or  $F_1(m_1^*)$ . This is true because  $\mathcal{G}'_{\text{enc}}$  uses  $F_1$  as a black box and evaluates it at  $m_0^*$  or  $m_1^*$  only when its input  $m$  is  $m_0^*$  or  $m_1^*$ . But when this happens,  $\mathcal{G}'_{\text{enc}}$  outputs  $\perp$  (by the check at line 1) no matter what the value of  $F_1$  is.

*Indistinguishability of  $H_1$  and  $H_2$ .* Let  $\mathcal{G}'_{\text{dec}}$  denote the circuit such that in  $H_1$ ,  $DK\{S\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}'_{\text{dec}})$ .  $H_2$  is obtained from  $H_1$  by replacing the PPRF key  $K_1$  in  $\mathcal{G}'_{\text{dec}}$  with a punctured key  $K_1\{m_0^*, m_1^*\}$ . Note that due to the check performed in line 2 of Figure 4.6,  $\mathcal{G}'_{\text{dec}}$  does not depend on the value of  $F_1(m_0^*)$  or  $F_1(m_1^*)$ : if  $m = m_b^*$ , then  $m \in S$ , so  $\mathcal{G}'_{\text{dec}}$  outputs  $\perp$  unconditionally. Then, since the punctured PRF scheme preserves functionality under puncturing, replacing  $K_1$  by  $K_1\{m_0^*, m_1^*\}$  is a functionality-preserving change. The indistinguishability of  $H_1$  and  $H_2$  then follows from the security of the indistinguishability obfuscator  $i\mathcal{O}$ .

*Indistinguishability of  $H_2$  and  $H_3$ .* Let  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  denote the circuits such that in  $H_2$ ,  $DK\{S\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}'_{\text{dec}})$ , and in  $H_3$ ,  $DK\{S\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}''_{\text{dec}})$ . The only difference between  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  is that  $\mathcal{G}''_{\text{dec}}$  outputs  $\perp$  if the input ciphertext  $c$  is equal to  $c_0^*$  or  $c_1^*$ . However, due to line 2 of Figure 4.6,  $\mathcal{G}'_{\text{dec}}$  already outputs  $\perp$  on  $c_0^*$  and  $c_1^*$ . Thus,  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  are functionally equivalent and the indistinguishability of  $H_2$  and  $H_3$  follows from the security of  $i\mathcal{O}$ .

*Indistinguishability of  $H_3$  and  $H_4$ .* This follows immediately from the security of the punctured PRF family  $\mathcal{F}_1$ . (Note that each ciphertext  $c_1, \dots, c_t$  can be generated using the punctured PRF key, because they are not encryptions of  $m_0^*$  or  $m_1^*$ .)

*Indistinguishability of  $H_4$  and  $H_5$ .* Note that with overwhelming probability, the random strings  $\alpha_b^*$  are not in the range of  $F_1$ . Therefore, except with negligible probability, there does not exist a message  $m$  such that  $F_1(m) = \alpha_b^*$  for  $b \in \{0, 1\}$ . Since the punctured PRF scheme preserves functionality under puncturing,  $\mathcal{G}'_{\text{enc}}$  (using  $K_2$ ) and  $\mathcal{G}''_{\text{enc}}$  (using  $K_2\{\alpha_0^*, \alpha_1^*\}$ ) behave identically on all input messages, except with negligible probability. The indistinguishability of  $H_4$  and  $H_5$  follows from the security of  $i\mathcal{O}$ .

*Indistinguishability of  $H_5$  and  $H_6$ .* Let  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  denote the circuits such that in  $H_5$ ,  $DK\{S\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}'_{\text{dec}})$ , and in  $H_6$ ,  $DK\{S\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}''_{\text{dec}})$ . The only difference between  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  is that  $\mathcal{G}''_{\text{dec}}$  explicitly outputs  $\perp$  whenever the input ciphertext  $c$  has  $\alpha_0^*$  or  $\alpha_1^*$  as a prefix, while  $\mathcal{G}'_{\text{dec}}$  does so only when  $c$  is exactly  $\alpha_0^*\|\beta_0^*$  or  $\alpha_1^*\|\beta_1^*$ . However, with overwhelming probability, each of the random strings  $\alpha_b^*$  is not in the image of  $F_1$ . Thus (because of the check that  $\alpha = F_1(m)$ )  $\mathcal{G}'_{\text{dec}}$  also (except with negligible probability) outputs  $\perp$  on every  $c$  with  $\alpha_0^*$  or  $\alpha_1^*$  as a prefix. As a consequence,  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  are functionally equivalent except with negligible probability, and the indistinguishability of  $H_5$  and  $H_6$  follows from the security of  $i\mathcal{O}$ .

*Indistinguishability of  $H_6$  and  $H_7$ .* Let  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  denote the circuits such that in  $H_6$ ,  $DK\{S\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}'_{\text{dec}})$ , and in  $H_7$ ,  $DK\{S\} \leftarrow i\mathcal{O}(1^\lambda, \mathcal{G}''_{\text{dec}})$ . Note that due to the check performed in line 2 (see Algorithm 4.8), line 3 is not executed in both  $\mathcal{G}'_{\text{dec}}$  and  $\mathcal{G}''_{\text{dec}}$  whenever the input ciphertext  $c$  is of the form  $\alpha^*\|\star$ . Then, since the punctured PRF scheme preserves functionality under puncturing,  $\mathcal{G}'_{\text{dec}}$  (using  $K_2$ ) and  $\mathcal{G}''_{\text{dec}}$  (using  $K_2\{\alpha_0^*, \alpha_1^*\}$ ) are functionally equivalent, and the indistinguishability of  $H_6$  and  $H_7$  follows from the security of  $i\mathcal{O}$ .



*Indistinguishability of  $H_7$  and  $H_8$ .* This follows immediately from the security of the punctured PRF family  $\mathcal{F}_2$  (note that with overwhelming probability, each ciphertext  $c_1, \dots, c_t$  can be generated with the punctured  $F_2$ , because the  $\tilde{\alpha}_b^*$  are chosen randomly).

*Finishing the proof.* Observe that in experiment  $H_8$ , every challenge ciphertext  $c_i^b$  consists of independent uniformly random strings  $\alpha_b^* \parallel \beta_b^*$  that information-theoretically hide the bit  $b$ . Further,  $EK\{U\}$  and  $DK\{S\}$  are also independent of bit  $b$ . Therefore, the adversary cannot guess the bit  $b$  with probability better than  $\frac{1}{2}$ .  $\square$

**4.2. Garbling Turing machines.** In this section, we describe a direct approach to randomized encodings for Turing machines. In section 4.3, we extend this approach to arbitrary RAM machines. We use the following natural construction.

To compute a randomized encoding of a Turing machine  $M$  and an input  $x$  which runs in space at most  $S$  and time at most  $T$ , our scheme first samples two authenticated encryption key pairs  $(EK_{\text{mem}}, DK_{\text{mem}})$  and  $(EK_{\text{state}}, DK_{\text{state}})$ . It then generates the randomized encoding in two parts: a garbled machine  $\tilde{M}$  and a garbled initial tape  $\tilde{T}_0$ , which we describe below. We assume that  $M$  has an *oblivious* access pattern—that is, there is some universal function  $\tau$  (independent of  $M$  and  $x$ ) so that  $\tau(i)$  is the  $i$ th address accessed by  $M$  on  $x$ . Furthermore,

- $\tau(i)$  is computable in time  $\text{poly}(\log i)$ ;
- the function  $\beta$ , where  $\beta(i) = \min\{i' : i' < i \wedge \tau(i') = \tau(i)\}$ , is also computable in time  $\text{poly}(\log i)$ .

This assumption is without loss of generality due to a transformation of Pippenger and Fischer [PF79], which shows how to compile any Turing machine into one with such an oblivious access pattern.

**The garbled machine.** Suppose that  $M$  has initial state  $q_0$  and transition function  $\delta$ . We define the garbled machine  $\tilde{M}$  by defining its initial state  $\tilde{q}_0$  and its transition function  $\tilde{\delta}$ , which crucially we give as an *obfuscated* circuit. We define  $\tilde{q}_0 = \text{Enc}(EK_{\text{state}}, (0, q_0))$ , and define  $\tilde{\delta} = i\mathcal{O}(1^\lambda, \hat{\delta})$ , where  $\hat{\delta}$  is described in Figure 4.9, but padded to a larger size  $p$ .

The padding parameter  $p$  is  $\text{poly}(\lambda) \cdot T^\epsilon$  for arbitrary  $\epsilon > 0$ , or  $\text{poly}(\lambda, \log T)$  if we assume subexponentially secure **iO** and one-way functions.

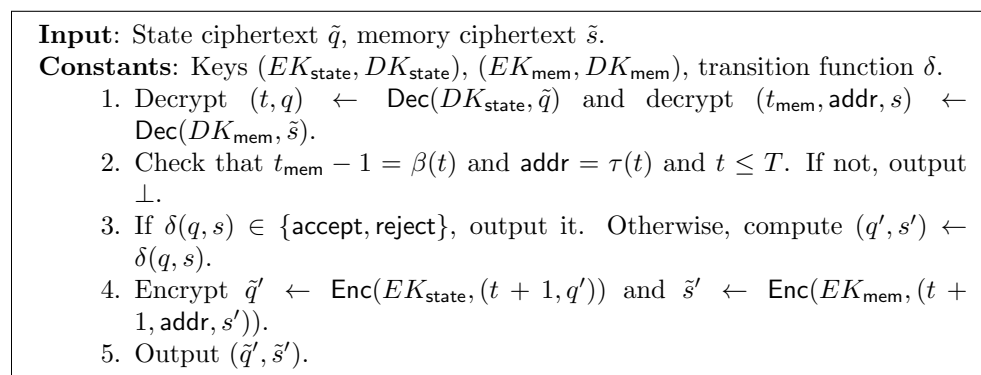


FIG. 4.9. Transition function  $\hat{\delta}$  to be obfuscated.

**The garbled initial tape.** We generate a garbled tape  $\tilde{\mathcal{T}}_0$  consisting of  $S$  ciphertexts. The  $i$ th ciphertext  $\tilde{\mathcal{T}}_0[i]$  is defined as

$$\tilde{\mathcal{T}}_0[i] = \begin{cases} \text{Enc}(EK_{\text{mem}}, (0, i, x_i)) & \text{if } i < |x|, \\ \text{Enc}(EK_{\text{mem}}, (0, i, \perp)) & \text{otherwise.} \end{cases}$$

**Evaluation.** On input  $\tilde{M} = (\tilde{q}_0, \tilde{\delta})$  and  $\tilde{\mathcal{T}}_0$ ,  $\text{Eval}(\tilde{M}, \tilde{\mathcal{T}}_0)$  is computed by running  $\tilde{M}$  as a Turing machine with initial tape contents  $\tilde{\mathcal{T}}_0$ . More precisely, one runs the procedure described in Figure 4.10.

**Input:** Garbled Turing machine  $\tilde{M} = (\tilde{q}_0, \tilde{\delta})$ , garbled initial tape  $\tilde{\mathcal{T}}_0 = (\tilde{x}_1, \dots, \tilde{x}_S)$ .

1. Initialize a mutable tape  $\tilde{\mathcal{T}} := \tilde{\mathcal{T}}_0$ .
2. Initialize a variable  $\tilde{q} := \tilde{q}_0$ .
3. For  $t = 1, 2, \dots, T$ :
  - (a) Compute  $\text{out} \leftarrow \tilde{\delta}(\tilde{q}, \tilde{\mathcal{T}}[\tau(t)])$ .
  - (b) If  $\text{out} \in \{\text{accept}, \text{reject}\}$ , output  $\text{out}$ .
  - (c) Otherwise, parse  $\text{out}$  as  $(\tilde{q}', \tilde{s}')$ .  
Set  $\tilde{q} := \tilde{q}'$  and set  $\tilde{\mathcal{T}}[\tau(t)] := \tilde{s}'$ .
4. Output  $\perp$ .

FIG. 4.10. Procedure for evaluating a garbled Turing machine.

*Remark 4.7.* In this section and the next, we achieve a stronger notion of garbling than the one presented in section 2.2. Specifically, we allow the garbling key **key** to be generated ahead of time, so that inputs and machines may syntactically be garbled separately.

**4.2.1. Security with virtual black box (VBB) obfuscation: A warmup and sanity check.** In this section we argue that the above garbling scheme is secure when instantiated with a VBB obfuscator and an authenticated encryption scheme satisfying existential unforgeability and semantic security. Precisely, we give a sketch of the following proposition.

**PROPOSITION 4.8.** *For any two machine-input pairs  $(M^{(0)}, x^{(0)})$  and  $(M^{(1)}, x^{(1)})$  satisfying  $M^{(0)}(x^{(0)}) = M^{(1)}(x^{(1)})$  and  $\text{TIME}(M^{(0)}, x^{(0)}) = \text{TIME}(M^{(1)}, x^{(1)})$ , it holds that for all PPT oracle algorithms  $\mathcal{A}$ ,*

$$\Pr \left[ \mathcal{A}^{\tilde{\delta}}(\tilde{\mathcal{T}}_0, \tilde{q}_0) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ M \leftarrow M^{(b)}, x \leftarrow x^{(b)} \\ \mathbf{key} \leftarrow \text{Gen}(1^\lambda) \\ (\tilde{\delta}, \tilde{q}_0) \leftarrow \text{Gb}(\mathbf{key}, M) \\ \tilde{\mathcal{T}}_0 \leftarrow \text{Encode}(\mathbf{key}, x) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

*Proof sketch.* Let  $(\tilde{q}_0, \tilde{s}_0), \dots, (\tilde{q}_{t^*}, \tilde{s}_{t^*})$  denote the queries made by  $\text{Eval}(\tilde{M}^{(b)}, \tilde{x})$  to  $\tilde{\delta}$ . First, we claim that any  $\mathcal{A}$  can be simulated by a  $\mathcal{B}$  which also queries only  $(\tilde{q}_0, \tilde{s}_0), \dots, (\tilde{q}_{t^*}, \tilde{s}_{t^*})$  but has nearly the same advantage. That is, in the above probability space,

$$\left| \Pr \left[ \mathcal{A}^{\tilde{\delta}}(\tilde{\mathcal{T}}_0, \tilde{q}_0) = b \right] - \Pr \left[ \mathcal{B}^{\tilde{\delta}}(\tilde{\mathcal{T}}_0, \tilde{q}_0) = b \right] \right| \leq \text{negl}(\lambda).$$

This follows because after having queried  $(\tilde{q}_0, \tilde{s}_0), \dots, (\tilde{q}_{i-1}, \tilde{s}_{i-1})$ , the set of ciphertexts seen by the adversary only contains one new pair which, when given to  $\tilde{\delta}$  as input, causes a non- $\perp$  output (by the checks on line 2 of  $\hat{\delta}$ ). That pair is  $(\tilde{q}_i, \tilde{s}_i)$ . The existential unforgeability of the authenticated encryption means that with high probability,  $\mathcal{A}$  cannot make any other queries for which  $\tilde{\delta}$  outputs a non- $\perp$  value.

Let  $q_0, \dots, q_{t^*-1}$  be the internal states of  $M$  in an honest execution on  $x$ , and let  $\mathcal{T}_0, \dots, \mathcal{T}_{t^*-1}$  be the tape configurations after every step of  $M$ . Then the view of  $\mathcal{B}$  is simply  $y (= M(x))$  along with an oracle which is simulatable given the following ciphertexts:

- $\text{Enc}(\text{sk}_{\text{mem}}, (0, i, x_i))$  for each  $i \in \{0, \dots, |x| - 1\}$ .
- $\text{Enc}(\text{sk}_{\text{mem}}, (0, i, \perp))$  for each  $i \in \{|x|, \dots, S\}$ .
- $\text{Enc}(\text{sk}_{\text{state}}, (i, q_i))$  for each  $i \in \{0, \dots, t^* - 1\}$ .
- $\text{Enc}(\text{sk}_{\text{mem}}, (i + 1, \tau(i), (\mathcal{T}_{i+1})_{\tau(i)}))$  for each  $i \in \{0, \dots, t^* - 1\}$ .

By the semantic security of the authenticated encryption scheme,  $\mathcal{B}$  cannot distinguish the case when these ciphertexts are generated from  $M = M^{(0)}$  and  $x = x^{(0)}$  from the case when  $M = M^{(1)}$  and  $x = x^{(1)}$ .  $\square$

**4.2.2. Security with iO: Technical overview.** We now undertake to prove the security of this construction assuming only indistinguishability obfuscation, following the same ideas as the proof of VBB security.

Recall our reasoning for a VBB obfuscator. We first argued (by existential unforgeability) that a computationally bounded adversary could query  $\hat{\delta}$  only using the set of ciphertexts it has already accumulated from the outputs of  $\hat{\delta}$ . In other words, there is a certain “reachable” set of inputs to  $\delta$ , and functionality differences outside that set are indistinguishable.

We then argued that the information available to an adversary is limited to the computation’s output, plus polynomially many ciphertexts of statically chosen messages, which computationally reveal no information.

There are two problems with applying this reasoning for an iO obfuscator. First, in contrast to oracle access, the obfuscated code of  $\hat{\delta}$  might leak information about the authenticated encryption secret key. An additional difficulty is that two obfuscated programs may be distinguishable even if it is computationally difficult to find an input on which they differ. We will see, nevertheless, that ACE is a particularly “iO-friendly” authenticated encryption scheme. We begin with some (much) simpler examples.

**iO-friendly authentication.** Let us first see how ACE and iO can hide functionality differences between two transition functions  $\delta_0$  and  $\delta_1$  outside of a “reachable” set of inputs. Transition functions are somewhat cumbersome because they take multiple inputs, but we can illustrate the symbiosis between ACE and IO with a couple of toy examples. The first toy example demonstrates how ACE is used together with IO to garble a circuit and an input, for the simple, non-iterated case where the circuit is applied only a single time. The second toy example demonstrates how ACE and IO are used for the extended case where the circuit is applied iteratively a fixed number of times, but the computation performed in all steps but the last one is trivial. Once these examples are in place, we move on to the actual construction and proof.

### Toy Example 1.

**INFORMAL CLAIM 4.9.** *For any circuit  $\delta : \{0, 1\}^n \rightarrow \{0, 1\}$  and input  $x^* \in \{0, 1\}^n$ ,  $(\tilde{x}^*, \tilde{\delta})$  reveals no more than  $(x^*, \delta(x^*))$ , where  $\tilde{x}^*$  and  $\tilde{\delta}$  are sampled according*

to the following procedure:

$$\begin{aligned} SK &\leftarrow ACE.Setup(1^\lambda, 1^n, 1^{O(n)}), \\ EK &\leftarrow ACE.GenEK(SK, \emptyset), \\ DK &\leftarrow ACE.GenDK(SK, \emptyset), \\ \tilde{x}^* &:= ACE.Enc(EK, x^*), \\ \hat{\delta} &:= \text{the circuit described in Figure 4.11,} \\ \tilde{\delta} &\leftarrow i\mathcal{O}(1^\lambda, \hat{\delta}). \end{aligned}$$

**Input:** Ciphertext  $\tilde{x}$ .

**Constants:** ACE decryption key  $DK$

1. Decrypt  $x := \text{Dec}(DK, \tilde{x})$  or output  $\perp$  if the output of Dec is  $\perp$ .
2. Output  $\delta(x)$ .

FIG. 4.11. The circuit  $\hat{\delta}$  to be obfuscated.

*Proof sketch.* Let  $\mathcal{A}$  be any PPT algorithm. We give a sequence of changes to the distribution of  $\tilde{x}, \tilde{\delta}$  which (assuming the security of  $ACE$  and  $i\mathcal{O}$ ) only negligibly affect the probability that  $\mathcal{A}$  outputs 1.

1. Instead of sampling

$$DK \leftarrow ACE.GenDK(SK, \emptyset),$$

sample

$$DK \leftarrow ACE.GenDK(SK, \{0, 1\}^n \setminus \{x^*\}).$$

Here we have written  $\{0, 1\}^n \setminus \{x^*\}$  as shorthand for the circuit which on input  $x \in \{0, 1\}^n$  outputs 1 if and only if  $x \neq x^*$ . The change to  $DK$  (and hence to  $(\tilde{x}, \tilde{\delta})$ ) is computationally indistinguishable by the security of constrained decryption.

2. Modify  $\hat{\delta}$  so that instead of *computing*  $\delta(x)$  when  $x = \perp$ , it outputs the hard-coded value  $\delta(x^*)$ . Pad  $\hat{\delta}$  so that its size is unchanged. This is a functionality-preserving change by the safety of constrained decryption, and so the change to  $(\tilde{x}, \tilde{\delta})$  is computationally indistinguishable by the security of  $i\mathcal{O}$ .

After applying the last change, the distribution of  $(\tilde{x}, \tilde{\delta})$  depends only on  $x$  and  $\delta(x)$ , and thus is a (computationally indistinguishable, efficiently sampleable) simulation of the “real-world” distribution of  $(\tilde{x}, \tilde{\delta})$ .  $\square$

Our next example illustrates how to handle circuits which need to decrypt (or in this case simply verify) their own encrypted (authenticated) outputs.

### Toy Example 2.

INFORMAL CLAIM 4.10. *For any circuit  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  and input  $x^* \in \{0, 1\}^n$ , let  $\hat{\delta}$  be the circuit, described formally in Figure 4.12, which “counts to  $T$ ” before computing  $f$ .  $(\tilde{x}^*, \tilde{\delta})$  reveals no more than  $(x^*, \delta(x^*))$ , where  $\tilde{x}^*$  and  $\tilde{\delta}$  are*

sampled according to the following procedure:

$$\begin{aligned} SK &\leftarrow ACE.Setup(1^\lambda, 1^n, 1^{O(n+\log(T))}), \\ EK &\leftarrow ACE.GenEK(SK, \emptyset), \\ DK &\leftarrow ACE.GenDK(SK, \emptyset), \\ \tilde{x}^* &:= ACE.Enc(EK, (0, x^*)), \\ \hat{\delta} &:= \text{the circuit described in Figure 4.12,} \\ \tilde{\delta} &\leftarrow i\mathcal{O}(1^\lambda, \hat{\delta}_b). \end{aligned}$$

**Input:** Ciphertext  $\tilde{x}$ .

**Constants:** ACE encryption, decryption key  $(EK, DK)$ .

1. Decrypt  $(i, x) := \text{Dec}(DK, \tilde{x})$ , or output  $\perp$  if the output of Dec is  $\perp$ .
2. If  $i = T$ , output  $f(x)$ . Otherwise, output  $\text{Enc}(EK, (i + 1, x))$ .

FIG. 4.12. The circuit  $\hat{\delta}$  to be obfuscated.

*Proof sketch.* The set of ciphertexts that should be “reachable” by an adversary given  $\tilde{x}^*$  and  $\tilde{\delta}$  is just the encryptions of  $(i, x^*)$  for  $0 \leq i < T$ . We would like, therefore, to imitate the steps of the proof in our first toy example:

1. Replace  $DK$  by the punctured key  $DK\{(i, x) : i \geq T \vee x \neq x^*\}$ .
2. Modify  $\hat{\delta}$  so that on an encryption of  $(T, x^*)$  it outputs the *hard-coded* value  $f(x^*)$ .

However, there is a difficulty with step 1.  $\hat{\delta}$  contains both a decryption key *and* an encryption key—thus we cannot indistinguishably puncture the decryption key without first puncturing the encryption key, but neither can we puncture the encryption key without first puncturing the decryption key.

