

# Limits of Extractability Assumptions with Distributional Auxiliary Input

Elette Boyle\*  
Technion Israel  
eboyle@alum.mit.edu

Rafael Pass†  
Cornell University  
rafael@cs.cornell.edu

October 21, 2014

## Abstract

Extractability, or “knowledge,” assumptions have recently gained popularity in the cryptographic community, leading to the study of primitives such as extractable one-way functions, extractable hash functions, succinct non-interactive arguments of knowledge (SNARKs), and differing-inputs obfuscation ( $di\mathcal{O}$ ), and spurring the development of a wide spectrum of new applications relying on these primitives. For most of these applications, it is required that the extractability assumption holds even in the presence of attackers receiving some *auxiliary information* that is sampled from some *fixed* efficiently computable distribution  $\mathcal{Z}$ .

We show that, assuming the existence of collision-resistant hash functions, there exist efficient distributions  $\mathcal{Z}, \mathcal{D}$  such that either

- extractable one-way functions w.r.t. auxiliary input  $\mathcal{Z}$  do not exist, or
- $di\mathcal{O}$  for a distribution of Turing machines and auxiliary input specified by  $\mathcal{D}$  does not exist.

A corollary of this result shows that assuming existence of fully homomorphic encryption with decryption in  $NC^1$ , there exist efficient distributions  $\mathcal{Z}, \mathcal{D}$  such that either

- SNARKs for NP w.r.t. auxiliary input  $\mathcal{Z}$  do not exist, or
- $di\mathcal{O}$  for a distribution of  $NC^1$  circuits and aux input specified by  $\mathcal{D}$  does not exist.

To achieve our results, we develop a “succinct punctured program” technique, mirroring the powerful punctured program technique of Sahai and Waters (STOC’14), and present several other applications of this new technique.

We additionally demonstrate that  $di\mathcal{O}$  w.r.t. any distribution  $\mathcal{D}$  of programs and bounded-length auxiliary input is directly implied by any obfuscator that satisfies the weaker indistinguishability obfuscation ( $i\mathcal{O}$ ) security notion and  $di\mathcal{O}$  for a slightly modified distribution  $\mathcal{D}'$  of programs (of slightly greater size) and no auxiliary input. As a consequence, we directly obtain negative results for  $di\mathcal{O}$  *in the absence of auxiliary input*.

---

\*This work was primarily completed while the first author was a postdoc at Cornell University, supported in part by AFOSR YIP Award FA9550-10-1-0093. This research has received funding from the European Union’s Tenth Framework Programme (FP10/ 2010-2016) under grant agreement no. 259426 ERC-CaC.

†Pass is supported in part by a Alfred P. Sloan Fellowship, 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. 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 US Government.

# 1 Introduction

**Extractability Assumptions.** Extractability, or “knowledge,” assumptions (such as the “knowledge-of-exponent” assumption), have recently gained in popularity, leading to the study of primitives such as extractable one-way functions, extractable hash-functions, SNARKs (succinct non-interactive arguments of knowledge), and differing-inputs obfuscation:

- **Extractable OWF:** An extractable family of one-way (resp. collision-resistant) functions [Dam91, HT98, CD09], is a family of one-way (resp. collision-resistant) functions  $\{f_i\}$  such that any attacker who outputs an element  $y$  in the range of a randomly chosen function  $f_i$  given the index  $i$  must “know” a pre-image  $x$  of  $y$  (i.e.,  $f_i(x) = y$ ). This is formalized by requiring for every adversary  $\mathcal{A}$ , the existence of an “extractor”  $\mathcal{E}$  that (with overwhelming probability) given the view of  $\mathcal{A}$  outputs a pre-image  $x$  whenever  $\mathcal{A}$  outputs an element  $y$  in the range of the function.

For example, the “knowledge-of-exponent” assumption of Damgard [Dam91] stipulates the existence of a particular such extractable one-way function.

- **SNARKs:** Succinct non-interactive arguments of knowledge (SNARKs) [Mic94, Val08, BCCT12] are communication-efficient (i.e., “short” or “succinct”) arguments for NP with the property that if a prover generates an accepting (short) proof, it must “know” a corresponding (potentially long) witness for the statement proved, and this witness can be efficiently “extracted” out from the prover.
- **Differing-inputs obfuscation:** [BGI<sup>+</sup>12, BCP14, ABG<sup>+</sup>13] A differing-inputs obfuscator  $\mathcal{O}$  for a class of programs  $\mathcal{C}$  is an efficient procedure which ensures if any efficient attacker  $\mathcal{A}$  can distinguish obfuscations  $\mathcal{O}(C_1)$  and  $\mathcal{O}(C_2)$  of any two programs  $C_1, C_2$  in the class  $\mathcal{C}$ , then it must “know” an input  $x$  such that  $C_1(x) \neq C_2(x)$ , and this input can be efficiently “extracted” (as above) from  $\mathcal{A}$ .

The above primitives have proven extremely useful in constructing cryptographic tools for which instantiations under complexity-theoretic hardness assumptions are not known (e.g., [HT98, BCCT12, GLR11, DFH12, BCP14, ABG<sup>+</sup>13]).

**Extraction with (Distribution-Specific) Auxiliary Input.** In all of these applications, we require a notion of an *auxiliary-input* extractable one-way function [HT98, CD09], where both the attacker and the extractor may receive an auxiliary input. The strongest formulation requires extractability in the presence of an *arbitrary* auxiliary input. Yet, as informally discussed already in the original work by Hada and Tanaka [HT98], extractability w.r.t. an arbitrary auxiliary input is an “overly strong” (or in the language of [HT98], “unreasonable”) assumption. Indeed, a recent result of Bitansky, Canetti, Rosen and Paneth [BCPR14] (formalizing earlier intuitions from [HT98, BCCT12]) demonstrates that assuming the existence of indistinguishability obfuscators for the class of polynomial-size circuits<sup>1</sup> there cannot exist auxiliary-input extractable one-way functions that remain secure for an arbitrary auxiliary input.

However, for most of the above applications, we actually do not require extractability to hold w.r.t. an arbitrary auxiliary input. Rather, as proposed by Bitansky et al [BCCT12, BCCT13], it

---

<sup>1</sup>The notion of indistinguishability obfuscation [BGI<sup>+</sup>12] requires that obfuscations  $\mathcal{O}(C_1)$  and  $\mathcal{O}(C_2)$  of any two *equivalent* circuits  $C_1$  and  $C_2$  (i.e., whose outputs agree on all inputs) from some class  $\mathcal{C}$  are computationally indistinguishable. A candidate construction for general-purpose indistinguishability obfuscation was recently given by Garg et al [GGH<sup>+</sup>13].

often suffices to consider extractability with respect to specific distributions  $\mathcal{Z}$  of auxiliary input.<sup>2</sup> More precisely, it would suffice to show that for every desired output length  $\ell(\cdot)$  and distribution  $\mathcal{Z}$  there exists a function family  $\mathcal{F}_{\mathcal{Z}}$  (which, in particular, may be tailored for  $\mathcal{Z}$ ) such that  $\mathcal{F}_{\mathcal{Z}}$  is a family of extractable one-way (or collision-resistant) functions  $\{0, 1\}^k \rightarrow \{0, 1\}^{\ell(k)}$  with respect to  $\mathcal{Z}$ . In fact, for some of these results (e.g., [BCCT12, BCCT13]), it suffices to just assume that extraction works for just for the *uniform* distribution.

In contrast, the result of [BCPR14] can be interpreted as saying that (assuming  $i\mathcal{O}$ ), there do not exist extractable one-way functions with respect to *every* distribution of auxiliary input: That is, for every candidate extractable one-way function family  $\mathcal{F}$ , there exists *some* distribution  $\mathcal{Z}_{\mathcal{F}}$  of auxiliary input that breaks it.

**Our Results.** In this paper, we show limitations of extractability primitives with respect to *distribution-specific* auxiliary input (assuming the existence of collision-resistant hash functions). Our main result shows a conflict between differing-inputs obfuscation for Turing machines [BCP14, ABG<sup>+</sup>13] and extractable one-way functions.

**Theorem 1.1** (Main Theorem – Informal). *Assume the existence of collision-resistant hash functions. Then for every polynomial  $\ell$ , there exist efficiently computable distributions  $\mathcal{Z}, \mathcal{D}$  such that one of the following two primitives does not exist:*

- *extractable one-way functions  $\{0, 1\}^k \rightarrow \{0, 1\}^{\ell(k)}$  w.r.t. auxiliary input from  $\mathcal{Z}$ .*
- *differing-inputs obfuscation for distribution  $\mathcal{D}$  of Turing machines and auxiliary inputs.*

By combining our main theorem with results from [BCCT12] and [BCP14], we obtain the following corollary:

**Theorem 1.2** (Informal). *Assume the existence of fully homomorphic encryption with decryption in  $NC^1$ .<sup>3</sup> Then there exist efficiently computable distributions  $\mathcal{Z}, \mathcal{D}$  such that one of the following two primitives does not exist:*

- *SNARKs w.r.t. auxiliary input from  $\mathcal{Z}$ .*
- *differing-inputs obfuscation for distribution  $\mathcal{D}$  of  $NC^1$  circuits and auxiliary inputs.*

To prove our results, we develop a new proof technique, which we refer to as the “succinct punctured program” technique, extending the “punctured program” paradigm of Sahai and Waters [SW14]; see Section 1.1 for more details. This technique has several other interesting applications, as we discuss in Section 1.3.

As a final contribution, we demonstrate that *even in the absence of auxiliary input*, care must be taken when making use of extractability assumptions. Specifically, we show that differing-inputs obfuscation ( $di\mathcal{O}$ ) for any distribution  $\mathcal{D}$  of programs and bounded-length auxiliary inputs, is directly implied by any obfuscator that satisfies a weaker indistinguishability obfuscation ( $i\mathcal{O}$ ) security notion (which is not an extractability assumption) and  $di\mathcal{O}$  security for a related distribution  $\mathcal{D}'$  of programs (of slightly greater size) which does *not contain auxiliary input*. Thus, negative results ruling out existence of  $di\mathcal{O}$  with bounded-length auxiliary input directly imply negative results for  $di\mathcal{O}$  in a setting without auxiliary input.

---

<sup>2</sup>As far as we know, the only exceptions are in the context of zero-knowledge simulation, where the extractor is used in the simulation (as opposed to being used as part of a reduction), and we require simulation w.r.t. arbitrary auxiliary inputs. Nevertheless, as pointed out in the works on zero-knowledge [HT98, GS12], to achieve “plain” zero-knowledge [GMR89, BLV06] (where the verifier does not receive any auxiliary input), weaker “bounded” auxiliary input assumptions suffice.

<sup>3</sup>As is the case for nearly all existing FHE constructions (e.g., [GSW13, BV14]).

**Theorem 1.3** (Informal). *Let  $\mathcal{D}$  be a distribution over pairs of programs and  $\ell$ -bounded auxiliary input information  $\mathcal{P} \times \mathcal{P} \times \{0, 1\}^\ell$ . There exists  $di\mathcal{O}$  with respect to  $\mathcal{D}$  if there exists an obfuscator satisfying  $i\mathcal{O}$  in addition to  $di\mathcal{O}$  with respect to a modified distribution  $\mathcal{D}'$  over  $\mathcal{P}' \times \mathcal{P}'$  for slightly enriched program class  $\mathcal{P}'$ , and no auxiliary input.*

As a consequence,  $di\mathcal{O}$  for the distributions  $\mathcal{D}$  in our theorems above can be directly replaced by  $i\mathcal{O}$  plus  $di\mathcal{O}$  for modified distributions  $\mathcal{D}'$  of programs that remain within the respective classes (i.e.,  $TM$  or  $NC^1$ ), but no longer include auxiliary input.

Our transformation further applies to a recent result of Garg *et al.* [GGHW14], which shows that based on a new assumption (pertaining to special-purpose obfuscation of Turing machines) general-purpose  $di\mathcal{O}$  w.r.t. auxiliary input cannot exist, by constructing a distribution over circuits and bounded-length auxiliary inputs for which no obfuscator can be  $di\mathcal{O}$ -secure. Our resulting conclusion is that, assuming such special-purpose obfuscation exists, then general-purpose  $di\mathcal{O}$  cannot exist, *even in the absence of auxiliary input.*

**Interpretation of Our Results.** Our results suggest that one must take care when making extractability assumptions, even in the presence of specific distributions of auxiliary inputs, and in certain cases even in the absence of auxiliary input. In particular, we must develop a way to distinguish “good” distributions of instances and auxiliary inputs (for which extractability assumptions may make sense) and “bad” ones (for which extractability assumptions are unlikely to hold). As mentioned above, for some applications of extractability assumptions, it in fact suffices to consider a particularly simple distribution of auxiliary inputs—namely the *uniform* distribution.<sup>4</sup> We emphasize that our results do not present any limitations of extractable one-way functions in the presence of uniform auxiliary input, and as such, this still seems like a plausible assumption at this point.

**Comparison to [GGHW14].** An interesting subsequent<sup>5</sup> work of Garg *et al.* [GGHW13, GGHW14] contains a related study of differing-inputs obfuscation. In [GGHW14], the authors propose a new “special-purpose” circuit obfuscation assumption, and demonstrate based on this assumption an auxiliary input distribution (whose size grows with the desired circuit size of circuits to be obfuscated) for which general-purpose  $di\mathcal{O}$  cannot exist. Using similar techniques of hashing and obfuscating Turing machines as in the current work, they further conclude that if the new obfuscation assumption holds also for *Turing machines*, then the “bad” auxiliary input distribution can have bounded length (irrespective of the circuit size).

Garg *et al.* [GGHW14] show the “special-purpose” obfuscation assumption is a falsifiable assumption (in the sense of [Nao03]) and is implied by virtual black-box obfuscation for the relevant restricted class of programs, but plausibility of the notion in relation to other primitives is otherwise unknown. In contrast, our results provide a direct relation between existing, studied topics (namely,  $di\mathcal{O}$ , EOWFs, and SNARKs). Even in the case that the special-purpose obfuscation assumption *does* hold, our primary results provide conclusions for  $di\mathcal{O}$  with natural “lightweight” auxiliary input distributions (e.g., descriptions of a collision-resistant hash function), whereas the auxiliary input considered in [GGHW14] is somewhat unlikely to occur naturally, itself consisting of an obfuscated program.

And, utilizing our final observation (which occurred subsequent to [GGHW14]), we show that based on their same special-purpose obfuscation assumption for Turing machines, we can in fact rule out general-purpose  $di\mathcal{O}$  for circuits even *in the absence of auxiliary input.*

---

<sup>4</sup>Note that this is not the case for all applications; for instance [HT98, GKP<sup>+</sup>13, BGI14, GS12] require considering more complicated distributions.

<sup>5</sup>A version of our paper has been on ePrint since October 2013 [BP13].

## 1.1 Proof Techniques

To explain our techniques, let us first explain earlier arguments against the plausibility of extractable one-way functions with auxiliary input. For simplicity of notation, we focus on extractable one-way function over  $\{0, 1\}^k \rightarrow \{0, 1\}^k$  (as opposed to over  $\{0, 1\}^k \rightarrow \{0, 1\}^{\ell(k)}$  for some polynomial  $\ell$ ), but emphasize that the approach described directly extends to the more general setting.

**Early Intuitions.** As mentioned above, already the original work of Hada and Tanaka [HT98], which introduced auxiliary input extractable one-way functions (EOWFs) (for the specific case of exponentiation), argued the “unreasonableness” of such functions, reasoning informally that the auxiliary input could contain a program that evaluates the function, and thus a corresponding extractor must be able to “reverse-engineer” *any* such program. Bitansky et al [BCCT12] made this idea more explicit: Given some candidate EOWF family  $\mathcal{F}$ , consider the distribution  $\mathcal{Z}_{\mathcal{F}}$  over auxiliary input formed by “obfuscating” a program  $\Pi^s(\cdot)$  for uniformly chosen  $s$ , where  $\Pi^s(\cdot)$  takes as input a function index  $e$  from the alleged EOWF family  $\mathcal{F} = \{f_i\}$ , applies a pseudorandom function (PRF) with hardcoded seed  $s$  to the index  $i$ , and then outputs the evaluation  $f_i(\text{PRF}_s(i))$ . Now, consider an attacker  $\mathcal{A}$  who, given an index  $i$ , simply runs the obfuscated program to obtain a “random” point in the range of  $f_i$ . If it were possible to obfuscate  $\Pi^s$  in a “virtual black-box (VBB)” way (as in [BGI<sup>+</sup>12]), then it easily follows that any extractor  $\mathcal{E}$  for this particular attacker  $\mathcal{A}$  can invert  $f_i$ . Intuitively, the VBB-obfuscated program hides the PRF seed  $s$  (revealing, in essence, only black-box access to  $\Pi^s$ ), and so if  $\mathcal{E}$  can successfully invert  $f_i$  on  $\mathcal{A}$ ’s output  $f_i(\text{PRF}_s(i))$  on a pseudorandom input  $\text{PRF}_s(i)$ , he must also be able to invert for a *truly* random input. Formally, given an index  $i$  and a random point  $y$  in the image of  $f_i$ , we can “program” the output of  $\Pi^s(i)$  to simply be  $y$ , and thus  $\mathcal{E}$  will be forced to invert  $y$ .

The problem with this argument is that (as shown by Barak et al [BGI<sup>+</sup>12]), for large classes of functions VBB program obfuscation simply does not exist.

**The Work of [BCPR14] and the “Punctured Program” Paradigm of [SW14].** Intriguingly, Bitansky, Canetti, Rosen and Paneth [BCPR14] show that by using a particular PRF and instead relying on indistinguishability obfuscation, the above argument still applies! To do so, they rely on the powerful “punctured-program” paradigm of Sahai and Waters [SW14] (and the closely related work of Hohenberger, Sahai and Waters [HSW14] on “instantiating random oracles”). Roughly speaking, the punctured program paradigm shows that if we use indistinguishability obfuscation to obfuscate a (function of) a special kind of “puncturable” PRF<sup>6</sup> [BW13, BGI14, KPTZ13], we can still “program” the output of the program on *one* input (which was used in [SW14, HSW14] to show various applications of indistinguishability obfuscation). Bitansky et al. [BCPR14] show that by using this approach, then from any alleged extractor  $\mathcal{E}$  we can construct a one-way function inverter  $Inv$  by “programming” the output of the program  $\Pi^s$  at the input  $i$  with the challenge value  $y$ . More explicitly, mirroring [SW14, HSW14], they consider a hybrid experiment where  $\mathcal{E}$  is executed with fake (but indistinguishable) auxiliary input, formed by obfuscating a “punctured” variant  $\Pi_{i,y}^s$  of the program  $\Pi^s$  that contains an  $i$ -punctured PRF seed  $s^*$  (enabling evaluation of  $\text{PRF}_s(j)$  for any  $j \neq i$ ) and directly outputs the hardcoded value  $y := f_i(\text{PRF}_s(i))$  on input  $i$ : indistinguishability of this auxiliary input follows by the security of indistinguishability obfuscation since the programs  $\Pi_{i,y}^s$  and  $\Pi^s$  are equivalent when  $y = f_i(\text{PRF}_s(i)) = \Pi^s(i)$ . In a second hybrid experiment, the “correct”

<sup>6</sup>That is, a PRF where we can surgically remove one point in the domain of the PRF, keeping the rest of the PRF intact, and yet, even if we are given the seed of the punctured PRF, the value of the original PRF on the surgically removed point remains computationally indistinguishable from random.

hardcoded value  $y$  is replaced by a *random* evaluation  $f_i(u)$  for uniform  $u$ ; here, indistinguishability of the auxiliary inputs follows directly by the security of the punctured PRF. Finally, by indistinguishability of the three distributions of auxiliary input in the three experiments, it must be that  $\mathcal{E}$  can extract an inverse to  $y$  with non-negligible probability in each hybrid; but, in the final experiment this implies the ability to invert a random evaluation, breaking one-wayness of the EOWF.

