On Necessary Padding with IO

Justin Holmgren*

Abstract

We show that the common proof technique of padding a circuit before IO obfuscation is sometimes necessary. That is, assuming indistinguishability obfuscation (IO) and one-way functions exist, we define samplers Sam_0, which outputs (aux_0, C_0), and Sam_1, which outputs (aux_1, C_1) such that:

- The distributions $(\mathsf{aux}_0, \mathsf{i}\mathcal{O}(C_0))$ and $(\mathsf{aux}_1, \mathsf{i}\mathcal{O}(C_1))$ are perfectly distinguishable.
- For padding $s = \mathsf{poly}(\lambda)$, the distributions $(\mathsf{aux}_0, \mathsf{i}\mathcal{O}(C_0\|0^s))$ and $(\mathsf{aux}_1, \mathsf{i}\mathcal{O}(C_1\|0^s))$ are computationally indistinguishable.

We note this refutes the recent "Superfluous Padding Assumption" of Brzuska and Mittelbach[BM15].

*Email: holmgren@csail.mit.edu. This work was done while the author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant #CNS-1523467. This research was also supported by NSF Eager CNS1347364, NSF Frontier CNS1413920, the Simons Foundation (agreement dated June 5, 2012), Air Force Laboratory FA875011-20225, and Lincoln Lab PO7000261954.

1 Introduction

Proofs of security for cryptographic constructions using indistinguishability obfuscation (IO) typically show that the obfuscations of two circuits are indistinguishable if those circuits are artificially padded to a larger size. This is a consequence of the fact that IO guarantees indistinguishability of two obfuscated circuits only if the input circuits are of equivalent sizes. But because of the ubiquitous hybrid argument style of proof, the two circuits must typically be padded not only to the size of larger circuit, but rather to the size of the largest circuit in the hybrid argument.

At this point, one may wonder whether this padding is an artifact of our security proof, or whether it is really necessary. In this work, we show that in the presence of arbitrary auxiliary information, there are distributions of equally-sized pairs of circuits, whose obfuscations are indistinguishable only if sufficiently padded.

1.1 Preliminaries

We assume familiarity with puncturable pseudorandom functions [BW13, BGI14]. In particular, we will use the fact that if one-way functions exist, then there is a puncturable PRF family $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda>0}$ in which each $f \in \mathcal{F}_\lambda$ maps $\{0,1\}^\lambda$ to $\{0,1\}$.

We also assume the existence of an indistinguishability obfuscator [GGH+13]. This is a p.p.t. Turing machine $\mathsf{i}\mathcal{O}$ such that:

- $\mathsf{i}\mathcal{O}(C, 1^\lambda)$ outputs a circuit which is functionally equivalent to C.
- If C and C' are two circuits of the same size and same functionality, then the advantage of any p.p.t. adversary in distinguishing $\mathsf{i}\mathcal{O}(C, 1^\lambda)$ from $\mathsf{i}\mathcal{O}(C', 1^\lambda)$ is negligible in λ .

We will frequently omit the security parameter 1^λ as an argument of $\mathsf{i}\mathcal{O}$.

We will write $C\|0^s$ to denote a *padded* version of the circuit C, which is of size $|C| + s$.

1.2 Techniques

We want to show a pair of distributions (aux_0, C_0) and (aux_1, C_1) such that:

- $(\mathsf{aux}_0, \mathsf{i}\mathcal{O}(C_0))$ is distinguishable from $(\mathsf{aux}_1, \mathsf{i}\mathcal{O}(C_1))$
- For some padding p, $(\mathsf{aux}_0, \mathsf{i}\mathcal{O}(C_0\|0^p))$ is indistinguishable from $(\mathsf{aux}_1, \mathsf{i}\mathcal{O}(C_1\|0^p))$.