We solve this problem by stratifying the message space by timestamp, and then alternately puncturing the encryption and decryption keys a little at a time. That is, instead of immediately replacing  $DK$  by  $DK\{(i, x) : i \geq T \vee x \neq x^*\}$ , we use a sequence of  $O(T)$  different indistinguishable changes to the distribution of  $(\tilde{x}^*, \tilde{\delta})$ , as follows:

- (1a) Change the sampling of  $EK$  to  $EK \leftarrow ACE.GenEK(SK, \{(i, x) : i \geq T \vee (i = 0 \wedge x \neq x^*)\})$ . This change does not affect the functionality of  $\hat{\delta}$  and hence is computationally indistinguishable.
- (1b) Change the sampling of  $DK$  to  $DK \leftarrow ACE.GenDK(SK, \{(i, x) : i \geq T \vee (i = 0 \wedge x \neq x^*)\})$ . By the security of constrained decryption, this is a computationally indistinguishable change to  $(\tilde{x}^*, DK, EK)$  and hence to  $(\tilde{x}^*, \tilde{\delta})$ .
- (1c) Repeat steps (a) and (b) alternately, first with “ $i = 0$ ” replaced by “ $i \leq 1$ ”, then by “ $i \leq 2$ ”, and so on, up until “ $i < T$ ”. Crucially (for ACE’s succinctness requirement on the punctured set), the punctured sets are decidable by a circuit of size  $O(n + \log T)$ .  $\square$

**iO-friendly encryption.** We now see how the ciphertext-indistinguishability of ACE allows the garbling of much more general “iterated functions.”

### Not-So-Toy Example 3.

**INFORMAL CLAIM 4.11.** *For any circuit  $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$  and input  $x^* \in \{0, 1\}^n$ ,  $(\tilde{x}^*, \tilde{\delta})$  reveals no more than  $f^T(x)$ , where  $\tilde{x}^*$  and  $\tilde{\delta}$  are sampled according to*

the following procedure:

$$\begin{aligned} SK &\leftarrow ACE.Setup(1^\lambda, 1^n, 1^{O(n)}), \\ EK &\leftarrow ACE.GenEK(SK, \emptyset), \\ DK &\leftarrow ACE.GenDK(SK, \emptyset), \\ \tilde{x}^* &:= ACE.Enc(EK, (0, x^*)), \\ \hat{\delta} &:= \text{the circuit described in Figure 4.13,} \\ \tilde{\delta} &\leftarrow i\mathcal{O}(1^\lambda, \hat{\delta}_b). \end{aligned}$$

**Input:** Ciphertext  $\tilde{x}$ .

**Constants:** ACE encryption, decryption key  $EK, DK$ , total number of iterations  $T$ .

1. Decrypt  $(t, x) := \text{Dec}(DK, \tilde{x})$ , or output  $\perp$  if the output of Dec is  $\perp$ .
2. If  $t = T$ , output  $x$ . Otherwise output  $\text{Enc}(EK, (t + 1, f(x)))$ .

FIG. 4.13. The circuit  $\hat{\delta}$  to be obfuscated.

*Proof sketch.* Suppose that  $\hat{\delta}$  were padded to be of size  $\Omega(T)$ . Then we could show indistinguishability from a simulated distribution by building on the techniques of Toy Example 2:

1. Change the sampling of  $DK$  to  $DK \leftarrow ACE.GenDK(SK, \{(i, x) : i \geq T \vee x \neq f^i(x^*)\})$ . This change is indistinguishable by the technique of Toy Example 2.
2. Change  $\hat{\delta}$  to a circuit that has  $f^T(x^*)$  and  $c_0, \dots, c_T$  hard-coded (but not  $EK$  or  $DK$ ), where  $c_i = \text{Enc}(EK, (i, f^i(x^*)))$ .  $\hat{\delta}$  now acts as follows: If the input is  $c_i$  for  $0 \leq i < T$ , output  $c_{i+1}$ . If the input is  $c_T$ , output  $f^T(x^*)$ . Otherwise, output  $\perp$ .

This change does not affect the functionality of  $\hat{\delta}$ , so indistinguishability follows directly from the security of  $i\mathcal{O}$ .

3. Change each ciphertext  $c_i$  to be  $\text{Enc}(EK, (i, 0))$ , and also change  $\tilde{x}^*$  to be  $\text{Enc}(EK, (i, 0))$ . The indistinguishability of this change follows from ACE's ciphertext indistinguishability.

Unfortunately, padding  $\hat{\delta}$  to size  $\Omega(T)$  would make this construction completely nonsuccinct, and therefore no better than a Yao garbled circuit.

We show a succinct version of this argument which step by step changes  $\hat{\delta}$  to a circuit which only does “dummy steps,” yielding  $T$  intermediate hybrid distributions  $H_1, \dots, H_T$ . In  $H_i$ ,  $\tilde{x}^*$  is defined as  $\text{Enc}(EK, (0, f^i(x^*)))$  and  $\hat{\delta}$  is defined as in Figure 4.14.

**Input:** Ciphertext  $\tilde{x}$ .

**Constants:** ACE encryption, decryption key  $EK, DK$ , total number of iterations  $T$ , hybrid number  $i$ .

1. Decrypt  $(t, x) := \text{Dec}(DK, \tilde{x})$ , or output  $\perp$  if the output of Dec is  $\perp$ .
2. If  $t < i$ , output  $\text{Enc}(EK, (t + 1, x))$ .
3. Otherwise, if  $i \leq t < T$ , output  $\text{Enc}(EK, (t + 1, f(x)))$ .
4. Otherwise (if  $i = T$ ), output  $x$ .

FIG. 4.14. The circuit  $\hat{\delta}$  to be obfuscated in hybrid  $H_i$ .

To show that  $H_i \approx H_{i+1}$ , we give a sequence of indistinguishable changes, proceeding through  $i$  intermediate hybrid distributions  $H_{i,i}, \dots, H_{i,0}$ .  $H_{i,j}$  is obtained from  $H_i$  by modifying  $\hat{\delta}$  to output  $\text{Enc}(EK, (j+1, f^{i+1}(x^*)))$  when  $t = j$ , as in Figure 4.15. Now we can indistinguishably change  $H_{i,j}$  to  $H_{i,j-1}$  as follows:

**Input:** Ciphertext  $\tilde{x}$ .

**Constants:** ACE encryption, decryption key  $EK, DK$ , total number of iterations  $T$ , hybrid number  $i, j$ , the value  $f^{i+1}(x^*)$ .

1. Decrypt  $(t, x) := \text{Dec}(DK, \tilde{x})$ , or output  $\perp$  if the output of Dec is  $\perp$ .
2. If  $t = j$ , output  $\text{Enc}(EK, (j+1, f^{i+1}(x^*)))$ .
3. Otherwise, if  $t < i$ , output  $\text{Enc}(EK, (t+1, x))$ .
4. Otherwise, if  $i \leq t < T$ , output  $\text{Enc}(EK, (t+1, f(x)))$ .
5. Otherwise (if  $i = T$ ), output  $x$ .

FIG. 4.15. The circuit  $\hat{\delta}$  to be obfuscated in hybrid  $H_{i,j}$ .

1. Change the sampling of  $DK$  to  $DK \leftarrow \text{ACE.GenDK}(SK, G_{i,j}^c)$ , where  $G_{i,j}$  is the set of “good” messages, namely,  $(t, x)$  for which the following hold:
  - If  $t \leq j$ , then  $x = f^i(x^*)$ .
  - If  $j < t \leq i+1$ , then  $x = f^{i+1}(x^*)$ . $G_{i,j}^c$  is just the set complement of  $G_{i,j}$ . The indistinguishability of this change follows from the techniques of Toy Example 2.
2. Change  $\tilde{\delta}$  so that if  $t = j-1$ , it outputs a hard-coded ciphertext  $c = \text{Enc}(EK, (j, f^i(x^*)))$ , and replace the “if  $t = j$ ” check with “if  $\tilde{x} = c$ ” (to be executed *before* decryption). Also, expand the set on which  $EK$  and  $DK$  are punctured to include all  $(t, x)$  for which  $t = j$ . These changes do not affect the functionality of  $\tilde{\delta}$ , so indistinguishability follows from  $i\mathcal{O}$ .
3. Change  $c$  to  $\text{Enc}(EK, (j, f^{i+1}(x^*)))$ . This is indistinguishable by ACE’s ciphertext indistinguishability.
4. Change  $\tilde{\delta}$  by removing the “if  $\tilde{x} = c$ ” special case. Change the special-case behavior when  $t = j-1$  to computing and outputting  $\text{Enc}(EK, (j, f(x)))$  rather than using the hard-coded ciphertext  $c$ . Remove the hard-coded ciphertext  $c$  altogether. Unpuncture  $EK$  and reduce the puncturing of  $DK$  to the set  $G_{i,j-1}^c$ . These changes do not affect the functionality of  $\hat{\delta}$ , so indistinguishability follows from  $i\mathcal{O}$ .
5. Apply the changes of step 1 in reverse to indistinguishably unpuncture  $DK$ .  $\square$

We note that when  $f$  is the function which maps one Turing machine configuration to the next, Not-So-Toy Example 3 already is a meaningfully succinct garbling for Turing machines. In particular, the size of the garbling is  $\text{poly}(S, \lambda)$ , and the evaluation time is  $T \cdot \text{poly}(S, \lambda)$ . Our main construction improves on this by reducing the evaluation time to  $T \cdot \text{poly}(\lambda)$  and the garbling size to  $S \cdot \text{poly}(\lambda)$ . The proof of security for the construction builds on our proof of security for Not-So-Toy Example 3, but with a finer-grained “per-cell” puncturing strategy.

#### 4.2.3. Security with $i\mathcal{O}$ : Proof.

**THEOREM 4.12.** *There is a PPT algorithm Sim such that for every Turing ma-*

chine  $M$  and input  $x$ , running in time  $T$  and space  $S$ , the distribution

$$(4.1) \quad \tilde{\delta}, \tilde{q}_0, \tilde{\mathcal{T}}_0 \left| \begin{array}{l} \mathbf{key} \leftarrow \text{Gen}(1^\lambda, T, S), \\ \tilde{\delta}, \tilde{q}_0 \leftarrow \text{Gb}(\mathbf{key}, M), \\ \tilde{\mathcal{T}}_0 \leftarrow \text{Encode}(\mathbf{key}, x) \end{array} \right.$$

is computationally indistinguishable from

$$\text{Sim}(1^\lambda, M(x), 1^T, 1^S).$$

*Proof.* We start with a generalization of the techniques given in the technical overview (specifically the technique used to show that  $H_{i,j}$  and  $H_{i,j-1}$  were indistinguishable in Not-So-Toy Example 3). Let  $\delta : Q \times \Sigma \rightarrow Q \times \Sigma \cup \{\text{accept}, \text{reject}, \perp\}$  be a circuit.

DEFINITION 4.13. Say that  $\delta$  is *timed* if for every  $q \in Q$  and  $s \in \Sigma$  there are associated timestamps  $q.t$  and  $s.t$  such that (i) if  $(q', s') = \delta(q, s)$ , then  $q'.t = s'.t = q.t + 1$ , and (ii) if  $s.t > q.t$ , then  $\delta(q, s) = \perp$ .

DEFINITION 4.14. Say that a timed circuit  $\delta$  is  $T$ -*bounded* if for all  $q$  and  $s$  with  $q.t > T$  it holds that  $\delta(q, s) = \perp$ .

DEFINITION 4.15. A pair of sets  $G_{\text{state}} \subseteq Q$  and  $G_{\text{mem}} \subseteq \Sigma$  is said to be  $\delta$ -invariant if

$$\delta(G_{\text{state}} \times G_{\text{mem}}) \cap (Q \times \Sigma) \subseteq G_{\text{state}} \times G_{\text{mem}}.$$

LEMMA 4.16. Let  $\delta_0, \delta_1 : Q \times \Sigma \rightarrow Q \times \Sigma \cup \{\text{accept}, \text{reject}, \perp\}$  be  $T$ -bounded timed circuits, and let  $(G_{\text{state}}, G_{\text{mem}})$  be  $\delta_0$ -invariant, with  $G_{\text{state}}$  and  $G_{\text{mem}}$  decidable by circuits of size, respectively,  $\text{poly}(\log |Q|)$  and  $\text{poly}(\log |\Sigma|)$ .

Suppose there exist  $q_0^* \in G_{\text{state}}$ ,  $q_1^* \in Q \setminus G_{\text{state}}$  and  $s_0^* \in G_{\text{mem}}$ ,  $s_1^* \in \Sigma \setminus G_{\text{mem}}$  such that for all  $q \in G_{\text{state}}$  and  $s \in G_{\text{mem}}$ ,  $(q', s') = \delta_0(q, s)$  if and only if

$$(\tau_{\text{state}}(q'), \tau_{\text{mem}}(s')) = \delta_1(\tau_{\text{state}}(q), \tau_{\text{mem}}(s)),$$


where  $\tau_{\text{state}} : Q \rightarrow Q$  denotes the transposition interchanging  $q_0^*$  and  $q_1^*$ , and  $\tau_{\text{mem}} : \Sigma \rightarrow \Sigma$  denotes the transposition interchanging  $s_0^*$  and  $s_1^*$ .

Then for any  $q_0 \in G_{\text{state}}$  and any  $s_0^{(0)}, \dots, s_0^{(S-1)} \in G_{\text{mem}}$ , it holds that  $(\tilde{c}_0, \tilde{\delta}_0)$  is computationally indistinguishable from  $(\tilde{c}_1, \tilde{\delta}_1)$  in the probability space defined by sampling

$$\begin{aligned} SK_{\text{state}} &\leftarrow \text{ACE.Setup}(1^\lambda, 1^{\log |Q|}, 1^{\text{poly}(\log |Q|)}), \\ EK_{\text{state}} &\leftarrow \text{ACE.GenEK}(SK_{\text{state}}, \emptyset), \\ DK_{\text{state}} &\leftarrow \text{ACE.GenDK}(SK_{\text{state}}, \emptyset), \\ SK_{\text{mem}} &\leftarrow \text{ACE.Setup}(1^\lambda, 1^{\log |\Sigma|}, 1^{\text{poly}(\log |\Sigma|)}), \\ EK_{\text{mem}} &\leftarrow \text{ACE.GenEK}(SK_{\text{mem}}, \emptyset), \\ DK_{\text{mem}} &\leftarrow \text{ACE.GenDK}(SK_{\text{mem}}, \emptyset), \\ \tilde{c}_0 &:= \left( \text{Enc}(EK_{\text{state}}, q_0), \left( \text{Enc}(EK_{\text{mem}}, s_0^{(i)}) \right)_{i=0}^{S-1} \right), \\ \tilde{c}_1 &:= \left( \text{Enc}(EK_{\text{state}}, \tau_{\text{state}}(q_0)), \left( \text{Enc}(EK_{\text{mem}}, \tau_{\text{mem}}(s_0^{(i)})) \right)_{i=0}^{S-1} \right), \\ \text{for each } b \in \{0, 1\}: \\ \tilde{\delta}_b &\leftarrow i\mathcal{O}(1^\lambda, \hat{\delta}_b), \end{aligned}$$

where  $\hat{\delta}_b$  is the circuit described in Figure 4.16, padded to size  $\text{poly}(\lambda, \log |Q|, \log |\Sigma|)$ .



FIG. 4.16. The circuit  $\hat{\delta}_b$  in Lemma 4.16.

*Proof.* We show this by a sequence of  $O(T)$  computationally indistinguishable changes which change  $(\tilde{c}_0, \tilde{\delta}_0)$  to  $(\tilde{c}_1, \tilde{\delta}_1)$ . First, we iteratively restrict  $EK_{\text{state}}, DK_{\text{state}}$  and  $EK_{\text{mem}}, DK_{\text{mem}}$  as follows:

1. Restrict  $EK_{\text{state}}$  (i.e., change the set at which it is punctured) so that

$$\text{Enc}(EK_{\text{state}}, q) \neq \perp \text{ if and only if } q.t > 0 \text{ or } q \in G_{\text{state}}.$$

Similarly, restrict  $EK_{\text{mem}}$  so that

$$\text{Enc}(EK_{\text{mem}}, s) \neq \perp \text{ if and only if } s.t > 0 \text{ or } s \in G_{\text{mem}}.$$

Because  $\hat{\delta}_0$  never encrypts messages with timestamp 0, these changes to  $EK_{\text{state}}$  and  $EK_{\text{mem}}$  are functionality-preserving, and thus by the security of  $i\mathcal{O}$  the overall change to  $\tilde{\delta}$  is computationally indistinguishable.

2. For  $j = 1, \dots, T$ , do the following:

- (a) Modify  $DK_{\text{state}}$  and  $DK_{\text{mem}}$ , such that

- if  $q$  is in the image of  $\text{Dec}(DK_{\text{state}}, \cdot)$ , then  $q.t > j - 1$  or  $q \in G_{\text{state}}$ ;
- if  $s$  is in the image of  $\text{Dec}(DK_{\text{mem}}, \cdot)$ , then  $s.t > j - 1$  or  $s \in G_{\text{mem}}$ .

The indistinguishability of this change follows from ACE's security of constrained decryption.

- (b) Restrict  $EK_{\text{state}}$  and  $EK_{\text{mem}}$  analogously to step 1, namely,

- $\text{Enc}(EK_{\text{state}}, q) \neq \perp$  if and only if  $q.t > j$  or  $q \in G_{\text{state}}$ ;
- $\text{Enc}(EK_{\text{mem}}, s) \neq \perp$  if and only if  $s.t > j$  or  $s \in G_{\text{mem}}$ .

The functional equivalence (and hence, by  $i\mathcal{O}$ , indistinguishability) of this change follows from the  $\delta_0$ -invariance of  $(G_{\text{state}}, G_{\text{mem}})$ .

By the  $T$ -boundedness of  $\delta_0$ ,  $EK_{\text{state}}$  and  $DK_{\text{state}}$  are now punctured at  $Q \setminus G_{\text{state}}$ , while  $EK_{\text{mem}}$  and  $DK_{\text{mem}}$  are now punctured at  $\Sigma \setminus G_{\text{mem}}$ . Let  $\tilde{q}_b^*$  denote  $\text{Enc}(EK_{\text{state}}, q_b^*)$ , and let  $\tilde{s}_b^*$  denote  $\text{Enc}(EK_{\text{mem}}, s_b^*)$ .

1. Modify  $\hat{\delta}_0$  so that instead of encrypting  $q_0^*$  (resp., decrypting  $\tilde{q}_0^*$ ), it uses the hard-coded correspondence  $q_0^* \leftrightarrow \tilde{q}_0^*$ . Similarly instead of encrypting  $s_0^*$  (resp. decrypting  $\tilde{s}_0^*$ ), it uses the hard-coded correspondence  $s_0^* \leftrightarrow \tilde{s}_0^*$ . Let  $EK_{\text{state}}$  and  $DK_{\text{state}}$  be punctured now on  $(Q \setminus G_{\text{state}}) \cup \{q_0^*\}$ , and let  $EK_{\text{mem}}$  and  $DK_{\text{mem}}$  be punctured on  $(\Sigma \setminus G_{\text{mem}}) \cup \{s_0^*\}$ . These changes are clearly functionality-preserving and hence indistinguishable by the security of  $i\mathcal{O}$ .
2. Change  $\tilde{q}_0^*$  to  $\text{Enc}(EK_{\text{state}}, q_1^*)$ , and similarly change  $\tilde{s}_0^*$  to  $\text{Enc}(EK_{\text{mem}}, s_1^*)$ . Similarly modify all ciphertexts in  $\tilde{c}$ . This is indistinguishable by ACE's ciphertext indistinguishability.
3. Replace  $\delta_0$  by  $\delta_1$  and remove the hard-coded correspondences  $q_0^* \leftrightarrow \tilde{q}_0^*$  and  $s_0^* \leftrightarrow \tilde{s}_0^*$ . Reduce the puncturing of  $EK_{\text{state}}$  and  $DK_{\text{state}}$  to  $Q \setminus \tau_{\text{state}}(G_{\text{state}})$ ,

and reduce the puncturing of  $EK_{\text{mem}}$  and  $DK_{\text{mem}}$  to  $\Sigma \setminus \tau_{\text{mem}}(G_{\text{mem}})$ . By assumption on the relation of  $\delta_1$  to  $\delta_0$ , this change is functionality-preserving and hence indistinguishable by the security of  $i\mathcal{O}$ .

Finally, remove all of the key restrictions and hard-codings. This change is indistinguishable by reversing the above hybrids.  $\square$

We now return to the proof of Theorem 4.12. We present a number of indistinguishable hybrid distributions  $H_0, \dots, H_T$ , each defining a different way of sampling  $\delta$ ,  $\tilde{q}_0$ , and  $\tilde{T}_0$ , where  $\tilde{\delta}$  is sampled as  $\tilde{\delta} \leftarrow i\mathcal{O}(1^\lambda, \hat{\delta})$  for some  $\hat{\delta}$ . In hybrid  $H_0$ ,  $(\tilde{\delta}, \tilde{q}_0, \tilde{T}_0)$  are sampled as in (4.1), while in hybrid  $H_T$  they are sampled given only  $M(x)$ ,  $T$ , and  $S$ . Thus, if we show that for all  $i \in \{0, \dots, T - 1\}$ ,  $H_i \approx H_{i+1}$ , then hybrid  $H_T$  defines the desired  $\text{Sim}$ .

All of our hybrids are designed around the transition function  $\hat{\delta}$ . We first describe the sorts of transition functions we will define.  $\hat{\delta}$  always takes the following as input:

- $\tilde{q}$ , expected to be an encryption of  $(t, q)$  for some timestamp  $t$  and state  $q$ .
- $\tilde{s}$ , expected to be an encryption of  $(t_{\text{mem}}, \text{addr}, s)$  for some different timestamp  $t_{\text{mem}}$ , an address tag  $\text{addr}$ , and an underlying memory symbol  $s$ .

On an input  $(\tilde{q}, \tilde{s})$ ,  $\hat{\delta}$  decrypts

$$\begin{aligned} (t, q) &\leftarrow \text{Dec}(DK_{\text{state}}, \tilde{q}), \\ (t_{\text{mem}}, \text{addr}, s') &\leftarrow \text{Dec}(DK_{\text{mem}}, \tilde{s}) \end{aligned}$$

and checks that  $t < T$ ,  $t_{\text{mem}} = \beta(t)$ , and  $\text{addr} = \tau(t)$ . If not, then  $\hat{\delta}$  outputs  $\perp$ . If so,  $\hat{\delta}$  does one of three things, depending on the hybrid distribution as well as on the timestamp  $t$ :

1. *Dummy step*: Only update the timestamps. That is, output  $(\tilde{q}', \tilde{s}')$ , where

$$\begin{aligned} \tilde{q}' &= \text{Enc}(EK_{\text{state}}, (t + 1, q)), \\ \tilde{s}' &= \text{Enc}(EK_{\text{mem}}, (t + 1, \text{addr}, s)). \end{aligned}$$

2. *Real step*: Perform one step of  $M$ 's computation. That is, compute  $\text{out} \leftarrow \delta(q, s)$ . If  $\text{out}$  is a "final answer," i.e.,  $\text{out} \in \{\text{accept}, \text{reject}\}$ , then output  $\text{out}$  in the clear. Otherwise, parse  $\text{out}$  as  $(q', s')$  and output  $(\tilde{q}', \tilde{s}')$ , where

$$\begin{aligned} \tilde{q}' &= \text{Enc}(EK_{\text{state}}, (t + 1, q')), \\ \tilde{s}' &= \text{Enc}(EK_{\text{mem}}, (t + 1, \text{addr}, s')). \end{aligned}$$

3. *Hard-coded output*: Output some hard-coded constant  $y$  in the clear.

With this terminology, we can crisply describe our hybrids. Let  $t^*$  denote the running time of  $M$  on  $x$ . Let  $q_t$  denote the  $t$ th internal state of  $M$  when executed on  $x$  for  $t$  steps. Similarly, let  $\mathcal{T}_t$  denote the tape of  $M$  when executed on  $x$  for  $t$  steps.

HYBRID  $H_i$  (IF  $i < t^*$ ):

- If  $t < i$ ,  $\hat{\delta}$  does a dummy step. If  $i \leq t \leq T$ ,  $\hat{\delta}$  does a real step.
- $\tilde{q}_0$  is defined as  $\text{Enc}(EK_{\text{state}}, (0, q_i))$ .
- $\tilde{T}_0$  is defined as  $(c_1, \dots, c_S)$ , where  $c_j = \text{Enc}(EK_{\text{mem}}, (0, j, \mathcal{T}_i[j]))$ .