**The Problem: Dependence on  $\mathcal{F}$ .** Note that in the above approach, the auxiliary input distribution is selected *as a function* of the family  $\mathcal{F} = \{f_j\}$  of (alleged) extractable one-way functions. Indeed, the obfuscated program  $\Pi^s$  must be able to evaluate  $f_j$  given  $j$ . One may attempt to mitigate this situation by instead obfuscating a universal circuit that takes as input both  $\mathcal{F}$  and the index  $j$ , and appropriately evaluates  $f_j$ . But here still the *size* of the universal circuit must be greater than the running time of  $f_j$ , and thus such an auxiliary input distribution would only rule out EOWFs with a-priori bounded running time. This does not suffice for what we aim to achieve: in particular, it still leaves open the possibility that for every distribution of auxiliary inputs, there may exist a family of extractable one-way functions that remains secure for that particular auxiliary input distribution (although the running time of the extractable one-way function needs to be greater than the length of the auxiliary input).

**A First Idea: Using Turing Machine Obfuscators.** At first sight, it would appear this problem could be solved if we could obfuscate *Turing machines*. Namely, by obfuscating a universal Turing machine in the place of a universal circuit in the construction above, the resulting program  $\Pi^s$  would depend only on the size of the PRF seed  $s$ , and not on the runtime of  $f_j \in F$ . Indeed, recently [BCP14] showed candidate constructions of obfuscators for Turing machines, based on the existence of differing-inputs obfuscators for  $NC^1$ , fully homomorphic encryption with decryption in  $NC^1$ , and succinct non-interactive arguments of knowledge (SNARKs).

But there is a catch. To rely on the punctured program paradigm, we must be able to obfuscate the program  $\Pi^s$  in such a way that the result is indistinguishable from an obfuscation of a related “punctured” program  $\Pi_{i,y}^s$ ; in particular, the *size* of the obfuscation must be at least as large as  $|\Pi_{i,y}^s|$ . Whereas the size of  $\Pi^s$  is now bounded by a polynomial in the size of the PRF seed  $s$ , the description of this punctured program must specify a punctured input  $i$  (corresponding to an index of the candidate EOWF  $\mathcal{F}$ ) and hardcoded output value  $y$ , and hence must grow with the size of  $\mathcal{F}$ . We thus run into a similar wall: even with obfuscation of Turing machines, the resulting auxiliary input distribution  $\mathcal{Z}$  would only rule out EOWF with a-priori bounded index length.

**Our “Succinct Punctured Program” Technique.** To deal with this issue, we develop a “succinct punctured program” technique. That is, we show how to make the size of the obfuscation be independent of the length of the input, while still retaining its usability as an obfuscator. The idea is two-fold: First, we modify the program  $\Pi^s$  to *hash* the input to the PRF, using a collision-resistant hash function  $h$ . That is, we now consider a program  $\Pi^{h,s}(j) = f_j(\text{PRF}_s(h(j)))$ . Second, we make use of *differing-inputs obfuscation*, as opposed to just indistinguishability obfuscation. Specifically, our constructed auxiliary input distribution  $\mathcal{Z}$  will sample a uniform  $s$  and a random hash function  $h$  (from some appropriate collection of collision-resistant hash functions) and then output a differing-inputs obfuscation of  $\Pi^{h,s}$ .

To prove that this “universal” distribution  $\mathcal{Z}$  over auxiliary input breaks *all* alleged extractable one-way functions over  $\{0, 1\}^k \rightarrow \{0, 1\}^k$ , we define a one-way function inverter  $Inv$  just as before, except that we now feed the EOWF extractor  $\mathcal{E}$  the obfuscation of the “punctured” variant  $\Pi_{i,y}^{h,s}$

which contains a PRF seed punctured at point  $h(i)$ . The program  $\Pi_{i,y}^{h,s}$  proceeds just as  $\Pi^{h,s}$  except on all inputs  $j$  such that  $h(j)$  is equal to this special value  $h(i)$ ; for those inputs it simply outputs the hardcoded value  $y$ . (Note that the index  $i$  is no longer needed to specify the function  $\Pi_{i,y}^{h,s}$ —rather, just its hash  $h(i)$ —but is included for notational convenience). As before, consider a hybrid experiment where  $y$  is selected as  $y := \Pi^{h,s}(i)$ .

Whereas before the punctured program was equivalent to the original, and thus indistinguishability of auxiliary inputs in the different experiments followed by the definition of indistinguishability obfuscation, here it is *no longer* the case that if  $y = \Pi^{h,s}(i)$ , then  $\Pi_{i,y}^{h,s}$  is equivalent to  $\Pi^{h,s}$ —in fact, they may differ on many points. More precisely, the programs may differ in all points  $j$  such that  $h(j) = h(i)$ , but  $j \neq i$  (since  $f_j$  and  $f_i$  may differ on the input  $PRF_s(h(i))$ ). Thus, we can no longer rely on indistinguishability obfuscation to provide indistinguishability of these two hybrids.

We resolve this issue by relying *differing-inputs obfuscation* instead of just indistinguishability obfuscation. Intuitively, if obfuscations of  $\Pi^{h,s}$  and  $\Pi_{i,y}^{h,s}$  can be distinguished when  $y$  is set to  $\Pi^{h,s}(i)$ , then we can efficiently recover some input  $j$  where the two programs differ. But, by construction, this must be some point  $j$  for which  $h(j) = h(i)$  (or else the two programs are the same), and  $j \neq i$  (since we chose the hardcoded value  $y = \Pi^{h,s}(i)$  to be consistent with  $\Pi^{h,s}$  on input  $i$ ). Thus, if the obfuscations can be distinguished, we can find a collision in  $h$ , contradicting its collision resistance.

To formalize this argument, we require a notion of differing-inputs obfuscation with auxiliary input. As it turns out, we only require extractability to hold with respect to the distribution that samples a random hash function  $h$  from the CRHF family. We have thus achieved our goal of demonstrating two distributions  $\mathcal{Z}, \mathcal{D}$  such that either extractable one-way functions  $\{0, 1\}^k \rightarrow \{0, 1\}^k$  do not exist w.r.t.  $\mathcal{Z}$ , or differing-inputs obfuscators do not exist w.r.t.  $\mathcal{D}$ . Finally, note that assuming the existence *public-coin* collision-resistant hash functions [HR04], then the auxiliary input portion of  $\mathcal{D}$  can be uniform.

## 1.2 Removing Auxiliary Input in $di\mathcal{O}$

We then turn our focus to differing-inputs obfuscation. We observe that in the case of  $di\mathcal{O}$ , the situation of “bad side information” cannot be circumvented simply by disallowing auxiliary input, but rather such information can appear in the *input-output behavior* of the programs to be obfuscated. Namely, we show that for any distribution  $\mathcal{D}$  over  $\mathcal{P} \times \mathcal{P} \times \{0, 1\}^\ell$  of programs  $\mathcal{P}$  and bounded-length auxiliary input, the existence of  $di\mathcal{O}$  w.r.t.  $\mathcal{D}$  is directly implied by the existence of any *indistinguishability* obfuscator ( $i\mathcal{O}$ ) that is  $di\mathcal{O}$ -secure for a slightly enriched distribution of programs  $\mathcal{D}'$  over  $\mathcal{P}' \times \mathcal{P}'$ , *without* auxiliary input.

Intuitively, this transformation works by embedding the “bad auxiliary input” into the *input-output behavior* of the circuits to be obfuscated themselves. That is, the new distribution  $\mathcal{D}'$  is formed by sampling first a triple  $(P_0, P_1, z)$  of programs and auxiliary input from the original distribution  $\mathcal{D}$ , and then instead considering the tweaked programs  $P_0^z, P_1^z$  that have a special additional input  $x^*$  (denoted later as “mode = \*”) for which  $P_0^z(x^*) = P_1^z(x^*)$  is defined to be  $z$ . This introduces no new differing inputs to the original program pair  $P_0, P_1$ , but now there is no hope of preventing the adversary from learning  $z$  without sacrificing correctness of the obfuscation scheme.

A technical challenge arises in the security reduction, however, in which we must modify the obfuscation of the  $z$ -embedded program  $P_b^z$  to “look like” an obfuscation of the original program  $P_b$ . Interestingly, this issue is solved by making use of a *second layer* of obfuscation, and is where the  $i\mathcal{O}$  security of the obfuscator is required. We refer the reader to Section 5 for details.

### 1.3 Other Applications of the “Succinct Punctured Program” Technique

As mentioned above, the “punctured program” paradigm of [SW14] has been used in multiple applications (e.g., [SW14, HSW14, GGHR, BZ14]). Many of them rely on punctured programs in an essentially identical way to the approach described above, and in particular follow the same hybrids within the security proof. Furthermore, for some of these applications, there are significant gains in making the obfuscation succinct (i.e., independent of the input size of the obfuscated program). Thus, for these applications, if we instead rely on differing-inputs obfuscation (and the existence of collision-resistant hash functions), by using our succinct punctured program technique, we can obtain significant improvements. For instance, relying on the same approach as above, we can show based on these assumptions:

- “Succinct” Perfect Zero-Knowledge Non-Interactive *Universal* Argument System (with communication complexity  $k^\epsilon$  for every  $\epsilon$ ), by relying on the non-succinct Perfect NIZK construction of [SW14].
- A *universal* instantiation of Random Oracles, for which the Full Domain Hash (FDH) signature paradigm [BR93] is (selectively) secure for *every* trapdoor (one-to-one) function (if hashing not only the message but also the index of the trapdoor function), by relying on the results of [HSW14] showing how to provide a trapdoor-function specific instantiation of the random oracle in the FDH.<sup>7</sup>

We explore these applications further in Section 4.

## 2 Preliminaries

### 2.1 Differing-Inputs Obfuscation

For each distribution  $\mathcal{D} = \{D_k\}$  over  $\mathcal{M}_k \times \mathcal{M}_k \times \{0, 1\}^*$  (specifying program pairs and auxiliary input), we define differing-inputs obfuscation secure with respect to  $\mathcal{D}$ . We present a definition as formalized in [BCP14], which is a variant of the notion of “differing-inputs” obfuscation of [BGI<sup>+</sup>12].<sup>8</sup>

**Definition 2.1** (Differing-Inputs Obfuscator w.r.t. Distribution  $\mathcal{D}$ ). A uniform PPT machine  $di\mathcal{O}$  is a *differing-inputs obfuscator w.r.t. distribution  $\mathcal{D}$*  over  $\mathcal{M}_k \times \mathcal{M}_k \times \{0, 1\}^*$  if it satisfies the following correctness and security properties:

- **Correctness:** There exists a negligible function  $\text{negl}(k)$  such that for every security parameter  $k \in \mathbb{N}$ , for all  $M \in \mathcal{M}_k$ , for all inputs  $x$ , we have

$$\Pr[\tilde{M} \leftarrow di\mathcal{O}(1^k, M) : \tilde{M}(x) = M(x)] = 1 - \text{negl}(k).$$

<sup>7</sup>That is, [HSW14] shows that for every trapdoor one-to-one function, there exists some way to instantiate the random oracle so that the resulting scheme is secure. In contrast, our results shows that there exists a single instantiation that works no matter what the trapdoor function is.

<sup>8</sup>Formally, our notion of differing-inputs obfuscation departs from differing-inputs obfuscation of [BGI<sup>+</sup>12] in three ways: First, we consider a setting in which the adversary (and extractor) may receive auxiliary input. Second, [BGI<sup>+</sup>12] require the extractor  $E$  to extract a differing input for  $M_0, M_1$  given *any* pair of programs  $M'_0, M'_1$  evaluating equivalent functions. Third, [BGI<sup>+</sup>12] consider also adversaries who distinguish with negligible advantage  $\epsilon(k)$ , and require that extraction still succeeds in this setting, but within time polynomial in  $1/\epsilon$ . In contrast, we restrict our attention only to adversaries who succeed with noticeable advantage.

- **Security:** For every non-uniform PPT adversary  $\mathcal{A}$  and polynomial  $p(k)$ , there exists a non-uniform PPT extractor  $E$  and polynomial  $q(k)$  such that for every  $k \in \mathbb{N}$ , it holds with overwhelming probability over  $(M_0, M_1, z) \leftarrow \mathcal{D}$  that

$$\begin{aligned} & \Pr \left[ b \leftarrow \{0, 1\}; \tilde{M} \leftarrow \text{di}\mathcal{O}(1^k, M_b) : \mathcal{A}(1^k, \tilde{M}, z) = b \right] \geq \frac{1}{2} + \frac{1}{p(k)} \\ \implies & \Pr \left[ w \leftarrow E(1^k, M_0, M_1, z) : M_0(w) \neq M_1(w) \right] \geq \frac{1}{q(k)}. \end{aligned}$$

**Definition 2.2** (Differing-Inputs Obfuscator for TM). A uniform PPT machine  $\text{di}\mathcal{O}_{\text{TM}}$  is called a *differing-inputs obfuscator for the class TM of polynomial-size Turing machines* if it satisfies the following. For each  $k$ , let  $\mathcal{M}_k$  be the class of Turing machines  $\Pi$  containing a description of a Turing machine  $M$  of size bounded by  $k$ , such that  $\Pi$  takes two inputs,  $(t, x)$ , with  $|t| = k$ , and the output of  $\Pi(t, x)$  is defined to be the the output of running the Turing machine  $M(x)$  for  $t$  steps. Then  $\text{di}\mathcal{O}_{\text{TM}}$  is a differing-inputs obfuscator for  $\{\mathcal{M}_k\}$ .<sup>9</sup>

## 2.2 Extractable One-Way Functions

For this work, we consider a slightly weakened version of EOWFs w.r.t. *distributional* auxiliary input information, from some distribution  $\mathcal{Z}$ . Namely, we require one-wayness and extractability to hold with overwhelming probability over auxiliary input  $z$  sampled from  $\mathcal{Z}$ . In contrast, typical definitions require these properties to hold for *any* auxiliary input  $z$  in some auxiliary input set  $\mathcal{Z}$  (e.g., *all* length-bounded values). For example, “standard” auxiliary-input-secure EOWFs [CD08] are required to be  $\mathcal{Z}$ -auxiliary-input EOWFs for *every* (possibly non-uniform) distribution  $\mathcal{Z}$ .

We present a non-uniform version of the definition, in which both one-wayness and extractability are with respect to *non-uniform* polynomial-time adversaries.

**Definition 2.3** ( $\mathcal{Z}$ -Auxiliary-Input EOWF). Let  $\ell, m$  be polynomially bounded length functions. An efficiently computable family of functions  $\mathcal{F} = \left\{ f_i : \{0, 1\}^k \rightarrow \{0, 1\}^{\ell(k)} \mid i \in \{0, 1\}^{m(k)}, k \in \mathbb{N} \right\}$ , associated with an efficient probabilistic key sampler  $\mathcal{K}_{\mathcal{F}}$ , is a  *$\mathcal{Z}$ -auxiliary-input extractable one-way function* if it satisfies:

- **One-wayness:** For non-uniform polynomial-time  $\mathcal{A}$  and sufficiently large  $k \in \mathbb{N}$ ,

$$\Pr \left[ z \leftarrow \mathcal{Z}_k; i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); x \leftarrow \{0, 1\}^k; x' \leftarrow \mathcal{A}(i, f_i(x); z) : f_i(x') = f_i(x) \right] \leq \text{negl}(k).$$

- **Extractability:** For any non-uniform polynomial-time adversary  $\mathcal{A}$ , there exists a non-uniform polynomial-time extractor  $\mathcal{E}$  such that, for sufficiently large security parameter  $k \in \mathbb{N}$ :

$$\Pr \left[ z \leftarrow \mathcal{Z}_k; i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); y \leftarrow \mathcal{A}(i; z); x' \leftarrow \mathcal{E}(i; z) : \exists x \text{ s.t. } f_i(x) = y \wedge f_i(x') \neq y \right] \leq \text{negl}(k).$$

---

<sup>9</sup>Note that applying the properties of differing-inputs obfuscation to this class of Turing machines  $\{\mathcal{M}_k\}$  implies that for programs  $\Pi_0, \Pi_1 \in \mathcal{M}_k$  defined above (corresponding to underlying size- $k$  Turing machines  $M_0, M_1$ ), efficiently distinguishing between obfuscations of  $\Pi_0$  and  $\Pi_1$  implies that one can efficiently extract an input *pair*  $(t', x')$  for which  $\Pi_0(t', x') \neq \Pi_1(t', x')$ . In particular, either  $M_0(x') \neq M_1(x')$  or  $\text{Runtime}(M_0, x') \neq \text{Runtime}(M_1, x)$ . Thus, if restricting attention to a subclass of  $\mathcal{M}_k$  for which each pair of programs satisfies  $\text{Runtime}(M_0, x) = \text{Runtime}(M_1, x)$  for each input  $x$ , then “standard” extraction is guaranteed (i.e., such that the extracted input contains  $x'$  satisfying  $M_0(x') \neq M_1(x')$ ). In the sequel, when referring to a differing-inputs obfuscation of a Turing machine  $M$ , we will implicitly mean the related program  $\Pi_M$  as above, but will suppress notation of the additional input  $t$ . For our application, it will be the case for the relevant class of Turing machines that every pair of programs  $M_0, M_1$  has same runtime per input (and thus we will achieve “standard” input-extraction guaranteed).

## 2.3 Succinct Non-Interactive Arguments (SNARGs and SNARKs)

We focus attention to *publicly verifiable* succinct arguments.

We consider succinct non-interactive arguments *of knowledge* (SNARKs) with adaptive soundness in Section 3.2, and consider the case of specific distributional auxiliary input.

**Definition 2.4** ( $\mathcal{Z}$ -Auxiliary Input Adaptive SNARK). A triple of algorithms (CRSGen, Prove, Verify) is a *publicly verifiable, adaptively sound succinct non-interactive argument of knowledge (SNARK)* for the relation  $\mathcal{R}$  if the following conditions are satisfied for security parameter  $k$ :

- **Completeness:** For any  $(x, w) \in \mathcal{R}$ ,

$$\Pr[\text{crs} \leftarrow \text{CRSGen}(1^k); \pi \leftarrow \text{Prove}(x, w, \text{crs}) : \text{Verify}(x, \pi, \text{crs}) = 1] = 1.$$

In addition,  $\text{Prove}(x, w, \text{crs})$  runs in time  $\text{poly}(k, |y|, t)$ .

