

Bootstrapping Obfuscators via Fast Pseudorandom Functions

Benny Applebaum*

October 26, 2013

Abstract

We show that it is possible to upgrade an obfuscator for a weak complexity class **WEAK** into an obfuscator for arbitrary polynomial size circuits, assuming that the class **WEAK** can compute pseudorandom functions. Specifically, under standard intractability assumptions (e.g., hardness of factoring, Decisional Diffie-Hellman, or Learning with Errors), the existence of obfuscators for \mathbf{NC}^1 or even \mathbf{TC}^0 implies the existence of general-purpose obfuscators for \mathbf{P} . Previously, such a bootstrapping procedure was known to exist under the assumption that there exists a fully-homomorphic encryption whose decryption algorithm can be computed in **WEAK**. Our reduction works with respect to virtual black-box obfuscators and relativizes to ideal models.

1 Introduction

General-purpose program obfuscation allows us to transform an arbitrary computer program into an “unintelligible” form while preserving its functionality. The latter property is formalized via the notion of *Virtual Black-Box* which asserts that the code of the obfuscated program reveals nothing more than what can be learned via oracle access to its input-output behavior. The seminal result of [6] shows that general purpose obfuscation is impossible in the standard model. Nevertheless, in a sequence of recent exciting works [9, 7, 5], it was shown that general-purpose obfuscation can be achieved in idealized models such as the Generic Colored Matrix Model, or the Generic Graded Encoding model.

All these works share a similar outline. First it is shown how to use the idealized model to obfuscate a weak complexity class such as \mathbf{NC}^1 , and then the weak obfuscator is bootstrapped into a general-purpose obfuscator for arbitrary polynomial-size circuits.¹ The bootstrapping step in all these works employs a fully homomorphic encryption [10] whose decryption algorithm can be implemented in \mathbf{NC}^1 . While recent constructions of FHE (e.g., [8]) make the latter assumption reasonable, the existence of FHE is still a strong public-key assumption and it is natural to ask whether it can be relaxed. Indeed, \mathbf{NC}^1 obfuscation already seems to put us at the “Heights of Cryptomania” [11], and so one may suspect whether the extra power of FHE is really needed for bootstrapping.

*School of Electrical Engineering, Tel-Aviv University, benny.applebaum@gmail.com. Supported by Alon Fellowship, ISF grant 1155/11, Israel Ministry of Science and Technology (grant 3-9094), and GIF grant 1152/2011.

¹The class \mathbf{NC}^1 is the class of polynomial-size circuits with logarithmic depth and bounded fan-in gates.

1.1 Our Results

In this note we show that bootstrapping can be based on a “Minicrypt” type assumption.

Theorem 1.1 (main theorem – informal). *Assume that the complexity class **WEAK** can compute a pseudorandom function (PRF). Then an obfuscator for **WEAK** (in some idealized model) can be bootstrapped into an obfuscator for every polynomial-size circuit family.*

(See Theorem 4.1 for a formal statement.) Since practical and theoretical implementations of PRFs tend to be highly efficient, we can instantiate the theorem with relatively low complexity classes. For example, by relying on PRFs constructions from [16, 4], we derive the following corollary.

Corollary 1.2. *Assuming the hardness of factoring, Decisional Diffie-Hellman, or Learning with Errors, the following holds. If \mathbf{TC}^0 can be obfuscated (in some idealized model), then every polynomial-size circuit family can be obfuscated as well.²*

Tightness. One may ask whether the PRF assumption in Theorem 1.1 can be further relaxed and replaced with the existence of a weaker primitive such as a one-way function or a pseudorandom generator. We note that the answer seems to be negative as such primitives can be computed in \mathbf{NC}^0 (under standard assumptions [3]), a class which is learnable and therefore trivially obfuscatable in the standard model. Therefore, such a strengthening of Theorem 1.1 would contradict the impossibility results of [6].