HYBRID  $H_i$  (IF  $t^* \leq i \leq T$ ):

- If  $t < t^*$ ,  $\hat{\delta}$  does a dummy step. If  $t = t^*$ ,  $\hat{\delta}$  outputs the hard-coded  $M(x)$ . If  $t^* < t \leq i$ ,  $\hat{\delta}$  does a dummy step. If  $i < t \leq T$ ,  $\hat{\delta}$  does a real step.
- $\tilde{q}_0$  is defined as  $\text{Enc}(EK_{\text{state}}, (0, \perp))$ .

- $\tilde{\mathcal{T}}_0$  is defined as  $(c_1, \dots, c_S)$ , where  $c_j = \text{Enc}(EK_{\text{mem}}, (0, j, \mathcal{T}_{t^*}[j]))$ .

HYBRID  $H_{T+1}$ :

- If  $t < t^*$ ,  $\hat{\delta}$  does a dummy step. If  $t = t^*$ ,  $\hat{\delta}$  outputs the hard-coded  $M(x)$ . If  $t^* < t$ ,  $\hat{\delta}$  does a dummy step.
- $\tilde{q}_0$  is defined as  $\text{Enc}(EK_{\text{state}}, (0, \perp))$ .
- $\tilde{\mathcal{T}}_0$  is defined as  $(c_1, \dots, c_S)$ , where  $c_j = \text{Enc}(EK_{\text{mem}}, (0, j, \perp))$ .

**Indistinguishability of hybrids  $H_i$  and  $H_{i+1}$  (for  $i < t^*$ ).** We will show that  $H_i$  and  $H_{i+1}$  are indistinguishable by following our earlier proof template using ACE. We give a sequence of subhybrids.

HYBRID  $H_{i,1}$ : This is the same as hybrid  $H_i$ , but  $EK_{\text{mem}}$ ,  $DK_{\text{mem}}$  and  $EK_{\text{state}}$ ,  $DK_{\text{state}}$  are, respectively, punctured at  $\Sigma \setminus G_{\text{mem}}^i$  and  $Q \setminus G_{\text{state}}^i$ , where

$$G_{\text{mem}}^i = \{(t, a, s) : \text{If } t \leq i \text{ and } a = \tau(i), \text{ then } s = \mathcal{T}_i[\tau(i)]\}$$

and

$$G_{\text{state}}^i = \{(t, q) : \text{If } t \leq i \text{ then } q = q_i\}.$$

HYBRID  $H_{i,2}$ : Add the following special-case branch to  $\hat{\delta}$ : If  $t = i$ , output  $(\tilde{q}^*, \tilde{s}^*)$ , where these are hard-coded constants

$$\tilde{q}^* = \text{Enc}(EK_{\text{state}}, (i + 1, q_{i+1}))$$

and

$$\tilde{s}^* = \text{Enc}(EK_{\text{mem}}, (i + 1, \tau(i), \mathcal{T}_{i+1}[\tau(i)])).$$

$\hat{\delta}$  is now described by Figure 4.17.

HYBRID  $H_{i,3}$ : In this step, we change the time at which the special-case branch above is activated. That is, for  $j$  ranging from  $i$  to 0, we have a hybrid  $H_{i,j}$  in which the transition function  $\hat{\delta}$  instead has the following special-case branches. Let  $\gamma(j)$  denote  $\min\{j' : j' \geq j \wedge \tau(j') = \tau(i)\}$ :

- If  $t = j$ , set  $\tilde{q}' = \tilde{q}^*$ , where  $\tilde{q}^*$  is the hard-coded value  $\text{Enc}(EK_{\text{state}}, (j + 1, q_{i+1}))$ . Otherwise, define  $\tilde{q}'$  as in a dummy step.
- If  $t = \gamma(j)$ , then set  $\tilde{s}' = \tilde{s}^*$ , where  $\tilde{s}^*$  is the hard-coded value

$$\text{Enc}(EK_{\text{mem}}, (j + 1, \tau(i), (\mathcal{T}_{i+1})[\tau(i)]).$$

Otherwise, define  $\tilde{s}'$  as in a dummy step.

CLAIM 4.17. *The distribution of  $(\hat{\delta}, \tilde{q}_0, \tilde{\mathcal{T}}_0)$  induced by hybrid  $H_{i,2}$  is indistinguishable from that induced by  $H_{i,3}$ .*

*Proof.* There are really two changes here: changing the time at which branch (a) is activated (together with the ciphertext  $c_{\text{state}}$ ), and changing the time at which branch (b) is activated (together with the ciphertext  $c_{\text{mem}}$ ). Although the two are treated analogously, we only illustrate with branch (a) for simplicity because its hard-coded time always changes from  $j$  to  $j - 1$ .

We first modify branch (a) to have two parts:

- If  $t = j - 1$ , set  $\tilde{q}' = c_{\text{state}}$ , where  $c_{\text{state}} = \text{Enc}(EK_{\text{state}}, (j, q_i))$ .
- If  $t = j$ , set  $\tilde{q}' = \tilde{q}^*$ , where  $\tilde{q}^* = \text{Enc}(EK_{\text{state}}, (j + 1, q_{i+1}))$ .

This change is indistinguishable by **IO**. Next, we restrict  $EK_{\text{state}}$  and  $DK_{\text{state}}$  to be unable to encrypt or decrypt anything with timestamp  $j$ . This is also indistinguishable by **IO** because

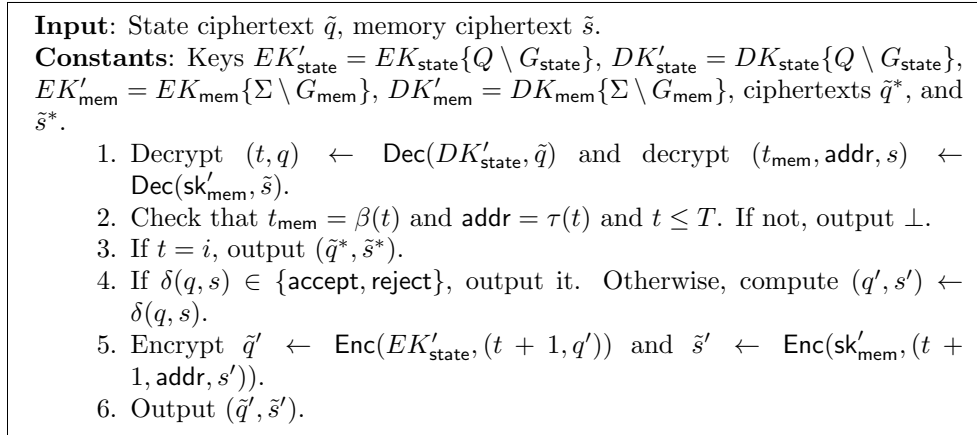


FIG. 4.17. Obfuscated transition function  $\hat{\delta}$  with punctured keys and hard-coded ciphertexts to output.

- $\hat{\delta}$  never encrypts something with timestamp  $j$ ; it outputs  $c_{\text{state}}$  instead;
- since  $DK_{\text{state}}$  was already restricted to  $G_{\text{state}}^i$ ,  $\text{Dec}(DK_{\text{state}}, c)$  would be of the form  $(j, \star)$  if and only if  $c = c_{\text{state}}$ , but  $\hat{\delta}$  doesn't ever decrypt  $c_{\text{state}}$ .

Next, we change the hard-coded value  $c_{\text{state}}$  to be equal to  $\text{Enc}(EK_{\text{state}}, (j, q_{i+1}))$ . This is indistinguishable due to ciphertext indistinguishability of ACE.

We now undo the hard-coding and key restriction we did to facilitate this change:

1.  $\hat{\delta}$  is modified so that when  $t = j$ , it computes  $\tilde{q}'$  as in a dummy step instead of hard-coding  $\tilde{q}^*$ . This is indistinguishable because, by the puncturing of  $DK_{\text{state}}$ , the only reachable case (mapping  $c_{\text{state}}$  to  $\tilde{q}^*$ ) is functionally equivalent to a dummy step.
2. We unrestrict  $EK_{\text{state}}$  and  $DK_{\text{state}}$  to allow encryption and decryption of  $(j, q_{i+1})$  and remove special-case branch (a.2). This is functionally equivalent, and thus indistinguishable by **iO**.
3. We then unrestrict  $DK_{\text{state}}$  to allow decryption of any message with timestamp  $j$ . This is indistinguishable by ACE's security of constrained decryption.
4. We then unrestrict  $EK_{\text{state}}$  to allow encryption of any message with timestamp  $j$ . This is indistinguishable by **iO** because  $\hat{\delta}$  only tries to encrypt  $(j, q)$  when  $q = q_{i+1}$ .

The result of these changes is hybrid  $H_{i,j-1}$ , so we have shown that for  $j > 1$ ,  $H_{i,j} \approx H_{i,j-1}$ . We must also show that  $H_{i,1} \approx H_{i+1}$ . The proof is exactly analogous, and hybrid  $H_{i+1}$  should be thought of as a degenerate hybrid  $H_{i,0}$ . Instead of modifying  $\hat{\delta}$  to have a hard-coded output when  $t = 0$  ( $\hat{\delta}$  is never run with  $t = 0$ —its first timestamp is  $t = 1$ ), one modifies the initial state  $\tilde{q}_0$  and the initial tape  $\tilde{T}_0$ .  $\square$

*Indistinguishability of hybrids  $H_{t^*-1}$  and  $H_{t^*}$ .* This proceeds analogously to the proof that  $H_{i-1} \approx H_i$  for  $i < t^*$ . The only difference is that the hard-coded behavior at time  $t^*$  is to output some hard-coded  $y \in \{\text{accept}, \text{reject}\}$  rather than to output two hard-coded ciphertexts  $(c_{\text{state}}, c_{\text{mem}})$ . There is also no need to go through hybrids  $H_{i,i}$  to  $H_{i,0}$  to change  $\tilde{q}_0$  (resp.,  $\tilde{T}_0$ ) from  $\text{Enc}(EK_{\text{state}}, (0, q_{t^*}))$  to  $\text{Enc}(EK_{\text{state}}, (0, q_{t^*+1}))$ ; since the running time of  $M$  on  $x$  is only  $t^*$ , there is no  $q_{t^*+1}$ .

*Indistinguishability of hybrids  $H_{t^*}$  and  $H_T$ .* We use the fact that in  $H_{t^*}$ 's definition of  $\hat{\delta}$ , the procedure  $\text{Enc}(EK_{\text{state}}, \cdot)$  is never invoked on a message of the form  $(t^* + 1, \star)$ . We then iteratively restrict  $EK_{\text{state}}$  and  $DK_{\text{state}}$  so that for any  $t \in \{t^* + 1, \dots, T\}$ , any message of the form  $(t, \star)$  is not in the image of  $\text{Dec}(DK_{\text{state}}, \cdot)$ . Then modifying  $\hat{\delta}$  to run a dummy step on all such  $t$  preserves functionality and is therefore indistinguishable.

*Indistinguishability of hybrids  $H_T$  and  $H_{T+1}$ .* We need to indistinguishably replace  $\tilde{q}_0$  and  $\tilde{\mathcal{T}}_0$  with null ciphertexts. This will require  $T \cdot S$  more hybrids. Due to its similarity to the lengthy argument described above, we only sketch these hybrids here.

We change the initial ciphertexts one at a time. For example, we might change  $\tilde{q}_0$  first, and then change  $\tilde{\mathcal{T}}_0$ , changing addresses from lowest to highest. But the exact order isn't important. Suppose that we want to change  $\tilde{\mathcal{T}}_0[0]$  to be  $\text{Enc}(EK_{\text{mem}}, (0, 0, \perp))$ . We go through hybrids  $H_{T,1}, \dots, H_{T,t^*}$ . In each of these hybrids,

$$\tilde{\mathcal{T}}_0[0] = \text{Enc}(EK_{\text{mem}}, (0, 0, \perp)),$$

but only in  $H_{T,T}$  will  $\hat{\delta}$  be independent of the old value of  $\tilde{\mathcal{T}}_0$ .

In hybrid  $H_{T,i}$ ,  $\hat{\delta}$  will do dummy steps everywhere, except for at  $t' = \min\{t' : t^* > t' > i \text{ and } \tau(t') = 0\}$ , if that set is nonempty. At this  $t'$ ,  $\hat{\delta}$  sets  $\tilde{s}' = \text{Enc}(EK_{\text{mem}}, (t' + 1, 0, \mathcal{T}_{t^*+1}[0]))$ .  $\square$

**4.3. Garbling RAM machines.** In this section we describe how to modify the construction of section 4.2 to work for RAM machines. For our purposes, a RAM machine consists of two components: an initial internal state  $q_0$ , and a transition function  $\delta$  which takes as input an internal state  $q$  and a memory symbol  $s$ , and outputs

1. an internal state  $q'$ , a memory symbol  $s'$ , and an address  $\text{addr}$ ; or
2. a "final output": either **accept** or **reject**.

An input to a RAM machine is a tape  $\mathcal{T}_0$ . To execute a RAM machine  $M = (q_0, \delta)$  on a tape  $\mathcal{T}_0$ , one does the following.

Define  $\text{addr}_0 = 0$ . For each  $i > 0$ , define  $(q_i, \text{addr}_i, s_i) = \delta(q_{i-1}, \mathcal{T}_{i-1}[\text{addr}_{i-1}])$  and define the  $i$ th memory configuration  $\mathcal{T}_i$  as

$$\mathcal{T}_i[\text{addr}] = \begin{cases} s_i & \text{if } s_i \neq \perp \text{ and } \text{addr} = \text{addr}_i, \\ \mathcal{T}_{i-1}[\text{addr}] & \text{otherwise.} \end{cases}$$

If  $\delta(q_{t-1}, \mathcal{T}_{t-1}[\text{addr}_{t-1}]) = y \in \{\text{accept}, \text{reject}\}$  for some  $t$ , then we say that  $M(\mathcal{T}_0) = y$ . If there is no such  $t$ , we say that  $M(\mathcal{T}_0) = \perp$ .

**Technical overview.** The difficulty with garbling RAM machines is the lack of an exact analog to the Pippenger–Fischer theorem. Although there is a vast literature on oblivious RAM (ORAM) transformations, these transformations (unlike the Pippenger–Fischer theorem) inherently produce randomized RAM machines, with access patterns from a fixed *distribution* on sequences of addresses. With *nonsuccinct* garbling schemes, this distinction is unimportant, because typically it is easier to construct a "weak" garbling scheme in which a garbling of  $(M, x)$  is simulatable given  $M(x)$  and the sequence of addresses accessed by  $M$ . Such a garbling scheme can be composed with any ORAM to yield a scheme with full security: a garbling of  $(M, x)$  is simulatable given only  $M(x)$ .

Weak garbling schemes typically preserve or nearly preserve the access pattern of  $M$  (indeed otherwise they would themselves be an ORAM, which runs counter to the goal of modularity). For *succinct* garbling schemes, such a weak garbling scheme is impossible: a simulator given an arbitrary access pattern cannot hope to compress it into a succinct  $(\tilde{M}, \tilde{x})$ , which reproduces that access pattern. Instead, one could hope for the natural weakening to an indistinguishability-based notion of security: if two RAM machines have *identical* access patterns, then their weak garblings should be indistinguishable. It is not clear how such a weak garbling scheme can be composed with an ORAM, but a subsequent work of Canetti and Holmgren [CH16] shows how to do it (for any ORAM satisfying a certain “localized randomness” property).

In this work, due to limitations of our techniques (which were later removed in [KLW15]), we don’t know a simpler way of obtaining such a weak garbling scheme. Indeed what we show is how to adapt the bounded space Turing machine garbling scheme of section 4.2 (and its security proof) to a garbling scheme for bounded space RAM machines.

Adapting the construction itself is quite straightforward. Starting with a RAM machine  $M = (q_0, \delta)$ , we transform it to obtain a different RAM machine  $\hat{M}$  as follows. First, we transform  $M$  to have an oblivious access pattern with a *predictably timed writes* property; i.e., just prior to reading any memory address, the last-written time for that address is efficiently computable from  $M$ ’s local state. Such a transformation is given in [GHRW14]. Then we construct a new RAM machine  $\hat{M}$  which simulates  $M$  underneath a basic memory authentication mechanism (starting with an initial tape  $\hat{\mathcal{T}}_0$ ): whenever  $M$  would perform a write “ $\mathcal{T}[a] := v$ ,’,  $\hat{M}$  instead performs the write “ $\hat{\mathcal{T}}[a] := \text{Enc}(t, a, v)$ .’. Whenever  $M$  would read  $\mathcal{T}[a]$ ,  $\hat{M}$  reads a ciphertext at  $\hat{\mathcal{T}}[a]$ , decrypts it to obtain  $(t, a', v)$ , checks that  $a = a'$  and  $t$  is the last time that  $M$  wrote to  $\mathcal{T}[a]$ , and if so returns  $v$  to  $M$ . Additionally,  $\hat{M}$  operates on an encrypted local state. Let  $(\hat{q}_0, \hat{\delta})$  denote the initial state and transition function of  $\hat{M}$ . Finally, the garbled RAM machine is  $(\hat{q}_0, \hat{\delta}, \hat{\mathcal{T}}_0)$ , where  $\hat{\delta} \leftarrow i\mathcal{O}(1^\lambda, \hat{\delta})$ .

The proof of security follows the same high-level structure as for the Turing machine garbling scheme, although several new technical difficulties arise. We give a sequence of computationally indistinguishable changes to  $(\hat{q}_0, \hat{\delta}, \hat{\mathcal{T}}_0)$ , proceeding through several hybrid distributions  $H_1, \dots, H_T$ , and culminating in a distribution which depends only on  $M(\mathcal{T}_0)$ . In hybrid  $H_i$ ,  $\hat{M}$  first performs  $i$  “dummy steps” in which it accesses memory at a sequence of addresses which simulate those accessed by the ORAM. At each of these addresses,  $\hat{M}$  only reencrypts the contents of each memory location it touches. After completing all of the  $i$  dummy steps,  $\hat{M}$  performs  $T - i$  steps of computation as before.

In showing that  $H_{i-1}$  is indistinguishable from  $H_i$ , we take two steps, each with their own challenges. First, we show that  $H_{i-1}$  and  $H_i$  are indistinguishable from distributions in which  $\hat{\delta}$  has hard-coded mappings for the input/output pairs that it would encounter in an honest execution at time  $i$  (and otherwise  $\hat{\delta}$  outputs  $\perp$  for inputs with timestamp  $i$ ). This is not easy (see Claim 4.23), but it leaves the hard-coded input/output mappings in  $\hat{\delta}$  as the only remaining difference between  $H_{i-1}$  and  $H_i$ . Within this, the most notable difference is that  $\hat{\delta}$ ’s hard-coded outputs contain (in the clear) different memory addresses. In  $H_{i-1}$ , these are real addresses accessed by the ORAM, while in  $H_i$  these are simulated addresses (part of a dummy step).

Next, we show that these hard-coded addresses are indistinguishable, even given the other auxiliary information which is present. In particular, we emphasize that in  $H_i$ , the initial state  $\hat{q}_0$  and  $\hat{\mathcal{T}}_0$  contain the entire ORAM internal state following the

first  $i$  accesses. We require the  $i$ th oblivious access to remain oblivious even given this information—this is a forward security property that to the best of our knowledge has not been previously studied for ORAMs. We show that an existing ORAM (namely, the ORAM of [CP13]) has this property. It seems likely that other tree-based ORAMs also have this property, but we have not verified this.

**4.3.1. Oblivious RAM.** We describe the ORAM of Chung and Pass [CP13], which is a simplification of [SCSL11], and we highlight the security property needed for our garbled RAM construction.

**Construction.** Starting with a RAM machine  $\Pi$  that uses  $N$  memory words, the construction transforms it into a machine  $\Pi'$  that uses  $N' = N \cdot \text{poly}(\log N, \lambda)$  memory words. While the eventual goal is to store  $O(1)$  words in the local state of  $\Pi'$ , Chung and Pass start with a “basic” construction in which the local state of  $\Pi'$  consists of a “position map”  $\text{pos} : [N/\alpha] \rightarrow [N/\alpha]$  for some constant  $\alpha > 1$ .

The  $N$  underlying memory locations are divided into  $N/\alpha$  “blocks,” each storing  $\alpha$  underlying memory words. The external memory is organized as a complete binary tree of  $N/\alpha$  leaves. The semantics of the position map is that the  $i$ th block of memory maps to the leaf labeled  $\text{pos}(i)$ . Let  $d = \log(N/\alpha)$ . The CP/SCSL invariant is that

*“Block  $i$  is stored in some node on the path from the root to the leaf labeled with  $\text{pos}(i)$ .”*

Each internal node of the tree stores a few memory blocks. In particular, each internal node, labeled by a string  $\gamma \in \{0, 1\}^{\leq d}$ , is associated with a “bucket” of  $\beta$  blocks for some  $\beta = \text{polylog}(N)$ .

The reads and writes to a location  $r \in [N]$  in the CP/SCSL ORAM proceed as follows:

- **Fetch:** Let  $b = \lceil r/\alpha \rceil$  be the block containing the memory location  $r$ , and let  $i = r \bmod \alpha$  be the component within block  $b$  containing the location  $r$ . We first look up the leaf corresponding to block  $b$  using the (locally stored) position map. Let  $p = \text{Pos}(b)$ .  
Next, we traverse the tree from the root to the leaf  $p$ , reading and writing the bucket associated to each internal node exactly once. In particular, we read the content once, and then we either write it back, or we erase a block once it is found, and write back the rest of the blocks.
- **Update Position Map:** Pick a uniformly random leaf  $p' \leftarrow [N/\alpha]$  and set (in the local memory)  $\text{Pos}(b) = p'$ .
- **Write Back:** In the case of a READ, add the tuple  $(b, p', v)$  to the root of the tree. In the case of a WRITE, add the tuple  $(b, p', v')$ , where  $v'$  is the new value to be written. If there is not enough space in the bucket associated with the root, output **overflow** and abort. (We note that [CP13, SCSL11] show that, setting the parameters appropriately, the probability that the **overflow** event happens is negligible.)
- **Flush the Block:** Pick a uniformly random leaf  $p^* \leftarrow [N/\alpha]$  and traverse the tree from the root to the leaf  $p^*$ , making exactly one read and one write operation for every memory cell associated with the nodes along the path so as to implement the following task: “push down” each tuple  $(\tilde{b}, \tilde{p}, \tilde{v})$  read in the nodes traversed as far as possible along the path to  $p^*$  while ensuring that the tuple is still on the path to its associated leaf  $\tilde{p}$  (i.e., maintaining the CP/SCSL invariant). In other words, the tuple ends up in the node  $\gamma =$  the longest common prefix of  $p^*$  and  $\tilde{p}$ . If at any point some bucket is about to

overflow, abort outputting overflow.

The following observation is central to the correctness and security of the CP/SCSL ORAM.

*Observation 1.* Each oblivious READ and WRITE operation traverses the tree along two randomly chosen paths, independent of the history of operations so far.

This key observation follows from the facts that (1) each position in the position map is used exactly once in a traversal (and before this traversal, this position is not used in determining what nodes to traverse), and (2) the flushing, by definition, traverses a random path, independent of the history.

**Security property.** Suppose an underlying access pattern is given; we will consider the randomized procedure of executing this sequence of accesses via this ORAM. We want a randomized “dummy” access algorithm *OSample* which on input  $j$  outputs a list of locations. This list should be distributed according to the real distribution of accesses corresponding to the  $j$ th underlying access.

We can now describe our desired security property. Let  $Q_s$  be a random variable for the entire ORAM state (both private registers and memory configuration) after the  $s$ th virtual step, and let  $I_j$  be a random variable for the physical addresses accessed on the  $j$ th virtual step.

LEMMA 4.18. *There exist PPT algorithms *Sim* and *OSample* such that for any  $s$ , any initial virtual memory configuration  $\mathcal{T}_0$ , any virtual access pattern  $\vec{a} = (a_1, \dots, a_t)$  (including both addresses and values written), and any possible (nonzero probability) values  $i_1, \dots, i_{s-1}$  of  $I_1, \dots, I_{s-1}$ , it holds that for all PPT adversaries  $\mathcal{A}$ ,*

$$\Pr \left[ \mathcal{A}(q_s^b, i_s^b) = b \left| \begin{array}{l} q_{s-1}^0 \leftarrow \text{Sim}(\mathcal{T}_0, \vec{a}, i_1, \dots, i_{s-1}) \\ q_s^0, i_s^0 \leftarrow \text{OAccess}(a_s; q_{s-1}) \\ i_s^1 \leftarrow \text{OSample}(s) \\ q_s^1 \leftarrow \text{Sim}(\mathcal{T}_0, \vec{a}, i_1, \dots, i_{s-1}, i_s^1) \\ b \leftarrow \{0, 1\} \end{array} \right. \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

This is a consequence of the following two claims.