- **Succinctness:** The length of the proof  $\pi$  output by  $\text{Prove}(x, w, \text{crs})$ , as well as the running time of  $\text{Verify}(x, \pi, \text{crs})$ , is bounded by  $p(k + |X|)$ , where  $p$  is a universal polynomial that does not depend on  $\mathcal{R}$ . In addition,  $\text{CRSGen}(1^k)$  runs in time  $\text{poly}(k)$ : in particular,  $\text{crs}$  is of length  $\text{poly}(k)$ .
- **Adaptive proof of knowledge:** For any non-uniform polynomial-size prover  $P^*$  there exists a non-uniform polynomial-size extractor  $\mathcal{E}_{P^*}$ , such that for all sufficiently large  $k \in \mathbb{N}$  and auxiliary input  $z \leftarrow \mathcal{Z}$ , it holds that

$$\Pr[z \leftarrow \mathcal{Z}; \text{crs} \leftarrow \text{CRSGen}(1^k); (x, \pi) \leftarrow P^*(z, \text{crs}); (x, w) \leftarrow \mathcal{E}_{P^*}(z, \text{crs}) : \text{Verify}(\text{crs}, x, \pi) = 1 \wedge w \notin R(x)] \leq \text{negl}(k).$$

We also consider the following notion of zero-knowledge (ZK) succinct non-interactive arguments (SNARGs) without the extraction property, in Section 4.1.

**Definition 2.5** (Perfect ZK-SNARG). A triple of algorithms (CRSGen, Prove, Verify) is a *publicly verifiable perfect zero-knowledge (ZK) succinct non-interactive argument (SNARG)* for the relation  $\mathcal{R}$  (corresponding to language  $L$ ) if it satisfies the Correctness and Succinctness properties as in Definition 2.4, in addition to the following properties:

- **(Non-Adaptive) Soundness:** For every PPT  $P^*$ , for every  $x \notin L$ , it holds that

$$\Pr[\text{crs} \leftarrow \text{CRSGen}(1^k); \pi \leftarrow P^*(1^k, x, \text{crs}) : \text{Verify}(x, \pi, \text{crs}) = 1] \leq \text{negl}(k).$$

- **Perfect Zero Knowledge:** There exist PPT algorithms  $\mathcal{S} = (\mathcal{S}^{\text{crs}}, \mathcal{S}^{\text{Proof}})$  such that for any polynomial collection  $(x_i, w_i) \in \mathcal{R}$ ,  $i \in [\ell(k)]$ , the following distributions are *identical*:

$$\{\text{crs} \leftarrow \text{CRSGen}(1^k); \pi_1 \leftarrow \text{Prove}(x_1, w_1, \text{crs}); \dots; \pi_\ell \leftarrow \text{Prove}(x_\ell, w_\ell, \text{crs}) : (\text{crs}, \pi_1, \dots, \pi_\ell)\}.$$

$$\{(\text{crs}^{\text{sim}}, \text{td}) \leftarrow \mathcal{S}^{\text{crs}}(1^k); \pi_1^{\text{sim}} \leftarrow \mathcal{S}^{\text{Proof}}(x_1, \text{crs}, \text{td}); \dots; \pi_\ell^{\text{sim}} \leftarrow \mathcal{S}^{\text{Proof}}(x_\ell, \text{crs}, \text{td}) : (\text{crs}^{\text{sim}}, \pi_1^{\text{sim}}, \dots, \pi_\ell^{\text{sim}})\}.$$

**Definition 2.6** ((Perfect ZK) Universal Arguments). [BG08] We say that (CRSGen, Prove, Verify) is a *(perfect zero-knowledge) universal argument* if it is a (Perfect ZK) SNARG for the universal relation  $R_U$ , defined to be the set of instance-witness pairs  $(y, w)$ , where  $y = (M, x, t)$ ,  $|w| \leq t$ , and  $M$  is a Turing machine, such that  $M$  accepts  $(x, w)$  after at most  $t$  steps.

Note that while the witness  $w$  for each instance  $y = (M, x, t)$  in the relation  $R_U$  is of size at most  $t$ , there is no a-priori polynomial bounding  $t$  in terms of  $|x|$ .

## 2.4 Puncturable PRFs

Our result makes use of puncturable PRFs, which are PRFs with an extra capability to generate keys that allow one to evaluate the function on all bit strings of a certain length, except for any polynomial-size set of inputs. We focus on the simple case of puncturing PRFs at a single point. The definition is formulated as in [SW14], following the specific exposition of [BCPR14].

**Definition 2.7** (Puncturable PRFs). Let  $m', \ell$  be polynomially bounded length functions. An efficiently computable family of functions

$$\mathcal{PRF} = \left\{ \text{PRF}_s : \{0, 1\}^{m'(k)} \rightarrow \{0, 1\}^{\ell(k)} \mid s \in \{0, 1\}^k, k \in \mathbb{N} \right\},$$

associated with an efficient (probabilistic) seed sampler  $\mathcal{K}_{\mathcal{PRF}}$ , is a puncturable PRF if there exists a puncturing algorithm  $\text{Punct}$  that takes as input a seed  $s \in \{0, 1\}^k$  and a point  $x^* \in \{0, 1\}^{m'(k)}$ , and outputs a punctured seed  $s_{x^*}$ , so that the following conditions are satisfied:

- **Functionality is preserved under puncturing:** For every  $x^* \in \{0, 1\}^{m'(k)}$ ,

$$\Pr \left[ s \leftarrow \mathcal{K}_{\mathcal{PRF}}(1^k); s_{x^*} \leftarrow \text{Punct}(k, x^*) : \forall x \neq x^*, \text{PRF}_s(x) = \text{PRF}_{s_{x^*}}(x) \right] = 1.$$

- **Indistinguishability at punctured points:** The following ensembles are computationally indistinguishable:

$$\begin{aligned} & \{s \leftarrow \mathcal{K}_{\mathcal{PRF}}(1^k), s_{x^*} \leftarrow \text{Punct}(s, x^*) : x^*, s_{x^*}, \text{PRF}_s(x^*)\}_{x^* \in \{0, 1\}^{m(k)}, k \in \mathbb{N}} \\ & \{s \leftarrow \mathcal{K}_{\mathcal{PRF}}(1^k), s_{x^*} \leftarrow \text{Punct}(s, x^*), u \leftarrow \{0, 1\}^{\ell(k)} : x^*, s_{x^*}, u\}_{x^* \in \{0, 1\}^{m(k)}, k \in \mathbb{N}}. \end{aligned}$$

Note that the definition is “selectively secure,” where  $x^*$  is specified before sampling the public parameters. For notational simplicity, (and without loss of generality), we will assume the punctured key  $s_{x^*}$  explicitly includes  $x^*$  in the clear.

As observed in [BW13, BGI14, KPTZ13], the GGM tree-based PRF construction [GGM86] yields puncturable PRFs as defined above, based on any one-way function. The size of such a punctured key  $s_{x^*}$  in this construction is  $O(m'(k) \cdot \ell(k))$  (specifically, a punctured key at input  $x^* = x_1 x_2 \cdots x_{m'(k)}$  can be attained by providing  $m'(k)$  size- $\ell(k)$  partial evaluations in the GGM tree, corresponding to prefixes  $(\bar{x}_1), (x_1 \bar{x}_2), \dots, (x_1 x_2 \cdots \bar{x}_{m'(k)})$ .)

**Theorem 2.8** ([BW13, BGI14, KPTZ13]). *If one-way functions exist, then for all efficiently computable functions  $m'(k)$  and  $\ell(k)$ , there exists a puncturable PRF family that maps  $m'(k)$  bits to  $\ell(k)$  bits, such that the size of a punctured key is  $O(m'(k) \cdot \ell(k))$ .*

## 3 Differing-Inputs Obfuscation or Extractable One-Way Functions

In this section, we present our main result: a conflict between differing-inputs obfuscation (“*diO*”) and extractable one-way functions (EOWF) w.r.t. particular fixed distributions of auxiliary information.

### 3.1 From $\mathcal{D}$ -Distributional $di\mathcal{O}$ to Impossibility of $\mathcal{Z}$ -Auxiliary-Input EOWF

We demonstrate bounded polynomial-time uniformly samplable distributions  $\mathcal{Z}$  (with bounded poly-size output length) and  $\mathcal{D}$  (over  $TM \times TM \times \{0, 1\}^{p(k)}$  for bounded poly-size  $p(k)$ ) such that if there exists differing-inputs obfuscation for Turing machines w.r.t. distribution  $\mathcal{D}$  of programs and auxiliary inputs, and there exist collision-resistant hash functions (CRHF), then there do *not* exist extractable one-way functions (EOWF) w.r.t. auxiliary information sampled from distribution  $\mathcal{Z}$ . In our construction,  $\mathcal{Z}$  consists of an obfuscated Turing machine, and the auxiliary input portion of  $\mathcal{D}$  is precisely the distribution of CRHF function descriptions. In particular, if there exist CRHFs whose indices are random (e.g., public-coin CRHFs), then the auxiliary input in  $\mathcal{D}$  can be *uniform*.

We emphasize that we provide a single distribution  $\mathcal{Z}$  of auxiliary inputs for which *all* candidate EOWF families  $\mathcal{F}$  with given output length will fail. This is in contrast to the result of [BCPR14], which show for each candidate family  $\mathcal{F}$  that there exists a tailored distribution  $\mathcal{Z}_{\mathcal{F}}$  (whose size grows with  $|\mathcal{F}|$ ) for which  $\mathcal{F}$  will fail.

**Theorem 3.1.** *For every polynomial  $\ell$ , there exists an efficient, uniformly samplable distributions  $\mathcal{Z}, \mathcal{D}$  such that, assuming the existence of collision-resistant hash functions and  $\mathcal{D}$ -distributional differing-inputs obfuscation for Turing machines, then there cannot exist  $\mathcal{Z}$ -auxiliary-input extractable one-way functions  $\{f_i : \{0, 1\}^k \rightarrow \{0, 1\}^{\ell(k)}\}$ .*

*Proof.* We construct an adversary  $\mathcal{A}$  and desired distribution  $\mathcal{Z}$  on auxiliary inputs, such that for any alleged EOWF family  $\mathcal{F}$ , there cannot exist an efficient extractor corresponding to  $\mathcal{A}$  given auxiliary input from  $\mathcal{Z}$  (assuming existence of the listed tools).

**The Universal Adversary  $\mathcal{A}$ .** We consider a universal PPT adversary  $\mathcal{A}$  that, given  $(i, z) \in \{0, 1\}^{\text{poly}(k)} \times \{0, 1\}^{n(k)}$ , parses  $z$  as a Turing machine and returns  $z(i)$ . Note that in our setting,  $i$  corresponds to the index of the selected function  $f_i \in \mathcal{F}$ , and (looking ahead) the auxiliary input  $z$  will contain an obfuscated program.

**The Auxiliary Input Distribution  $\mathcal{Z}$ .** Let  $\mathcal{PRF} = \{\text{PRF}_s : \{0, 1\}^{m(k)} \rightarrow \{0, 1\}^k\}_{s \in \{0, 1\}^k}$  be a puncturable pseudorandom function family, and  $\mathcal{H} = \{\mathcal{H}_k\}$  a collision-resistant hash function family with  $h : \{0, 1\}^* \rightarrow \{0, 1\}^{m(k)}$  for each  $h \in \mathcal{H}_k$ . (Note that by Theorem 2.8, punctured PRFs for these parameters exist based on OWFs, which are implied by CRHF). We begin by defining two classes of *Turing machines*:

$$\begin{aligned} \mathcal{M} &= \left\{ \Pi^{h,s} \mid s \in \{0, 1\}^k, h \in \mathcal{H}_k, k \in \mathbb{N} \right\}, \\ \mathcal{M}^* &= \left\{ \Pi_{i,y}^{h,s} \mid s \in \{0, 1\}^k, y \in \{0, 1\}^{\ell(k)}, h \in \mathcal{H}_k, k \in \mathbb{N} \right\}, \end{aligned}$$

which we now describe. We assume without loss of generality for each  $k$  that the corresponding collection of Turing machines  $\Pi^{h,s} \in \mathcal{M}_k$ ,  $\Pi_{i,y}^{h,s} \in \mathcal{M}_k^*$  are of the *same size*; this can be achieved by padding. (We address the size bound of each class of machines below). In a similar fashion, we may further assume that for each  $k$  the runtime of each  $\Pi^{h,s}$  and  $\Pi_{i,y}^{h,s}$  on any given input  $f_i$  is equal (see discussion in Section 2.1).

At a high level, each machine  $\Pi^{h,s}$  accepts as input a poly-size circuit description of a function  $f_i$  (with canonical description, including a function index  $i$ ), computes the hash of the corresponding index  $i$  w.r.t. the hardcoded hash function  $h$ , applies a PRF with hardcoded seed  $s$  to the hash, and then evaluates the circuit  $f_i$  on the resulting PRF output value  $x$ : that is,  $\Pi_{i,y}^{h,s}(f_i)$  outputs  $U_k(f_i, \text{PRF}_s(h(i)))$ , where  $U_k$  is the universal Turing machine. See Figure 1. Note that each  $\Pi^{h,s}$  can be described by a Turing machine of size  $O(|s| + |h| + |U_k|)$ , which is bounded by  $p(k)$  for some fixed polynomial  $p$ .

Turing Machine  $\Pi^{h,s}$ :

**Hardwired:** Hash function  $h : \{0, 1\}^* \rightarrow \{0, 1\}^{m(k)}$ , PRF seed  $s \in \{0, 1\}^k$ .

**Inputs:** Circuit description  $f_i$

1. Hash the index:  $v = h(i)$ .
2. Compute the PRF on this hash:  $x = \text{PRF}_s(v)$ .
3. Output the evaluation of the universal Turing machine on inputs  $f_i, x$ : i.e.,  $y = U_k(f_i, x)$ .

**Figure 1:** Turing machines  $\Pi^{h,s} \in \mathcal{M}$ .

Auxiliary Input Distribution  $\mathcal{Z}_k$ :

1. Sample a hash function  $h \leftarrow \mathcal{H}_k$  and PRF seed  $s \leftarrow \mathcal{K}_{\mathcal{PRF}}(1^k)$ .
2. Output an obfuscation  $\tilde{\Pi} \leftarrow \text{diO}(\Pi^{h,s})$ .

**Figure 2:** The auxiliary input distribution  $\mathcal{Z}_k$ .

The machines  $\Pi_{i,y}^{h,s}$  perform a similar task, except that instead of having the entire PRF seed  $s$  hardcoded, they instead only have a *punctured* seed  $s^*$  derived from  $s$  by puncturing it at the point  $h(i)$  (i.e., enabling evaluation of the PRF on all points except  $h(i)$ ). In addition, it has hardcoded an output  $y$  to replace the punctured result. More specifically, on input a circuit description  $f_j$  (with explicitly specified index  $j$ ), the program  $\Pi_{i,y}^{h,s}$  first computes the hash  $h = h(j)$ , continues computation as usual for any  $h \neq h(i)$  using the punctured PRF key, and for  $h = h(i)$ , it skips the PRF and  $U_k$  evaluation steps and directly outputs  $y$ . Note that because  $h$  is not injective, this puncturing may change the value of the program on multiple inputs  $f_j$  (corresponding to functions  $f_j \in \mathcal{F}$  with  $h(j) = h(i)$ ). When the hardcoded value  $y$  is set to  $y = f_i(\text{PRF}_s(h(i)))$ , then  $\Pi_{i,y}^{h,s}$  agrees with  $\Pi^{h,s}$  additionally on the input  $f_i$ , but not necessarily on the other inputs  $f_j$  for which  $h(j) = h(i)$ . (Indeed, whereas the hash of their indices collide, and thus their corresponding PRF outputs,  $\text{PRF}(h(j))$ , will agree, the final step will apply *different* functions  $f_j$  to this value). We remark that indistinguishability obfuscation arguments will thus not apply to this scenario, since we are modifying the computed functionality. In contrast, differing-inputs obfuscation will guarantee that the two obfuscated programs are indistinguishable, otherwise we can efficiently *find* one of the disagreeing inputs, which will correspond to a collision in the CRHF.

Note that each  $\Pi_{i,y}^{h,s}$  can be described by a Turing machine of size  $O(|s^*| + |h| + |y| + |U_k|)$ . Recall by Theorem 2.8 the size of the punctured PRF key  $|s^*| \in O(m'(k)\ell(k))$ , where the PRF has input and output lengths  $m'(k)$  and  $\ell(k)$ . In our application, note that the input to the PRF is not the function index  $i$  itself (in which case the machine  $\Pi_{i,y}^{h,s}$  would need to grow with the size of the alleged EOWF family), but rather the *hashed* index  $h(i)$ , which is of fixed polynomial length. Thus, collectively, we have  $|\Pi_{i,y}^{h,s}|$  is bounded by a fixed polynomial  $p'(k)$ , and finally that there exists a single fixed polynomial bound on the size of *all* programs  $\Pi^{h,s} \in \mathcal{M}, \Pi_{i,y}^{h,s} \in \mathcal{M}^*$ . This completely determines the defined auxiliary input distribution  $\mathcal{Z} = \{\mathcal{Z}_k\}$ , described in full in Figure 2. (Note that the size of the auxiliary output generated by  $\mathcal{Z}$ , which corresponds to an obfuscation of an appropriately padded program  $\Pi^{h,s}$  is thus also bounded by a fixed polynomial in  $k$ ).

**$\mathcal{A}$  Has No Extractor.** We show that, based on the assumed security of the underlying tools, the constructed adversary  $\mathcal{A}$  given auxiliary input from the constructed distribution  $\mathcal{Z} = \{\mathcal{Z}_k\}$ , cannot have an extractor  $\mathcal{E}$  satisfying Definition 2.3:

Turing Machine  $\Pi_{i,y}^{h,s}$ :

**Hardwired:** Hash function  $h : \{0, 1\}^* \rightarrow \{0, 1\}^{m(k)}$ , punctured PRF seed  $s^* \in \{0, 1\}^k$ , punctured point  $h(i)$ , bit string  $y \in \{0, 1\}^{\ell(k)}$ .

**Input:** Circuit description  $f_j$  (containing index  $j$ )

1. Hash the index:  $v = h(j)$ .
2. If  $v \neq h(i)$ , compute  $x = \text{PRF}_{s^*}(v)$ , and output  $U_k(f_j, x)$ .
3. If  $v = h(i)$ , output  $y$ .

**Figure 3:** “Punctured” Turing machines  $\Pi_{i,y}^{h,s} \in \mathcal{M}^*$ .

Auxiliary Input Distribution  $\mathcal{Z}_k(i, y)$ :

1. Sample a hash function  $h \leftarrow \mathcal{H}_k$  and PRF seed  $s \leftarrow \mathcal{K}_{\mathcal{PRF}}(1^k)$ .
2. Sample a punctured PRF seed  $s^* \leftarrow \text{Punct}(s, h(i))$ , punctured at point  $h(i)$ .
3. Compute the “correct” punctured evaluation:  $y = f_i(\text{PRF}_s(h(i)))$ .
4. Output an obfuscation  $\tilde{M} \leftarrow \text{di}\mathcal{O}(\Pi_{i,y}^{h,s})$ , where  $\Pi_{i,y}^{h,s}$  is defined from  $(h, s^*, y)$ , as in Figure 10.