In fact, we do not expect to prove a statement like Corollary 1.2 for classes lower than \mathbf{TC}^0 , as \mathbf{TC}^0 is the lowest complexity class for which the impossibility results of [6] apply. More generally, [6] essentially show that if a complexity class **WEAK** contains a PRF then it cannot be obfuscated in the standard model. Our results complement the picture by showing that if such a class can be obfuscated in some idealized model, then everything can be obfuscated. This suggests the existence of a zero-one law: if an idealized model admits non-trivial obfuscation (i.e., for some class which is non-obfuscatable in the standard model) then it admits general purpose obfuscation for all (polynomial-size) circuits.

Virtual Black-Box vs. Indistinguishability Obfuscation. Our results hold with respect to the strongest security notion of *Virtual Black-Box* (VBB) obfuscation [6] relative to some ideal model. An alternative (weaker) notion of security is *indistinguishability Obfuscation*. The latter notion is highly attractive as it may be achievable in the standard model (no impossibility results are known), and, quite surprisingly, it suffices for a wide range of applications [17]. In this work we focus on VBB obfuscators, as we believe that when *constructing* obfuscators it is best to strive for the strongest form of security. (See detailed discussion in [5].) We do not know whether our results apply in the indistinguishability setting, and leave this question for future research. Interestingly, the FHE-based bootstrapping works in both settings [9, 7].

²The class \mathbf{TC}^0 is the class of all Boolean circuits with constant depth and polynomial size, containing only unbounded-fan in AND gates, OR gates, and majority gates. This class is a subclass of \mathbf{NC}^1 .

1.2 Techniques

Our main technical tool is randomized encoding (RE) of functions [15, 3]. Intuitively, a function $f(x)$ is encoded by a randomized function $\hat{f}(x; r)$ if the distribution $\hat{f}(x)$, induced by a random choice of r , reveals nothing but $f(x)$. Formally, a sample from $\hat{f}(x)$ can be *decoded* to $f(x)$, and vice versa, given $f(x)$ one can efficiently *simulate* the distribution $\hat{f}(x)$ – so the functions are essentially “equivalent”. REs become non-trivial (and useful) when their complexity is smaller than the complexity of f . This “equivalence” between a complicated function f to a simpler encoding \hat{f} , was exploited in various applications to reduce a complex task to a simpler one. (See the surveys [1] and [14].) We can adopt a similar approach in our context as well.

In order to obfuscate a function f taken from a family \mathcal{F} let us obfuscate its low-complexity encoding \hat{f} and release the latter obfuscated program composed with the decoder algorithm. For this to work, let us assume the existence of universal decoder and universal simulator that work uniformly for all functions in \mathcal{F} . (This can be guaranteed by encoding the evaluation function of the collection \mathcal{F} which maps a circuit of f and an input x to the value $f(x)$.)

Unfortunately, this approach is somewhat problematic as the encoding \hat{f} employs internal randomness. One potential solution is to treat the randomness as an additional input and let the user of the obfuscated program choose it. While decoding still succeeds, this solution fails to be secure. Once the randomness r is revealed, the function f can be fully recovered. (Technically, universal simulation cannot be achieved anymore.) Another (flawed) solution, is to fix some secret randomness r and obfuscate the mapping $x \mapsto \hat{f}(x; r)$. Unfortunately, the privacy of the encoding holds only when fresh randomness is being used, and one can easily recover the circuit of f when r is fixed.

The problem is solved via the extra use of a PRF. Specifically, we choose a PRF $h \xleftarrow{R} \mathcal{H}$ and obfuscate the function $\hat{f}(x; h(x))$. The security of the PRF ensures that this function behaves essentially as a standard encoding whose internal randomness is freshly chosen in each invocation, and so simulation succeeds. By using the low-complexity encoding from [2], the complexity of \hat{f} is dominated by the complexity of the PRF $\mathcal{H} \in \mathbf{WEAK}$ (assuming that the class \mathbf{WEAK} satisfies some basic closure properties), and one can prove Theorem 1.1.