CLAIM 4.19.  *$I_s$  is independent of  $I_1, \dots, I_{s-1}$  and  $a_1, \dots, a_s$ , and is efficiently sampleable.*

*Proof.* This follows immediately from Observation 1. □

This lets us define *OSample*( $s$ ) as the efficient algorithm which samples  $I_s$ . Specifically, in the CP/SCSL ORAM, *OSample* samples and outputs two uniformly random paths in each tree.

CLAIM 4.20. *The conditional distribution  $Q_s | I_1, \dots, I_s$  is efficiently sampleable for all values of  $I_1, \dots, I_s$  (that jointly occur with nonzero probability).*

*Proof.* Recall that  $Q_s$  has two parts: a position map *Pos* and memory contents  $\tilde{D}_s$ , which are structured as a tree. We first give the sampling procedure for the basic construction, and then extend it to the recursive case. It is easy to verify that this procedure samples the correct distribution.

To sample  $Q_s$  given a sequence of physical memory accesses  $(i_1, \dots, i_s)$ , do the following. For every memory block  $b \in [N/\alpha]$ , let  $\tau_b \leq s$  be the last time when block  $b$  was accessed. Let  $I_j = (I_j^{\text{read}}, I_j^{\text{flush}})$  be the pair of paths that comprise each  $I_j$ .

- For each block  $b$ , pick a uniformly random leaf  $p_b \leftarrow [N/\alpha]$ . Compute the unique internal node  $\gamma_b$  such that  $\gamma_b$  is the largest common prefix between  $p_b$  and any of  $I_{\tau_b}^{\text{flush}}, \dots, I_s^{\text{flush}}$ .



- Construct  $\text{Pos}$  by letting  $\text{Pos}(b) = p_b$ .
- Construct  $\tilde{D}_s$  by writing each memory block  $b$  together with its value at time  $s$  to the internal node  $\gamma_b$ .

To sample  $Q_s$  for the recursive construction, we note that this basic sampler doesn't need to know  $\vec{x}$  and the entire access pattern; it only needs to know each  $\tau_b$  described above, as well as the memory contents at time  $s$ . This information  $(\{\tau_{b'}\}, \vec{x}'_s)$  for the next smaller recursive case is readily computable.  $\vec{x}'_s$  for the next level ORAM is just  $\text{Pos}$ , which we have already computed.  $\tau_{b'}$  is the maximum of  $\tau_b$  over all  $b$  corresponding to  $b'$ . So we can run the basic sampler repeatedly until we have sampled  $Q_s$  for the whole recursive construction.  $\square$

This allows us to define  $\text{Sim}$  as the above procedure for efficiently sampling  $Q_s$  conditioned on  $I_1 = i_1, \dots, I_s = i_s$ .

We can now prove a stronger (statistical) version of Lemma 4.18, knowing that  $\text{Sim}$  and  $\text{OSample}$  output the correct conditional distributions.

CLAIM 4.21.  $\Pr[q_s^0, i_s^0 = q_s, i_s] \approx \Pr[q_s^1, i_s^1 = q_s, i_s]$ .

*Proof.*

$$\begin{aligned} \Pr[q_s^0, i_s^0 = q_s, i_s] &\approx \mathbb{E}_{q_{s-1}}[\Pr[q_s, i_s | q_{s-1}, i_1, \dots, i_{s-1}] | i_1, \dots, i_{s-1}] \\ &= \Pr[q_s, i_s | i_1, \dots, i_{s-1}] \\ &= \Pr[q_s | i_1, \dots, i_s] \Pr[i_s | i_1, \dots, i_{s-1}] \\ &= \Pr[q_s | i_1, \dots, i_s] \Pr[i_s] \\ &\approx \Pr[q_s^1, i_s^1 = q_s, i_s]. \end{aligned}$$

1. The first approximate equality follows from  $\text{Sim}$  approximately sampling  $q_{s-1}$  given  $i_1, \dots, i_{s-1}$  and from  $\text{OAccess}$  (exactly) sampling  $q_s, i_s$  given  $q_{s-1}$ .
2. The second equality is just marginalization over  $q_{s-1}$ .
3. The third equality is the chain rule for probabilities.
4. The fourth equality is Claim 4.19, namely, that the locations accessed at time  $s$  are independent of previously accessed locations.
5. The fifth approximate equality follows from  $\text{OSample}$  approximately sampling  $i_s$ .  $\square$

**4.3.2. Garbling construction.** As for Turing machines, our garbling scheme for RAM machines consists of three parts:

- an initial state  $\tilde{q}_0$ ;
- an obfuscated transition function  $i\mathcal{O}(\hat{\delta})$ ;
- an initial tape  $\tilde{\mathcal{T}}_0$ .

We generate  $(\tilde{q}_0, \hat{\delta}, \tilde{\mathcal{T}}_0)$  by transforming  $(q_0, \hat{\delta}, \mathcal{T}_0)$  into three steps:

1. We apply the ORAM transformation of [CP13]. Additionally, we give the transformed machine a “predictably timed writes” property [GHRW14]. Informally, this states that whenever it accesses a memory address, it already knows the last time at which that address was written to. To supply the randomness needed by the ORAM, we give the transition function a hard-coded puncturable PRF  $F_{\text{ORAM}}$ . The “randomness” used at time  $t$  is given by  $F_{\text{ORAM}}(t)$ .
2. We augment the internal states to contain a timestamp which starts at 0 and is incremented on every step. We also augment the contents of each memory cell to store its address, as well the timestamp at which it was last written.

Finally, we modify  $\delta$  to check, given its internal state, that the accessed memory value is labeled with the expected timestamp and address (checking the expected timestamp uses the predictably timed writes property).

3. We modify the transition function so that its inputs and outputs are all ciphertexts, with the exception of the address  $\text{addr}$ , which it outputs in the clear. We hard-code two secret keys,  $\text{sk}_{\text{state}}$  and  $\text{sk}_{\text{mem}}$ , which it uses to decrypt its inputs and encrypt its outputs using our special authenticated encryption scheme (ACE).

We will let  $M' = (q'_0, \delta')$  and  $\mathcal{T}'_0$  refer to the result of step 1 above. Internal states  $q'$  of  $M'$  will be tuples  $(q, q_{\text{ORAM}})$ , where  $q$  is a state of  $M$  and  $q_{\text{ORAM}}$  is an ORAM state. Our hybrids below will primarily modify  $M'$ , and then will apply the same steps 2 and 3 above.

**4.3.3. Security proof.** Given a RAM machine  $M = (q_0, \delta)$  and an input  $\mathcal{T}_0$ , we give a sequence of hybrid distributions  $\mathbf{H}_0, \dots, \mathbf{H}_T$  such that  $\mathbf{H}_0$  is distributed as our garbling of  $(M, \mathcal{T}_0)$  constructed above, and  $\mathbf{H}_T$  is efficiently sampleable given  $M(\mathcal{T}_0)$ .

**Definition of hybrids.** To better describe our hybrids, we will introduce a bit of terminology. Suppose the ORAM has an overhead of  $\eta$ —so one step of executing  $M$  on  $\mathcal{T}_0$  naturally corresponds to a chunk of  $\eta$  steps of  $\hat{M}$  on  $\hat{\mathcal{T}}_0$ . We refer to one such chunk as a *virtual step*. When we wish to emphasize that we are referencing a single step in the execution of  $\hat{M}$  on  $\hat{\mathcal{T}}_0$ , we use the term *physical step*.

In the description of our hybrids, the greatest complexity arises in describing the transition function  $\hat{\delta}$ . For simplicity, we describe transition functions by how they act at the level of virtual steps. Precisely, we will consider transition functions  $\hat{\delta}$  which on input  $(\tilde{q}, \tilde{s})$  first decrypt  $(t, q) \leftarrow \text{Dec}(\text{sk}_{\text{state}}, \tilde{q})$  and decompose  $t$  as  $\eta \cdot t_{\text{vrt}} + t_{\text{phys}}$ .  $\hat{\delta}$  then does one of four things:

1. *Real virtual step:*
  - (a) Decrypt  $(t_{\text{mem}}, \text{addr}, s) \leftarrow \text{Dec}(\text{sk}_{\text{mem}}, \tilde{s})$ .
  - (b) Check that  $t_{\text{mem}}$  and  $\text{addr}$  are as expected given  $(t, q)$  (i.e.,  $t_{\text{mem}}$  matches the “expected timestamp” given by predictably timed writes). Otherwise output  $\perp$ .
  - (c) Compute  $\text{out} \leftarrow \delta'(q, s)$ .
  - (d) If  $\text{out} \in \{\text{accept}, \text{reject}\}$ , output it. Otherwise, parse out as  $(q', s', \text{addr}')$ .
  - (e) Let  $\tilde{q} = \text{Enc}(\text{sk}_{\text{state}}, (t + 1, q'))$  and let  $\tilde{s} = \text{Enc}(\text{sk}_{\text{mem}}, (t + 1, \text{addr}, s'))$ .
  - (f) Output  $(\tilde{q}, \tilde{s}, \text{addr}')$ .
2. *Dummy virtual step:*
  - (a) Decrypt  $(t_{\text{mem}}, \text{addr}, s) \leftarrow \text{Dec}(\text{sk}_{\text{mem}}, \tilde{s})$ .
  - (b) Check that  $t_{\text{mem}}$  and  $\text{addr}$  are as expected, given  $(t, q)$ . Otherwise output  $\perp$ .
  - (c) Compute  $(\text{addr}_1, \dots, \text{addr}_\eta) \leftarrow \text{OSample}(i; F_{\text{dummy}}(i))$ .
  - (d) Let  $\tilde{q} = \text{Enc}(\text{sk}_{\text{state}}, (t + 1, q))$  and let  $\tilde{s} = \text{Enc}(\text{sk}_{\text{mem}}, (t + 1, \text{addr}, s))$ .
  - (e) Output  $(\tilde{q}, \tilde{s}, \text{addr}_{t_{\text{phys}}})$ .
3. *Half-dummy virtual step:* Do the same thing as in a dummy virtual step, except use hard-coded values for  $(\text{addr}_1, \dots, \text{addr}_\eta)$  instead of sampling them from  $\text{OSample}(i)$ .
4. *Hard-coded virtual step:* There are  $\eta$  different hard-coded output values  $w_1, \dots, w_\eta$ , where each  $w_i$  is a triple  $(c_i^{\text{state}}, c_i^{\text{mem}}, \text{addr}_i)$ . Output  $w_{t_{\text{phys}}}$ .

**Definition of hybrid  $\mathbf{H}_i$ .** In the  $i$ th hybrid,  $\tilde{M}$  is a RAM program  $M''$  which performs  $i$  dummy virtual steps, and then continues for up to  $T - i$  real virtual steps.

$\mathcal{T}_0''$  and  $q_0''$  are generated to be consistent with  $M''$ 's  $i$  dummy steps by using the ORAM's Sim procedure. Namely, let  $\vec{a} = \text{addr}_1, \dots, \text{addr}_{i^*}$  be the virtual addresses accessed by  $M$  on  $x$ . Let  $\vec{a}'' = \text{OSample}(0; F_{\text{dummy}}(0)), \dots, \text{OSample}(i; F_{\text{dummy}}(i))$  be the physical addresses accessed in the first  $i$  virtual steps of  $M''$ . We sample

$$\mathcal{T}_0'', q_{\text{ORAM}}'' \leftarrow \text{Sim}_{\text{ORAM}}(\mathcal{T}_0, \vec{a}, \vec{a}'').$$

$\tilde{\mathcal{T}}_0$  is defined such that

$$\tilde{\mathcal{T}}_0[\text{addr}] = \text{Enc}(\text{sk}_{\text{mem}}, (0, \text{addr}, \mathcal{T}_0''[\text{addr}])),$$

and  $\tilde{q}_0$  is defined such that

$$\tilde{q}_0 = \text{Enc}(\text{sk}_{\text{state}}, (0, (q_i, q_{\text{ORAM}}''))).$$

*Indistinguishability of  $H_{i-1}$  and  $H_i$ .* To show that hybrid  $i - 1$  is indistinguishable from hybrid  $i$ , we introduce additional hybrids  $H_{i-1,1}$  through  $H_{i-1,3}$ . These hybrids' transition functions differ from that of hybrid  $H_{i-1}$  only in how  $\hat{\delta}$  behaves on the  $i$ th virtual step.

HYBRID  $H_{i-1,1}$ : The  $i$ th virtual step is a hard-coded step, with hard-coded values chosen so that the transcript of execution is identical to that of hybrid  $H_{i-1}$ .

HYBRID  $H_{i-1,2}$ : The  $i$ th virtual step is a half-dummy step, with the addresses hard-coded as in hybrid  $H_{i-1,1}$ .

In this hybrid, the initial state  $\tilde{q}_0$  is defined as  $\text{Enc}(\text{sk}_{\text{state}}, (0, q_i))$ , and the initial tape  $\tilde{\mathcal{T}}_0$  is defined such that

$$\tilde{\mathcal{T}}_0[\text{addr}] = \text{Enc}(\text{sk}_{\text{mem}}, (0, \text{addr}, \mathcal{T}_i[\text{addr}])).$$

HYBRID  $H_{i-1,3}$ : The  $i$ th virtual step is again a half-dummy step, but the hard-coded addresses  $(\text{addr}_1, \dots, \text{addr}_\eta)$  are now obtained by sampling

$$(\text{addr}_1, \dots, \text{addr}_\eta) \leftarrow \text{OSample}(i).$$

Additionally,  $\mathcal{T}_0''$  and  $q_{\text{ORAM}}''$  are now sampled as

$$\mathcal{T}_0'', q_{\text{ORAM}}'' \leftarrow \text{Sim}_{\text{ORAM}}(\mathcal{T}_0, \vec{a}, \vec{a}''),$$

where  $\vec{a}''$  denotes the physical addresses accessed by the first  $i$  virtual steps, rather than the first  $i - 1$ .

CLAIM 4.22. *Hybrid  $H_{i-1} \approx H_{i-1,1}$ .*

*Proof.* We apply the template from section 4.2 of iteratively restricting the encryption/decryption keys to ensure that for each  $t_{\text{phys}} \in \{0, \dots, \eta - 1\}$ , conditioned on  $t = \eta i + t_{\text{phys}}$ , there is only one possible output of  $\hat{\delta}$ . Once this unique output condition is satisfied, hard-coding the output values when  $t_{\text{vrt}} = i$  preserves the functionality of  $\hat{\delta}$  and is hence indistinguishable by the security of  $iO$ .

In order to guarantee the above uniqueness of possible outputs, we define the good set  $G_{\text{state}}$  as all  $(t, q)$  satisfying the following:

- If  $t < \eta i$ , then  $q$  differs from  $q_{\eta i}$  only in the parts of the state required for predictably timed writes. In other words,  $q$  essentially doesn't change during the dummy steps.

- If  $\eta i \leq t < (\eta + 1)i$ , then  $q$  is “correct,” i.e., its value is the same as in an honest execution. Note that this only requires hard-coding  $\eta$  values into the description of  $G_{\text{state}}$ .

And we define the good set  $G_{\text{mem}}$  as all  $(t_{\text{mem}}, \text{addr}, s)$  satisfying the following:

- If  $\text{addr}$  is one of the physical addresses accessed in the  $i$ th virtual step, and if  $t_{\text{mem}} < (\eta + 1)i$ , then  $s$  is “correct”; i.e., its value is the same as in an honest execution. Since  $\mathcal{T}[\text{addr}]$  is modified at most  $\eta$  times before virtual step  $i + 1$ , and there are at most  $\eta$  such addresses, the description size of this constraint is  $O(\eta^2)$ .
- If  $\text{addr}$  is the parent of a physical address  $\text{addr}'$  accessed in the  $i$ th virtual step (with memory viewed as a tree according to the predictably timed writes structure), let  $\tau_{\text{addr}'}$  denote the last time that a dummy step accesses  $\text{addr}$ . If  $\tau_{\text{addr}'} \leq t_{\text{mem}} < \eta i$ , then  $s$  encodes that the expected timestamp for  $\text{addr}'$  is exactly  $\tau_{\text{addr}'}$ .

$EK_{\text{state}}, DK_{\text{state}}$  and  $EK_{\text{mem}}, DK_{\text{mem}}$  can be restricted and unrestricted to  $G_{\text{state}}$  and  $G_{\text{mem}}$ , respectively, because  $(G_{\text{state}}, G_{\text{mem}})$  is  $\delta$ -invariant.  $\square$

CLAIM 4.23. *Hybrid*  $H_{i-1,1} \approx H_{i-1,2}$ .

*Proof.* This claim is analogous to Claim 4.17 in our Turing machine construction, but with additional difficulties.

**The challenges.** The heart of the proof of Claim 4.17 replaced a conditional “if  $t = t^*$ ” with a different condition—“if  $\tilde{q} = c_{\text{state}}$ ” or “if  $\tilde{s} = c_{\text{mem}}$ ”—and then argued that these conditions were functionally equivalent with suitably (indistinguishably) punctured keys. In the case of  $\tilde{s}$ , it suffices to show that  $t = t^*$  if and only if  $t_{\text{mem}} = \beta(t^*)$ , where  $\beta(t^*)$  is the time before  $t^*$  that the same address was accessed. For Turing machines, these checks were equivalent, because  $\hat{\delta}$  directly computed and checked that  $t_{\text{mem}} = \beta(t)$ .  $\beta(t)$  was the most recent time the  $t$ th address of the Pippenger–Fischer access pattern was accessed.

However, for an oblivious RAM, there is no such fixed access pattern, and  $\beta$  is therefore not computable by the real-world  $\hat{\delta}$ .  $\hat{\delta}$  computes  $\beta$  instead via the predictably timed write transformation of [GHRW14]. In this transformation, each access to memory is transformed into several ( $\log S$ ) accesses, such that if all previous accesses were performed correctly, then the RAM machine’s local state will know what timestamp to expect in the address presently being accessed.

**The approach.** We will define more elaborate sets on which  $\hat{\delta}$ ’s ACE keys are punctured. These sets must be restrictive enough to ensure that the predictably timed write steps yield the correct answer at time  $t^*$ , but not so restrictive that they have too large of a description size. For example, we cannot define the sets to be “all of the values in a correct execution.” To construct the appropriate sets, we will need to look at the structure of the predictably timed write transformation.

We want to emulate virtual random access to some database  $D$ , such that on every physical memory access to an address  $\text{addr}$ , we know the timestamp at which  $\text{addr}$  was last written to (and therefore the timestamp we should expect to see upon reading  $\text{addr}$ ). We structure physical memory as a complete binary tree whose leaves are the cells of  $D$ . This tree obeys the invariant that, at the beginning of any virtual access, every node stores the times at which each of its children were most recently accessed. To virtually access  $D[i]$ , we first access the root, and then access each successive child on the path to  $D[i]$ . When accessing a node, we first read it, to see when the next child was last written. We then write back the node with an updated

timestamp for that child, knowing that the child is about to be accessed.

**The proof.** To show that  $H_{i-1,1} \approx H_{i-1,2}$ , we use a hybrid argument, which we describe as a sequence of indistinguishable changes.

1. First, we separate the hard-coded state ciphertexts from the hard-coded memory ciphertexts. That is, instead of

If  $t = \eta i + j$ , output  $w_j$

for some hard-coded tuple of outputs  $w_j$ ,  $\hat{\delta}$  will instead have three conditions:

If  $t = \eta i + j$ , then set  $\tilde{q}' = c_j^{\text{state}}$

and

If  $t = \eta i + j$ , then set  $\tilde{s}' = c_j^{\text{mem}}$

and

If  $t = \eta i + j$ , then set  $\text{addr}' = \text{addr}_j$ .

Clearly, this change preserves the functionality of  $\hat{\delta}$ .

2. Next, for every address, we remove all but the last time at which the address is written. So,  $\hat{\delta}$  contains a line of code saying

If  $t = \eta i + \eta - 1$ , then set  $\tilde{q}' = c_{\eta-1}^{\text{state}}$

and also, for every address  $\text{addr}_j$ , a line which says

If  $t = \eta i + j$ , then set  $\tilde{s}' = c_j^{\text{mem}}$ .

We can do such a thing because all but the last hard-coded value will never be decrypted, so the others may as well be dummy steps. The good sets required to enforce this are similar to good sets we have seen already.

3. Finally, for each of these lines, we go through a sequence of hybrids whose result is to (i) remove the line and (ii) modify  $\tilde{T}_0$  or  $\tilde{q}_0$  instead to incorporate the hard-coded  $c_j^{\text{mem}}$  or  $c_{\eta-1}^{\text{state}}$ , respectively. This is the focus of the paragraph below. For brevity, we include only the (more involved) case of removing the lines which give a hard-coded value of  $\tilde{s}'$ .

Let  $\text{addr}^*$  be the address to which  $\tilde{s}'$  is being written, let  $c^* = \text{Enc}(\text{sk}_{\text{mem}}, m^*)$  be the hard-coded value of  $\tilde{s}'$ , and let  $t^*$  be the time at which it is written. We will replace  $t^*$  by  $\beta(t^*)$  until  $\beta(t^*) = 0$ —that is, until  $t^*$  is the first time at which  $\text{addr}^*$  is written. Recall that  $\beta(t^*)$  is the last time before  $t^*$  that  $\text{addr}^*$  was written. We similarly change the timestamp of  $m^*$  from  $t^*$  to  $\beta(t^*)$ . We write this as  $m^*[t_{\text{mem}} \mapsto \beta(t^*)]$ .

If  $\beta(t^*)$  is 0, then we completely remove the line of code and replace  $\tilde{T}_0[\text{addr}]$  by  $\text{Enc}(\text{sk}_{\text{mem}}, m^*)$  (but with the timestamp of  $m^*$  set to 0).

We now must show how  $t^*$  can be replaced by  $\beta(t^*)$ . Again, this is a sequence of indistinguishable changes.

1. We first add another branch to  $\hat{\delta}$ , saying

If  $t = \beta(t^*)$ , then set  $\tilde{s}' = c'^*$

with hard-coded  $\beta(t^*)$  and hard-coded  $c'^*$  chosen to be the “correct” value. This change is indistinguishable for the same reason that  $H_{i-1}$  was indistinguishable from  $H_{i-1,1}$ .

2. Next, we change the condition “if  $t = t^*$ ” into “if  $\tilde{s} = c'^*$ .” With appropriately defined sets  $G_{\text{state}}$  and  $G_{\text{mem}}$ , this change preserves the functionality of  $\hat{\delta}$  and is hence indistinguishable.

3. We replace  $c'_{\text{mem}}^*$  by  $\text{Enc}(\text{sk}_{\text{mem}}, m^*[t_{\text{mem}} \mapsto \beta(t^*)])$ . With appropriately restricted ACE keys (no memory ciphertext with a timestamp of  $\beta(t^*)$  needs to be encrypted or decrypted by  $\hat{\delta}$ ), this is indistinguishable by semantic security.

4. We remove the “if  $\tilde{s} = c'^*$ ” branch. With the newly modified value of  $c'^*$ , its action of setting  $\tilde{s}' = c^*$  is the same as what a dummy step would do, so

removing the branch has no effect on the functionality of  $\hat{\delta}$ .