**Figure 4:** The “punctured” distribution  $\mathcal{Z}_k(i, y)$ .

**Proposition 3.2.** *For any non-uniform polynomial-time candidate extractor  $\mathcal{E}$  for  $\mathcal{A}$ , it holds that  $\mathcal{E}$  fails with overwhelming probability: i.e.,*

$$\Pr \left[ z \leftarrow \mathcal{Z}_k; i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); y \leftarrow \mathcal{A}(i; z); x' \leftarrow \mathcal{E}(i; z) : \exists x \text{ s.t. } f_i(x) = y \wedge f_i(x') \neq y \right] \geq 1 - \text{negl}(k).$$

*Proof.* First note that given auxiliary input  $z \leftarrow \mathcal{Z}_k$ ,  $\mathcal{A}$  produces an element in the image of the selected  $f_i$  with high probability. That is,

$$\Pr \left[ z \leftarrow \mathcal{Z}_k; i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); y \leftarrow \mathcal{A}(i; z) : \exists x \text{ s.t. } f_i(x) = y \right] \geq 1 - \text{negl}(k).$$

Indeed, by the definition of  $\mathcal{A}$  and  $\mathcal{Z}_k$ , and the correctness of the obfuscator  $\text{di}\mathcal{O}$ , then we have with overwhelming probability

$$\mathcal{A}(i; z) = \tilde{M}(f_i) = \Pi^{h,s}(f_i) = f_i(\text{PRF}_s(h(i))),$$

where  $z = \tilde{M}$  is an obfuscation of  $\Pi^{h,s} \in \mathcal{M}$ ; i.e.,  $z = \tilde{M} \leftarrow \text{di}\mathcal{O}(\Pi^{h,s})$ .

Now, suppose for contradiction that there exists a non-negligible function  $\epsilon(k)$  such that for all  $k \in \mathbb{N}$  the extractor  $\mathcal{E}$  successfully outputs a preimage corresponding to the output  $\mathcal{A}(i; z) \in \text{Range}(f_i)$  with probability  $\epsilon(k)$ : i.e.,

$$\Pr \left[ z \leftarrow \mathcal{Z}_k; i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); x' \leftarrow \mathcal{E}(i; z) : f_i(x') = \mathcal{A}(i; z) = f_i(\text{PRF}_s(h(i))) \right] \geq \epsilon(k).$$

where as before,  $s, h$  are such that  $z = \text{di}\mathcal{O}(\Pi^{h,s})$ . We show that this cannot be the case, via three steps.

**Step 1: Replace  $\mathcal{Z}$  with “punctured” distribution  $\mathcal{Z}(i, y)$ .** For every index  $i$  of the EOWF family  $\mathcal{F}$  and  $k \in \mathbb{N}$ , consider an alternative distribution  $\mathcal{Z}_k(i, y)$  that, instead of sampling and obfuscating a Turing machine  $\Pi^{h,s}$  from the class  $\mathcal{M}$ , as is done for  $\mathcal{Z}$ , it does so with a Turing machine  $\Pi_{i,y}^{h,s} \in \mathcal{M}^*$  as follows. First, it samples a hash function  $h \leftarrow \mathcal{H}_k$  and PRF seed  $s$  as usual. It then generates a *punctured* PRF key  $s^* \leftarrow \text{Punct}(s, h(i))$  that enables evaluation of the PRF on all points except the value  $h(i)$ . For the specific index  $i$ , it computes the correct full evaluation  $y := f_i(\text{PRF}_s(h(i)))$ . Finally,  $\mathcal{Z}_k(i, y)$  outputs an obfuscation of the constructed program  $\Pi_{i,y}^{h,s}$  as specified in Figure 10 from the values  $(h, s^*, y)$ : i.e.,  $\tilde{M} \leftarrow \text{diO}(\Pi_{i,y}^{h,s})$ . See Figure 4 for a full description of  $\mathcal{Z}(i, y)$ .

We now argue that the extractor  $\mathcal{E}$  must also succeed in extracting a preimage when given a value  $z^* \leftarrow \mathcal{Z}_k(i, y)$  from this modified distribution instead of  $\mathcal{Z}_k$ . At a high level, the argument runs as follows. If  $\mathcal{E}$ 's extraction success drops by a non-negligible amount when instead given a sample from this new distribution, then he can be used to distinguish between obfuscations of corresponding pair of programs  $\Pi^{h,s}, \Pi_{i,y}^{h,s}$ . By the differing-inputs obfuscation property, this implies that we can efficiently extract an input  $j$  on which the two programs differ. But, by construction, the only such inputs  $j$  are those that collide with  $i$  with respect to the collision-resistant hash function. Thus, such an extractor cannot exist, as it can be used to efficiently find collisions in the CRHF.

We now formalize this intuition.

**Lemma 3.3.** *It holds that*

$$\Pr \left[ i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); z^* \leftarrow \mathcal{Z}_k(i, y); x' \leftarrow \mathcal{E}(i; z^*) : \right. \\ \left. f_i(x') = \mathcal{A}(i; z^*) = f_i(\text{PRF}_s(h(i))) \right] \geq \epsilon(k) - \text{negl}(k). \quad (1)$$

*Proof.* Suppose, to the contrary, there exists a non-negligible function  $\alpha(k)$  for which the probability in Equation (1) is less than  $\epsilon(k) - \alpha(k)$ . In particular, there exists a polynomial  $p(k)$  such that for infinitely many values of  $k$ , there exists an index  $i_k$  in the range of  $\mathcal{K}_{\mathcal{F}}(1^k)$  for which

$$\Pr \left[ z^* \leftarrow \mathcal{Z}_k(i_k, y^*); x' \leftarrow \mathcal{E}(i_k; z^*) : f_{i_k}(x') = \mathcal{A}(i_k; z^*) = f_{i_k}(\text{PRF}_s(h(i_k))) \right] < \epsilon(k) - \frac{1}{p(k)}. \quad (2)$$

Denote by  $I = \{i_k\}$  this ensemble of function indices.

Consider the following non-uniform obfuscation adversary  $\mathcal{A}_{\text{obf}}^I$  making use of  $\mathcal{E}$ , hardcoded with the ensemble of “good” index values  $I$ , and given auxiliary input a hash function description  $h$  (note that we eventually wish to turn  $\mathcal{E}$  into an adversary for the CRHF family).

Obfuscation adversary  $\mathcal{A}_{\text{obf}}^I(1^k, \tilde{M}, h)$ :

1. Execute the EOWF extractor  $\mathcal{E}$ , giving  $i_k$  (from the hardcoded ensemble  $I$ ) as the target index, and the obfuscated program  $\tilde{M}$  as auxiliary input. That is,  $x' \leftarrow \mathcal{E}(i_k; \tilde{M})$ .
2. Output 0 if  $\mathcal{E}$  succeeded in extracting a preimage: i.e., if  $f_{i_k}(x') = \tilde{M}(i_k)$ . Otherwise, output a random bit  $b' \leftarrow \{0, 1\}$ .

Now, expanding out the sampling procedures of the distributions  $\mathcal{Z}_k, \mathcal{Z}_k(i_k, y)$ , Equation (2) implies that for the hardcoded ensemble  $E = \{i_k\}$ , it holds for infinitely many values of  $k$  that

$$\Pr \left[ h \leftarrow \mathcal{H}_k; (M_0, M_1) \leftarrow \text{ProgSample}(1^k, i_k, h); \right. \\ \left. b \leftarrow \{0, 1\}; \tilde{M} \leftarrow \text{diO}(M_b) : \mathcal{A}_{\text{obf}}^I(1^k, \tilde{M}, h) = b \right] \geq \frac{1}{2} + \frac{1}{2p(k)}, \quad (3)$$

where  $\text{ProgSample}(1^k, i_k, h)$  is given by:

1.  $s \leftarrow \mathcal{K}_{\mathcal{P}\mathcal{R}\mathcal{F}}(1^k)$ .
2.  $s^* \leftarrow \text{Punct}(s, h(i_k))$ .
3. Output  $M_0 := \Pi^{h,s}$ ,  $M_1 := \Pi_{i,y}^{h,s}$ .

Indeed, note that if  $b = 0$  then the obfuscated program  $\tilde{M}$  given to  $\mathcal{A}_{\text{obf}}^I$  in Equation (3) is distributed as  $\tilde{M} \leftarrow \mathcal{Z}_k$ , in which case  $\mathcal{E}$  extracts a preimage with probability at least  $\epsilon(k)$ . In contrast, if  $b = 1$ , then  $\tilde{M}$  given to  $\mathcal{A}_{\text{obf}}^E$  is distributed as  $\tilde{M} \leftarrow \mathcal{Z}_k(i_k, y)$ , in which case (for infinitely many  $k$ )  $\mathcal{E}$  extracts a preimage with probability  $\epsilon(k) - 1/p(k)$ .

Thus, with probability at least  $1/4p(k)$  over  $h \leftarrow \mathcal{H}_k$ ;  $(M_0, M_1) \leftarrow \text{ProgSample}(1^k, i_k, h)$ , it holds that

$$\Pr \left[ b \leftarrow \{0, 1\}; \tilde{M} \leftarrow \text{di}\mathcal{O}(M_b) : \mathcal{A}_{\text{obf}}^I(1^k, \tilde{M}, h) = b \right] \geq \frac{1}{2} + \frac{1}{4p(k)}. \quad (4)$$

Now, by the differing-inputs obfuscation security of  $\text{di}\mathcal{O}$  (with respect to distribution  $\mathcal{D}$  of programs and auxiliary input corresponding to sampled  $\Pi^{h,s}$ ,  $\Pi_{i,y}^{h,s}$  and auxiliary input  $h$ ) there exists a non-uniform PPT algorithm  $\mathcal{E}_{\text{obf}}^{I'}$  (with some non-uniform advice ensemble  $I'$ ) and a polynomial  $q(k)$  corresponding to  $\mathcal{A}_{\text{obf}}^I$  and  $p(k)$ , such that for any pair of programs  $M_0, M_1$  for which  $\mathcal{A}_{\text{obf}}^I$  successfully distinguishes between the obfuscations  $\text{di}\mathcal{O}(M_0)$  and  $\text{di}\mathcal{O}(M_1)$  with advantage  $1/4p(k)$  (when given auxiliary input  $h \leftarrow \mathcal{H}_k$ ), then  $\mathcal{E}_{\text{obf}}^{I'}$  (given the pair  $M_0, M_1$ ) will extract a disagreeing input  $j$  such that  $M_0(j) \neq M_1(j)$  with non-negligible probability  $1/q(k)$  when given the same auxiliary input  $h \leftarrow \mathcal{H}_k$ .

By (4), this means that for infinitely many  $k$ , with probability  $1/4p(k)$  over  $h \leftarrow \mathcal{H}_k$ ;  $(M_0, M_1) \leftarrow \text{ProgSample}(1^k, i_k, h)$ , we have

$$\Pr \left[ j \leftarrow \mathcal{E}_{\text{obf}}^{I'}(1^k, M_0, M_1, h) : M_0(j) \neq M_1(j) \right] \geq \frac{1}{q(k)}. \quad (5)$$

We now argue that this contradicts the collision resistance of the hash family  $\mathcal{H}$ . Namely, consider the following non-uniform PPT collision-finding adversary  $\mathcal{A}_{\text{CR}}^{I, I'}$ , who is given both advice ensembles  $I, I'$ .

Collision-finding adversary  $\mathcal{A}_{\text{CR}}^{I, I'}(1^k)$ :

1. Receive a random hash function  $h \leftarrow \mathcal{H}_k$  from the CRHF challenger.
2. Sample a corresponding pair of programs  $(M_0, M_1) \leftarrow \text{ProgSample}(1^k, i_k, h)$  where  $i_k$  is specified by the advice ensemble  $I$ . Recall that  $M_0$  is a program  $\Pi^{h,s}$  and  $M_1$  is a corresponding “ $i_k$ -punctured” program  $\Pi_{i,y}^{h,s}$ .
3. Execute the (non-uniform) extractor  $\mathcal{E}_{\text{obf}}^{I'}$  from above on this pair of programs, hash function descriptor  $h$ , and advice from  $I'$ : that is,  $j \leftarrow \mathcal{E}_{\text{obf}}^{I'}(1^k, M_0, M_1, h)$ .
4. Output the pair  $(i_k, j)$  as the alleged collision for  $h$ .

By Equation (5), we have that with probability at least  $1/4p(k)$  over the sampling of the challenge function  $h \leftarrow \mathcal{H}_k$  and  $\mathcal{A}_{\text{CR}}^{I, I'}$ 's sampling of  $(M_0, M_1) \leftarrow \text{ProgSample}(1^k, i_k, h)$  that the extractor  $\mathcal{E}_{\text{obf}}^{I'}$  succeeds in extracting a disagreeing input with probability  $1/q(k)$ . Thus, with probability  $1/4p(k)q(k)$ , the value  $j$  output by  $\mathcal{A}_{\text{CR}}^{I, I'}(1^k)$  satisfies  $M_0(j) \neq M_1(j)$ ; that is,  $\Pi^{h,s}(j) \neq \Pi_{i,y}^{h,s}(j)$ . But, by construction of  $\Pi_{i,y}^{h,s}$ , the set of such inputs consists exactly of the values  $\{j \neq i_k \mid h(j) = h(i_k)\}$ . (Indeed, recall that for all inputs  $j$  with  $h(j) = h(i_k)$ , we hardcoded their final output to the value corresponding to  $\Pi^{h,s}(i_k)$ ). Therefore, with non-negligible probability  $(1/4p(k)q(k))$  for infinitely many  $k$ ,  $\mathcal{A}_{\text{CR}}^{I, I'}$  has found a collision in  $h$ , yielding a contradiction to the security of  $\mathcal{H}$ . The lemma follows.  $\square$

**Step 2: Replace “correct” hardcoded  $y$  in  $\mathcal{Z}(i, y)$  with random  $f_i$  evaluation.** Next, we consider another experiment where  $\mathcal{Z}_k(i, y)$  is altered to a nearly identical distribution  $\mathcal{Z}_k(i, u)$  where, instead of hardcoding the “correct”  $i$ -evaluation value  $y = f_i(\text{PRF}_s(h(i)))$  in the generated “punctured” program  $\Pi_{i,y}^{h,s}$ , the distribution  $\mathcal{Z}_k(i, u)$  now simply samples a random  $f_i$  output  $y = f_i(u)$  for an independent random  $u \leftarrow \{0, 1\}^k$ . We claim that the original EOWF extractor  $\mathcal{E}$  still succeeds in finding a preimage when given this new auxiliary input distribution:

**Lemma 3.4.** *It holds that*

$$\Pr \left[ i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); z^{**} \leftarrow \mathcal{Z}_k(i, u); x' \leftarrow \mathcal{E}(i; z^{**}) : \right. \\ \left. f_i(x') = \mathcal{A}(i; z^{**}) = f_i(u) \right] \geq \epsilon(k) - \text{negl}(k). \quad (6)$$

*Proof.* This follows from the fact that  $\text{PRF}_s(h(i))$  is pseudorandom, even given the  $h(i)$ -punctured key  $s^*$ .

Formally, consider an algorithm  $\mathcal{A}_{\text{PRF}}^0$  which, on input the security parameter  $1^k$ , a pair of values  $i, h$ , and a pair  $s^*, x$  (that will eventually correspond to a challenge punctured PRF key, and either  $\text{PRF}_s(h(i))$  or random  $u$ ), performs the following steps.

Algorithm  $\mathcal{A}_{\text{PRF}}^0(1^k, i, h, s^*, x)$ :

1. Take  $y = f_i(x)$ , and obfuscate the associated program  $\Pi_{i,y}^{h,s}$ : i.e.,  $z^{**} \leftarrow \text{diO}(1^k, \Pi_{i,y}^{h,s})$ .
2. Run the EOWF extractor given index  $i$  and auxiliary input  $z^{**}$ :  $x' \leftarrow \mathcal{E}(i; z^{**})$ .
3. Output 0 if  $\mathcal{E}$  succeeds in extracting a valid preimage: i.e., if  $f_i(x') = y^* = f_i(x)$ . Otherwise, output a random bit  $b \leftarrow \{0, 1\}$ .

Now, suppose Lemma 3.4 does not hold: i.e., the probability in Equation (6) differs by some non-negligible amount from  $\epsilon(k)$ . Then, expanding out the sampling procedure of  $\mathcal{Z}_k(i, y)$  and  $\mathcal{Z}_k(i, u)$ , we have for some non-negligible function  $\alpha(k)$  that

$$\Pr \left[ i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); h \leftarrow \mathcal{H}_k; s \leftarrow \mathcal{K}_{\mathcal{PRF}}(1^k); s^* \leftarrow \text{Punct}(s, h(i)); u \leftarrow \{0, 1\}^k; \right. \\ \left. b \leftarrow \{0, 1\} : \mathcal{A}_{\text{PRF}}^0(1^k, i, h, x_b) = b \right] \geq \frac{1}{2} + \alpha(k), \quad (7)$$

where  $x_0 := \text{PRF}_s(h(i))$  and  $x_1 := u$ . Indeed, in the case  $b = 0$ , the auxiliary input  $z^{**}$  generated by  $\mathcal{A}_{\text{PRF}}$  and given to  $\mathcal{E}$  has distribution exactly  $\mathcal{Z}(i, y)$ , whereas in the case  $b = 1$ , the generated  $z^{**}$  has distribution exactly  $\mathcal{Z}(i, u)$ .

In particular, there exists a polynomial  $p(k)$  such that for infinitely many  $k$ , there exists an index  $i_k$  and hash function  $h_k \in \mathcal{H}_k$  with

$$\Pr \left[ s \leftarrow \mathcal{K}_{\mathcal{PRF}}(1^k); s^* \leftarrow \text{Punct}(s, h(i_k)); u \leftarrow \{0, 1\}^k; \right. \\ \left. b \leftarrow \{0, 1\} : \mathcal{A}_{\text{PRF}}^0(1^k, i_k, h, x_b) = b \right] \geq \frac{1}{2} + \frac{1}{p(k)}, \quad (8)$$

where  $x_0, x_1$  are as before.