2 Preliminaries

Circuit families. A family of polynomial-size boolean circuits \mathcal{F} is an infinite sequence of circuit families $\{\mathcal{F}_n\}_{n \in \mathbb{N}}$ where for every $n \in \mathbb{N}$ the family \mathcal{F}_n consists of boolean circuits with n inputs, $m(n)$ outputs, and circuit size $\ell(n)$ where m, ℓ are polynomials in n . For such a family there is always a universal efficient evaluator F such that for every length parameter n , circuit $f \in \mathcal{F}_n$ and input $x \in \{0, 1\}^n$, we have that $F(f, x) = f(x)$.

Pseudorandom functions [12]. Let $\mathcal{H} = \{\mathcal{H}_n\}_{n \in \mathbb{N}}$ be a family of polynomial-size boolean circuits and let \mathcal{K} be a PPT sampling algorithm that on input 1^n samples a circuit in \mathcal{H}_n . (The probability distribution induced by \mathcal{K} is not necessarily uniform.) We say that \mathcal{H} is a *pseudorandom function family* (PRF) if for every (non-uniform) efficient oracle-aided adversary \mathcal{A} we have that

$$\left| \Pr_{h \leftarrow \mathcal{K}(1^n)} [\mathcal{A}^h(1^n) = 1] - \Pr_{R_n} [\mathcal{A}^{R_n}(1^n) = 1] \right| \leq \text{neg}(n),$$

where R_n is a uniformly chosen function with the same input and output lengths as the functions in \mathcal{H}_n . We say that the PRF is secure relative to a probabilistic oracle \mathcal{M} (an “idealized model”) if the above holds when the adversary is given an additional oracle access to \mathcal{M} . To simplify notation, we will typically make the sampler implicit and write $h \stackrel{R}{\leftarrow} \mathcal{H}_n$ with the understanding that the distribution is induced by some efficient sampler.

2.1 Randomized Encoding of Functions

Let F be a polynomial-time computable function that maps n bits to $m(n)$ bits. Intuitively, a randomized function \hat{F} is an “encoding” of F if for every input x the distribution $\hat{F}(x)$ reveals the value of $F(x)$ but no other additional information. We formalize this via the notion of *computationally private randomized encoding* from [2].

Definition 2.1 (Computational randomized encoding). *Let $F : \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}$ be an efficiently computable function and let \hat{F} be an efficiently computable randomized function. We say that \hat{F} is a computational randomized encoding of F (or encoding for short), if there exist an efficient decoder algorithm D and an efficient probabilistic simulator algorithm Sim that satisfy the following:*

- **Perfect correctness.** *For every n and every input $x \in \{0, 1\}^n$,*

$$D(1^n, \hat{F}(x)) = F(x).$$

- **Computational privacy.** *For every non-uniform efficient oracle-aided adversary \mathcal{A} we have*

$$\left| \Pr[\mathcal{A}^{\hat{F}_n}(1^n) = 1] - \Pr[\mathcal{A}^{\text{Sim}(1^n, F_n(\cdot))}(1^n) = 1] \right| \leq \text{neg}(n), \quad (1)$$

where \hat{F}_n and F_n are the restrictions of \hat{F} and F to n -bit inputs, and both oracles are probabilistic functions (and fresh randomness is used in each invocation).

We say that the encoding is secure relative to a probabilistic oracle \mathcal{M} (an “idealized model”) if privacy holds when the adversary is given an additional oracle access to \mathcal{M} .