In our construction, C_0 and C_1 are defined simply as circuits which evaluate a (puncturable) PRF. aux_0 and aux_1 are (sufficiently padded) obfuscated circuits, each of which have this same PRF inside. aux_b takes a "small" circuit as input, and checks whether this circuit agrees with the PRF on a "large" set of test inputs. If it does, then aux_b outputs b. Otherwise, aux_b outputs 0.

"Small" and "large" are chosen so that the obfuscated C_0 or C_1 is small, but not large.

The first bullet is easy; the slightly tricky part is to show that when C_0 and C_1 are padded to be large (even before obfuscation is applied), then $(\mathsf{aux}_0, i\mathcal{O}(C_0))$ is indistinguishable from $(\mathsf{aux}_1, i\mathcal{O}(C_1))$. We show this in a three-step hybrid argument, starting with $(\mathsf{aux}_b, i\mathcal{O}(C_b))$ for arbitrary b.

1. The PRF in aux_b and in C_b is punctured on the whole set of test inputs, and the corresponding test values are hard-coded. This is indistinguishable by $i\mathcal{O}$.
2. These test values are replaced by truly random bits. This is indistinguishable by the puncturable PRF security.
3. Now there statistically does not exist any small circuit which agrees with all the test values. So aux_b is replaced by the obfuscated all zero function, which is indistinguishable by $i\mathcal{O}$.

This last hybrid distribution on $(\mathsf{aux}_b, i\mathcal{O}(C_b))$ is independent of the bit b.

1.3 Related Work

We note a similarity between our work and the previous work of Goldwasser and Kalai [GK05], which amongst other things shows the impossibility of VBB-obfuscating pseudorandom functions given auxiliary information. They use the fact that oracle access to a pseudo-entropic circuit C, does not reveal how to find a small circuit that agrees with C. In particular, suppose one is given an obfuscated circuit which tests whether a (small) input circuit agrees with C, and if so outputs a secret bit b. Now one can compute the bit b given any obfuscation of C, but not given black-box access to C.

Our result modifies this argument to assume only that the underlying obfuscators are indistinguishability obfuscators instead of VBB obfuscators. We show that one can compute the bit b given any sufficiently small obfuscation of C (i.e. $i\mathcal{O}(C)$), but not given an obfuscation of a functionally equivalent C' which has been padded to be larger (i.e. $i\mathcal{O}(C\|0^s)$).

In personal communication, Nir Bitansky points out that the results of [BCC+14] may yield an alternative proof of our result, and that ordinary PRFs (not necessarily puncturable) and witness encryption (which is implied by IO) suffice to prove our main result.

2 Main Result

We now prove our main result more formally. Let q be a polynomial such that $q(\lambda)$ bounds the size of $i\mathcal{O}(f, 1^\lambda)$, when f is sampled from the punctural PRF family $\mathcal{F}_\lambda$.

We now define a pair of algorithms $(\mathsf{Sam}_0, \mathsf{Sam}_1)$. Sam_b will output a pair (aux_b, C_b). The algorithm $\mathsf{Sam}_b(1^\lambda)$ first samples a puncturable PRF $f \leftarrow \mathcal{F}_\lambda$.

Sam_b first computes $\mathsf{aux}_b \leftarrow \mathsf{i}\mathcal{O}(A_{b,f})$, where $A_{b,f}$ is described in Algorithm 1, and Sam_b then computes $C_b = B_f$, where B_f is described in Algorithm 2.

Input: Circuit C of size $q(\lambda)$
Data: Bit b, PPRF f
if $C(i) = f(i)$ *for all* $i \in \{0, \ldots, q(\lambda) + \lambda\}$ **then**
  return b
else
  return 0
end

Algorithm 1: Circuit $A_{b,f}$

Input: $x \in \{0,1\}^\lambda$
Data: PPRF f
return $f(x)$.