Consider a non-uniform punctured-PRF adversary  $\mathcal{A}_{\text{PRF}}^I$  (with the ensemble  $I = \{i_k, h_k\}$  hardcoded) that first selects the challenge point  $h_k(i_k)$ ; receives the PRF challenge information  $(s^*, x)$  for this point; executes  $\mathcal{A}_{\text{PRF}}^0$  on input  $(1^k, i_k, h_k, s^*, x)$ , and outputs the corresponding bit  $b$  output by  $\mathcal{A}_{\text{PRF}}^0$ . Then by (8), it follows that  $\mathcal{A}_{\text{PRF}}^I$  breaks the security of the punctured PRF.  $\square$

**Step 3: Such an extractor breaks one-wayness of EOWF.** Finally, we observe that this means that  $\mathcal{E}$  can be used to break the one-wayness of the original function family  $\mathcal{F}$ . Indeed, given a random key  $i$  and a challenge output  $y = f_i(u)$ , an inverter can simply sample a hash function  $h$  and  $h(i)$ -punctured PRF seed  $s^*$  on its own, construct the program  $\Pi_{i,y}^{h,s}$  with its challenge  $y$  hardcoded in, and sample an obfuscation  $z^{**} \leftarrow \text{diO}(\Pi_{i,y}^{h,s})$ . Finally, it runs  $\mathcal{E}(i, z^{**})$  to invert  $y^*$ , with the same probability  $\epsilon(k) - \text{negl}(k)$ . □

This concludes the proof of Theorem 3.1. □

## 3.2 Differing-Inputs Obfuscation or SNARKs

We link the existence of differing-inputs obfuscation for  $NC^1$  and the existence of succinct non-interactive arguments of knowledge (SNARKs), via an intermediate step of *proximity* extractable one-way functions (PEOWFs), a notion related to EOWFs, introduced in [BCCT12]. Namely, building upon the results of the previous subsection, and results of [BCP14, BCCT12], we show:

1. Assuming the existence of fully homomorphic encryption (FHE) with decryption in  $NC^1$  and SNARKs for NP, there exist efficient distributions  $\mathcal{Z}, \mathcal{D}$  such that  $\mathcal{D}$ -distributional differing-inputs obfuscation for  $NC^1$  implies that there *cannot* exist PEOWFs  $\{f : \{0, 1\}^k \rightarrow \{0, 1\}^k\}$  w.r.t.  $\mathcal{Z}$ .
2. PEOWFs  $\{f : \{0, 1\}^k \rightarrow \{0, 1\}^k\}$  w.r.t. this auxiliary input distribution  $\mathcal{Z}$  are *implied by* the existence of SNARKs for NP secure w.r.t. a second efficient auxiliary input distribution  $\mathcal{Z}''$  (and collision-resistant hash functions), as shown in [BCCT12].
3. Thus, one of these conflicting hypotheses must be false. That is, there exist efficient distributions  $\mathcal{Z}'', \mathcal{D}$  such that assuming existence of FHE with decryption in  $NC^1$  and collision-resistant hash functions, then either: (1)  $\mathcal{D}$ -distributional differing-inputs obfuscation for  $NC^1$  does not exist, or (2) SNARKs for NP w.r.t.  $\mathcal{Z}''$  do not exist.

Note that we focus on the specific case of PEOWFs with  $k$ -bit inputs and  $k$ -bit outputs, as this suffices to derive the desired contradiction; however, the theorems following extend also to the more general case of PEOWF output length (demonstrating an efficient distribution  $\mathcal{Z}$  to rule out each potential output length  $\ell(k)$ ).

### 3.2.1 Proximity EOWFs

We begin by defining Proximity EOWFs.

**Proximity Extractable One-Way Functions (PEOWFs).** In a Proximity EOWF (PEOWF), the extractable function family  $\{f_i\}$  is associated with a “proximity” equivalence relation  $\sim$  on the range of  $f_i$ , and the one-wayness and extractability properties are modified with respect to this relation. The one-wayness is strengthened: not only must it be hard to find an exact preimage of  $v$ , but it is also hard to find a preimage of any equivalent  $v \sim v'$ . The extractability requirement is weakened accordingly: the extractor does not have to output an exact preimage of  $v$ , but only a preimage of some equivalent value  $v' \sim v$ .

As an example, consider functions of the form  $f : x \mapsto (f_1(x), f_2(x))$  and equivalence relation on range elements  $(a, b) \sim (a, b')$  whose first components agree. Then the proximity extraction property requires for any adversary  $\mathcal{A}$  who outputs an image element  $(a, b) \in \text{Range}(f)$  that there exists an extractor  $\mathcal{E}$  who finds an input  $x$  for which  $f(x) = (a, b')$  for some  $b'$  not necessarily equal to  $b$ .

In this work, we allow the relation  $\sim$  to depend on the function index  $i$ , but require that the relation  $\sim$  is *publicly* (and efficiently) testable. We further consider non-uniform adversaries and extraction algorithms, and (in line with this work) auxiliary inputs coming from a specified distribution  $\mathcal{Z}$ .

**Definition 3.5** ( $\mathcal{Z}$ -Auxiliary-Input Proximity EOWFs). Let  $\ell, m$  be polynomially bounded length functions. An efficiently computable family of functions

$$\mathcal{F} = \left\{ f_i : \{0, 1\}^k \rightarrow \{0, 1\}^{\ell(k)} \mid i \in \{0, 1\}^{m(k)}, k \in \mathbb{N} \right\},$$

associated with an efficient probabilistic key sampler  $\mathcal{K}_{\mathcal{F}}$ , is a  $\mathcal{Z}$ -auxiliary-input proximity extractable one-way function if it satisfies the following (strong) one-wayness, (weak) extraction, and public testability properties:

- **(Strengthened) One-wayness:** For non-uniform polynomial-time  $\mathcal{A}$  and sufficiently large security parameter  $k \in \mathbb{N}$ ,

$$\Pr \left[ z \leftarrow \mathcal{Z}_k; i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); x \leftarrow \{0, 1\}^k; x' \leftarrow \mathcal{A}(i, f_i(x); z) : f_i(x') \sim f_i(x) \right] \leq \text{negl}(k).$$

- **(Weakened) Extractability:** For any non-uniform polynomial-time adversary  $\mathcal{A}$ , there exists a non-uniform polynomial-time extractor  $\mathcal{E}$  such that, for sufficiently large security parameter  $k \in \mathbb{N}$ ,

$$\Pr \left[ z \leftarrow \mathcal{Z}_k; i \leftarrow \mathcal{K}_{\mathcal{F}}(1^k); y \leftarrow \mathcal{A}(i; z); x' \leftarrow \mathcal{E}(i; z) : \exists x \text{ s.t. } f_i(x) = y \wedge f_i(x') \not\sim y \right] \leq \text{negl}(k).$$

- **Publicly Testable Relation:** There exists a deterministic polytime machine  $\mathcal{T}$  such that, given the function index  $i$ ,  $\mathcal{T}$  accepts  $y, y' \in \{0, 1\}^{\ell(k)}$  if and only if  $y \sim_k y'$ .

### 3.2.2 ( *diO* for $NC^1$ + FHE + SNARK ) $\implies$ No $\mathcal{Z}$ -PEOWF

We now show that, assuming the existence of fully homomorphic encryption (FHE) with decryption in  $NC^1$ ,<sup>10</sup> then for some efficiently computable distributions  $\mathcal{D}_{\text{obf}}, \mathcal{Z}_{\text{SNARK}}, \mathcal{Z}_{\text{PEOWF}}$ , if there exist differing-inputs obfuscators w.r.t distribution  $\mathcal{D}_{\text{obf}}$  of  $NC^1$  circuits and auxiliary inputs, and SNARKs w.r.t. auxiliary input  $\mathcal{Z}_{\text{SNARK}}$ , then there *cannot* exist PEOWFs w.r.t. auxiliary input  $\mathcal{Z}_{\text{PEOWF}}$ . This takes place in two steps.

First, we remark that an identical proof to that of Theorem 3.1 rules out the existence of  $\mathcal{Z}$ -auxiliary-input *proximity EOWFs* in addition to standard EOWFs, based on the same assumptions: namely, assuming differing-inputs obfuscation w.r.t.  $\mathcal{D}$  for Turing machines and collision-resistant hash functions (which are implied by FHE [IKO05]). Indeed, assuming the existence of a PEOWF extractor  $\mathcal{E}$  for the adversary  $\mathcal{A}$  and auxiliary input distribution  $\mathcal{Z}$  (who extracts a “related” preimage to the target value), the same procedure yields a PEOWF inverter who similarly extracts a “related” preimage to any challenge output. In the reduction, it is merely required that the success of  $\mathcal{E}$  is efficiently and publicly testable (this is used to construct a distinguishing adversary for the differing-inputs obfuscation scheme, in Step 1). However, this is directly implied by the public testability of the PEOWF relation  $\sim$ , as specified in Definition 3.5.

**Theorem 3.6.** *There exist efficient, uniformly samplable distributions  $\mathcal{Z}, \mathcal{D}$  such that, assuming the existence of collision-resistant hash functions and differing-inputs obfuscation w.r.t.  $\mathcal{D}$  for polynomial-size Turing machines, there cannot exist (publicly testable)  $\mathcal{Z}$ -auxiliary-input PEOWFs  $\{f_i : \{0, 1\}^k \rightarrow \{0, 1\}^k\}$ .*

<sup>10</sup>As is the case for nearly all existing FHE constructions (e.g., [GSW13, BV14]).

Now, in [BCP14], it was shown that differing-inputs obfuscation for all polynomial-size Turing machines can be achieved by bootstrapping up from differing-inputs obfuscation for  $NC^1$ , assuming the existence of FHE with decryption in  $NC^1$  and SNARKs.<sup>11</sup> The resulting Turing machine obfuscator will be secure w.r.t.  $\mathcal{D}$  if the underlying SNARK scheme is secure w.r.t. auxiliary input  $\mathcal{D}$ , and the original  $NC^1$  obfuscator is secure w.r.t. an augmented program and auxiliary input distribution  $\mathcal{D}_{\text{obf}} = \mathcal{D}_{\text{obf}} \parallel \mathcal{Z}_{\text{crs}}$ , formed by concatenating a sample from  $\mathcal{D}'$  with additional auxiliary input composed of a CRS generated for the SNARK scheme.

Putting this together with Theorem 3.6, we thus have the following corollary.

**Corollary 3.7.** *There exist efficient, uniformly samplable distributions  $\mathcal{Z}, \mathcal{D}_{\text{obf}}$  (and polynomial  $p$ ) such that, assuming the existence of SNARKs and FHE with decryption in  $NC^1$ , then assuming the existence of differing-inputs obfuscation w.r.t.  $\mathcal{D}_{\text{obf}}$  over  $NC^1 \times NC^1 \times \{0, 1\}^{p(k)}$ , there cannot exist PEOWFs  $\{f_i : \{0, 1\}^k \rightarrow \{0, 1\}^k\}$  w.r.t. auxiliary input  $\mathcal{Z}$ .*

### 3.2.3 ( SNARK + CRHF ) $\implies$ $\mathcal{Z}$ -PEOWF

As shown in [BCCT12], Proximity EOWFs (PEOWFs) with respect to an auxiliary input distribution  $\mathcal{Z}$  are *implied by* collision-resistant hash functions (CRHF) and SNARKs secure with respect to a related auxiliary input distribution.<sup>12</sup>

Loosely, the transformation converts any collision-resistant hash function family  $\mathcal{F}$  into a PEOWF by appending to the output of each  $f \in \mathcal{F}$  a succinct SNARK argument  $\pi_x$  that there exists a preimage  $x$  yielding output  $f(x)$ . (If the Prover algorithm of the SNARK system is randomized, then the function is also modified to take an additional input, which is used as the random coins for the SNARK generation). The equivalence relation on outputs is defined by  $(y, \pi) \sim (y', \pi')$  if  $y = y'$  (note that this relation is publicly testable). More explicitly, consider the new function family  $\mathcal{F}'$  composed of functions

$$f'_{\text{crs}}(x, r) = \left( f(x), \text{Prove}(1^k, \text{crs}, f(x), x; r) \right),$$

where a function  $f'_{\text{crs}} \in \mathcal{F}'$  is sampled by first sampling a function  $f \leftarrow \mathcal{F}$  from the original CRHF family, and then sampling a CRS for the SNARK scheme,  $\text{crs} \leftarrow \text{CRSGen}(1^k)$ .

Now (as proved in [BCCT12]), the resulting function family will be a PEOWF with respect to auxiliary input  $\mathcal{Z}$  if the underlying SNARK system is secure with respect to an augmented auxiliary input distribution  $\mathcal{Z}_{\text{SNARK}} := (\mathcal{Z}, \mathcal{Z}')$ , formed by concatenating a sample from  $\mathcal{Z}$  with a function index sampled from the collision-resistant hash function family  $\mathcal{F}$ .

**Theorem 3.8** ([BCCT12]). *There exist efficient, uniformly samplable distributions  $\mathcal{Z}, \mathcal{Z}_{\text{SNARK}}$  such that, assuming the existence of collision-resistant hash functions and SNARKs for NP secure w.r.t. auxiliary input distribution  $\mathcal{Z}_{\text{SNARK}}$ , then there exist PEOWFs  $\{f_i : \{0, 1\}^k \rightarrow \{0, 1\}^k\}$  w.r.t.  $\mathcal{Z}$ .*

### 3.2.4 Reaching a Standoff

Observe that the conclusions of Corollary 3.7 and Theorem 3.8 are in direct contradiction. Thus, it must be that one of the two sets of assumptions is false. Namely,

<sup>11</sup>Explicitly, [BCP14] prove that bootstrapping can be based on the weaker assumption of  $P$ -certificates (succinct arguments for languages in  $P$ ) in the case where there is a polynomial bound on the input size. A slight modification to their construction, in which the user instead generates a SNARK with respect to the *hash* of the input, admits bootstrapping for unbounded input length based on SNARKs. Indeed, the extraction property of the SNARK guarantees that the entire input will be recovered, as required.

<sup>12</sup>[BCCT12] consider the setting of arbitrary auxiliary input; however, their construction directly implies similar results for specific auxiliary input distributions.

**Corollary 3.9.** *Assuming the existence of fully homomorphic encryption with decryption in  $NC^1$ , there exist efficiently samplable distributions  $\mathcal{Z}_{\text{SNARK}}, \mathcal{D}_{\text{obf}}$  such that one of the following two objects cannot exist:*

- *SNARKs w.r.t. auxiliary input distribution  $\mathcal{Z}_{\text{SNARK}}$ .*
- *Differing-inputs obfuscation w.r.t. distribution  $\mathcal{D}_{\text{obf}}$  of  $NC^1$  circuits and auxiliary input.*

More explicitly, we have that  $\mathcal{Z}_{\text{SNARK}} = (\mathcal{Z}, \mathcal{Z}')$  and  $\mathcal{D}_{\text{obf}} = \mathcal{D} \parallel \mathcal{Z}_{\text{crs}}$ , where  $\mathcal{Z}$  is composed of an obfuscated program,  $\mathcal{Z}'$  consists of a randomly sampled index from a CRHF family,  $\mathcal{Z}_{\text{crs}}$  corresponds to a randomly sampled CRS from a SNARK system, and  $\mathcal{D} \parallel \mathcal{Z}_{\text{crs}}$  denotes concatenation of additional auxiliary input to the distribution  $\mathcal{D}$ .

## 4 Applications of “Succinct Punctured Programs” Technique

We now demonstrate a variety of applications of our “succinct punctured programs” technique.

### 4.1 Perfect NIZK Universal Arguments

In [SW14], Sahai and Waters demonstrated a construction of a Non-Interactive Zero-Knowledge (NIZK) argument system with perfect zero knowledge from indistinguishability obfuscation, supporting fixed NP languages with statements (and witnesses) up to an a-priori bounded size. Using our succinct punctured programs technique, we achieve a non-interactive *universal argument system*, also perfectly zero knowledge, for languages and statements of unbounded polynomial size (see Definition 2.6).

**Theorem 4.1.** *Assume the existence of auxiliary-input-secure differing-inputs obfuscation for TM and collision-resistant hash functions. Then for any constant  $\epsilon > 0$  there exists a perfect zero-knowledge universal argument system, as in Definition 2.6.*

As with the construction of [SW14], the obfuscation scheme we use must be secure with respect to a particular distribution  $\mathcal{Z}$  consisting of a second obfuscated program. For simplicity, however, we state our result with respect to obfuscation secure against arbitrary auxiliary input.

**NIZK of [SW14].** We first recall the NIZK construction of [SW14]. The system consists of two obfuscated circuits (serving as the CRS):

- A Prove circuit, which has hardcoded a PRF seed  $s$ , takes as input a statement and witness pair  $(x, w)$  and outputs the PRF evaluation  $\text{PRF}_s(x)$  on  $x$  if  $w$  is a valid witness (i.e.,  $R(x, w) = 1$ ).
- A Verify circuit, which has hardcoded the same PRF seed  $s$  and a description of a one-way function  $f$ , takes as input a statement and alleged proof  $(x, \pi)$ , and outputs 1 exactly if  $f(\pi)$  is equal to  $f(\text{PRF}_s(x))$ . (The introduction of the one-way function  $f$  is not needed for correctness, but rather to argue security by use of the “punctured programs” technique).

Sahai and Waters show this scheme satisfies perfect zero knowledge and (non-adaptive) soundness assuming the obfuscation scheme used is a secure indistinguishability obfuscator with respect to arbitrary auxiliary input [SW14].

Note that here the size of the Prove and Verify circuits must grow with the size of potential statements  $x$ , thus inherently fixing an upper bound on the handled statement size at the time of CRS generation.

**Succinct ZK Universal Arguments.** Following our technique of succinct punctured programs, we modify the construction of [SW14] in two ways.

First, we replace the obfuscated Prove and Verify circuits (which depend on a fixed NP relation  $R$ ) with obfuscated *Turing machines*, which can accept arbitrary polynomial-size relations as part of their input. That is, we now consider instances of the universal relation  $R_U$ , as described in Definition 2.6. To obtain obfuscation of Turing machines, we use differing-inputs obfuscation in the place of indistinguishability obfuscation. On the surface, this modification almost appears to suffice for our goal. However, one problem remains: to prove soundness of the resulting scheme for an instance  $(M, x, t) \notin R_U$ , we must argue that the obfuscated Turing machines are indistinguishable from obfuscations of corresponding  $(M, x, t)$ -punctured programs, whose size must grow with  $|(M, x, t)|$ . However, this requires the size of the obfuscated Turing machines to grow with  $|(M, x, t)|$ , thus annihilating our universality.

To solve this problem, we incorporate a second modification to the [SW14] construction, by first computing a *collision-resistant hash* of the input statement, and proceeding with this hashed value  $h(M, x, t)$ . Now we need only that the obfuscated Prove and Verify programs are indistinguishable from an obfuscation of a program that is punctured at  $h(M, x, t)$ , whose size can now be made independent of  $|(M, x, t)|$ . As in the previous section, indistinguishability of these programs will follow from the security of the *differing-inputs* obfuscation together with the collision resistance of  $h$ .

Formally, consider the following tools:

1.  $\mathcal{PRF} = \{\text{PRF}_s : \{0, 1\}^{m(k)} \rightarrow \{0, 1\}^k\}_{s \in \{0, 1\}^k}$  a puncturable PRF family.  
Note that by Theorem 2.8, punctured PRFs for these parameters exist based on OWFs, which are implied by CRHF.
2.  $\mathcal{H} = \{\mathcal{H}_k\}$  a CRHF family with  $h : \{0, 1\}^* \rightarrow \{0, 1\}^{m(k)}$  for each  $h \in \mathcal{H}_k$ .
3.  $\mathcal{F} = \{f : \{0, 1\}^k \rightarrow \{0, 1\}^{k^\epsilon}\}$  be a OWF family.