Encoding Collections. We encode a family of polynomial-size boolean circuits \mathcal{F} by encoding its evaluation algorithm $F(f, x)$. Specifically, for every $f \in \mathcal{F}$ we define the randomized function $\hat{f}(x; r) = \hat{F}(f, x; r)$ where \hat{F} is the encoder of F . By definition, the decoder D and the simulator Sim of F apply universally for all $f \in \mathcal{F}$. Formally, for every n and $f \in \mathcal{F}_n$:

$$D(1^{|x|}, \hat{f}(x)) = f(x) \quad \forall x \in \{0, 1\}^n.$$

Also for every non-uniform efficient oracle-aided adversary \mathcal{A} and every sequence of functions $\{f_n\}$ and their encodings $\{\hat{f}_n\}$ we have that

$$\left| \Pr[\mathcal{A}^{\hat{f}_n}(1^n) = 1] - \Pr[\mathcal{A}^{\text{Sim}(1^n, f_n(\cdot))}(1^n) = 1] \right| \leq \text{neg}(n). \quad (2)$$

(The oracles in (2) are simply the restriction of the oracles in (1) to inputs of the form (f_n, \cdot) and so (2) follows immediately from (1).)

2.2 Obfuscation

Definition 2.2 (Virtual Black-Box Obfuscator [6]). *Let $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ be a family of polynomial-size boolean circuits. An obfuscator \mathcal{O} for \mathcal{F} is a PPT algorithm which maps a circuit $f \in \mathcal{F}_n$ to a new circuit $[f]$ (not necessarily in \mathcal{F}) such that the following properties hold:*

1. **Preserving Functionality.** *For every $n \in \mathbb{N}$, every $f \in \mathcal{F}_n$ and every input $x \in \{0, 1\}^n$*

$$\Pr_{[f] \stackrel{R}{\leftarrow} \mathcal{O}(f)} [[f](x) \neq f(x)] \leq \text{neg}(n).$$

2. **Polynomial slowdown.** *There exists a polynomial p such that for every $n \in \mathbb{N}$ and $f \in \mathcal{F}_n$ the circuit $\mathcal{O}(f)$ is of size at most $p(|f|)$.*
3. **Virtual Black-Box.** *For every (non-uniform) efficient adversary \mathcal{A} there exists a (non-uniform) efficient simulator Sim such that for every n and every $f \in \mathcal{F}_n$:*

$$\left| \Pr[\mathcal{A}(\mathcal{O}(f)) = 1] - \Pr[\text{Sim}^f(1^{|f|}, 1^n) = 1] \right| \leq \text{neg}(n).$$

We define obfuscation with respect to a probabilistic oracle \mathcal{M} (an “idealized model”) by granting \mathcal{O} and \mathcal{A} an oracle access to \mathcal{M} . (The circuit f cannot have oracle gates to \mathcal{M} .) Properties (1) and (2) should hold for every possible coins of \mathcal{M} . A complexity class \mathbf{C} is obfuscatable if there exists an efficiently computable mapping that maps every efficiently computable function family $\mathcal{F} \in \mathbf{C}$ (represented by its evaluator F) to an obfuscator \mathcal{O} for \mathcal{F} .

3 Our Reduction

Let \mathcal{F} be a family of polynomial-size boolean circuits with an evaluator F . We will construct an obfuscator \mathcal{O} for \mathcal{F} based on the following ingredients. (1) An encoding \hat{F} of F ; (2) a pseudorandom function family \mathcal{H} where the output length of functions in \mathcal{H}_n equals to the length of the random input of \hat{F}_n ; and (3) a weak obfuscator $\text{weak}\mathcal{O}$ for the circuit family $\mathcal{G} = \{\mathcal{G}_n\}$ where \mathcal{G}_n contain all circuits of the form

$$g_{f,h} : x \mapsto \hat{F}(f, x; h(x)), \quad \forall f \in \mathcal{F}_n, h \in \mathcal{H}_n.$$

Construction 3.1. *Given a circuit $f \in \mathcal{F}_n$ the obfuscator \mathcal{O} does the following:*

- *Sample a random $h \stackrel{R}{\leftarrow} \mathcal{H}$ and obfuscate $g_{f,h}$ by $[g] := \text{weak}\mathcal{O}(g_{f,h})$.*
- *Output the circuit $[f] = \text{D} \circ [g]$ where D is the RE decoder and \circ denotes function composition.*

It is easy to verify that $[f]$ preserves the functionality of f .