The main thing now is to show how to make the conditions “if  $t = t^*$ ” and “if  $\tilde{s} = c'^*$ ” functionally equivalent. It suffices to define the good sets such that the following hold:

- $G_{\text{state}}$  is the set of  $(t, q)$  satisfying the following:
  - If  $t \leq t^*$ , then  $q$  differs from  $q_{t^*}$  only in the substate of the predictably timed write mechanism.
  - If  $t = t^*$ , then  $q = q_{t^*}$ .
- $G_{\text{mem}}$  is the set of  $(t_{\text{mem}}, \text{addr}, s)$  satisfying the following:
  - If  $\text{addr}$  is an internal node of the predictably timed write tree on the path from the root to  $\text{addr}^*$ , and if  $t_{\text{mem}}^* \leq t_{\text{mem}} < t^*$ , then  $s$  encodes that the next node on the path to  $\text{addr}^*$  should also have a timestamp which is exactly  $t_{\text{mem}}^*$ .  $\square$

CLAIM 4.24. *Hybrid  $H_{i-1,2} \approx H_{i-1,3}$ .*

*Proof.* The indistinguishability of this change is exactly our ORAM security property, which is stated and proved in Lemma 4.18 of section 4.3.1.  $\square$

CLAIM 4.25. *Hybrid  $H_{i-1,3} \approx H_i$ .*

*Proof.* The difference between  $H_{i-1,3}$  and  $H_i$  is that  $H_{i-1,3}$  has a hard-coded list of addresses  $(\text{addr}_1, \dots, \text{addr}_\eta)$  to access on virtual step  $i$ , while  $H_i$  computes the addresses as  $(\text{addr}_1, \dots, \text{addr}_\eta) = \text{OSample}(i; F_{\text{dummy}}(i))$ .

We give a sequence of indistinguishable changes which transform  $H_{i-1,3}$  into  $H_i$ :

1. Replace the PRF key  $F_{\text{dummy}}$  in  $\hat{\delta}$  with a punctured PRF key  $F_{\text{dummy}}\{i\}$ . This preserves the functionality of  $\hat{\delta}$  and is thus indistinguishable.
2. Instead of hard-coding the addresses as

$$(\text{addr}_1, \dots, \text{addr}_\eta) \leftarrow \text{OSample}(i)$$

(i.e., with true randomness), hard-code them as

$$(\text{addr}_1, \dots, \text{addr}_\eta) = \text{OSample}(i; F_{\text{dummy}}(i)).$$

This change is indistinguishable because  $F_{\text{dummy}}(i)$  is pseudorandom *even to a distinguisher which is given the punctured key  $F_{\text{dummy}}\{i\}$ .*

3. Change  $\hat{\delta}$  to be as in  $H_i$ . This preserves the functionality of  $\hat{\delta}$  and is hence indistinguishable.  $\square$

CLAIM 4.26. *There is a PPT simulator  $\text{Sim}$  such that hybrid  $H_T$  is computationally indistinguishable from  $\text{Sim}(M(\mathcal{T}_0), 1^S, 1^T)$ .*

*Proof.* This is not immediately clear, since in our definition of  $H_T$ ,  $\tilde{\mathcal{T}}_0$  still contains an encryption of  $\mathcal{T}_T$ , so we need to argue that  $\mathcal{T}_T$  is not revealed. Before describing how to do this, we first remark that one can get around this in a black-box way with no asymptotic overhead if restricted to RAM machines whose running time is longer than their database size. Instead of garbling a RAM machine  $M$  directly, one simply garbles an  $M'$  which runs  $M$ , and then erases its memory.

However, since the RAM model of computation is arguably most useful when applied to sublinear computations, we would like to achieve something more efficient.

We use a sequence of  $S + 1$  hybrids  $H_{T,0}, \dots, H_{T,S}$ , where hybrid  $H_{T,j}$  differs from  $H_T$  only in that, for all  $0 \leq k < j$ ,

$$\tilde{\mathcal{T}}_0[k] = \text{Enc}(\text{sk}_{\text{mem}}, (0, k, \perp)).$$

It therefore suffices to show how to change  $\tilde{\mathcal{T}}_0[j]$  from  $\text{Enc}(\text{sk}_{\text{mem}}, (0, j, \mathcal{T}_T[j]))$  to  $\text{Enc}(\text{sk}_{\text{mem}}, (0, j, \perp))$ .

For this, we give another sequence of hybrids  $H_{T,j-1,0}, \dots, H_{T,j-1,N_j}$ , where the number of hybrids  $N_j$  is equal to the number of times that  $\hat{\delta}$  accesses address  $j$ . The structure of these hybrids resembles those used in Claim 4.23, but in reverse.

Let  $t_k$  denote the  $k$ th time that  $\hat{\delta}$  accesses address  $j$ . Hybrid  $H_{T,j-1,k}$  is defined to differ from  $H_{T,j}$  only in that  $\hat{\delta}$  hard-codes  $\tilde{s}' = \text{Enc}(\text{sk}_{\text{mem}}, (t_k + 1, j, \mathcal{T}_T[j]))$  at time  $t_k$ . Or, if  $k = 0$ ,  $\tilde{\mathcal{T}}[j]$  has the hard-coded value  $\text{Enc}(\text{sk}_{\text{mem}}, (0, j, \mathcal{T}_T[j]))$  instead.

Just as in Claim 4.23, these hybrids are shown to be indistinguishable by introducing another hard-coded branch at  $t_{k+1}$ , and then switching the hard-coded value of  $\tilde{s}'$  so that the “if  $t = t_k$ ” branch is functionally equivalent to a dummy step. All of this is done with analogously defined good sets.  $\square$

**4.4. The bounded space requirement and subsequent work.** In this section we describe why the size of our randomized encodings grows proportional to the space complexity of the underlying computation, and give a taste of the techniques used by [KLW15] to remove this dependence.

The dependence on the space complexity appears to come from our proof technique rather than being inherent in the construction. Consider for simplicity our Turing machine construction. Recall that only one part of the construction grows with  $S$ —the initial garbled tape  $\tilde{\mathcal{T}}_0$ , defined as

$$\tilde{\mathcal{T}}_0[i] = \begin{cases} \text{Enc}(\text{sk}_{\text{mem}}, (0, i, x_i)) & \text{if } i < |x|, \\ \text{Enc}(\text{sk}_{\text{mem}}, (0, i, \perp)) & \text{otherwise} \end{cases}$$

for  $i = 0, \dots, S - 1$ . In other words,  $\tilde{\mathcal{T}}_0$  is a bit-by-bit encryption of the input  $x$ , padded with  $S - |x|$  dummy ciphertexts. One natural approach to making the randomized encoding fully succinct is to give out the dummy ciphertexts implicitly (and succinctly) instead of explicitly, e.g., via a small obfuscated circuit that on input  $i$  (for  $|x| \leq i < S$ ) outputs  $\text{Enc}(\text{sk}_{\text{mem}}, (0, i, \perp))$ . One may conjecture the security of this variant, and we do not know of any attacks. However, major difficulties arise in reducing the security of this variant to the security of the underlying primitives (**IO** for circuits and ACE).

Indeed, recall that we prove security via a hybrid argument, where in the  $i$ th hybrid  $\tilde{\mathcal{T}}_0$  encrypts the (padded)  $i$ th memory configuration  $\mathcal{T}_i$  of  $M$  when executed on  $x$ . In order for these hybrids to be indistinguishable from one another,  $\tilde{\mathcal{T}}_0$  (in all hybrids) needs to be as large as  $S$ , the size of the largest memory configuration of  $M$  when executed on  $x$ . This hybrid structure was no accident: due to the succinctness requirements on the sets to which ACE keys are constrainable, we were only able to change the behavior of  $\hat{\delta}$  on the  $i + 1$ st step once  $\hat{\delta}$  already performs dummy actions on the first  $i$  steps. This behavior of  $\hat{\delta}$  necessitated hard-coding the  $i$ th configuration in  $\tilde{\mathcal{T}}_0$  for correctness.

[KLW15] get around this difficulty by developing a new way of authenticating the inputs to  $\hat{\delta}$ : Rather than using an authenticated encryption scheme directly on encrypted memory words, they maintain a signed Merkle tree of the entire tape contents. This hash function is designed so that the Merkle tree is “somewhere statistically binding” as in [HW15]. Namely, the “local opening” property of the Merkle tree can be strengthened: for any index  $j \in [S]$ , the hash function can be indistinguishably sampled so that for any Merkle root, there is only one local opening at  $j$ . Thus the problem of authenticating the memory inputs to  $\hat{\delta}$  on step  $i$  is reduced to maintaining

a small authenticated local state (the TM state and Merkle tree root).

To authenticate this state, [KLW15] uses a special “splittable signature” scheme with constrainability properties resembling those of ACE. They use a puncturable PRF to generate one signing/verification key pair per time step. As in our security proof, the signing and verification keys are alternately constrained until the verification key used on the  $i$ th time step will only accept signatures of the *correct*  $i$ th local state. The main difficulty now is how to “forget” (for the sake of succinctness) that keys of the first  $i - 1$  time steps are constrained—otherwise the number of hard-coded keys or messages in  $\hat{\delta}$  will grow linearly with  $T$ . [KLW15] show how this can be accomplished by augmenting each local state  $q_i$  with a special hash of  $q_{i-1}$ .

**5. Applications.** In this section, we address three of our main applications of succinct garbling schemes: succinct **iO**, publicly verifiable delegation, and  $\mathcal{SNARG}$ s (succinct noninteractive argument systems). (The rest of the applications outlined in the introduction follow directly by plugging in our succinct garbling into previous work.) In fact, all of our applications can be instantiated with succinct randomized encodings; namely, they do not require separate input encoding. We first recall the syntax and properties of randomized encodings.

**Randomized encodings.** A randomized encoding scheme  $\mathcal{RE} = (\text{REnc}, \text{Dec})$  for  $\{\mathcal{AL}_\lambda\}$  consists of a randomized encoding algorithm  $\text{REnc}$  and a decoding algorithm  $\text{Dec}$ .  $\text{REnc}(1^\lambda, AL, x)$ , given any function  $AL \in \mathcal{AL}_\lambda$  and input  $x$ , returns the encoded computation  $\widehat{AL}(x)$ . Given such an encoding,  $\text{Dec}$  can decode the result  $AL(x)$ . Any garbling scheme  $\mathcal{GS} = (\text{Garb}, \text{Encode}, \text{Eval})$  for  $\{\mathcal{AL}_\lambda\}$  can be projected to a corresponding randomized encoding where  $\text{REnc} = \text{Garb} \circ \text{Encode}$  is given by

$$(\widehat{AL}, \hat{x}) \stackrel{\$}{\leftarrow} \text{REnc}(1^\lambda, AL, x), \text{ where } (\widehat{AL}, \text{key}) \stackrel{\$}{\leftarrow} \text{Garb}(1^\lambda, AL, x), \hat{x} = \text{Encode}(\text{key}, x)$$

and the evaluation algorithm  $\text{Eval}$  is the decoding algorithm  $\text{Dec}$ .

In accordance, the correctness, security, and efficiency properties are all defined similarly to garbling schemes, as defined in section 2.2 (in particular, it will be convenient to consider randomized encodings that, like garbling schemes, also guarantee the privacy of the program and not just the input). When projecting a garbling scheme to a randomized encoding scheme as above, the randomized encoding inherits the corresponding efficiency properties of the garbling scheme.

**5.1. Overview.** We briefly sketch the main ideas behind our main applications of succinct randomized encodings.

**Succinct iO.** The construction of succinct **iO** from randomized encoding and exponential **iO** for circuits is a natural instantiation of the bootstrapping approach suggested by Applebaum [App14]. There, the goal is to reduce obfuscation of general circuits to obfuscation of  $\text{NC}^1$  circuits; our goal is to reduce obfuscation of programs with large running time (but bounded space) to obfuscation of significantly smaller circuits. To obfuscate a succinct program  $\Pi$  with respect to inputs of size at most  $n$ , we obfuscate a small circuit  $C^{\Pi, K}$  that has a hardwired seed  $K$  for a PRF and, given input  $x$ , applies the PRF to  $x$  to derive randomness and then computes a succinct randomized encoding of  $\widehat{\Pi}(x)$ . The obfuscated  $i\mathcal{O}(\Pi)$ , given input  $x$ , computes the encoding, decodes it, and returns the result.

The analysis in [App14] establishes security in case the circuit obfuscator  $i\mathcal{O}$  is virtually black-box secure. We show that if  $i\mathcal{O}$  has  $2^{-\lambda^\epsilon}$ -security for security parameter  $\lambda \gg n^\epsilon$ , and the PRF is puncturable and is also  $2^{-\lambda^\epsilon}$ -secure, then a similar



result holds for **iO** (rather than virtual black box). The proof is based on a general *probabilistic iO* argument, an abstraction recently made by Canetti et al. [CLTV15].

**Succinct FE.** The construction of succinct functional encryption (FE) follows rather directly by plugging our randomized encodings into previous constructions of nonsuccinct functional encryption. Concretely, starting with the scheme of Gentry et al. [GHRW14], we can replace the nonsuccinct randomized encodings for RAM in their construction with our succinct randomized encodings and obtain selectively secure FE.<sup>13</sup> Alternatively, starting from the scheme of Garg et al. [GGHZ16], we can replace randomized encodings for circuits in their construction with our succinct randomized encodings and get an adaptively secure succinct FE scheme. (Here we also need to rely on the fact that succinct randomized encodings can be computed in low depth, which is required in their construction.) We note that in both cases, our succinct randomized encodings already satisfy the required security for their security proof to go through, and only the succinctness features change.

**Publicly verifiable delegation.** Finally, we sketch the basic ideas behind the delegation scheme. The delegation scheme is pretty simple and similar in spirit to previous delegation schemes (in a weaker processing model) [AIK10, GGP10, PRV12, GKP<sup>+</sup>13]. To delegate a computation, given by  $\Pi$  and  $x$ , the verifier simply sends the prover a randomized encoding  $\widehat{\Pi}'(x, r)$ , where  $\Pi'$  is a machine that returns  $r$  if and only if it accepts  $x$ , and  $r$  is a sufficiently long random string. The security of the randomized encoding implies that the prover learns nothing of  $r$ , unless the computation is accepting. The scheme can be easily made publicly verifiable by publishing  $f(r)$  for some one-way function  $f$ . Furthermore, the scheme ensures input-privacy for the verifier.

We then propose a simple transformation that can be applied to any delegation scheme in order to make the first verifier message reusable. The idea is natural: we let the verifier's first message be an obfuscation of a circuit  $C^K$  that has a hardwired key  $K$  for a puncturable PRF and, given a computation  $(\Pi, x)$ , applies the PRF to derive randomness and generates a first message for the delegation scheme. Thus, for each computation, a first message is effectively sampled afresh. Relying on **iO** and the security of the puncturable PRF, we can show that (nonadaptive) soundness is guaranteed. The transformation can also be applied to privately verifiable delegation schemes, such as the one of [KRR14], and maintains soundness, even if the prover has a verification oracle.

**5.2. From randomized encodings to iO.** We present a generic transformation from a garbling scheme for an algorithm class  $\{\mathcal{AL}_\lambda\}$  to an indistinguishability obfuscator for  $\{\mathcal{AL}_\lambda\}$ , assuming subexponentially indistinguishability obfuscators for circuits. We require that the algorithm class have the property that for any  $\lambda < \lambda' \in \mathbb{N}$ , it holds that every algorithm  $AL \in \mathcal{AL}_\lambda$  is also contained in  $\mathcal{AL}_{\lambda'}$ —we say that such a class is “monotonically increasing.” For instance, the class of Turing machines TM and RAM machines RAM are all monotonically increasing.

**PROPOSITION 5.1.** *Let  $\{\mathcal{AL}_\lambda\}$  be any monotonically increasing class of deterministic algorithms. It holds that if there are*

- (i) *a subexponentially indistinguishable **iO**,  $i\mathcal{O}^C$ , for circuits, and*
- (ii) *a subexponentially indistinguishable randomized encoding  $\mathcal{RE}$  for  $\{\mathcal{AL}_\lambda\}$ ,*

<sup>13</sup>Formally, their construction is given in terms of garbling for RAM rather than randomized encodings, but these are actually used as randomized encodings, without making special use of independent input encoding.

then there is an indistinguishability obfuscator  $i\mathcal{O}^A$  for  $\{\mathcal{AL}_\lambda\}$ .

Furthermore, the following efficiency preservation holds:

- If  $\mathcal{RE}$  has optimal efficiency or I/O-dependent complexity,  $i\mathcal{O}^A$  has I/O-dependent complexity.
- If  $\mathcal{RE}$  has space-dependent complexity, so does  $i\mathcal{O}^A$ .
- If  $\mathcal{RE}$  and  $i\mathcal{O}^C$  have linear-time-dependent complexity, so does  $i\mathcal{O}^A$ .

Before moving to the proof of the proposition, we first note that combining Proposition 5.1 with constructions of garbling schemes for TM and RAM in section 3, we directly obtain  $\mathbf{iO}$  for TM and RAM with space-dependent complexity.

**THEOREM 5.2.** *Assume a subexponentially indistinguishable  $\mathbf{iO}$  for circuits and subexponentially secure OWF. There is an indistinguishability obfuscator for TM and RAM with space-dependent complexity.*

**Proof of Proposition 5.1.** This result relies on the following natural way of obfuscating probabilistic circuits, abstracted in [CLTV15].

**Probabilistic  $\mathbf{iO}$ .** Let  $i\mathcal{O}$  and  $F$  be  $2^{\lambda^\varepsilon}$ -indistinguishable  $\mathbf{iO}$  and puncturable PRF. Given a probabilistic circuit  $C$ , obfuscate it in the following way: Consider another circuit  $\Pi^{C,k}$  that on input  $x$  computes  $C$  using pseudorandom coins  $F(k, x)$  generated with a hardwired PRF key  $k$ , that is,  $\Pi^{C,k}(x) = C(x; F(k, x))$ . The obfuscation of  $C$ , denoted by  $pi\mathcal{O}(1^\lambda, C)$ , is an  $\mathbf{iO}$  obfuscation of  $\Pi^{C,k}$  for a randomly sampled key  $C$ , that is,

$$\widehat{C} \stackrel{\$}{\leftarrow} pi\mathcal{O}(1^\lambda, C), \text{ where } k \stackrel{\$}{\leftarrow} \text{PRF.Gen}(1^{\lambda'}); \widehat{C} \stackrel{\$}{\leftarrow} i\mathcal{O}(1^{\lambda'}, \Pi^{C,k}),$$

where  $\lambda' = (\lambda + n)^{1/\varepsilon}$  for  $n = C.n$ , so that  $\mathbf{iO}$  and  $F$  are  $\text{negl}(\lambda)2^n$ -indistinguishable. The work of [CLTV15] showed that the above obfuscations are indistinguishable for circuits whose output distributions are strongly indistinguishable for every input. More specifically, circuits  $C_1$  and  $C_2$  with the same input length  $n$  are strongly indistinguishable (with respect to auxiliary input  $z$ ) if for every input  $x \in \{0, 1\}^n$  the outputs  $C_1(x)$  and  $C_2(x)$  are  $\text{negl}(\lambda)2^{-n}$ -indistinguishable (given  $z$ ). The following lemma summarizes this.

**LEMMA 5.3** ( $pi\mathcal{O}$  for circuits [CLTV15]). *Assume subexponentially indistinguishable  $\mathbf{iO}$  for circuits  $i\mathcal{O}^C$ , and subexponentially indistinguishable OWF. Then, for every class  $\{C_\lambda\}$  of polynomial-sized circuits, and every nonuniform PPT sampleable distribution  $\mathcal{D}$  over the support of  $\{C_\lambda \times C_\lambda \times \{0, 1\}^{\text{poly}(\lambda)}\}$ , if it holds that for every nonuniform PPT adversary  $\mathcal{R}$ , and input  $x$ , that*

$$\begin{aligned} & \left| \Pr[(C_1, C_2, z) \stackrel{\$}{\leftarrow} \mathcal{D}_\lambda, y \stackrel{\$}{\leftarrow} C_1(x) : \mathcal{R}(C_1, C_2, x, y, z) = 1] \right. \\ & \left. - \Pr[(C_1, C_2, z) \stackrel{\$}{\leftarrow} \mathcal{D}_\lambda, y \stackrel{\$}{\leftarrow} C_2(x) : \mathcal{R}(C_1, C_2, x, y, z) = 1] \right| \leq \text{negl}(\lambda) \cdot 2^{-n}, \end{aligned}$$

then the following ensembles are computationally indistinguishable:

$$\{C_1, C_2, pi\mathcal{O}(1^\lambda, C_1), z\}_\lambda \approx \{C_1, C_2, pi\mathcal{O}(1^\lambda, C_2), z\}_\lambda.$$

For completeness, we include a proof sketch of the lemma.

*Proof sketch of Lemma 5.3.* The lemma essentially follows from complexity leveling. To see the proof, first consider a simpler case, where the two circuits  $C_1$  and  $C_2$  have identical implementation on all but one input  $x^*$ , and the outputs on  $x^*$ ,  $C_1(x^*)$  and  $C_2(x^*)$ , are indistinguishable. In this case, it follows directly from the security

**iO** that obfuscation of  $C_b, \hat{C}_b \stackrel{s}{\leftarrow} pi\mathcal{O}(1^\lambda, C_1)$  is indistinguishable to the obfuscation of  $C'_b \stackrel{s}{\leftarrow} i\mathcal{O}(C'_b)$ , where  $C'_b$  has a punctured key  $k(x^*)$  and  $C_b(x^*; F(k, x^*))$  hardwired in; then it follows from the pseudorandomness of puncturable PRF and the indistinguishability of  $C_1(x^*)$  and  $C_2(x^*)$  that  $i\mathcal{O}(C'_b)$  and  $i\mathcal{O}(C'_1)$  are indistinguishable. Therefore, overall obfuscation of  $C_1$  and  $C_2$  are indistinguishable.

Now consider the case where  $C_1$  and  $C_2$  are sampled from  $\mathcal{D}(1^\lambda)$ , and their output distributions for every input are  $\text{negl}(\lambda)2^{-n}$ -indistinguishable. To show that their **pIO** obfuscation are indistinguishable, consider an exponential,  $2^n$ , number of hybrids where, in each hybrid, a circuit  $C_i$  is obfuscated, which outputs  $C_2(x)$  for every input  $x \leq i$  and outputs  $C_1(x)$  for every input  $x > i$ . Since in every two neighboring hybrids,  $C_i$  and  $C_{i+1}$  are the same except on one input,  $x^* = i + 1$ . By the argument above, neighboring hybrids have a distinguishing gap  $O(\text{negl}(\lambda)2^{-n})$ . Thus, by a hybrid argument, obfuscations of  $C_1$  and  $C_2$  are indistinguishable. This concludes the lemma.  $\square$

**Construction of iO for general algorithms.** Using Lemma 5.3, we now prove Proposition 5.1.

Given  $2^{-\lambda^\epsilon}$ -indistinguishable **iO**  $i\mathcal{O}^C$  and  $2^{-\lambda^\epsilon}$ -indistinguishable randomized encoding  $\mathcal{RE}$ , let  $pi\mathcal{O}$  be the obfuscator for probabilistic circuits constructed from  $i\mathcal{O}^C$ . Our **iO** for the a general algorithm class  $\{\mathcal{AL}_\lambda\}$  is defined as follows:

$$\widehat{AL}(\cdot) \stackrel{s}{\leftarrow} i\mathcal{O}^A(1^\lambda, AL), \text{ where } \widehat{AL}(\cdot) \stackrel{s}{\leftarrow} pi\mathcal{O}(\lambda, \text{REnc}(1^{\lambda'}, AL, \cdot)),$$

where the security parameter  $\lambda' = (\lambda + n)^{1/\epsilon}$  for  $n = AL.n$  so that  $\text{REnc}$  is  $\text{negl}(\lambda)2^n$ -indistinguishable. (Note that the reason that we can use the security parameter  $\lambda' > \lambda$  is because the algorithm class is monotonically increasing, and thus  $AL \in \mathcal{AL}_{\lambda'}$  also belongs to  $\mathcal{AL}_{\lambda'}$ .) The correctness of  $i\mathcal{O}^A$  follows from the correctness of  $\mathcal{RE}$  and  $i\mathcal{O}^C$  underlying  $pi\mathcal{O}$ . Next, we show the security of  $i\mathcal{O}^A$ .