We now present our construction.

Succinct ZK Universal Argument Construction (CRSGen, Prove, Verify):

**CRSGen( $1^k$ ):** on input the security parameter, the CRS generation procedure samples a PRF seed  $s \leftarrow \mathcal{K}_{\mathcal{PRF}}(1^k)$ , a hash function  $h \leftarrow \mathcal{H}_k$ , and a OWF  $f \leftarrow \mathcal{F}_k$ . It then generates obfuscations of the corresponding Turing machines  $P^{h,s}$  and  $V^{h,s,f}$ , as defined in Figures 5 and 7. That is,  $\tilde{P} \leftarrow \text{diO}(1^k, P^{h,s})$  and  $\tilde{V} \leftarrow \text{diO}(1^k, V^{h,s,f})$ . Output the pair of obfuscated programs  $\text{crs} = (\tilde{P}, \tilde{V})$ .

**Prove( $\text{crs}, \vec{x}, w$ ):** Evaluate the obfuscated program  $\tilde{P} \in \text{crs}$  on input  $(\vec{x}, w)$ : i.e., output  $\tilde{P}(\vec{x}, w)$ .

**Verify( $\text{crs}, \vec{x}, \pi$ ):** Evaluate the obfuscated program  $\tilde{V} \in \text{crs}$  on input  $(\vec{x}, \pi)$ : i.e., output  $\tilde{V}(\vec{x}, \pi)$ .

*Proof of Theorem 4.1.* Perfect zero knowledge holds as the distribution of proofs can be perfectly simulated given the PRF seed  $s$  (without knowledge of any witness).

The size and time complexities follow from a straightforward analysis of the programs  $P^{h,s}, V^{h,s,f}$ . In particular, for any input  $(M, x, t)$ , the proof size is  $k^\epsilon$  bits (independent of  $|M|, t$ ) corresponding to the output of the OWF, and verification of a proof on  $(M, x, t)$  requires only computing a hash and PRF, and not executing  $M$ .

(Non-adaptive) soundness follows an analogous sequence of hybrids as Theorem 3.1, mirroring the approach of [SW14]. Namely, for  $\vec{x} = (M, x, t) \notin R_U$ , the obfuscated programs  $P^{h,s}, V^{h,s,f}$  are sequentially replaced with their “ $\vec{x}$ -punctured” counterparts  $P_{\vec{x}}^{h,s}$  and  $V_{\vec{x},y}^{h,s,f}$  (as defined in Figures 6

**Turing Machine  $P^{h,s}$ :****Hardwired:** Hash function  $h : \{0,1\}^* \rightarrow \{0,1\}^{m(k)}$ , PRF seed  $s \in \{0,1\}^k$ .**Inputs:** Universal statement  $\vec{x} = (M, x, t)$ , alleged witness  $w$ .

1. Check validity of  $w$  as witness for  $\vec{x} = (M, x, t)$  by executing the Turing machine  $M$  on input  $(x, w)$  for  $t$  steps. If  $M(x, w) \neq 1$ , output  $\perp$  and terminate. Otherwise, continue.
2. Hash the statement description:  $v = h(\vec{x})$ .
3. Compute the PRF on this hash:  $y = \text{PRF}_s(v)$ .
4. Output  $y$ .

**Figure 5:** Turing machines  $P^{h,s}$ .**Punctured Turing Machine  $P_{\vec{x}}^{h,s}$ :****Hardwired:** Hash function  $h : \{0,1\}^* \rightarrow \{0,1\}^{m(k)}$ , *punctured* PRF seed  $s^* \in \{0,1\}^k$ , punctured PRF input  $h(\vec{x})$ .**Inputs:** Universal statement  $\vec{x}' = (M', x', t')$ , alleged witness  $w$ .

1. Check validity of  $w$  as witness for  $\vec{x}' = (M', x', t')$  by executing the Turing machine  $M'$  on input  $(x', w)$  for  $t'$  steps. If  $M'(x', w) \neq 1$ , output  $\perp$  and terminate. Otherwise, continue.
2. Hash the statement description:  $v = h(\vec{x})$ .
3. Compute the PRF on this hash:  $y = \text{PRF}_{s^*}(v)$ .
4. Output  $y$ .

**Figure 6:** Punctured Turing machines  $P_{\vec{x}}^{h,s}$ . Note for  $(M, x, t) \notin R_U$ , the punctured output is  $\perp$ .

and 8), and then replacing the hardcoded output  $y$  in the  $V_{\vec{x},y}^{h,s,f}$  program with a random output  $f(u)$  of the OWF  $f$  for uniform  $u \in \{0,1\}^k$ . By the security of the differing-inputs obfuscator, an adversary's success probability in generating a proof on  $\vec{x}$  cannot decrease by too much, or such adversary can be used to find collisions in the hash function  $h$ . Producing a verifying proof on  $\vec{x}$  within this final hybrid then corresponds to inverting a random evaluation of the OWF, contradicting its one-wayness. □

**Turing Machine  $V^{h,s,f}$ :****Hardwired:** Hash function  $h : \{0,1\}^* \rightarrow \{0,1\}^{m(k)}$ , PRF seed  $s \in \{0,1\}^k$ , function  $f \in \mathcal{F}$  from the OWF family.**Inputs:** Universal statement  $\vec{x} = (M, x, t)$ , alleged proof  $\pi$ .

1. Compute the “correct” proof for  $\vec{x}$ : i.e.,  $\pi' = \text{PRF}_s(h(\vec{x}))$ .
2. Verify whether  $f(\pi) = f(\pi')$ . If so, output 1; otherwise output 0.

**Figure 7:** Turing machines  $V^{h,s,f}$ .

**Punctured Turing Machine**  $V_{\vec{x},y}^{h,s,f}$ :

**Hardwired:** Hash function  $h : \{0,1\}^* \rightarrow \{0,1\}^{m(k)}$ , punctured PRF seed  $s \in \{0,1\}^k$ , punctured PRF input  $h(\vec{x})$ , punctured output  $y$ , OWF  $f \in \mathcal{F}$ .

**Inputs:** Universal statement  $\vec{x}' = (M', x', t')$ , alleged proof  $\pi$ .

1. If  $h(\vec{x}') \neq h(\vec{x})$ , then compute the “correct” proof for  $x'$  as  $\pi' = \text{PRF}_s(h(\vec{x}'))$ , and output 1 if and only if  $f(\pi) = f(\pi')$ .
2. If  $h(\vec{x}') = h(\vec{x})$ , then output 1 if and only if  $f(\pi) = y$ , where  $y$  is hardcoded.

**Figure 8:** Punctured Turing machines  $V_{\vec{x},y}^{h,s,f}$ .

## 4.2 Universal Instantiation of Full-Domain Hash

The full-domain hash (FDH) signature paradigm, first proposed by Bellare and Rogaway [BR93, BR96], provides a means of building a signature scheme from any trapdoor permutation, within the heuristic random oracle model. Specifically, a signature on a message  $m$  is generated by evaluating the random oracle at input  $m$ , and then computing the inverse of the trapdoor permutation on this value  $\text{RO}(m)$ . This work has been very influential and formed the foundation for part of the PKCS#1 standard [KS98]. However, negative results in later years called into question the rigorous implications of security proofs in the random oracle model [CGH04, GK03, BBP04], showing e.g. that for some such applications *no* concrete instantiation of the random oracle can yield security.

In a recent work, Hohenberger, Sahai, and Waters [HSW14] presented a methodology for instantiating the random oracle that provides (selective) security for full-domain hash signatures in the standard model, building on recent advances in indistinguishability obfuscation (in particular, using the punctured programs paradigm of [SW14]). They demonstrate for every trapdoor permutation  $f$  that there exists a hash function  $\mathcal{R}_f$  (tailored to  $f$ ) such that the FDH signature scheme is (selectively) secure in the standard model when using  $f$  and instantiating the random oracle by  $\mathcal{R}_f$ .

We show that our succinct punctured programs technique yields a *universal* instantiation of the random oracle providing security for full-domain hash signatures. That is, we provide a single family of Turing machines  $\mathcal{R} = \{\mathcal{R}_k\}$  such that, for *any* injective trapdoor function  $f$ , the Bellare-Rogaway Full-Domain Hash signature scheme [BR93, BR96] using  $f$  and instantiating the random oracle by  $\mathcal{R}$  is selectively secure in the standard model. This construction involves a tweak to the FDH signature structure, in which the random oracle takes as input a description of the trapdoor permutation  $f$  in addition to the message to be signed: i.e.,  $\text{Sign}(m) = f^{-1}(\text{RO}(m, f))$ .

**Our construction.** Intuitively, to instantiate the random oracle, we would like for each message  $m$  to be able to sample a random image of a given trapdoor function  $f$ , without revealing information about the corresponding preimage. In [HSW14], this is done by providing an (indistinguishability) obfuscation of a circuit that computes a PRF on the input  $m$  and then evaluates  $f$  on this outcome. Using the technique of punctured programs [SW14], Hohenberger *et al.* [HSW14] show that given this obfuscated circuit, no information is revealed about the evaluation of the PRF on the (pre-selected) forgery challenge message  $m$ , so that forging on message  $m$  implies inverting a random(-looking) output  $f(\text{PRF}_s(m))$  of the trapdoor function, contradicting its assumed security. In this construction, however, the obfuscated circuit (i.e., the instantiation of the random oracle) inherently depends on the trapdoor function  $f$ .

We avoid this dependency by obfuscating a *Turing machine* that takes as input a message  $m$  and description of the desired trapdoor function  $f$  (described as a poly-size circuit), and outputs a

**Program  $\Pi^{h,s}$ :**

**Hardwired:** Collision-resistant hash function  $h : \{0, 1\}^* \rightarrow \{0, 1\}^{m(k)}$ , PRF seed  $s \in \{0, 1\}^k$ .

**Inputs:** Message  $m$ , circuit description  $f$

1. Hash the input:  $v = h(m, f)$ .
2. Compute the PRF on this hash:  $x = \text{PRF}_s(v)$ .
3. Evaluate the universal Turing machine on inputs  $f, x$ : i.e.,  $y = U_k(f, x)$ .
4. Output  $y$ .

**Figure 9:** Program  $\Pi^{h,s}$  that is obfuscated to form  $R^{h,s} \in \mathcal{R}_k$ .

seemingly random evaluation of  $f$ . As in our constructions from the previous sections, in order to allow the size of the obfuscated program to be independent of the size of the input trapdoor function (while still maintaining security), we first *hash* the input  $(m, f)$ , and then proceed with this hashed value. Indeed, security of the construction will be shown by replacing the obfuscated program with a corresponding obfuscated “punctured” program, with the punctured (challenge) input and output hardcoded; this hardcoded input corresponds to a challenge pair  $(m, f)$  in the naïve case, forcing the obfuscated program to grow in size with  $|f|$ , but is reduced to a short hashed value  $h(m, f)$  in the latter case. Again, to provide security in this modified setting, we rely on differing-inputs obfuscation. Signature queries  $m'$  with  $h(m', f) \neq h(m, f)$  can be simulated in the security reduction given the punctured PRF key; any query of the form  $m'$  with  $h(m', f) = h(m, f)$  yields a collision in  $h$  (and thus will not occur).

More explicitly, each program  $R^{h,s} \in \mathcal{R}_k$  in our random oracle instantiation is a (differing-inputs) obfuscation of a program  $\Pi^{h,s}$ , indexed by a PRF seed  $s$  and collision-resistant hash function  $h$ .  $\Pi^{h,s}$  accepts as input a message  $m$ , and a description of the trapdoor function  $f$  (modeled as a polynomial-size circuit), and it functions by: (1) computing the hash of  $(m, f)$  with respect to  $h$  (using a Merkle-Damgard hash tree approach), (2) applying a PRF with hardcoded seed  $s$  to this hash value, and then (3) evaluating  $f$  on the resulting PRF output value  $x$ . That is,

$$\Pi^{h,s}(m, f) = U_k(f, \text{PRF}_s(h(m, f))),$$

where  $U_k$  is the universal Turing machine, and

$$R^{h,s} \leftarrow \text{diO}(1^k, \Pi^{h,s}).$$

(See Figure 9 for details). Denote by  $\mathcal{M}$  the class of Turing machines

$$\mathcal{M} = \left\{ \Pi^{h,s} \mid s \in \{0, 1\}^k, h \in \mathcal{H}_k, k \in \mathbb{N} \right\}.$$

To prove security of the resulting FDH signature scheme, we consider a second, related class of (“punctured”) Turing machines

$$\mathcal{M}^* = \left\{ \Pi_{m,y}^{h,s} \mid s \in \{0, 1\}^k, h \in \mathcal{H}_k, k \in \mathbb{N} \right\},$$

where each Turing machine  $\Pi_{m,y}^{h,s}$  is defined as in Figure 10. The obfuscator we use is with respect to the class of Turing machines  $\mathcal{M} \cup \mathcal{M}^*$ .

Now, consider the instantiation of the full-domain hash signature scheme using a hash function  $R^{h,s} \in \mathcal{R}_k$  in the place of a random oracle, and with an injective trapdoor function family  $\mathcal{F}$ .

**Program  $\Pi_{m,y}^{h,s}$ :**

**Hardwired:** Hash function  $h : \{0,1\}^* \rightarrow \{0,1\}^{m(k)}$ , punctured PRF seed  $s^* \in \{0,1\}^k$  (punctured at point  $h(m, f)$ ), bit string  $y \in \{0,1\}^{\ell'(k)}$ .

**Inputs:** Message  $m'$ , circuit description  $f'$

1. Hash the input:  $v' = h(m', f')$ .
2. If  $v' \neq h(m, f)$ , compute  $x = \text{PRF}_{s^*}(h)$ , and output  $U_k(f', x)$ .
3. If  $v' = h(m, f)$ , output  $y$ .

**Figure 10:** “Punctured” program  $\Pi_{m,y}^{h,s}$ , used within the security proof.

**Setup( $1^k$ ) :** On input the security parameter  $1^k$ , the setup algorithm samples a program  $R^{h,s} \leftarrow \mathcal{R}_k$  to be used as a random oracle. That is, it samples a random seed for a puncturable PRF  $s \leftarrow \mathcal{F}_{\mathcal{PRF}}(1^k)$ , and an underlying hash function  $h \leftarrow \mathcal{H}_k$  from the collision-resistant hash function family. This uniquely defines a program  $\Pi^{h,s}$ . Then, the algorithm obfuscates this program as  $R^{h,s} \leftarrow \text{diO}(1^k, \Pi^{h,s})$ .

The setup algorithm then samples  $(f, f^{-1}) \leftarrow \text{TDFSetup}(1^k)$  that produces a public index  $f$  and trapdoor  $f^{-1}$  (that allows inversion)

The verification key for the signature scheme is set to  $\text{vk} = (f, R^{h,s})$ , consisting of the TDF description  $f$  and the “random oracle”  $R^{h,s}$ . The secret key  $\text{sk}$  is the trapdoor  $f^{-1}$  and  $R^{h,s}$ .

**Sign( $\text{sk}, m$ ) :** The signature algorithm outputs  $\sigma = f^{-1}(R^{h,s}(m, f))$ .

**Verify( $\text{vk}, m, \sigma$ ) :** The verification algorithm tests if  $f(\sigma) \stackrel{?}{=} (R^{h,s}(m, f))$  and outputs accept if and only if this holds.

**Theorem 4.2.** *Assuming the existence of collision-resistant hash functions, and differing-inputs obfuscation for the class TM, then for any injective trapdoor function family  $\mathcal{F}$ , the scheme described above is a selectively secure signature scheme.*

*Proof.* The proof follows an analogous sequence of hybrids as in Theorem 3.1, mirroring the approach of [HSW14]. Namely, for challenge forgery message  $m$  and injective trapdoor function  $f$ , the hybrids are (loosely) as follows:

Hybrid 0: The real (selective) security experiment.

Hybrid 1: The obfuscated program  $R^{h,s} \leftarrow \text{diO}(1^k, \Pi^{h,s})$  is replaced by an obfuscation of the corresponding “ $(m, f)$ -punctured” program  $\Pi_{m,y}^{h,s}$  with the “correct” hardcoded output  $y = \Pi^{h,s}(m, f)$ . Indistinguishability follows by the security of the differing-inputs obfuscator, together with the collision resistance of  $h$ . (Note that in the reduction, signature queries can be simulated using the inversion trapdoor to  $f$ ).

Hybrid 2: The obfuscated program is replaced by an obfuscation of  $\Pi_{m,y}^{h,s}$ , with hardcoded output  $y$  set to a *random* evaluation of  $f$ : i.e.,  $f(u)$  for uniform  $u$ . Indistinguishability follows by the pseudo randomness of PRF. (Note that in the reduction, signature queries can again be simulated using the inversion trapdoor to  $f$ ).

By the indistinguishability of hybrids, it follows that a forging adversary must continue to successfully forge in this final experiment, Hybrid 2. But, this implies the adversary can invert a random output of the trapdoor function, yielding the desired contradiction.  $\square$

## 5 Removing Auxiliary Input in $di\mathcal{O}$

We now demonstrate that *even in the absence of auxiliary input*, care must be taken when making use of differing-inputs obfuscation ( $di\mathcal{O}$ ). Specifically, we observe a general transformation obtaining  $di\mathcal{O}$  with respect to any distribution  $\mathcal{D}$  over  $\mathcal{P} \times \mathcal{P} \times \{0, 1\}^{p(k)}$  of program pairs and polynomial bounded-length auxiliary input from any  $\mathcal{O}$  that satisfies a weaker notion of *indistinguishability* obfuscation ( $i\mathcal{O}$ ) security (which is not an extractability assumption; see Section 5.1 below), and is  $di\mathcal{O}$ -secure with respect to distribution  $\mathcal{D}'$  over slightly enriched program pairs  $\mathcal{P}' \times \mathcal{P}'$ , with *no* auxiliary input.<sup>13</sup>

Intuitively, we demonstrate that “bad side information” cannot be avoided by disallowing auxiliary input, but rather can appear embedded within the *input-output behavior* of the programs to be obfuscated. Since any candidate obfuscator  $\mathcal{O}$  must preserve the input-output functionality of the underlying programs, this “bad” information is unavoidably released.

**Theorem 5.1** (Removing Auxiliary Input in  $di\mathcal{O}$ ). *For any polynomial  $p$  and distribution  $\mathcal{D}$  over  $\mathcal{P} \times \mathcal{P} \times \{0, 1\}^{p(k)}$ , there exists a polynomial  $q$ , a class of programs  $\mathcal{P}'$  with  $\max_{P' \in \mathcal{P}'} |P'| \leq q(\max_{P \in \mathcal{P}} |P|)$ , and a distribution  $\mathcal{D}'$  over  $\mathcal{P}' \times \mathcal{P}'$  for which the following holds. The existence of  $di\mathcal{O}$  w.r.t.  $\mathcal{D}$  (with auxiliary input) is implied by the existence of obfuscation that is  $i\mathcal{O}$  secure and  $di\mathcal{O}$  secure w.r.t.  $\mathcal{D}'$  (without auxiliary input).*

In the following two subsections, we present a formal definition of  $i\mathcal{O}$ , and then present our transformation and proof for Theorem 5.1.