Lemma 3.2. *The obfuscator \mathcal{O} is functionality preserving.*

Proof. Fix some x and $h \in \mathcal{H}$, and let us condition on the event that the weak obfuscator preserves the functionality, namely, $[g](x) = g_{f,h}(x)$. Then, by the correctness of the encoding, we have

$$[f](x) = \text{D}([g](x)) = \text{D}(g_{f,h}(x)) = \text{D}(\hat{F}(f, x; h(x))) = f(x).$$

Since the weak obfuscator is correct with all but negligible probability the claim follows. □

In the next section, we will prove that the obfuscator is secure.

Lemma 3.3. *The obfuscator \mathcal{O} satisfies the Virtual Black-Box property. Furthermore, the proof relativizes, namely, if the weak obfuscator, the PRF, and the randomized encoding are secure relative to an oracle \mathcal{M} , then so is the obfuscator \mathcal{O} .*

3.1 Security (Proof of Lemma 3.3)

Let \mathcal{A} be an efficient adversary for the new obfuscator \mathcal{O} . Our goal is to construct a simulator Sim that simulates \mathcal{A} . For this aim, let us first define an adversary \mathcal{B} for the weak obfuscator $\text{weak}\mathcal{O}$ as follows. Given an obfuscated circuit $[f]$, the adversary \mathcal{B} applies \mathcal{A} to the circuit $\text{D} \circ [f]$ where D is the (universal) decoder of the encoding. Let weakSim be the simulator of $\text{weak}\mathcal{O}$ which simulates the adversary \mathcal{B} , and let reSim be the (universal) RE simulator.

We define the simulator $\text{Sim}^f(1^n)$ for the adversary \mathcal{A} as follows. Invoke the oracle-aided weak simulator $\text{weakSim}^g(1^n)$, and whenever weakSim makes a new oracle query $x \in \{0, 1\}^n$, answer it with $\text{reSim}(1^n, f(x))$, where the latter is computed via the help of the oracle f . If x was previously queried, respond with the same answer as before.

Fix some $f \in \mathcal{F}_n$. We will prove that

$$\left| \Pr[\mathcal{A}(\mathcal{O}(f)) = 1] - \Pr[\text{Sim}^f(1^n) = 1] \right| \leq \varepsilon_1(n) + \varepsilon_2(n) + \varepsilon_3(n) \quad (3)$$

where ε_1 (resp., $\varepsilon_2, \varepsilon_3$) upper-bounds the distinguishing advantage of efficient adversaries against the weak obfuscator (resp., against the PRF, against the encoding). Through the proof, we simplify notation by omitting the unary input 1^n ; also for a binary random variable Y we write $\Pr[Y]$ to denote $\Pr[Y = 1]$.

Let $\hat{f}(x; r)$ be the circuit that computes the encoding of $f(x)$, i.e., $\hat{f}(x; r) = \hat{F}(f, x; r)$. For a function $h \in \mathcal{H}$, let \hat{f}_h denote the circuit that computes $\hat{f}(x; h(x))$. Then, by definition, for every h , we have that

$$\Pr[\mathcal{A}(\mathcal{O}(f; h))] = \Pr[\mathcal{B}(\text{weak}\mathcal{O}(\hat{f}_h))]. \quad (4)$$

Also, by the VBB property of the weak obfuscator, for every h we have that

$$\left| \Pr[\mathcal{B}(\text{weak}\mathcal{O}(\hat{f}_h))] - \Pr[\text{weakSim}^{\hat{f}_h}] \right| \leq \varepsilon_1(n). \quad (5)$$

By relying on the security of the PRF we will prove the following claim.