**Security.** Fix a polynomial  $T$ , a *nonuniform* PPT sampleable distribution  $\mathcal{D}$  over the support  $\{\mathcal{AL}_\lambda^T \times \mathcal{AL}_\lambda^T \times \{0, 1\}^{\text{poly}(\lambda)}\}$ , such that, with overwhelming probability,  $(AL_1, AL_2, z) \leftarrow \mathcal{D}(1^\lambda)$  satisfies that  $AL_1$  and  $AL_2$  are functionally equivalent and have matching parameters. We want to show that the following distributions are indistinguishable:

$$\left\{ (AL_1, AL_2, z) \stackrel{s}{\leftarrow} \mathcal{D}(1^\lambda) : (i\mathcal{O}^A(1^\lambda, AL_1), z) \right\}_\lambda, \\ \left\{ (AL_1, AL_2, z) \stackrel{s}{\leftarrow} \mathcal{D}(1^\lambda) : (i\mathcal{O}^A(1^\lambda, AL_2), z) \right\}_\lambda.$$

By construction of  $i\mathcal{O}^A$ , this is equivalent to showing

$$\left\{ (AL_1, AL_2, z) \stackrel{s}{\leftarrow} \mathcal{D}(1^\lambda) : (pi\mathcal{O}(1^\lambda, \text{REnc}(1^{\lambda'}, AL_1, \cdot)), z) \right\}_\lambda, \\ \left\{ (AL_1, AL_2, z) \stackrel{s}{\leftarrow} \mathcal{D}(1^\lambda) : (pi\mathcal{O}(1^\lambda, \text{REnc}(1^{\lambda'}, AL_2, \cdot)), z) \right\}_\lambda.$$

Consider the sampler  $\mathcal{D}'(1^\lambda)$  that outputs  $C'_1, C'_2, z$  by sampling  $(AL_1, AL_2, z) \stackrel{s}{\leftarrow} \mathcal{D}(1^\lambda)$  and setting  $C'_b = \text{REnc}(1^{\lambda'}, AL_b, \cdot)$ . It follows from the security of  $\mathcal{RE}$  that for every nonuniform adversary  $\mathcal{R}$  and every input  $x$ , the output distributions of  $C'_1(x)$  and  $C'_2(x)$  are  $\text{negl}(\lambda)2^n$ -indistinguishable, given  $x, z, C'_1, C'_2$ . Thus, it follows from Lemma 5.3 that the above two ensembles are indistinguishable, as are the obfuscations of  $AL_1$  and  $AL_2$ .

**Efficiency.** Finally, we analyze the efficiency of  $i\mathcal{O}^A$ . It is easy to see that  $pi\mathcal{O}(1^\lambda, C)$  runs in time  $T_{pIO}(\lambda, C.n, |C|)$ , where  $T_{pIO}$  is a polynomial depending on the running time of the underlying  $\mathbf{iO}$  and PRF as well as the parameters of their subexponential security; moreover, if the underlying  $\mathbf{iO}$  has linear-time-dependent complexity,  $p_{pIO}$  also depends linearly in  $|C|$  (still polynomially in  $\lambda$  and  $C.n$ ). Let  $T_{REnc}(\lambda', |AL|, n, m, S, T)$  be the running time of  $REnc(1^{\lambda'}, AL, x)$ . Overall, the running time of  $i\mathcal{O}^A(1^\lambda, AL)$  is

$$T_{pIO}(\lambda, n, T_{REnc}(\lambda', |AL|, n, m, S, T)), \text{ where } \lambda' = \text{poly}(\lambda, n).$$

Therefore,

- if  $\mathcal{RE}$  has optimal efficiency (that is,  $T_{REnc}$  depends only on  $m$ ) or I/O-dependent complexity (that is,  $T_{REnc}$  does not depend on  $S, T$ ),  $i\mathcal{O}^A$  has I/O-dependent complexity;
- if  $\mathcal{RE}$  has space-dependent complexity (that is,  $T_{REnc}$  does not depend on  $T$ ), so does  $i\mathcal{O}^A$ ;
- if  $\mathcal{RE}$  and the underlying  $\mathbf{iO}$  has linear-time-dependent complexity (that is,  $T_{REnc}$  depends linearly on  $T$  and  $T_{pIO}$  depends linearly on  $|C|$ ), so does  $i\mathcal{O}^A$ .

This concludes the proof of Proposition 5.1.

**A corollary: Output-independence.** The size of the randomized encodings (or garbling schemes) described in previous sections depends (linearly) on the output of the encoded computation; accordingly, so does the succinct  $\mathbf{iO}$  construction described in this section. We start by noting that in the succinct  $\mathbf{iO}$  construction this dependence can be easily removed. Concretely, rather than considering the machine  $AL(x)$  that for any input  $x$  might have an  $m$ -bit output  $y$ , we can consider a new single-bit machine  $AL'(x, i)$  that, given additional input  $i \in \{0, 1\}^{\log m}$ , outputs  $y_i$ . Observe that if  $AL_0$  and  $AL_1$  compute the same function, then clearly so do their single-bit versions,  $AL'_0$  and  $AL'_1$ . The overhead is only polylogarithmic in the output size  $m$ . Thus, we directly obtain succinct  $\mathbf{iO}$  that is output-independent. We note that this does not involve making any additional computational assumptions. (Note that we only increase the input size logarithmically, and thus the exponential loss in the input incurred by the transformation given by Theorem 5.1 is only polynomial.)

Next, we observe that this directly implies indistinguishability-based succinct randomized encodings that are output-independent. The encoding of  $AL, x$  simply consists of an (output-independent) obfuscation of a machine that has no input and output  $AL(x)$ . (Note that this only requires polynomial  $\mathbf{iO}$ , since the exponential blowup in the input size of the transformation given by Theorem 5.1 is completely avoided.)

**More efficient construction.** Evaluating the  $\mathbf{iO}$  for TM and RAM obtained in Theorem 5.2 on input  $x$  involves evaluating the obfuscated program on  $x$  once to obtain a randomized encoding  $\widehat{AL}(x)$ , and then decoding it. When relying on an arbitrary randomized encoding with space-dependent complexity, the overall evaluation takes time  $T_{AL}(x) \times \text{poly}(\lambda, |AL|, S)$ . When the space is large, the overhead on the running time is large.

We now improve the evaluation efficiency by combining Proposition 5.1 with the specific RAM garbling scheme of Theorem 3.5.

**THEOREM 5.4.** *Assume a subexponentially indistinguishable  $\mathbf{iO}$  for circuits and subexponentially secure OWFs. There is an indistinguishability obfuscator for TM and RAM, where obfuscation of a machine  $R$  takes linear time in the space complexity*

$\text{poly}(\lambda, n, |R|) \times S$ , and evaluation of the obfuscated program on input  $x$  takes time  $\text{poly}(\lambda, n, |R|) \times (T_R(x) + S)$ , with  $n = R.n$  and  $S = R.S$ .

Toward the above theorem, consider instantiating Proposition 5.1 using the RAM garbling scheme of Theorem 3.5 and a subexponentially secure **iO** scheme with quasi-linear complexity (implied by subexponentially secure **iO** and OWF as shown in section 5.4). Recall that the RAM garbling scheme of Theorem 3.5 has linear-space-dependent complexity  $\text{poly}(\lambda, |R|) \times S$  and evaluation time  $\text{poly}(\lambda, |R|) \times (T^* + S)$  with  $T^* = T_R(x)$ ; such a garbling scheme leads to a randomized encoding algorithm **REnc** with the same encoding and decoding complexity. By the same efficiency analysis as in Proposition 5.1, this instantiation yields an **iO** for RAM with linear-space-dependent complexity, namely,  $\text{poly}(\lambda, |R|, n)S$ . Therefore, its evaluation time is now  $\text{poly}(\lambda, |R|, n) \times S + \text{poly}(\lambda, |R|) \times (T^* + S)$ , which is  $\text{poly}(\lambda, |R|, n) \times (S + T^*)$ .

*Remark 5.5* (indistinguishability obfuscation with linear complexity in the program size). In section 3.3.1, we showed that the efficiency of our RAM garbling scheme can be improved to depend only linearly in program description size  $|R|$ , namely, it has garbling complexity of  $\text{poly}(\lambda)(|R| + S)$  and evaluation complexity of  $\text{poly}(\lambda)(|R| + S + T^*)$ . When using such a RAM garbling scheme as the underlying scheme in our construction of **iO** scheme for RAM, we obtain an **iO** scheme with complexity  $\text{poly}(\lambda, n)(|R| + S)$  and evaluation time  $\text{poly}(\lambda, n)(|R| + S + T^*)$ .

**5.3. Publicly verifiable delegation,  $\mathcal{SNARG}$ s for  $\mathbf{P}$ , and succinct NIZKs for  $\mathbf{NP}$ .** We now present the publicly verifiable delegation scheme for bounded space computations, following from our succinct randomized encodings, as well as a general transformation from delegation schemes to succinct noninteractive arguments. We also note the implications to succinct NIZKs as a corollary of our succinct **iO** and the work of [SW14].

**5.3.1. P-delegation.** A delegation system for  $\mathbf{P}$  is a 2-message protocol between a verifier and a prover. The verifier consists of two algorithms  $(\mathcal{G}, \mathcal{V})$ : given a (well-formed) algorithm, input, and security parameter  $z = (AL, x, \lambda)$ ,  $\mathcal{G}$  generates a message  $\sigma$ . The prover, given  $(z, \sigma)$ , produces a proof  $\pi$  attesting that  $AL$  accepts  $x$  within  $AL.T$  steps.  $\mathcal{V}$  then verifies the proof. In a privately verifiable system, the  $\mathcal{G}$  produces, in addition to the (public) message  $\sigma$ , a secret verification state  $\tau$ , and verification by  $\mathcal{V}$  requires  $(z, \sigma, \tau, \pi)$ . In a publicly verifiable scheme,  $\tau$  can be published (together with  $\sigma$ ), without compromising soundness.

We shall require that the running time of  $(\mathcal{G}, \mathcal{V})$  will be significantly smaller than  $AL.T$ , and that the time to prove is polynomially related to  $AL.T$ .

**DEFINITION 5.6 (P-delegation).** A prover and verifier  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$  constitute a delegation scheme for  $\mathbf{P}$  if it satisfies the following:

1. **Completeness:** For any  $z = (AL, x, \lambda)$ , such that  $AL$  accepts  $x$  within  $AL.T$  steps,

$$\Pr \left[ \mathcal{V}(z, \sigma, \tau, \pi) = 1 \mid \begin{array}{l} (\sigma, \tau) \leftarrow \mathcal{G}(z) \\ \pi \leftarrow \mathcal{P}(z, \sigma) \end{array} \right] = 1.$$

2. **Soundness:** For any polynomial-sized prover  $\mathcal{P}^*$ , polynomial  $T(\cdot)$ , there exists a negligible  $\alpha(\cdot)$  such that for any  $z = (AL, x, \lambda)$ , such that  $AL.T \leq T(\lambda)$ , and  $AL$  does not accept  $x$  within  $AL.T$  steps,

$$\Pr \left[ \mathcal{V}(z, \sigma, \tau, \pi) = 1 \mid \begin{array}{l} (\sigma, \tau) \leftarrow \mathcal{G}(z) \\ \pi \leftarrow \mathcal{P}^*(z, \sigma) \end{array} \right] \leq \alpha(\lambda).$$

- 3. Optimal verification and instance-based prover efficiency: *There exists a (universal) polynomial  $p$  such that for every  $z = (AL, x, \lambda)$ ,*
  - *the verifier algorithms  $(\mathcal{G}, \mathcal{V})$  run in time  $p(\lambda, |AL|, |x|, \log AL.T)$ ;*
  - *the prover  $\mathcal{P}$  runs in time  $p(\lambda, |AL|, |x|)T_{AL}(x)$ .*
- 3'. Space-dependent verification complexity: *The scheme has space-dependent verification complexity if the running time of the  $(\mathcal{G}, \mathcal{V})$  may also depend on space; concretely, there exists a (universal) polynomial  $p$  such that for every  $z = (AL, x, \lambda)$ ,*
  - *the verifier algorithms  $(\mathcal{G}, \mathcal{V})$  run in time  $p(\lambda, |AL|, \log AL.T, AL.S)$ .*

The system is said to be publicly verifiable if soundness is maintained when the malicious prover  $\mathcal{P}^*$  is also given the verification state  $\tau$ .

*Remark 5.7 (input privacy).* Our construction achieves an additional property of input privacy which states that the first message of the delegation scheme  $\sigma$  leaks no information about the input  $x$  on which the computation of  $AL$  is being delegated, beyond the output  $AL(x)$ . This ensures that, in the outsourcing computation application, the server performing the computation learns no more than is necessary about the input to the computation.

We next present a publicly verifiable delegation with fast verification based on any succinct randomized encoding, and one-way functions.

**The scheme.** Let  $f$  be a one-way function, and let  $(REnc, Dec)$  be a randomized encoding scheme. We describe  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$  as follows: Let  $z = (AL, x, \lambda)$  be a tuple consisting of an algorithm, input, and security parameter.

**Generator  $\mathcal{G}(z)$ .** For  $r \leftarrow \{0, 1\}^\lambda$ , let  $AL'(x, r)$  be the machine that returns  $r$  if  $AL(x) = 1$ , and  $\perp$  otherwise.  $\mathcal{G}$  generates and outputs  $\sigma \leftarrow REnc(1^\lambda, \Pi', (x, r))$  and  $\tau = f(r)$ .

**Prover  $\mathcal{P}(z, \sigma)$ .**  $\mathcal{P}$  simply runs  $\pi \leftarrow Dec(\sigma)$  and outputs  $\pi$ .

**Verifier  $\mathcal{V}(z, \sigma, \tau, \pi)$ .**  $\mathcal{V}$  outputs 1 if and only if  $f(\pi) = \tau$ .

We prove that  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$  is a **P**-delegation scheme as follows.

**THEOREM 5.8.** *If  $(REnc, Dec)$  is a randomized encoding scheme with optimal complexity (resp., space-dependent complexity), then  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$  as described above is a publicly verifiable **P**-delegation scheme with optimal verification (resp., space-dependent verification).*

*Proof.* The completeness of  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$  follows directly from the correctness of  $(REnc, Dec)$ . Also, note that the running time of the verifier algorithms  $(\mathcal{G}, \mathcal{V})$  is related to the running time of  $REnc$ . Therefore, it also follows directly that if  $(REnc, Dec)$  has optimal complexity (resp., space-dependent complexity), then  $(\mathcal{G}, \mathcal{V})$  satisfies the property of optimal verification (resp., space-dependent verification), and the instance-based prover efficiency follows from the fact the randomized encoding has instance efficiency. It remains to show the soundness of  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$ .

To show soundness, we will rely on the security of  $(REnc, Dec)$  and the one-wayness of  $f$ . Assume for contradiction there exist polynomial-size prover  $\mathcal{P}^*$  and polynomial  $p(\cdot)$  such that for infinitely many  $z = (AL, x, \lambda)$ , where  $AL$  does not accept  $x$  and  $AL.T \leq p(\lambda)$ , we have that

$$\Pr \left[ \mathcal{V}(z, \sigma, \tau, \pi) = 1 \mid \begin{array}{l} (\sigma, \tau) \leftarrow \mathcal{G}(z) \\ \pi \leftarrow \mathcal{P}^*(z, \sigma, \tau) \end{array} \right] \geq \frac{1}{p(\lambda)} .$$

Let  $\mathcal{Z}$  be the sequence of such  $z = (AL, x, \lambda)$  and consider any  $z \in \mathcal{Z}$ . Recall that  $\mathcal{G}(z)$  samples  $r \leftarrow \{0, 1\}^\lambda$  and outputs  $\sigma \leftarrow \text{REnc}(1^\lambda, AL', (x, r))$  and  $\tau \leftarrow f(r)$ . Since  $AL$  does not accept  $x$ , we have that  $AL'(x, r)$  outputs  $\perp$ . By the security of  $(\text{REnc}, \text{Dec})$ , there exists a PPT simulator  $\text{Sim}$  such that the ensembles  $\{\text{REnc}(1^\lambda, AL', (x, r))\}_{r \in \{0, 1\}^\lambda, z \in \mathcal{Z}}$  and  $\{\text{Sim}(1^\lambda, \perp, AL', 1^{|x|+|r|})\}_{r \in \{0, 1\}^\lambda, z \in \mathcal{Z}}$  are indistinguishable. Therefore, given a simulated  $\sigma \leftarrow \text{Sim}(1^\lambda, \perp, AL', 1^{|x|+|r|})$  we have that  $\mathcal{P}^*$  still convinces  $\mathcal{V}$  with some noticeable probability. More formally, for infinitely many  $z \in \mathcal{Z}$ , we have that

$$\Pr \left[ \mathcal{V}(z, \sigma, \tau, \pi) = 1 \mid \begin{array}{l} r \leftarrow \{0, 1\}^\lambda \\ \sigma \leftarrow \text{Sim}(1^\lambda, \perp, AL', 1^{|x|+|r|}) \\ \tau \leftarrow f(r) \\ \pi \leftarrow \mathcal{P}^*(z, \sigma, \tau) \end{array} \right] \geq \frac{1}{p(\lambda)} - \alpha(\lambda)$$

for some negligible function  $\alpha(\cdot)$ .

Recall that  $\mathcal{V}$  outputs 1 if and only if  $f(\pi) = \tau$ . Therefore  $\mathcal{V}(z, \sigma, \tau, \pi) = 1$  implies that  $\mathcal{P}^*$ , when given  $\tau = f(r)$ , outputs  $\pi$ , which is in the preimage of  $f(r)$ . Hence  $\mathcal{P}^*$  can be used to break the one-wayness of  $f$ , and we have a contradiction. This completes the proof of the theorem.  $\square$

**5.3.2.  $\mathcal{SNARG}$ s for  $\mathbf{P}$ .** A succinct noninteractive argument system ( $\mathcal{SNARG}$ ) for  $\mathbf{P}$  is a delegation system where the first message  $\sigma$  is *reusable*, independent of any specific computation, and can be used to verify an unbounded number of computations. In a privately verifiable  $\mathcal{SNARG}$ , soundness might not be guaranteed if the prover learns the result of verification on different inputs, which can be seen as certain leakage on the private state  $\tau$  (this is sometimes referred to as the *verifier rejection problem*). Accordingly, in this case, we shall also address a strong soundness requirement, which says that soundness holds even in the presence of a verification oracle.

**DEFINITION 5.9 ( $\mathcal{SNARG}$ ).** A  $\mathcal{SNARG}$   $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$  is defined as a delegation scheme, with the following change to the syntax of  $\mathcal{G}$ : the generator  $\mathcal{G}$  now gets as input a security parameter, time bound, and input bound  $\lambda, T, n \in \mathbb{N}$ , and does not get  $AL, x$  as before. We require the following:

1. Completeness: For any  $z = (AL, x, \lambda)$ , such that  $AL.T \leq T$  and  $|AL, x| \leq n$ , and  $AL$  accepts  $x$ ,

$$\Pr \left[ \mathcal{V}(z, \sigma, \tau, \pi) = 1 \mid \begin{array}{l} (\sigma, \tau) \leftarrow \mathcal{G}(\lambda, T, n) \\ \pi \leftarrow \mathcal{P}(z, \sigma) \end{array} \right] = 1 .$$

2. Soundness: For any polynomial-sized prover  $\mathcal{P}^*$ , polynomials  $T(\cdot), n(\cdot)$ , there exists a negligible  $\alpha(\cdot)$  such that for any  $z = (AL, x, \lambda)$ , where  $AL.T \leq T(\lambda)$ ,  $|AL, x| \leq n(\lambda)$ , and  $AL$  does not accept  $x$ ,

$$\Pr \left[ \mathcal{V}(z, \sigma, \tau, \pi) = 1 \mid \begin{array}{l} (\sigma, \tau) \leftarrow \mathcal{G}(\lambda, T(\lambda), n(\lambda)) \\ \pi \leftarrow \mathcal{P}^*(z, \sigma) \end{array} \right] \leq \alpha(\lambda) .$$

- 2\*. Strong soundness: For any polynomial-sized oracle-aided prover  $\mathcal{P}^*$ , polynomials  $T(\cdot), n(\cdot)$ , there exists a negligible  $\alpha(\cdot)$  such that for any  $z = (AL, x, \lambda)$ , where  $AL.T \leq T(\lambda)$ ,  $|AL, x| \leq n(\lambda)$ , and  $AL$  does not accept  $x$ ,

$$\Pr \left[ \mathcal{V}(z, \sigma, \tau, \pi) = 1 \mid \begin{array}{l} (\sigma, \tau) \leftarrow \mathcal{G}(\lambda, T(\lambda), n(\lambda)) \\ \pi \leftarrow \mathcal{P}^{\mathcal{V}(\cdot, \sigma, \tau, \cdot)}(z, \sigma) \end{array} \right] \leq \alpha(\lambda) .$$

3. Optimal verification and instance-based prover efficiency: *There exists a (universal) polynomial  $p$  such that for every  $z = (AL, x, \lambda)$ ,*
  - *the verifier algorithms  $(\mathcal{G}, \mathcal{V})$  run in time  $p(\lambda, n(\lambda), \log AL.T)$ ;*
  - *the prover  $\mathcal{P}$  runs in time  $p(\lambda, |AL|, |x|)T_{AL}(x)$ .*

As before, the system is said to be publicly verifiable if soundness is maintained when the malicious prover is also given the verification state  $\tau$ . (In this case, strong soundness follows from standard soundness.) Also, we can naturally extend the definition for the case of semisuccinctness, in which case,  $\mathcal{G}$  will also get a space bound  $S$ , and the running time of algorithms  $(\mathcal{G}, \mathcal{V})$  may also depend on  $S$ .

*Remark 5.10* (nonadaptive soundness). Note that in the definition above and in our construction, we will consider only *nonadaptive* soundness, as opposed to *adaptive* soundness, where the malicious prover  $\mathcal{P}^*$  can pick the statement  $z$  after seeing the first message  $\sigma$ .

We now show a simple transformation, based on **iO**, that takes any 2-message delegation scheme (e.g., the one constructed above) and turns it into a *SNARG* for **P**. The transformation works in either the public or private verification setting. Furthermore, it always results in a *SNARG* with strong soundness; even the delegation we start with does not have strong soundness (such as the scheme of [KRR14]).

**The scheme.** Let  $(\mathcal{P}_d, (\mathcal{G}_d, \mathcal{V}_d))$  be a **P**-delegation scheme,  $(\text{PRF.Gen}, \text{PRF.Punc}, \text{F})$  a puncturable PRF scheme, and **iO** an indistinguishability obfuscator. We describe a *SNARG*  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$  as follows.

Let  $z = (AL, x, \lambda)$  be a tuple consisting of an algorithm, input, and security parameter such that  $|AL, x| \leq n$  and  $AL.T \leq T$ . For notational convenience, we decompose  $\mathcal{G}_d$  into  $(\mathcal{G}_{\sigma_d}, \mathcal{G}_{\tau_d})$ , where  $\mathcal{G}_{\sigma_d}(z)$  only outputs the message  $\sigma_d$  and  $\mathcal{G}_{\tau_d}(z)$  only outputs the secret verification state  $\tau_d$ .

**Generator  $\mathcal{G}(\lambda, T, n)$ .**

1.  $\mathcal{G}$  samples a puncturable PRF key  $K \leftarrow \text{PRF.Gen}(1^\lambda)$ .
2. Let  $C_\sigma[K]$  be a circuit that on input  $z$  runs  $r \leftarrow \text{F}(K, z)$  and outputs  $\sigma_d \leftarrow \mathcal{G}_{\sigma_d}(z; r)$ . That is,  $C_\sigma$  runs  $\mathcal{G}_{\sigma_d}$  to generate a first message of the delegation scheme, using randomness from the PRF key  $K$ . Similarly,  $C_\tau[K]$  on input  $z$  runs  $r \leftarrow \text{F}(K, z)$  and outputs  $\tau_d \leftarrow \mathcal{G}_{\tau_d}(z; r)$ .  $\mathcal{G}$  generates the circuits  $C_\sigma[K]$  and  $C_\tau[K]$  and pads them to be of size  $\ell_\sigma$  and  $\ell_\tau$ , respectively, which will be specified precisely later in the analysis. For now, we mention that if we use a delegation scheme with optimal verification, then  $\ell_\sigma, \ell_\tau \leq \text{poly}(\lambda, n, \log T)$ . We subsequently assume the circuits  $C_\sigma$  and  $C_\tau$  are padded.
3.  $\mathcal{G}$  runs  $\sigma \leftarrow i\mathcal{O}(1^\lambda, C_\sigma[K])$ ,  $\tau \leftarrow i\mathcal{O}(1^\lambda, C_\tau[K])$  and outputs  $(\sigma, \tau)$ .