### 5.1 Indistinguishability Obfuscation

We begin by defining the notion of indistinguishability obfuscation ( $i\mathcal{O}$ ), which we rely on as a tool in our transformation. Note that in general,  $i\mathcal{O}$  is weaker than (and implied by)  $di\mathcal{O}$ . However, in our setting, we consider general-purpose  $i\mathcal{O}$ , versus  $di\mathcal{O}$  that is secure with respect to a particular distribution. In such case, the two notions are incomparable.

**Definition 5.2** (Indistinguishability Obfuscation). A uniform PPT machine  $i\mathcal{O}$  is an *indistinguishability obfuscator* ( $i\mathcal{O}$ ) (w.r.t. worst-case inputs) for a class of Turing machines  $\{\mathcal{M}_k\}_{k \in \mathbb{N}}$  if the following conditions are satisfied:

- **Correctness:** There exists a negligible function  $\text{negl}(k)$  such that for every security parameter  $k \in \mathbb{N}$ , for all  $M \in \mathcal{M}_k$ , for all inputs  $x$ , we have

$$\Pr[M' \leftarrow di\mathcal{O}(1^k, M) : M'(x) = M(x)] = 1 - \text{negl}(k).$$

- **Security:** For every non-uniform PPT adversary  $\mathcal{A}$  and polynomial  $p(k)$ , for every  $k \in \mathbb{N}$ , every pair of Turing machines  $M_0, M_1 \in \mathcal{M}_k$ ,

$$\begin{aligned} \Pr \left[ b \leftarrow \{0, 1\}; \tilde{M} \leftarrow i\mathcal{O}(1^k, M_b) : \mathcal{A}(1^k, \tilde{M}) = b \right] &\geq \frac{1}{2} + \frac{1}{p(k)} \\ \implies \exists x \text{ s.t. } M_0(x) &\neq M_1(x). \end{aligned}$$

**Remark 5.3** ( $i\mathcal{O}$  and Auxiliary Input). In contrast to  $di\mathcal{O}$  obfuscation, indistinguishability obfuscation directly implies security against arbitrary auxiliary input. Indeed, for any non-uniform adversary  $\mathcal{A}$  and auxiliary input  $z$ , we can simply consider a modified non-uniform adversary  $\mathcal{A}_z$

<sup>13</sup>Note that standard  $i\mathcal{O}$  security (for  $P/poly$ ) is not implied by  $di\mathcal{O}$  for a *specific distribution*.

who has  $z$ , and the same conclusion holds. (Recall the issue in the more demanding  $di\mathcal{O}$  setting is that one must demonstrate an *extractor* algorithm  $\mathcal{E}$  for the original adversary  $\mathcal{A}$  such that the same extractor works for any possible auxiliary input  $z$ ).

## 5.2 $di\mathcal{O}$ with Auxiliary Input from $i\mathcal{O}$ + $di\mathcal{O}$ without Auxiliary Input

We now present our “auxiliary input removal” transformation, constituting the proof of Theorem 5.1. Formally, for class of programs  $\mathcal{P}$ , and distribution  $\mathcal{D}$  over  $\mathcal{P} \times \mathcal{P} \times \{0, 1\}^{p(k)}$ , we define generic program modifiers, and then corresponding “enriched” variants  $\mathcal{P}', \mathcal{D}'$  below. (Note that the  $p$ -notation will mostly be suppressed).

**Definition 5.4** (Program Modifiers). For program class  $\mathcal{P}$ , define the following program “wrappers”:

- **Expand( $P, z$ ):** For  $P \in \mathcal{P}, z \in \{0, 1\}^{p(k)}$ , the program  $\text{Expand}(P, z)$  takes as input a pair  $(x, \text{mode})$  for  $x$  a valid input to  $P$  and  $\text{mode} \in \{\emptyset, *\}$  (corresponding to “standard” and “special” modes), and embeds  $z$  into the program input-output as follows:
  1. If  $\text{mode} = \emptyset$ : Output  $P(x)$ .
  2. If  $\text{mode} = *$ : Output  $z$ .

We will sometimes use shorthand notation and denote by  $P^z$  the program  $\text{Expand}(P, z)$

- **ztoZero( $P^z$ ):** For  $P^z$  for some  $P \in \mathcal{P}, z \in \{0, 1\}^{p(k)}$ , the program  $\text{ztoZero}(P^z)$  takes as input a *pair*  $(x, \text{mode})$  but “zeroes” out the special embedded output of  $P^z$ , as follows:
  1. If  $\text{mode} = \emptyset$ : Output  $P^z(x, \text{mode})$ .
  2. If  $\text{mode} = *$ : Output 0.
- **pad( $P, \text{size}$ ):** For any program  $P$  (or  $P^z$ ) with  $|P| \leq \text{size}$ , the program  $\text{pad}(P, \text{size})$  has identical input-output behavior to  $P$ , but satisfies  $|\text{pad}(P, \text{size})| = \text{size}$ .
- **Restrict( $P^z, \text{mode}^*$ ):** For  $P^z$  for some  $P \in \mathcal{P}, z \in \{0, 1\}^{p(k)}$ , and  $\text{mode}^* \in \{\emptyset, *\}$ , the program  $\text{Restrict}(P^z, \text{mode}^*)$  takes as input a *single* value  $x$ , and outputs  $P^z(x, \text{mode}^*)$ , with  $\text{mode}^*$  hardcoded.

**Definition 5.5** (Enriched Program Classes). Let  $\mathcal{P} = \{P\}$  be a class of programs, and  $\mathcal{D}$  be a distribution over  $\mathcal{P} \times \mathcal{P} \times \{0, 1\}^{p(k)}$  for polynomial  $p$ . We define the ( $p$ )-enriched variants  $\mathcal{P}', \mathcal{D}'$  as:

- $\mathcal{P}' := \{\text{Expand}(P, z) : P \in \mathcal{P}, z \in \{0, 1\}^{p(k)}\}$ , where  $\text{Expand}(P, z)$  is as above.
- $\mathcal{D}'$  is sampled over  $\mathcal{P}' \times \mathcal{P}'$  as follows:
  1. Sample  $(P_0, P_1, z) \leftarrow \mathcal{D}$ .
  2. Output the pair  $(\text{Expand}(P_0, z), \text{Expand}(P_1, z)) \in \mathcal{P}' \times \mathcal{P}'$ .

*Proof of Theorem 5.1.* Suppose  $\mathcal{O}'$  is a program obfuscation algorithm such that: (1)  $\mathcal{O}'$  is  $i\mathcal{O}$ -secure, as in Definition 5.2, and (2)  $\mathcal{O}'$  is  $di\mathcal{O}$ -secure w.r.t. program distribution  $\mathcal{D}'$  over  $\mathcal{P}' \times \mathcal{P}'$  (with *no* auxiliary input), as in Definition 2.1. We define a new obfuscation algorithm  $\mathcal{O}$  and show that it is a differing-inputs obfuscator secure w.r.t.  $\mathcal{D}$  over  $\mathcal{P} \times \mathcal{P} \times \{0, 1\}^{p(k)}$  (which includes auxiliary input).

*The Obfuscator  $\mathcal{O}$ :* To obfuscate a program  $P \in \mathcal{P}$ , the algorithm  $\mathcal{O}$  first “expands”  $P$  to include a new input (“ $\text{mode} = *$ ”) with hardcoded output  $0 \in \{0, 1\}^{p(k)}$ . It then pads the expanded program to an appropriate size, and runs the obfuscator  $\mathcal{O}'$  on the padded program. Finally, it “wraps” the result to accept a *single* input  $x$  and run the obfuscated program with the second input  $\text{mode}$  hardcoded

to  $\emptyset$  (i.e., “standard” evaluation mode). More formally, let  $\text{size} = |\mathcal{O}'(\text{ztoZero}(\text{Expand}(P, z)))|$  for  $P \in \mathcal{P}, z \in \{0, 1\}^{p(k)}$ . Then we define

$$\mathcal{O}(1^k, P; r) := \text{Restrict}(\mathcal{O}'(\text{pad}(\text{Expand}(P, 0), \text{size}); r), \emptyset).$$

*Correctness:* Follows directly by the correctness of  $\mathcal{O}'$ . Namely,

$$\begin{aligned} & \left[ \text{Restrict}(\mathcal{O}'(\text{pad}(\text{Expand}(P, 0), \text{size}); r), \emptyset) \right](x) \\ &= \mathcal{O}'(\text{pad}(\text{Expand}(P, 0), \text{size}); r)(x, \emptyset) \quad \text{by defn of Restrict} \\ &= \text{pad}(\text{Expand}(P, 0), \text{size})(x, \emptyset) \quad \text{whp, by correctness of } \mathcal{O}' \\ &= \text{Expand}(P, 0)(x, \emptyset) \quad \text{since pad does not affect input-output} \\ &= P(x) \quad \text{by the defn of Expand (in “standard” mode } \emptyset). \end{aligned}$$

*Security:* Consider a (non-uniform) PPT adversary  $\mathcal{A}$  and polynomial  $q_{\mathcal{A}}$ . To prove the theorem, we must construct an extractor algorithm  $\mathcal{E}$ , polynomial  $q_{\mathcal{E}}$ , and negligible function  $\nu$ , and show that for all but negligible probability  $\nu(k)$  over  $(P_0, P_1, z) \leftarrow \mathcal{D}$ , if

$$\Pr \left[ b \leftarrow \{0, 1\}; \tilde{P} \leftarrow \mathcal{O}(1^k, P_b) : \mathcal{A}(1^k, \tilde{P}, z) = b \right] \geq \frac{1}{2} + \frac{1}{q_{\mathcal{A}}(k)}$$

for these  $P_0, P_1, z$ , then  $\mathcal{E}$  succeeds with probability  $1/q_{\mathcal{E}}(k)$  in extracting a differing input for  $P_0, P_1$ , given input  $(1^k, P_0, P_1, z)$ . Roughly we will do this in three steps: (1) We first show for any  $\mathcal{O}$ -adversary  $\mathcal{A}$  that distinguishes as above, there is an  $\mathcal{O}'$ -adversary that distinguishes in the *diO* challenge w.r.t.  $\mathcal{D}'$  with similar success (making use of the *iO* security of  $\mathcal{O}'$  in the process). (2) By the assumed *diO* security of  $\mathcal{O}'$  w.r.t.  $\mathcal{D}'$ , this guarantees existence of an extractor  $\mathcal{E}'$  who succeeds in extracting differing inputs from programs sampled as in  $\mathcal{D}'$ . (3) As the final step, we tweak the extractor  $\mathcal{E}'$  to instead work for programs sampled as in  $\mathcal{D}$ , yielding the desired extractor algorithm.

From the  $\mathcal{O}$ -adversary  $\mathcal{A}$ , we begin by constructing a (non-uniform) PPT adversary  $\mathcal{A}'$  for the original obfuscator  $\mathcal{O}'$  for *diO* security w.r.t. program distribution  $\mathcal{D}'$ . Recall the distribution  $\mathcal{D}'$  corresponds to program pairs  $(P_0^z, P_1^z)$  where  $z$  has been “embedded” into programs  $P_0, P_1$  (where  $(P_0, P_1, z)$  are as in  $\mathcal{D}$ ), and with no auxiliary input.

At a high level, there are two challenges we must solve in order for  $\mathcal{A}'$  to make use of the adversary  $\mathcal{A}$ :

- First,  $\mathcal{A}'$  must generate from his challenge  $\mathcal{O}'(P_b^z)$  an appropriate value of auxiliary input  $z$  to give to  $\mathcal{A}$ . This will be done by reading the embedded value  $z$  at the “special input” of  $\mathcal{O}'(P_b^z)$ .
- Second (and slightly more complicated), he must convert his  $\mathcal{O}'$ -obfuscated program  $\mathcal{O}'(P_b^z)$  to “look like” an  $\mathcal{O}$ -obfuscation  $\mathcal{O}(P_b)$ , where recall  $\mathcal{O}(P_b) = \text{Restrict}(\mathcal{O}'(\text{pad}(\text{Expand}(P_b, 0), \text{size})), \emptyset)$ . Indeed, it is possible that  $\mathcal{A}$  may successfully distinguish given  $z$  and obfuscations of the latter type, yet may not succeed when given  $z$  and obfuscations of the former. In particular, note that the program  $\mathcal{O}'(P_b^z)$  contains information  $z$  that agrees with the auxiliary input, whereas  $\mathcal{O}(P_b)$  has only 0 in this role; since we make no assumptions on the structure of the obfuscator  $\mathcal{O}'$ , there is no direct way to “strip off” this extra information from the obfuscation  $\mathcal{O}'(P_b^z)$ . Instead, to achieve this goal, we make use of a *second layer* of obfuscation. Namely, given the (obfuscated) code  $\mathcal{O}'(P_b^z)$  in the  $\mathcal{O}'$  challenge, we can generate a second program  $M_b^0$  that has the *same input-output behavior* and size as the desired underlying program  $\text{pad}(\text{Expand}(P_b, 0), \text{size})$  that appears within  $\mathcal{O}'$  in the obfuscation  $\mathcal{O}(P_b)$ . Explicitly,

$$\left[ \text{pad}(\text{Expand}(P_b, 0), \text{size}) \right] \equiv \left[ \text{ztoZero}(\mathcal{O}'(P_b^z)) \right], \quad (9)$$

(i.e., are functionally equivalent), where recall the second program bypasses the code of  $\mathcal{O}'(P_b^z)$  on the special input  $\text{mode} = *$  and directly outputs 0 (see Definition 5.4). Indeed, both programs evaluate to 0 on any input of the form  $(x, *)$  and to  $P_b(x)$  on any input of the form  $(x, \emptyset)$ . Define  $M_b^0 := \text{ztoZero}(\mathcal{O}'(P_b^z))$ . Note that the value of  $\text{size}$  for  $\text{pad}$  was precisely chosen to be comparable with the size of the program  $M_b^0$  (see Definition 5.4).

Now, because  $M_b^0$  and  $\text{pad}(\text{Expand}(P_b, 0), \text{size})$  have identical input-output behavior (and size), by the *indistinguishability* obfuscation security of  $\mathcal{O}'$  (which, recall, supports auxiliary input; see Remark 5.3), it must be that their  $\mathcal{O}'$ -obfuscations are indistinguishable, even given  $z$ :

$$\left(\mathcal{O}'(M_b^0), z\right) \stackrel{c}{\cong} \left(\mathcal{O}'(\text{pad}(\text{Expand}(P_b, 0), \text{size})), z\right).$$

This indistinguishability must further be preserved when applying the  $\text{Restrict}(\cdot, \emptyset)$  operation to both programs:

$$\left(\text{Restrict}(\mathcal{O}'(M_b^0), \emptyset), z\right) \stackrel{c}{\cong} \left(\text{Restrict}(\mathcal{O}'(\text{pad}(\text{Expand}(P_b, 0), \text{size})), \emptyset), z\right).$$

But, note that the right-hand distribution is now precisely  $(\mathcal{O}(P_b), z)$ . We have thus succeeded in converting challenges between  $\mathcal{O}'$  and  $\mathcal{O}$ : if the original  $\mathcal{O}$ -adversary  $\mathcal{A}$  succeeded in predicting  $b$  from  $(1^k, \mathcal{O}(P_b), z)$ , he *must also succeed* when given  $(1^k, \text{Restrict}(\mathcal{O}'(M_b^0), \emptyset), z)$ , where  $\text{Restrict}(\mathcal{O}'(M_b^0), \emptyset) = \text{Restrict}(\mathcal{O}'(\text{ztoZero}(\mathcal{O}'(P_b^z))), \emptyset)$  can be directly simulated using the  $\mathcal{O}'$ -challenge  $\mathcal{O}'(P_b^z)$ .

We now formally define the new  $\mathcal{O}'$ -adversary  $\mathcal{A}'$ . On input the security parameter  $1^k$  and an obfuscated program  $\mathcal{O}'(P_b^z)$  (where the underlying program  $P_b^z$  is sampled via  $(P_0^z, P_1^z) \leftarrow \mathcal{D}'$ ;  $b \leftarrow \{0, 1\}$ ),  $\mathcal{A}'$  performs the following:

1. Evaluate  $\mathcal{O}'(P_b^z)$  on the special “ $\text{mode} = *$ ” input to learn  $z := \mathcal{O}'(P_b^z)(x, *)$ , for arbitrary  $x$ .
2. Construct a *new* program to simulate the  $\mathcal{O}$ -challenge program distribution. Namely,
  - 1: Define  $M_b^0 := \text{ztoZero}(\mathcal{O}'(P_b^z))$ .
  - 2: Sample an obfuscation  $\tilde{M}_b^0 \leftarrow \mathcal{O}'(1^k, M_b^0)$ , using the original obfuscator  $\mathcal{O}'$ .
  - 3: Define  $M_b := \text{Restrict}(\tilde{M}_b^0, \emptyset)$ .
3. Execute the adversary  $\mathcal{A}$  on input  $(1^k, M_b, z)$ , and output the resulting bit  $b'$ .

We now show that  $\mathcal{A}'$  as constructed succeeds in the  $di\mathcal{O}$  security challenge for  $\mathcal{O}'$  w.r.t.  $\mathcal{D}'$  with essentially as much advantage as  $\mathcal{A}$  in the security challenge for  $\mathcal{O}$ . Let  $\nu(\cdot)$  be the negligible function provided in the  $i\mathcal{O}$  security guarantee for  $\mathcal{O}$  (see Definition 5.2).

**Claim 5.6.** *By the  $i\mathcal{O}$  security of  $\mathcal{O}'$ , for every triple  $(P_0, P_1, z) \in \mathcal{P} \times \mathcal{P} \times \{0, 1\}^{p(k)}$  in the support of  $\mathcal{D}$ , for  $P_b^z := \text{Expand}(P_b, z)$ , and for the adversary  $\mathcal{A}'$  constructed above, it holds that*

$$\begin{aligned} & \Pr \left[ b \leftarrow \{0, 1\}; \mathcal{O}'(P_b^z) \leftarrow \mathcal{O}'(1^k, P_b^z); \text{guess}' \leftarrow \mathcal{A}'(1^k, \mathcal{O}'(P_b^z)) : \text{guess}' = b \right] \\ & \geq \Pr \left[ b \leftarrow \{0, 1\}; \mathcal{O}(P_b) \leftarrow \mathcal{O}(1^k, P_b); \text{guess} \leftarrow \mathcal{A}(1^k, \mathcal{O}(P_b), z) : \text{guess} = b \right] - \nu(k). \end{aligned}$$

*Proof.* Denote these two probabilities by  $\text{win}_{\mathcal{A}'}$  and  $\text{win}_{\mathcal{A}}$ . Suppose, to the contrary, there is a triple  $(P_0, P_1, z)$  for which  $\mathcal{A}'$  predicts  $b$  in the  $\mathcal{O}'$ -challenge with probability strictly less than  $\nu(k)$  below the probability that  $\mathcal{A}$  succeeds in predicting  $b$  in the  $\mathcal{O}$ -challenge: i.e., that  $\text{win}_{\mathcal{A}'} < \text{win}_{\mathcal{A}} - \nu(k)$ . From this, we construct a (non-uniform) adversary  $\mathcal{A}_{i\mathcal{O}}^z$  who has  $z$  hardcoded and breaks the  $i\mathcal{O}$  security of  $\mathcal{O}'$  for the pair of functionally equivalent programs constructed in Equation 9, where one

program appears within the obfuscation  $\mathcal{O}(P_b)$  as per the definition of  $\mathcal{O}$ , and the second appears within the strategy of  $\mathcal{A}'$  given his challenge  $\mathcal{O}'(P_b^z)$ .