Claim 3.4.

$$\left| \Pr_{h \xleftarrow{R} \mathcal{H}_n} [\text{weakSim}^{\hat{f}_h}] - \Pr[\text{weakSim}^{\hat{f}}] \right| \leq \varepsilon_2(n),$$

where $\hat{f}(\cdot)$ is viewed as a randomized function and repeated queries to this function are answered consistently based on the first answer.

Proof. If the claim does not hold, we can break the PRF as follows. Let M^R be an oracle aided adversary which calls weakSim and whenever weakSim makes a query x to its oracle, M answers the query with $\hat{f}(x; R(x))$. Observe that if R is a truly random function then M accepts with probability exactly $\Pr[\text{weakSim}^{\hat{f}}]$. On the other hand, if the oracle $R \xleftarrow{R} \mathcal{H}_n$ then the acceptance probability is exactly $\Pr_{h \xleftarrow{R} \mathcal{H}_n} [\text{weakSim}^{\hat{f}_h}]$. It follows that M breaks the security of the PRF, and the claim follows. \square

Finally, we rely on the privacy of the randomized encoding to prove the following claim.

Claim 3.5.

$$\left| \Pr[\text{weakSim}^{\hat{f}}] - \Pr[\text{Sim}^f] \right| \leq \varepsilon_3(n),$$

where $\hat{f}(\cdot)$ is viewed as a randomized function and repeated queries to this function are answered consistently based on the first answer.

Proof. Recall that $\text{Sim}^{f(\cdot)}$ is simply $\text{weakSim}^{\text{reSim}(f(\cdot))}$, where repeated queries are answered consistently. Therefore, if the claim does not hold, we can use weakSim to distinguish between the randomized functions $\text{reSim}(f(\cdot))$ and $\hat{f}(\cdot)$, in contradiction to the privacy of the RE (Eq. 2). (The distinguisher will simply invoke weakSim and will answer repeated queries based on the first answer.) \square

The lemma (i.e., Eq. 3) now follows from Eqs. 4 and 5 and Claims 3.4 and 3.5. The “furthermore” part follows by noting that the proof relativizes. \square

4 Main Result

Fix some circuit complexity class **WEAK** with the following basic properties: (1) The class **WEAK** contains the class \mathbf{NC}^0 ; (2) (Closure under concatenation) If each output bit of a multi-output function f is computable in **WEAK** then so is f ; (3) (Closure under composition) if $f \in \mathbf{WEAK}$ and $g \in \mathbf{WEAK}$ then $g \circ f$ is in **WEAK**, where \circ denotes function composition.

We can now prove our main theorem.

Theorem 4.1. *Assume that there exists a one-way function and a pseudorandom function \mathcal{H} in **WEAK**. If **WEAK** can be obfuscated then every family of polynomial-size circuits can be obfuscated.*

Proof. First, we claim that, under the above assumption, any efficiently computable function F has an encoding \hat{F} computable in **WEAK**. In [2] it was shown, based on Yao’s garbled circuit technique, that $F(x)$ can be encoded by an encoding $\hat{F}(x; r)$ which is reducible to a minimal-stretch pseudorandom generator $G : \{0, 1\}^\kappa \rightarrow \{0, 1\}^{\kappa+1}$ via a non-adaptive \mathbf{NC}^0 reduction. Namely, $\hat{F}(x; r)$ can be written as $g(x, r, G(r_1), \dots, G(r_t))$ where g is an \mathbf{NC}^0 function and the r_i ’s are sub-blocks of the random string r . Such a PRG is also reducible to a one-way function via a (non-adaptive) \mathbf{NC}^0 reduction by [3, 13]. As a result, under the theorem’s assumption, the resulting encoding can be written as a composition of an \mathbf{NC}^0 function with a function in **WEAK** which results in a **WEAK** function.