**Prover  $\mathcal{P}(z, \sigma)$ .**  $\mathcal{P}$  runs  $\sigma$  on input  $z$  to get  $\sigma_d \leftarrow \sigma(z)$ . Note that  $\sigma_d$  is a first message of the underlying delegation scheme  $(\mathcal{P}_d, (\mathcal{G}_d, \mathcal{V}_d))$ . Next,  $\mathcal{P}$  generates the corresponding proof of the delegation scheme  $\pi \leftarrow \mathcal{P}_d(z, \sigma_d)$  and outputs  $\pi$ .

**Verifier  $\mathcal{V}(z, \sigma, \tau, \pi)$ .**  $\mathcal{V}$  runs  $\sigma_d \leftarrow \sigma(z)$ ,  $\tau_d \leftarrow \tau(z)$  and outputs the result of  $\mathcal{V}_d(z, \sigma_d, \tau_d, \pi)$ .

**THEOREM 5.11.** *Assume the existence of an indistinguishability obfuscator  $i\mathcal{O}$ . If  $(\mathcal{P}_d, (\mathcal{G}_d, \mathcal{V}_d))$  is a privately verifiable (resp., publicly verifiable) **P**-delegation scheme, then  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$  as described above is a privately verifiable (resp., publicly verifiable) *SNARG* with strong soundness. Moreover, if the delegation scheme has optimal or space-dependent verification and relative prover efficiency, then so does the *SNARG*.*



*Proof.* The completeness of  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$  follows directly from that of  $(\mathcal{P}_d, (\mathcal{G}_d, \mathcal{V}_d))$  and the correctness of  $\mathbf{iO}$ . The running time of  $\mathcal{G}(\lambda, T, n)$  is polynomial in  $\lambda$  and the maximum running time of  $\mathcal{G}_d$  on inputs  $z = (AL, x, \lambda)$ , where  $|AL, x| \leq n$  and  $AL.T \leq T$ . Similarly, the running times of  $\mathcal{P}$  and  $\mathcal{V}$  are polynomial in  $\lambda$  and the running times of  $\mathcal{P}_d$  and  $\mathcal{V}_d$ , respectively. Therefore, the optimal (or space-dependent) verification and prover efficiency of  $(\mathcal{P}_d, (\mathcal{G}_d, \mathcal{V}_d))$  implies that the same properties hold for  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$ .

To show strong soundness of  $(\mathcal{P}, (\mathcal{G}, \mathcal{V}))$ , we will rely on the soundness of  $(\mathcal{P}_d, (\mathcal{G}_d, \mathcal{V}_d))$  and the security of  $\mathbf{iO}$  and the punctured PRF (PRF.Gen, PRF.Punc, F). We will first consider the privately verifiable setting. Assume for contradiction that there exist polynomial-sized oracle-aided prover  $\mathcal{P}^*$ , polynomials  $T(\cdot), n(\cdot), p(\cdot)$  such that for infinitely many  $z = (AL, x, \lambda)$ , where  $AL.T \leq T(\lambda)$ ,  $|AL, x| \leq n(\lambda)$ , and  $AL$  does not accept  $x$ ,

$$\Pr \left[ \mathcal{V}(z, \sigma, \tau, \pi) = 1 \mid \begin{array}{l} (\sigma, \tau) \leftarrow \mathcal{G}(\lambda, T(\lambda), n(\lambda)) \\ \pi \leftarrow \mathcal{P}^{*\mathcal{V}(\cdot, \sigma, \tau, \cdot)}(z, \sigma) \end{array} \right] \geq \frac{1}{p(\lambda)} .$$

We will refer to the above probability as the advantage  $\mathcal{A}(z, \mathcal{P}^*)$ . We will now construct a malicious prover  $\mathcal{P}_d^*$  to break the soundness of the delegation scheme.  $\mathcal{P}_d^*$  gets as input  $z$  and  $\sigma_d$ , which is some first message of the delegation scheme.  $\mathcal{P}^*$  runs a subroutine  $\mathcal{D}$  described in the following paragraph, on input  $(z, \sigma_d)$ , to obtain a “fake”  $\mathcal{SNARG}$  message and verification state  $(\sigma, \tau)$ , which it will then use to run  $\mathcal{P}^*$  and answer its queries. That is,  $\mathcal{P}_d^*$  runs  $(\sigma, \tau) \leftarrow \mathcal{D}(z, \sigma_d)$ ,  $\pi \leftarrow \mathcal{P}^{*\mathcal{V}(\cdot, \sigma, \tau, \cdot)}(z, \sigma)$  and outputs  $\pi$ . The subroutine  $\mathcal{D}$  is defined as follows.

**Subroutine  $\mathcal{D}(z, \sigma_d)$ .**

1.  $\mathcal{D}$  samples a puncturable PRF key  $K \leftarrow \text{PRF.Gen}(1^\lambda)$  and punctures it at the input  $z$  to obtain a punctured key  $K_z \leftarrow \text{PRF.Punc}(K, z)$ .
2. Let  $C_\sigma^*[K_z, \sigma_d]$  be a circuit that on input  $z^*$  behaves as follows: If  $z^* = z$ , then  $C_\sigma^*$  simply outputs the hardwired value  $\sigma_d$ . Otherwise,  $C_\sigma^*$  runs  $r \leftarrow \text{F}(K_z, z^*)$  and outputs the result of  $\mathcal{G}_{\sigma_d}(z^*; r)$ .
3. Similarly, let  $C_\tau^*[K_z]$  be a circuit that on input  $z^*$  behaves as follows: If  $z^* = z$ , then  $C_\tau^*$  simply outputs  $\perp$ . Otherwise,  $C_\tau^*$  runs  $r \leftarrow \text{F}(K_z, z^*)$  and outputs the result of  $\mathcal{G}_{\tau_d}(z^*; r)$ .
4.  $\mathcal{D}$  generates the circuits  $C_\sigma^*[K_{z^*}, \sigma_d]$  and  $C_\tau^*[K_{z^*}]$  and pads them to sizes  $\ell_\sigma$  and  $\ell_\tau$ , respectively, where  $\ell_\sigma$  is the maximum size of the circuits  $C_\sigma^*[K_{z^*}, \sigma_d]$  and  $C_\sigma^*[K]$  and  $\ell_\tau$  is the maximum size of the circuits  $C_\tau^*[K_{z^*}]$  and  $C_\tau^*[K]$ . We subsequently assume the circuits  $C_\sigma^*$  and  $C_\tau^*$  are padded.
5.  $\mathcal{D}$  generates  $\sigma \leftarrow i\mathcal{O}(1^\lambda, C_\sigma^*[K_z, \sigma_d])$ ,  $\tau \leftarrow i\mathcal{O}(1^\lambda, C_\tau^*[K_z])$  and outputs  $(\sigma, \tau)$ .

Note that when  $\mathcal{P}_d^*$  uses  $\tau$ , as generated by  $\mathcal{D}$  above, to answer  $\mathcal{P}^*$ 's verification oracle queries on the input  $z$ , then, unlike a “real” verification state,  $\tau$  simply outputs  $\perp$ . In this case,  $\mathcal{P}_d^*$  answers the query with the bit 0 (rejecting the proof submitted in the query).

We now analyze the success probability of  $\mathcal{P}_d^*$ . We want to show there exists a polynomial  $p'$  such that for infinitely many  $z = (AL, x, \lambda)$ , where  $AL$  does not accept  $x$ , the following holds:

$$\mathcal{A}_d(z, \mathcal{P}_d^*) = \Pr \left[ \mathcal{V}_d(z, \sigma_d, \tau_d, \pi) = 1 \mid \begin{array}{l} (\sigma_d, \tau_d) \leftarrow \mathcal{G}_d(z) \\ \pi \leftarrow \mathcal{P}_d^*(z, \sigma_d) \end{array} \right] \geq \frac{1}{p'(\lambda)} .$$

Let  $\mathcal{Z}$  be the sequence of such  $z = (AL, x, \lambda)$ .

To show  $\mathcal{P}_d^*$  succeeds with noticeable probability, we will consider a hybrid malicious prover  $\mathcal{P}_d^{\text{Hyb}}$  that is very similar to  $\mathcal{P}_d^*$  except that it also gets the secret verification state  $\tau_d$  as input and uses it in a different subroutine  $\mathcal{D}^{\text{Hyb}}$ . We will first show that for every  $z \in \mathcal{Z}$ ,  $\mathcal{A}_d(z, \mathcal{P}_d^*) = \mathcal{A}_d(z, \mathcal{P}_d^{\text{Hyb}})$ . Next, we show that relying on the security of the indistinguishability obfuscator and the puncturable PRF,  $\mathcal{A}_d(z, \mathcal{P}_d^{\text{Hyb}})$  is negligibly close to  $\mathcal{A}(z, \mathcal{P}^*)$  for all  $z \in \mathcal{Z}$ . By assumption,  $\mathcal{A}(z, \mathcal{P}^*)$  is noticeable, and hence we have that  $\mathcal{A}_d(z, \mathcal{P}_d^*)$  is noticeable, contradicting the soundness of the  $\mathbf{P}$ -delegation scheme.

We now describe the hybrid malicious prover  $\mathcal{P}_d^{\text{Hyb}}$ .  $\mathcal{P}_d^{\text{Hyb}}$  gets as input  $z$  and both  $\sigma_d$  and  $\tau_d$ . It uses the hybrid subroutine  $\mathcal{D}^{\text{Hyb}}$  on input  $(z, \sigma_d, \tau_d)$  to generate a hybrid “fake”  $(\sigma, \tau)$  to run  $\mathcal{P}^*$  and answer its queries. However, unlike  $\mathcal{P}_d^*$ , it uses  $\tau$  to answer *all* of  $\mathcal{P}^*$ ’s queries (including those on input  $z$ ).  $\mathcal{D}^{\text{Hyb}}$  is defined as follows.

**Subroutine  $\mathcal{D}^{\text{Hyb}}(z, \sigma_d, \tau_d)$ .**

1.  $\mathcal{D}^{\text{Hyb}}$  samples  $K_z$  and generates  $\sigma$  exactly as in  $\mathcal{D}$ . The only difference is in the generation of  $\tau$ .
2. Let  $C_\tau^*[K_z, \tau_d]$  be a circuit that on input  $z^*$  behaves as follows: If  $z^* = z$ , then  $C_\tau^*$  simply outputs the hardwired value  $\tau_d$ . Otherwise,  $C_\tau^*$  runs  $r \leftarrow F(K_z, z^*)$  and outputs the result of  $\mathcal{G}_{\tau_d}(z^*; r)$ .
3.  $\mathcal{D}^{\text{Hyb}}$  generates  $C_\tau^*[K_z, \tau_d]$ , pads it to the maximum size of  $C_\tau^*[K_z, \tau_d]$  and  $C_\tau[K]$ , and generates  $\tau \leftarrow i\mathcal{O}(1^\lambda, C_\tau^*[K_z, \tau_d])$ .  $\mathcal{D}^{\text{Hyb}}$  outputs  $(\sigma, \tau)$ .

We now observe that for every  $z \in \mathcal{Z}$ ,  $\mathcal{A}_d(z, \mathcal{P}_d^*) = \mathcal{A}_d(z, \mathcal{P}_d^{\text{Hyb}})$ . The only difference in the two experiments is in the view of  $\mathcal{P}^*$ : When run by  $\mathcal{P}_d^*$ , its oracle responses are answered using  $\tau$  as generated by  $\mathcal{D}$ , and when run by  $\mathcal{P}_d^{\text{Hyb}}$ , its oracle responses are answered using  $\tau$  as generated by  $\mathcal{D}^{\text{Hyb}}$ . However, we claim that the responses are distributed identically in both cases. They could only potentially differ on queries on the input  $z$ , but since  $z$  is a “false” input, i.e.,  $AL$  does not accept  $x$ , in both cases the verification oracle response on such queries is 0 (reject).

Next we show that there is a negligible function  $\alpha(\cdot)$  such that for every  $z \in \mathcal{Z}$ ,

$$|\mathcal{A}_d(z, \mathcal{P}_d^{\text{Hyb}}) - \mathcal{A}(z, \mathcal{P}^*)| \leq \alpha(\lambda).$$

We first observe that in the experiment corresponding to  $\mathcal{A}_d(z, \mathcal{P}_d^{\text{Hyb}})$ , the event  $\mathcal{V}_d(z, \sigma_d, \tau_d, \pi) = 1$  is equivalent to the event  $\mathcal{V}(z, \sigma, \tau, \pi) = 1$ , where  $(\sigma, \tau) \leftarrow \mathcal{D}^{\text{Hyb}}(z, \sigma_d, \tau_d)$ . This follows directly from the construction of  $\mathcal{V}$  and the fact that  $\sigma$  and  $\tau$  are hardwired to output  $\sigma_d$  and  $\tau_d$  on input  $z$ . Hence we have that

$$\mathcal{A}_d(z, \mathcal{P}_d^{\text{Hyb}}) = \Pr \left[ \mathcal{V}(z, \sigma, \tau, \pi) = 1 \mid \begin{array}{l} (\sigma_d, \tau_d) \leftarrow \mathcal{G}_d(z) \\ (\sigma, \tau) \leftarrow \mathcal{D}^{\text{Hyb}}(z, \sigma_d, \tau_d) \\ \pi \leftarrow \mathcal{P}^{\mathcal{V}(\cdot, \sigma, \tau, \cdot)}(z, \sigma) \end{array} \right].$$

Viewed this way, we can now observe that the only difference between the above experiment and that of  $\mathcal{A}(z, \mathcal{P}^*)$  is in how  $(\sigma, \tau)$  are generated. In the above experiment,  $(\sigma, \tau)$  comes from  $\mathcal{D}^{\text{Hyb}}$  and  $\mathcal{G}_d$ , whereas in the experiment for  $\mathcal{A}(z, \mathcal{P}^*)$ ,  $(\sigma, \tau)$  comes from  $\mathcal{G}$ . It suffices to show the following claim.

CLAIM 5.12. *The following ensembles are computationally indistinguishable:*

$$(5.1) \quad \{(\sigma, \tau) : (\sigma, \tau, \pi) \leftarrow \mathcal{D}^{\text{Hyb}}(z, \sigma_d, \tau_d), (\sigma_d, \tau_d) \leftarrow \mathcal{G}_d(z)\}_{z \in \mathcal{Z}}$$

$$(5.2) \quad \approx_c \{(\sigma, \tau) : (\sigma, \tau) \leftarrow \mathcal{G}(\lambda, T(\lambda), n(\lambda))\}_{z \in \mathcal{Z}}.$$

*Proof.* Recall that in ensemble (5.1),  $\sigma \leftarrow i\mathcal{O}(C_\sigma^*[K_z, \sigma_d])$ , where  $K_z$  is a PRF key punctured at input  $z$  and  $C_\sigma^*$  on all input  $z$  outputs  $\sigma_d$  and on all other inputs  $z^*$  outputs  $\mathcal{G}_d(z^*; F(K_z, z^*))$ . However, in ensemble (5.2),  $\sigma \leftarrow i\mathcal{O}(C_\sigma[K])$ , where  $C$  on input  $z^*$  outputs  $\mathcal{G}_d(z^*; F(K, z^*))$ . The difference between  $\tau$  in ensembles (5.1) and (5.2) is the same. Indistinguishability follows from the security of  $\mathbf{iO}$  and that of (PRF.Gen, PRF.Punc, F) in the standard way. We provide a brief overview.

Consider a hybrid ensemble that is identical to ensemble (5.1) except that instead of uniform randomness,  $\mathcal{G}_d$  uses randomness from  $F(K, z)$ , where  $K$  is a PRF key.  $K$  is then punctured at input  $z$  and given to  $\mathcal{D}^{\text{Hyb}}$  to use as  $K_z$ . By the security of the punctured PRF, this hybrid ensemble is indistinguishable from ensemble (5.1). Furthermore, the circuits obfuscated as  $\sigma$  and  $\tau$  in this hybrid ensemble and in ensemble (5.2) are functionally equivalent. Hence, by the security of  $\mathbf{iO}$ , ensemble (5.2) is indistinguishable from the hybrid ensemble. A hybrid argument completes the proof of the claim.  $\square$

This completes the proof of strong soundness in the privately verifiable setting. Proving strong soundness in the publicly verifiable setting is very similar. The malicious prover for the  $\mathcal{SNARG} \mathcal{P}^*$  now also requires  $\tau$  as input to generate the convincing proof  $\pi$ . On the other hand the prover we want to construct for the delegation scheme  $\mathcal{P}_d^*$  gets  $\tau_d$  as input from the challenger.  $\mathcal{P}_d^*$  uses the same strategy as  $\mathcal{P}_d^{\text{Hyb}}$  to generate  $\tau$  and simply gives it to  $\mathcal{P}^*$ . Using the same proof as above, we have that if  $\mathcal{P}^*$  succeeds with noticeable probability, then so does  $\mathcal{P}_d^*$ .  $\square$

**5.3.3. Succinct perfect NIZK for NP.** In this section we briefly note that by using the succinct indistinguishability obfuscator from section 5.2 in the construction of [SW14], we can obtain an NIZK argument scheme for any NP language  $\mathcal{R}_L$  that is perfect zero-knowledge and additionally *succinct* in the following sense: Let  $\Pi^{\mathcal{R}}$  be a uniform program that computes the NP relation  $\mathcal{R}(x, w)$ , and let  $\tau(n)$  and  $s(n)$  be, respectively, bounds on the length of witness and space needed by  $\Pi^{\mathcal{R}}$  for  $n$ -bit statements. The length of the common reference string (CRS) of the scheme for proving  $n$ -bit statements grows polynomially with  $n$ ,  $\tau(n)$ , and  $s(n)$  (and is essentially independent of the verification time of the language). Below we provide a brief overview of the [SW14] construction and how it can be made succinct using succinct  $\mathbf{iO}$ .

In [SW14], the NIZK scheme relies on  $\mathbf{iO}$  for circuits as follows: The CRS contains an obfuscation of two circuits that contain the same PRF key. The first obfuscation is used by the Prover to generate proofs: the circuit takes as input a statement and witness  $(x, w)$  of lengths  $n$  and  $\tau(n)$ , and outputs the image of the input under the PRF as the proof if the witness is valid, that is,  $\Pi^{\mathcal{R}}(x, w) = 1$ . The second obfuscation is used by the Verifier to check if this proof is valid. [SW14] shows how to use this idea relying on  $\mathbf{iO}$  and puncturable PRFs. In their construction, the length of the proof is succinct: it depends only on the security parameter. However, the length of the CRS is related to the size of the circuits obfuscated in the CRS, which is related to the verification time. We note that by obfuscating the pair of Turing machines that perform the above functionality, and using our succinct indistinguishability obfuscator instead, the length of the CRS can be made to depend on the statement and witness lengths, as well as the space complexity of the verification program, independent of the verification time.

Note that this succinct construction relies on our succinct indistinguishability obfuscator, which in turn relies on subexponentially secure  $\mathbf{iO}$  for circuits (as opposed to standard  $\mathbf{iO}$  for circuits that the [SW14] construction is based on).

**THEOREM 5.13** (follows from [SW14]). *Assuming subexponentially secure  $\mathbf{iO}$  for circuits and subexponentially secure OWFs, there exists an NIZK argument scheme for every NP language determined by a uniform polynomial-time program  $\Pi^{\mathcal{R}}$  with the following properties:*

1. *The scheme is perfectly zero-knowledge.*
2. *The scheme has adaptive soundness.<sup>14</sup>*
3. *There are universal polynomials  $p$ ,  $p'$ , and  $p''$ , such that the length of the CRS of the scheme for verifying statements  $x$  of length  $n$  is  $p(\lambda, n, \tau(n), s(n))$ , where  $\lambda$  is the security parameter,  $\tau(n)$  is a bound on the length of witness, and  $s(n)$  is the space complexity of  $\Pi^{\mathcal{R}}$  for verifying  $n$ -bit statements. The length of the proof is  $p'(\lambda)$ . The running time of the prover for statement and witness  $(x, w)$  is  $p''(\lambda, n, \tau(n), s(n))T$ , where  $T$  is the running time of  $\Pi^{\mathcal{R}}(x, w)$ .*

*Remark 5.14* (improved efficiency of delegation and  $\mathcal{SNARG}$ s). When plugging in the more efficient garbling scheme of section 3.3.1, we directly obtain a publicly verifiable delegation scheme that the verification complexity is  $\text{poly}(\lambda)(|AL| + AL \cdot S + |x|)$  and the prover complexity for a  $T$ -time computation  $AL(x)$  is  $\text{poly}(\lambda)(|AL| + AL \cdot S + T)$ . Furthermore, combining this delegation scheme with Theorem 5.11, while applying the trick of “obfuscating in a piecemeal fashion” as in sections 3.3 and 5.2, we obtain a  $\mathcal{SNARG}$  for  $\mathbf{P}$  with the same verification and prover efficiency.

Finally, using the more efficiency succinct  $\mathbf{iO}$  of Remark 5.5 in Theorem 5.13, the size of the CRS can be improved to  $\text{poly}(\lambda, n, \tau(n)) \cdot s(n)$ , and the prover efficiency can be improved to  $\text{poly}(\lambda, n, \tau(n))(s(n) + T)$ .

**5.4. A new bootstrapping theorem.** In general, when considering  $\mathbf{iO}$  for circuits, the size of an obfuscation  $|i\mathcal{O}(C)|$  (or more generally the time required to obfuscate) is allowed to be an arbitrary polynomial in the original circuit size  $|C|$ . In known candidate constructions (e.g., [GGH<sup>+</sup>13]) the blowup is quadratic (see the discussion in [GHRW14]). In this section, we show how to bootstrap  $\mathbf{iO}$  for circuits (with arbitrary polynomial blowup) to obtain  $\mathbf{iO}$  for circuits where the blowup is quasi-linear. The transformation has the additional feature that the underlying obfuscation scheme is only used for (simple) circuits whose size depends only on the security parameter and input.

**The high-level idea.** The transformation relies on ideas similar to those used in section 5.2 to construct succinct  $\mathbf{iO}$  from succinct randomized encodings, which in turn go back to the bootstrapping technique of Applebaum [App14]. Concretely, we rely on plain randomized encodings [IK02, AIK06] for circuits known to have the following basic locality property: Given a circuit  $C$  with  $s$  gates and  $n$ -bit input  $x$ , computing a randomized encoding  $\widehat{C}(x; R)$  can be decomposed into  $s$  computations  $\widehat{C}_1(x; R), \dots, \widehat{C}_s(x; R)$ , each of fixed size  $\ell$  independent of the circuit size  $|C|$ . In particular, each such computation  $\widehat{C}_i(x; R)$  involves at most  $\ell$  bits of the shared randomness  $R$ .

Similarly to the transformation in section 5.2, the transformation here is based on the basic idea of obfuscating the circuit that computes the randomized encoding  $\widehat{C}(x; r)$  for any input  $x$ , while deriving the randomness  $R$ , by applying a puncturable PRF to the input  $x$ . The only difference is that, rather than obfuscating this circuit

<sup>14</sup>The perfect NIZK construction of [SW14] only satisfies nonadaptive soundness, but by a standard complexity leveraging trick it can be made to satisfy adaptive soundness. Since we assume subexponential security of the  $\mathbf{iO}$  anyway, this comes at no cost to us.

as a whole, we separately obfuscate  $s$  smaller circuits computing the corresponding  $\widehat{C}_i(x; R)$ . To make sure that deriving the randomness is also local, we associate  $r = |R|$  PRF keys  $K_1, \dots, K_r$  with each of the bits of the shared randomness  $R$ . Each one of the  $s$  obfuscated circuits only includes the PRF keys required for its local computation. The corresponding bits of randomness are again derived by applying the corresponding PRFs to the input  $x$ .