Recall that by the definition of  $\mathcal{O}$  and of  $\mathcal{A}'$ , we have:

- $\mathcal{O}(P_b) \leftarrow \mathcal{O}(1^k, P_b)$  is equivalent to  $[\text{Restrict}(\tilde{P}, \emptyset) : \tilde{P} \leftarrow \mathcal{O}'(1^k, [\text{pad}(\text{Expand}(P_b, 0), \text{size}))]]$ .
- $\text{guess}' \leftarrow \mathcal{A}'(1^k, \mathcal{O}'(P_b^z))$  is equivalent to

$$[\text{guess}' \leftarrow \mathcal{A}(1^k, \tilde{P}, z) : z = \mathcal{O}'(P_b^z)(x, *); \tilde{P} \leftarrow \mathcal{O}'(1^k, [\text{ztoZero}(\mathcal{O}'(P_b^z))])].$$

Thus,  $\text{win}_{\mathcal{A}'}$  and  $\text{win}_{\mathcal{A}}$  correspond to the probabilities of the following experiments:

$\text{win}_{\mathcal{A}'}$ :		$\text{win}_{\mathcal{A}}$ :
$b \leftarrow \{0, 1\};$ $\mathcal{O}'(P_b^z) \leftarrow \mathcal{O}'(1^k, P_b^z);$ $z := \mathcal{O}'(P_b^z)(x, *);$ $M'_b := \text{ztoZero}(\mathcal{O}'(P_b^z));$ $\tilde{M}'_b \leftarrow \mathcal{O}'(1^k, M'_b);$ $\text{guess}' \leftarrow \mathcal{A}(1^k, \tilde{M}'_b, z);$ $: \text{guess}' = b$		$b \leftarrow \{0, 1\};$  $M_b := \text{pad}(\text{Expand}(P_b, 0), \text{size});$ $\tilde{M}_b \leftarrow \mathcal{O}'(1^k, M_b);$ $\text{guess} \leftarrow \mathcal{A}(1^k, \tilde{M}_b, z);$ $: \text{guess} = b$

If  $\text{win}_{\mathcal{A}'} < \text{win}_{\mathcal{A}} - \nu(k)$ , then in particular there must exist  $b \in \{0, 1\}$  and  $\mathcal{O}'(P_b^z) \in \mathcal{O}'(1^k, P_b^z)$  such that for the pair of programs  $M'_b := \text{ztoZero}(\mathcal{O}'(P_b^z))$  and  $M_b := \text{pad}(\text{Expand}(P_b, 0), \text{size})$ ,

$$\begin{aligned} & \Pr[\tilde{M}'_b \leftarrow \mathcal{O}'(1^k, M'_b); \text{guess}' \leftarrow \mathcal{A}(1^k, \tilde{M}'_b, z) : \text{guess}' = b] \\ & < \Pr[\tilde{M}_b \leftarrow \mathcal{O}'(1^k, M_b); \text{guess} \leftarrow \mathcal{A}(1^k, \tilde{M}_b, z) : \text{guess} = b] - \nu(k). \end{aligned}$$

But, recall that the programs  $M'_b, M_b$  are *functionally equivalent* (and of comparable size). Indeed, both accept a pair of inputs  $(x, \text{mode})$ , and satisfy  $M'_b(x, *) = M_b(x, *) = 0$  and  $M'_b(x, \emptyset) = M_b(x, \emptyset) = P_b(x)$  (in addition,  $M_b$  was padded precisely to the appropriate value of  $\text{size}$ ; see Definition 5.4). Thus, the algorithm  $\mathcal{A}$  directly gives us an adversary who breaks the  $i\mathcal{O}$  security of  $\mathcal{O}'$  for the program pair  $M'_b, M_b$ . That is, consider the (non-uniform) PPT adversary  $\mathcal{A}_{i\mathcal{O}}^z$  who has hardcoded the program pair  $M'_b, M_b$ , the bit  $b$ , and the value  $z$  from the original triple  $(P_0, P_1, z)$ . In the  $i\mathcal{O}$  security challenge,  $\mathcal{A}_{i\mathcal{O}}^z$  is given as input the security parameter  $1^k$  and a  $\mathcal{O}'$ -obfuscated program  $\tilde{M}$ , which is either generated as  $\tilde{M} \leftarrow \mathcal{O}'(1^k, M'_b)$  or as  $\tilde{M} \leftarrow \mathcal{O}'(1^k, M_b)$ .

The adversary  $\mathcal{A}_{i\mathcal{O}}^z(1^k, \tilde{M})$ :

- 1: Evaluate  $\text{guess} \leftarrow \mathcal{A}(1^k, \tilde{M}, z)$ .
- 2: If  $\text{guess} = b$ , output 1 (i.e., guess of  $M_b$ ); otherwise, output 0.

By construction, in combination with the above, the probability that  $\mathcal{A}_{i\mathcal{O}}^z$  outputs 1 in the case that  $\tilde{M} \leftarrow \mathcal{O}'(1^k, M'_b)$  is *strictly smaller* than  $\nu(k)$  less than the probability that  $\mathcal{A}_{i\mathcal{O}}^z$  outputs 1 in the case that  $\tilde{M} \leftarrow \mathcal{O}'(1^k, M_b)$ . This contradicts the  $i\mathcal{O}$  security of  $\mathcal{O}'$ . The claim follows.  $\square$

Now, suppose for some  $(P_0, P_1, z) \in \text{supp}(\mathcal{D})$  that  $\mathcal{A}$  succeeds in guessing  $b$  in the  $\mathcal{O}$ -obfuscation challenge:

$$\Pr \left[ b \leftarrow \{0, 1\}; \tilde{P} \leftarrow \mathcal{O}(1^k, P_b) : \mathcal{A}(1^k, \tilde{P}, z) = b \right] \geq \frac{1}{2} + \frac{1}{q_{\mathcal{A}}(k)}.$$

By Claim 5.6, it must thus be that  $\mathcal{A}'$  similarly succeeds in guessing  $b$  in the  $\mathcal{O}'$ -obfuscation challenge for the corresponding pair  $(P_0^z, P_1^z) \in \text{supp}(\mathcal{D}')$ :

$$\Pr \left[ b \leftarrow \{0, 1\}; \tilde{P} \leftarrow \mathcal{O}'(1^k, P_b^z) : \mathcal{A}'(1^k, \tilde{P}) = b \right] \geq \frac{1}{2} + \frac{1}{q_{\mathcal{A}}(k)} - \nu(k) \quad (10)$$

$$\geq 1/2 + 1/[2q_{\mathcal{A}}(k)]. \quad (11)$$

By the  $di\mathcal{O}$  security of  $\mathcal{O}'$ , we know that for this adversary  $\mathcal{A}'$  and polynomial  $2q_{\mathcal{A}}$ , there exists an extractor  $\mathcal{E}'$  and polynomial  $q_{\mathcal{E}'}$  such that, with all but negligible probability over  $(P_0^z, P_1^z) \leftarrow \mathcal{D}'$ , if (11) holds, then with probability at least  $1/q_{\mathcal{E}'}(k)$ , running the extractor  $(x^*, \text{mode}) \leftarrow \mathcal{E}'(1^k, P_0^z, P_1^z)$  succeeds in finding a differing input  $(x^*, \text{mode})$  for which  $P_0^z(x^*) \neq P_1^z(x^*, \text{mode})$ .

As the final step, we wish to construct from  $\mathcal{E}'$  a new extractor  $\mathcal{E}$  that instead, given a tuple  $(1^k, P_0, P_1, z)$ , outputs a differing input  $x^*$  for which the programs  $P_0$  and  $P_1$  disagree. Define the algorithm  $\mathcal{E}'$  that, on input  $(1^k, P_0, P_1, z)$ , constructs the programs  $P_0^z := \text{Expand}(P_0, z)$  and  $P_1^z := \text{Expand}(P_1, z)$ , executes the original extractor algorithm  $\mathcal{E}'$  as  $(x^*, \text{mode}) \leftarrow \mathcal{E}'(1^k, P_0^z, P_1^z)$ , and outputs the value  $x^*$ . Since by construction,  $P_0^z(x, *) = P_1^z(x, *) = z$  for every value of  $x$  (corresponding to “special” input mode), it must be that the differing input  $(x^*, \text{mode})$  generated by  $\mathcal{E}'$  satisfies  $\text{mode} = \emptyset$  (i.e., “standard” mode evaluation). That is, it must be that  $P_0^z(x^*, \emptyset) \neq P_1^z(x^*, \emptyset)$ , which means precisely that  $P_0(x^*) \neq P_1(x^*)$ .

Theorem 5.1 follows. □

**Applying the theorem to [GGHW14].** In [GGHW14], the authors propose a new “special-purpose” Turing machine obfuscation assumption, and demonstrate based on this assumption a distribution of circuits and bounded-length auxiliary input for which  $di\mathcal{O}$  cannot exist. Applying Theorem 5.1, we conclude that, assuming such special-purpose obfuscation exists, there exists a distribution of circuits that is  $di\mathcal{O}$ -unobfuscatable for any  $i\mathcal{O}$  obfuscator, *even when no auxiliary input is present*.

Intuitively, their counterexample is for  $di\mathcal{O}$  obfuscation of circuits that differ exactly on inputs corresponding to valid message-signature pairs (with respect to a secure digital signature scheme). Their “bad” auxiliary input distribution amounts to a “special” obfuscated program with side information on the signing key: the program takes as input a circuit description  $C$ , signs a hash of the circuit (to produce unique messages), and outputs the bit obtained by running  $C$  on the resulting message-signature pair. Access to this program allows an adversary to distinguish obfuscations of circuits as above, but (by their special obfuscation assumption) does not enable any efficient algorithm to extract a message-signature pair.

More formally, let  $\mathcal{H}$  be a collision-resistant hash function family, and  $(\text{Gen}, \text{Sign}, \text{Verify})$  a signature scheme (with keys denoted  $(\text{sk}, \text{vk})$ ). The “special-purpose” obfuscation assumption of [GGHW14] is given below, in Definition 5.7 and Conjecture 5.8.

**Definition 5.7** (Special-Purpose TM Class  $\{P_{H, \text{sk}}^*\}$ ). [GGHW14].

Program  $P_{H, \text{sk}}^*$ : (for hardcoded  $H \in \mathcal{H}$ , signing key  $\text{sk}$ )

Input: circuit description  $C$  of size  $s(k)$  with 1-bit output.

- 1: Compute  $m = H(C)$  and  $\sigma = \text{Sign}(\text{sk}, m)$ .
- 2: Output the bit  $C(m, \sigma, 0)$ .

**Conjecture 5.8** (Special-Purpose TM Obfuscation [GGHW14]). *There exists a deterministic signature scheme  $(\text{Gen}, \text{Sign}, \text{Verify})$ , collision-resistant hash function family  $\mathcal{H}$ , and an obfuscator  $\text{sp}\mathcal{O}$ , such that  $\text{sp}\mathcal{O}$  satisfies correctness and the following security property. There exists a negligible function  $\nu$  such that for any PPT attacker  $\mathcal{A}$ ,*

$$\Pr \left[ (\text{sk}, \text{vk}) \leftarrow \text{Gen}(1^k); H \leftarrow \mathcal{H}(1^k); \tilde{P} \leftarrow \text{sp}\mathcal{O}(1^k, P_{H, \text{sk}}^*); (m, \sigma) \leftarrow \mathcal{A}(1^k, \text{vk}, \tilde{P}) \right. \\ \left. : \text{Verify}(\text{vk}, m, \sigma) = 1 \right] \leq \nu(k),$$

where  $P_{H, \text{sk}}^*$  is the Turing machine defined in Definition 5.7.

Garg *et al.* [GGHW14] prove the following theorem.

**Theorem 5.9** ([GGHW14]). *Under the special-purpose Turing machine obfuscation conjecture (Conjecture 5.8), there exists a distribution  $\mathcal{D}$  over  $P/\text{poly} \times P/\text{poly} \times \mathcal{Z}$  of circuit pairs and bounded-length auxiliary input information for which there cannot exist secure differing-inputs obfuscation.*

Combining this with Theorem 5.1, we thus obtain the following corollary.

**Corollary 5.10.** *Under the Special-Purpose TM Obfuscation Conjecture 5.8, there exists a distribution of circuit pairs  $\mathcal{D}'$  over  $P/\text{poly} \times P/\text{poly}$  with respect to which no  $i\mathcal{O}$ -secure obfuscator can be  $di\mathcal{O}$ -secure, even in the absence of auxiliary input.*

In particular, this implies that if the special-purpose obfuscator conjecture holds, then general-purpose  $di\mathcal{O}$  cannot exist, even without auxiliary input.

**Corollary 5.11.** *Under the Special-Purpose TM Obfuscation Conjecture 5.8, there cannot exist  $di\mathcal{O}$  for  $P/\text{poly}$ , even in the absence of auxiliary input.*

*Proof.* Follows from Corollary 5.10, since any  $di\mathcal{O}$  obfuscator for  $P/\text{poly}$  is also  $i\mathcal{O}$  secure for  $P/\text{poly}$  and  $di\mathcal{O}$ -secure for any efficiently samplable distribution of circuits  $\mathcal{D}'$  over  $P/\text{poly} \times P/\text{poly}$ .  $\square$

## Acknowledgements

The authors would like to thank Kai-Min Chung for several insightful discussions.

## References

- [ABG<sup>+</sup>13] Prabhanjan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013.
- [BBP04] Mihir Bellare, Alexandra Boldyreva, and Adriana Palacio. An uninstantiable random-oracle-model scheme for a hybrid-encryption problem. In *EUROCRYPT*, pages 171–188, 2004.
- [BCCT12] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *ITCS*, pages 326–349, 2012.
- [BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for snarks and proof-carrying data. In *STOC*, pages 111–120, 2013.
- [BCP14] Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC*, pages 52–73, 2014.
- [BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 505–514, 2014.
- [BG08] Boaz Barak and Oded Goldreich. Universal arguments and their applications. *SIAM J. Comput.*, 38(5):1661–1694, 2008.
- [BGI<sup>+</sup>12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6, 2012.

- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In *Public-Key Cryptography - PKC 2014 - 17th International Conference on Practice and Theory in Public-Key Cryptography*, pages 501–519, 2014.
- [BLV06] Boaz Barak, Yehuda Lindell, and Salil P. Vadhan. Lower bounds for non-black-box zero knowledge. *J. Comput. Syst. Sci.*, 72(2):321–391, 2006.
- [BP13] Elette Boyle and Rafael Pass. Limits of extractability assumptions with distributional auxiliary input. Cryptology ePrint Archive, Report 2013/703, 2013.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM Conference on Computer and Communications Security*, pages 62–73. ACM, 1993.
- [BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures - how to sign with rsa and rabin. In *EUROCRYPT*, pages 399–416, 1996.
- [BV14] Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based FHE as secure as PKE. In *Innovations in Theoretical Computer Science, ITCS'14, Princeton, NJ, USA, January 12-14, 2014*, pages 1–12, 2014.
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security*, pages 280–300, 2013.
- [BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference*, pages 480–499, 2014.
- [CD08] Ran Canetti and Ronny Ramzi Dakdouk. Extractable perfectly one-way functions. In *ICALP (2)*, pages 449–460, 2008.
- [CD09] Ran Canetti and Ronny Ramzi Dakdouk. Towards a theory of extractable functions. In Omer Reingold, editor, *TCC*, volume 5444 of *Lecture Notes in Computer Science*, pages 595–613. Springer, 2009.
- [CGH04] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. *J. ACM*, 51(4):557–594, 2004.
- [Dam91] Ivan Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In *CRYPTO*, pages 445–456, 1991.
- [DFH12] Ivan Damgård, Sebastian Faust, and Carmit Hazay. Secure two-party computation with low communication. In *TCC*, pages 54–74, 2012.
- [GGH<sup>+</sup>13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, 2013.
- [GGHR] Sanjam Garg, Craig Gentry, Shai Halevi, and Mariana Raykova. Two-round secure MPC from indistinguishability obfuscation. In *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC*.
- [GGHW13] Sanjam Garg, Craig Gentry, Shai Halevi, and Daniel Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. Cryptology ePrint Archive, Report 2013/860, 2013.

- [GGHW14] Sanjam Garg, Craig Gentry, Shai Halevi, and Daniel Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference*, pages 518–535, 2014.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- [GK03] Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the fiat-shamir paradigm. In *FOCS*, pages 102–, 2003.
- [GKP<sup>+</sup>13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. How to run turing machines on encrypted data. In *CRYPTO (2)*, pages 536–553, 2013.
- [GLR11] Shafi Goldwasser, Huijia Lin, and Aviad Rubinfeld. Delegation of computation without rejection problem from designated verifier cs-proofs. *IACR Cryptology ePrint Archive*, 2011:456, 2011.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [GS12] Divya Gupta and Amit Sahai. On constant-round concurrent zero-knowledge from a knowledge assumption. Cryptology ePrint Archive, Report 2012/572, 2012. <http://eprint.iacr.org/>.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO (1)*, pages 75–92, 2013.
- [HR04] Chun-Yuan Hsiao and Leonid Reyzin. Finding collisions on a public road, or do secure hash functions need secret coins? In *CRYPTO*, pages 92–105, 2004.
- [HSW14] Susan Hohenberger, Amit Sahai, and Brent Waters. Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In *Advances in Cryptology - EURO-CRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 201–220, 2014.
- [HT98] Satoshi Hada and Toshiaki Tanaka. On the existence of 3-round zero-knowledge protocols. In Hugo Krawczyk, editor, *CRYPTO*, volume 1462 of *Lecture Notes in Computer Science*, pages 408–423. Springer, 1998.
- [IKO05] Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. Sufficient conditions for collision-resistant hashing. In *TCC*, pages 445–456, 2005.
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13*, pages 669–684, 2013.
- [KS98] B. Kaliski and J. Staddon. PKCS #‘: RSA cryptography specifications version 2.0, 1998.
- [Mic94] Silvio Micali. Cs proofs (extended abstracts). In *FOCS*, pages 436–453, 1994.
- [Nao03] Moni Naor. On cryptographic assumptions and challenges. In *CRYPTO*, pages 96–109, 2003.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 475–484, 2014.

- [Val08] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In Ran Canetti, editor, *TCC*, volume 4948 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2008.