Next, we observe that, since **WEAK** is closed under concatenation, the assumption implies the existence of a PRF in **WEAK** whose output length is equal to the randomness complexity of the encoding. (By using direct product, one can transform a **WEAK** PRF with a single output bit into a new **WEAK** PRF with arbitrary polynomial number of output bits.) It follows that the circuit family \mathcal{G} (from Construction 3.1) can be written as a composition of two **WEAK** functions, and the theorem follows from Lemmas 3.2 and 3.3. \square

Remark. We assume nothing on the complexity of the *sampler* of the PRF, and therefore the sampler can preprocess the function $h \in \mathcal{H}$. (This preprocessing is exploited in fast implementations of PRFs [16, 4].) As a result, the additional one-wayness assumption does not seem to follow from the fact that \mathcal{H} is in **WEAK**.³

Under standard intractability assumptions, one can instantiate **WEAK** with \mathbf{NC}^1 or even \mathbf{TC}^0 . Indeed, the existence of PRFs in \mathbf{TC}^0 can be based on the hardness of factoring, the DDH assumption or the intractability of lattice/learning problems [16, 4], and the existence of one-way functions in \mathbf{TC}^0 (or even in \mathbf{NC}^0) follow from these assumptions as well [3]. Therefore Corollary 1.2 follows from Theorem 4.1.

Acknowledgement. We thank Alon Rosen and Zvika Brakerski for useful discussions.

References

- [1] B. Applebaum. Randomly encoding functions: A new cryptographic paradigm - (invited talk). In *ICITS*, pages 25–31, 2011.
- [2] B. Applebaum, Y. Ishai, and E. Kushilevitz. Computationally private randomizing polynomials and their applications. *Computational Complexity*, 15(2):115–162, 2006.
- [3] B. Applebaum, Y. Ishai, and E. Kushilevitz. Cryptography in \mathbf{NC}^0 . *SIAM J. Comput.*, 36(4):845–888, 2006.
- [4] A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *Advances in Cryptology - EUROCRYPT*, 2012.
- [5] B. Barak, S. Garg, Y. T. Kalai, O. Paneth, and A. Sahai. Protecting obfuscation against algebraic attacks. Cryptology ePrint Archive, Report 2013/631, 2013. <http://eprint.iacr.org/>.
- [6] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. *J. of the ACM*, 59(2), 2012.
- [7] Z. Brakerski and G. N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. Cryptology ePrint Archive, Report 2013/563, 2013. <http://eprint.iacr.org/>.
- [8] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS*, 2011.
- [9] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *IEEE 54th Annual Symposium on Foundations of Computer Science, FOCS*, 2013.

³The one-wayness assumption becomes redundant if the mapping $(k, x) \mapsto h_k(x)$ (where $h_k = \mathcal{K}(1^n; k)$) is computable in **WEAK**.

- [10] C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC*, pages 169–178, 2009.
- [11] C. Gentry. Encrypted messages from the heights of cryptomania. In *TCC*, pages 120–121, 2013.
- [12] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *J. of the ACM*, 33:792–807, 1986.
- [13] I. Haitner, O. Reingold, and S. P. Vadhan. Efficiency improvements in constructing pseudo-random generators from one-way functions. *SIAM J. Comput.*, 42(3):1405–1430, 2013.
- [14] Y. Ishai. Randomization techniques for secure computation. In M. Prabhakaran and A. Sahai, editors, *Secure Multi-Party Computation*, volume 10 of *Cryptology and Information Security Series*, pages 222–248. IOS press, Amsterdam, 2012.
- [15] Y. Ishai and E. Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *IEEE 41st Annual Symposium on Foundations of Computer Science, FOCS*, pages 294–304, 2000.
- [16] M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. *J. of the ACM*, 51(2):231–262, 2004.
- [17] A. Sahai and B. Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. Cryptology ePrint Archive, Report 2013/454, 2013. <http://eprint.iacr.org/>.