The gain is that the size of the resulting obfuscated circuit is thus  $s \cdot \text{poly}(\lambda, n)$ , as required. The proof of security relies on a variant of the probabilistic **iO** argument invoked in section 5.2, with the difference that puncturing is performed simultaneously across all  $r$  PRF keys. Accordingly, it incurs a  $2^n$  security loss in the input length  $n$  (which is polynomial when considering circuits with logarithmic-sized inputs, as is often the case in this work).

We next describe the transformation in more detail and sketch the proof of security. We start by defining the required notion of locality for the randomized encoding.

**DEFINITION 5.15** (locality of randomized encodings). *A randomized encoding  $\mathcal{RE} = (\text{REnc}, \text{Dec})$  for circuits is said to be local if*

$$\text{REnc}(1^\lambda, C, x; R) = \widehat{C}_1(x; R|_{S_1}), \dots, \widehat{C}_s(x; R|_{S_s}) ,$$

where  $s = \Theta(|C|)$ ,  $S_i \subseteq \{1, \dots, |R|\}$ ,  $R|_{S_i}$  is the restriction of  $R$  to  $S_i$ , and the following properties are satisfied:

- $\widehat{C}_i$  is a circuit of fixed size  $\ell(\lambda, |x|) = \text{poly}(\lambda, |x|)$ , independent of  $|C|$ .
- The circuits  $\{\widehat{C}_i\}$  and sets  $\{S_i\}$  can be computed from  $C$  in time  $|C| \cdot \text{poly}(\lambda, |x|)$ .
- Decoding can be done in time  $|C| \cdot \text{poly}(\lambda, |x|)$ .

Such randomized encodings can be constructed based on any one-way function [Yao86, AIK06].

**A quasi-linear obfuscator  $i\mathcal{O}^*$ .** We now describe the new obfuscator. The obfuscator relies on the following building blocks:

- A randomized encoding  $\mathcal{RE} = (\text{REnc}, \text{Dec})$  for circuits that is local and which used randomness of length at most  $r = r(|C|, \lambda)$ .
- An indistinguishability obfuscator  $i\mathcal{O}$  for circuits (with arbitrary polynomial blowup).
- A puncturable PRF ( $\text{PRF.Gen}, \text{PRF.Punc}, \text{F}$ ).

All building blocks are assumed to be  $2^{-n+\omega(\log \lambda)}$ -secure.

The obfuscator  $i\mathcal{O}^*(1^\lambda, C)$  proceeds as follows:

1. Compute the circuits  $\widehat{C}_1(\cdot; \cdot), \dots, \widehat{C}_s(\cdot; \cdot)$  and sets  $S_1, \dots, S_s$ .
2. Sample PRF keys  $K_1, \dots, K_r \leftarrow \text{PRF.Gen}(1^\lambda)$ .
3. For each  $i \in [s]$ , obfuscate using  $i\mathcal{O}$  the circuit  $\mathbb{C}_i$  that has hardwired  $\{K_j : j \in S_i\}$  and given  $x$  operates as follows:
  - Derive randomness  $R|_{S_i}$  by invoking  $\text{F}_{K_j}(x)$  for  $j \in S_i$ .
  - Output  $\widehat{C}_i(x, R|_{S_i})$ .

The circuit is further padded to be of total size  $\ell(\lambda, x)$ , where  $\ell$  is determined in the analysis.

4. Output the obfuscations  $i\mathcal{O}(\mathbb{C}_1), \dots, i\mathcal{O}(\mathbb{C}_s)$ .

To evaluate  $i\mathcal{O}^*(1^\lambda, C)$  on input  $x$ , first evaluate each  $i\mathcal{O}(\mathbb{C}_i)$  on  $x$ , obtain the randomized encoding

$$\widehat{C}(x) = \widehat{C}_1(x; R|_{S_1}), \dots, \widehat{C}_s(x; R|_{S_s}) ,$$

and decode to obtain the result  $C(x)$ .

PROPOSITION 5.16.  $i\mathcal{O}^*$  is a circuit obfuscator with quasi-linear blowup.

*Proof sketch.* The functionality of  $i\mathcal{O}^*$  follows directly from the functionality of the underlying  $\mathbf{iO}$  and the correctness of decoding. The fact that the size  $|i\mathcal{O}^*(1^\lambda, C)|$  obfuscated circuit is  $|C| \cdot \text{poly}(\lambda, |x|)$  follows from the locality of the randomized encoding. We next sketch the security.

We use a probabilistic  $\mathbf{iO}$  argument similar to the one used in section 5.2. Concretely, given two circuits  $C, C'$  of the same size and functionality, we consider  $2^n + 1$  hybrids that transition from  $i\mathcal{O}^*(1^\lambda, C)$  to  $i\mathcal{O}^*(1^\lambda, C')$ . In the  $j$ th hybrid, the  $s$  obfuscations are with respect to hybrid circuits  $\mathbb{C}_1^j, \dots, \mathbb{C}_s^j$ , where  $\mathbb{C}_i^j$  uses  $\widehat{C}_i$  for all inputs  $x < j$  and  $\widehat{C}'_i$  for all inputs  $x \geq j$ . Each two consecutive hybrids differ on only a single point,  $j$ . Similarly to section 5.2, we puncture the underlying PRFs at this point  $j$  and hardwire the values  $\widehat{C}_1(x; R|_{S_1}), \dots, \widehat{C}_s(x; R|_{S_s})$  (or  $\widehat{C}'_1(x; R|_{S_1}), \dots, \widehat{C}'_s(x; R|_{S_s})$ , respectively), using true randomness instead of pseudorandomness. Then we can rely on the security of the randomized encodings to switch between the two.

The padding parameter  $\ell(\lambda, |x|)$  is chosen to account for the above hybrids (and only induces quasi-linear blowup).  $\square$

**Acknowledgments.** We thank Boaz Barak and Guy Rothblum for their input regarding the plausibility of interactive proofs with fast verification (relevant to the plausibility of perfectly private succinct randomized encodings). We thank Daniel Wichs for discussing several aspects of [GHRW14]. We thank Stefano Tessaro for many delightful discussions at the early stage of the project. Finally, we thank the anonymous reviewers of STOC and SICOMP for their valuable comments.

## REFERENCES

- [ACC<sup>+</sup>15] P. ANANTH, Y.-C. CHEN, K.-M. CHUNG, H. LIN, AND W.-K. LIN, *Delegating RAM computations with adaptive soundness and privacy*, in Theory of Cryptography – TCC 2016, M. Hirt and A. Smith, eds., Lecture Notes in Comput. Sci. 9986, Springer, 2016, pp. 3–30.
- [AIK04] B. APPLEBAUM, Y. ISHAI, AND E. KUSHILEVITZ, *Cryptography in NC<sup>0</sup>*, in 45th Annual Symposium on Foundations of Computer Science (Rome, Italy, 2004), IEEE Computer Society Press, 2004, pp. 166–175.
- [AIK06] B. APPLEBAUM, Y. ISHAI, AND E. KUSHILEVITZ, *Computationally private randomizing polynomials and their applications*, Comput. Complexity, 15 (2006), pp. 115–162.
- [AIK10] B. APPLEBAUM, Y. ISHAI, AND E. KUSHILEVITZ, *From secrecy to soundness: Efficient verification via secure computation*, in ICALP 2010: 37th International Colloquium on Automata, Languages and Programming, Part I (Bordeaux, France, 2010), Lecture Notes in Comput. Sci. 6198, S. Abramsky, C. Gavaille, C. Kirchner, F. Meyer auf der Heide, and P. G. Spirakis, eds., Springer, 2010, pp. 152–163.
- [AIK11] B. APPLEBAUM, Y. ISHAI, AND E. KUSHILEVITZ, *How to garble arithmetic circuits*, in 52nd Annual Symposium on Foundations of Computer Science (Palm Springs, CA, 2011), R. Ostrovsky, ed., IEEE Computer Society Press, 2011, pp. 120–129.
- [AJL<sup>+</sup>12] G. ASHAROV, A. JAIN, A. LÓPEZ-ALT, E. TROMER, V. VAIKUNTANATHAN, AND D. WICHS, *Multiparty computation with low communication, computation and interaction via threshold FHE*, in Advances in Cryptology – EUROCRYPT 2012 (Cambridge, UK, 2012), Lecture Notes in Comput. Sci. 7237, D. Pointcheval and T. Johansson, eds., Springer, 2012, pp. 483–501.
- [AJS15] P. ANANTH, A. JAIN, AND A. SAHAI, *Indistinguishability Obfuscation with Constant Size Overhead*, IACR Cryptology ePrint Archive, 2015:1023, 2015.
- [AJS17] P. ANANTH, A. JAIN, AND A. SAHAI, *Indistinguishability Obfuscation for Turing Machines: Constant Overhead and Amortization*, in Advances in Cryptology – CRYPTO 2017, J. Katz and H. Shacham, eds., Lecture Notes in Comput. Sci. 10402, Springer, 2017, pp. 252–279.
- [App11a] B. APPLEBAUM, *Key-dependent message security: Generic amplification and completeness*, in Advances in Cryptology – EUROCRYPT 2011 (Tallinn, Estonia, 2011),

- Lecture Notes in Comput. Sci. 6632, K. G. Paterson, ed., Springer, 2011, pp. 527–546.
- [App11b] B. APPLEBAUM, *Randomly encoding functions: A new cryptographic paradigm (invited talk)*, in Proceedings of the 5th International Conference on Information Theoretic Security, ICITS 2011 (Amsterdam, The Netherlands), 2011, pp. 25–31.
- [App14] B. APPLEBAUM, *Bootstrapping obfuscators via fast pseudorandom functions*, in Advances in Cryptology - ASIACRYPT 2014 - Proceedings of the 20th International Conference on the Theory and Application of Cryptology and Information Security (Kaoshiung, Taiwan, R.O.C.), 2014, pp. 162–172.
- [BCCT12] N. BITANSKY, R. CANETTI, A. CHIESA, AND E. TROMER, *From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again*, in ITCS 2012: 3rd Innovations in Theoretical Computer Science (Cambridge, MA, 2012), S. Goldwasser, ed., Association for Computing Machinery, 2012, pp. 326–349.
- [BCCT13] N. BITANSKY, R. CANETTI, A. CHIESA, AND E. TROMER, *Recursive composition and bootstrapping for SNARKS and proof-carrying data*, in 45th Annual ACM Symposium on Theory of Computing (Palo Alto, CA, 2013), D. Boneh, T. Roughgarden, and J. Feigenbaum, eds., ACM Press, 2013, pp. 111–120.
- [BCC<sup>+</sup>17] N. BITANSKY, R. CANETTI, A. CHIESA, S. GOLDWASSER, H. LIN, A. RUBINSTEIN, AND E. TROMER, *The hunting of the SNARK*, J. Cryptology, 30 (2017), pp. 989–1066.
- [BCP14] E. BOYLE, K.-M. CHUNG, AND R. PASS, *On extractability obfuscation*, in TCC 2014: 11th Theory of Cryptography Conference (San Diego, CA, 2014), Lecture Notes in Comput. Sci. 8349, Y. Lindell, ed., Springer, 2014, pp. 52–73.
- [BGI<sup>+</sup>01] B. BARAK, O. GOLDREICH, R. IMPAGLIAZZO, S. RUDICH, A. SAHAI, S. VADHAN, AND K. YANG, *On the (im)possibility of obfuscating programs*, in Advances in Cryptology CRYPTO 2001, Springer, 2001, pp. 1–18.
- [BGI14] E. BOYLE, S. GOLDWASSER, AND I. IVAN, *Functional signatures and pseudorandom functions*, in PKC, 2014, pp. 501–519.
- [BGL<sup>+</sup>15] N. BITANSKY, S. GARG, H. LIN, R. PASS, AND S. TELANG, *Succinct randomized encodings and their applications*, in Proceedings of the 47th Annual ACM on Symposium on Theory of Computing, STOC 2015 (Portland, OR), ACM, 2015, pp. 439–448.
- [BHHI10] B. BARAK, I. HAITNER, D. HOFHEINZ, AND Y. ISHAI, *Bounded key-dependent message security*, in Advances in Cryptology – EUROCRYPT 2010 (French Riviera, 2010), Lecture Notes in Comput. Sci. 6110, H. Gilbert, ed., Springer, 2010, pp. 423–444.
- [BHK13] M. BELLARE, V. T. HOANG, AND S. KEELVEEDHI, *Instantiating random oracles via UCEs*, in Advances in Cryptology – CRYPTO 2013, Part II (Santa Barbara, CA, 2013), Lecture Notes in Comput. Sci. 8043, R. Canetti and J. A. Garay, eds., Springer, pp. 398–415.
- [BHR12a] M. BELLARE, V. T. HOANG, AND P. ROGAWAY, *Adaptively secure garbling with applications to one-time programs and secure outsourcing*, in Advances in Cryptology – ASIACRYPT 2012 (Beijing, China, 2012), Lecture Notes in Comput. Sci. 7658, X. Wang and K. Sako, eds., Springer, 2012, pp. 134–153.
- [BHR12b] M. BELLARE, V. T. HOANG, AND P. ROGAWAY, *Foundations of garbled circuits*, in The ACM Conference on Computer and Communications Security, CCS’12 (Raleigh, NC), 2012, pp. 784–796.
- [BLMR13] D. BONEH, K. LEWI, H. W. MONTGOMERY, AND A. RAGHUNATHAN, *Key homomorphic PRFs and their applications*, in Advances in Cryptology – CRYPTO 2013, Part I (Santa Barbara, CA, 2013), Lecture Notes in Comput. Sci. 8042, R. Canetti and J. A. Garay, eds., Springer, 2013, pp. 410–428.
- [BW13] D. BONEH AND B. WATERS, *Constrained pseudorandom functions and their applications*, in ASIACRYPT (2), 2013, pp. 280–300.
- [CCC<sup>+</sup>16] Y.-C. CHEN, S. S. M. CHOW, K.-M. CHUNG, R. W. F. LAI, W.-K. LIN, AND H.-S. ZHOU, *Cryptography for Parallel RAM from Indistinguishability Obfuscation*, IACR Cryptology ePrint Archive, Report 2015/406, 2015, <https://eprint.iacr.org/2015/406>.
- [CCHR15] R. CANETTI, Y. CHEN, J. HOLMGREN, AND M. RAYKOVA, *Succinct Adaptive Garbled RAM*, Cryptology ePrint Archive, Report 2015/1074, 2015, <https://eprint.iacr.org/2015/1074>.
- [CH16] R. CANETTI AND J. HOLMGREN, *Fully succinct garbled RAM*, in Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science (Cambridge, MA, 2016), MIT, 2016, pp. 169–178, <http://doi.acm.org/10.1145/2840728.2840765>.
- [CHJV15] R. CANETTI, J. HOLMGREN, A. JAIN, AND V. VAIKUNTANATHAN, *Succinct garbling and*

- indistinguishability obfuscation for RAM programs*, in Proceedings of the 47th Annual ACM Symposium on Theory of Computing, STOC 2015 (Portland, OR), 2015, pp. 429–437.
- [CIJ+13] A. DE CARO, V. IOVINO, A. JAIN, A. O’NEILL, O. PANETH, AND G. PERSIANO, *On the achievability of simulation-based security for functional encryption*, in Advances in Cryptology - CRYPTO 2013, Part II (Santa Barbara, CA, 2013), Springer, 2013, pp. 519–535.
- [CLP13] K.-M. CHUNG, H. LIN, AND R. PASS, *Constant-round concurrent zero knowledge from  $p$ -certificates*, in 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013 (Berkeley, CA), 2013, pp. 50–59.
- [CLTV15] R. CANETTI, H. LIN, S. TESSARO, AND V. VAIKUNTANATHAN, *Obfuscation of probabilistic circuits and applications*, in Theory of Cryptography - 12th Theory of Cryptography Conference, Part II, TCC 2015 (Warsaw, Poland), 2015, pp. 468–497.
- [CP13] K.-M. CHUNG AND R. PASS, *A Simple ORAM*, IACR Cryptology ePrint Archive, 2013:243, 2013.
- [DFH12] I. DAMGÅRD, S. FAUST, AND C. HAZAY, *Secure two-party computation with low communication*, in TCC 2012: 9th Theory of Cryptography Conference (Sicily, Italy, 2012), Lecture Notes in Comput. Sci. 7194, R. Cramer, ed., Springer, 2012, pp. 54–74.
- [DHRW16] Y. DODIS, S. HALEVI, R. D. ROTHBLUM, AND D. WICHS, *Spooky encryption and its applications*, in Advances in Cryptology - CRYPTO 2016, Part III (Santa Barbara, CA), Springer, 2016, pp. 93–122.
- [Gen09] C. GENTRY, *Fully homomorphic encryption using ideal lattices*, in 41st Annual ACM Symposium on Theory of Computing (Bethesda, MD, 2009), M. Mitzenmacher, ed., ACM Press, pp. 169–178.
- [GGH+13] S. GARG, C. GENTRY, S. HALEVI, M. RAYKOVA, A. SAHAI, AND B. WATERS, *Candidate indistinguishability obfuscation and functional encryption for all circuits*, in 54th Annual Symposium on Foundations of Computer Science (Berkeley, CA, 2013), IEEE Computer Society Press, 2013, pp. 40–49.
- [GGHR14] S. GARG, C. GENTRY, S. HALEVI, AND M. RAYKOVA, *Two-round secure MPC from indistinguishability obfuscation*, in TCC 2014: 11th Theory of Cryptography Conference (San Diego, CA, 2014), Lecture Notes in Comput. Sci. 8349, Y. Lindell, ed., Springer, 2014, pp. 74–94.
- [GGHW13] S. GARG, C. GENTRY, S. HALEVI, AND D. WICHS, *On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input*, in Advances in Cryptology - CRYPTO 2014 (Santa Barbara, CA, 2014), Springer, 2014, pp. 518–535.
- [GGHZ16] S. GARG, C. GENTRY, S. HALEVI, AND M. ZHANDRY, *Functional encryption without obfuscation*, in Theory of Cryptography - 13th Theory of Cryptography Conference, Part II, Lecture Notes in Comput. Sci. 9563, Springer, 2016, pp. 480–511.
- [GGM86] O. GOLDBREICH, S. GOLDWASSER, AND S. MICALI, *How to construct random functions*, J. ACM, 33 (1986), pp. 792–807.
- [GGP10] R. GENNARO, C. GENTRY, AND B. PARNO, *Non-interactive verifiable computing: Outsourcing computation to untrusted workers*, in Advances in Cryptology – CRYPTO 2010 (Santa Barbara, CA, 2010), Lecture Notes in Comput. Sci. 6223, T. Rabin, ed., Springer, 2010, pp. 465–482.
- [GHL+14] C. GENTRY, S. HALEVI, S. LU, R. OSTROVSKY, M. RAYKOVA, AND D. WICHS, *Garbled RAM revisited*, in Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques (Copenhagen, Denmark), 2014, pp. 405–422.
- [GHRW14] C. GENTRY, S. HALEVI, M. RAYKOVA, AND D. WICHS, *Outsourcing private RAM computation*, in 55th Annual Symposium on Foundations of Computer Science, 2014.
- [GKP+13] S. GOLDWASSER, Y. T. KALAI, R. A. POPA, V. VAIKUNTANATHAN, AND N. ZELDOVICH, *Reusable garbled circuits and succinct functional encryption*, in 45th Annual ACM Symposium on Theory of Computing (Palo Alto, CA, 2013), D. Boneh, T. Roughgarden, and J. Feigenbaum, eds., ACM Press, 2013, pp. 555–564.
- [GLOS15] S. GARG, S. LU, R. OSTROVSKY, AND A. SCAFURO, *Garbled RAM from one-way functions*, in 47th Annual ACM Symposium on Theory of Computing, ACM Press, 2015.
- [HW15] P. HUBÁČEK AND D. WICHS, *On the communication complexity of secure function evaluation with long output*, in ITCS, ACM Press, 2015, pp. 163–172.



- [IK00] Y. ISHAI AND E. KUSHILEVITZ, *Randomizing polynomials: A new representation with applications to round-efficient secure computation*, in 41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000), IEEE Computer Society Press, 2000, pp. 294–304.
- [IK02] Y. ISHAI AND E. KUSHILEVITZ, *Perfect constant-round secure computation via perfect randomizing polynomials*, in Automata, Languages and Programming, 29th International Colloquium, ICALP 2002 (Malaga, Spain), 2002, pp. 244–256.
- [KLW15] V. KOPPULA, A. B. LEWKO, AND B. WATERS, *Indistinguishability obfuscation for Turing machines with unbounded memory*, in 47th Annual ACM Symposium on Theory of Computing, ACM Press, 2015.
- [KPTZ13] A. KIAYIAS, S. PAPAPOPOULOS, N. TRIANOPOULOS, AND T. ZACHARIAS, *Delegatable pseudorandom functions and applications*, in CCS, 2013, pp. 669–684.
- [KRR14] Y. T. KALAI, R. RAZ, AND R. D. ROTHBLUM, *How to delegate computations: The power of no-signaling proofs*, in 46th Annual ACM Symposium on Theory of Computing, D. B. Shmoys, ed., ACM Press, 2014, pp. 485–494.
- [Kv99] A. KLIVANS AND D. VAN MELKEBEEK, *Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses*, in 31st Annual ACM Symposium on Theory of Computing (Atlanta, GA, 1999), ACM Press, 1999, pp. 659–667.
- [LO13] S. LU AND R. OSTROVSKY, *How to garble RAM programs*, in Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques (Athens, Greece), 2013, pp. 719–734.
- [LTV12] A. LÓPEZ-ALT, E. TROMER, AND V. VAIKUNTANATHAN, *On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption*, in 44th Annual ACM Symposium on Theory of Computing, H. J. Karloff and T. Pitassi, eds., ACM Press, 2012, pp. 1219–1234.
- [Mic00] S. MICALI, *Computationally sound proofs*, SIAM J. Comput., 30 (2000), pp. 1253–1298, <https://doi.org/10.1137/S0097539795284959>.
- [MV99] P. B. MILTERSEN AND N. V. VINODCHANDRAN, *Derandomizing Arthur-Merlin games using hitting sets*, in 40th Annual Symposium on Foundations of Computer Science (New York, 1999), IEEE Computer Society Press, 1999, pp. 71–80.
- [MW16] P. MUKHERJEE AND D. WICHS, *Two round multiparty computation via multi-key FHE*, in Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Part II (Vienna, Austria), 2016, pp. 735–763.
- [Nao03] M. NAOR, *On cryptographic assumptions and challenges (invited talk)*, in Advances in Cryptology - CRYPTO 2003 (Santa Barbara, CA, 2003), Lecture Notes in Comput. Sci. 2729, D. Boneh, ed., Springer, 2003, pp. 96–109.
- [PF79] N. PIPPENGER AND M. J. FISCHER, *Relations among complexity measures*, J. ACM, 26 (1979), pp. 361–381.
- [PRV12] B. PARNO, M. RAYKOVA, AND V. VAIKUNTANATHAN, *How to delegate and verify in public: Verifiable computation from attribute-based encryption*, in TCC 2012: 9th Theory of Cryptography Conference (Sicily, Italy, 2012), Lecture Notes in Comput. Sci. 7194, R. Cramer, ed., Springer, 2012, pp. 422–439.
- [Rog91] P. ROGAWAY, *The Round Complexity of Secure Protocols*, Ph.D. thesis, Massachusetts Institute of Technology, 1991.
- [SCSL11] E. SHI, T.-H. H. CHAN, E. STEFANOV, AND M. LI, *Oblivious RAM with  $O((\log N)^3)$  worst-case cost*, in ASIACRYPT, 2011, pp. 197–214.
- [SW14] A. SAHAI AND B. WATERS, *How to use indistinguishability obfuscation: Deniable encryption, and more*, in Proceedings of the 46th ACM Symposium on Theory of Computing, STOC 2014, D. B. Shmoys, ed., ACM Press, 2014, pp. 475–484.
- [Wat14] B. WATERS, *A Punctured Programming Approach to Adaptively Secure Functional Encryption*, Cryptology ePrint Archive, Report 2014/588, 2014.
- [Yao86] A. C.-C. YAO, *How to generate and exchange secrets (extended abstract)*, in 27th Annual Symposium on Foundations of Computer Science (Toronto, Ontario, Canada, 1986), IEEE Computer Society Press, 1986, pp. 162–167.