Algorithm 2: Circuit B_f

Claim 1. *The distributions* $(\mathsf{aux}_0, \mathsf{i}\mathcal{O}(C_0))$ *and* $(\mathsf{aux}_1, \mathsf{i}\mathcal{O}(C_1))$ *are perfectly distinguishable when* $(\mathsf{aux}_b, C_b) \leftarrow \mathsf{Sam}_b(1^\lambda)$.

Proof. aux_b computes the same function as $A_{b,f}$ and C_b is B_f. The claim follows simply because $A_{b,f}(\mathsf{i}\mathcal{O}(B_f)) = A_{b,f}(B_f) = b$. □

Claim 2. *For some padding* p_b*, the distributions* $(\mathsf{aux}_0, \mathsf{i}\mathcal{O}(C_0 \| 0^{p_b}))$ *and* $(\mathsf{aux}_1, \mathsf{i}\mathcal{O}(C_1 \| 0^{p_b}))$ *are computationally indistinguishable when* $(\mathsf{aux}_b, C_b) \leftarrow \mathsf{Sam}_b(1^\lambda)$.

Proof. Let p_b be padding so that $\left|B_f \| 0^{p_b}\right| = |B_f^1|$, where B_f^1 is described in Algorithm 4. We define three indistinguishable hybrid distributions H_b^1 through H_b^3 such that:

- H_b^1 is indistinguishable from $(\mathsf{aux}_b, \mathsf{i}\mathcal{O}(C_b \| 0^{p_b}))$ when $(\mathsf{aux}_b, C_b) \leftarrow \mathsf{Sam}_b(1^\lambda)$.
- H_b^3 is independent of b.

Hybrid H_b^1: Hybrid H_b^1 is sampled by first sampling a PPRF $f \leftarrow \mathcal{F}_\lambda$, and puncturing it on the set $\{0, \ldots, q(\lambda) + \lambda\}$ to obtain the punctured PRF f'. Let $y_i = f(i)$ for $i \in \{0, \ldots, q(\lambda) + \lambda\}$.

H_b^1 then consists of $(\mathsf{i}\mathcal{O}(A_{b,f}^1), \mathsf{i}\mathcal{O}(B_f^1))$. The circuit $A_{b,f}^1$ is described in Algorithm 3, with the values y_i hard-coded, and is padded to be as large as $A_{b,f}$. The circuit B_f^1 is described in Algorithm 4, with the PPRF f' and the values y_i hard-coded.

Hybrid H_b^2: Hybrid H_b^2 is sampled identically to H_b^1, but each y_i is sampled uniformly at random.

Input: Circuit C of size $q(\lambda)$
Data: Bit b, values y_i
if $C(i) = y_i$ *for all* $i \in \{0, \ldots, q(\lambda) + \lambda\}$ **then**
  return b
else
  return 0
end

Algorithm 3: Circuit $A^1_{b,f}$

Input: $x \in \{0,1\}^\lambda$
Data: Punctured PPRF f', values y_i for $i \in \{0, \ldots, q(\lambda) + \lambda\}$
if $x \in \{0, \ldots, q(\lambda) + \lambda\}$ **then**
  return y_x;
else
  Return $f'(x)$;
end

Algorithm 4: Circuit B^1_f

Hybrid H^3_b: In hybrid H^3_b, the circuit $A^1_{b,f}$ is replaced with the constant zero function, appropriately padded.

Claim 3. $H^1_b \approx \mathsf{Sam}_b(1^\lambda)$.

Proof. This follows from the security of $\mathsf{i}\mathcal{O}$: the obfuscated circuits have the same functionality and size in both H^1_b and $\mathsf{Sam}_b(1^\lambda)$. □

Claim 4. *Hybrid* $H^2_b \approx H^1_b$.

Proof. This follows from the pseudorandomness of the punctured PRF f' at the (selectively) punctured set $\{0, \ldots, q(\lambda) + \lambda\}$. □

Claim 5. *Hybrid* $H^3_b \approx H^2_b$.

Proof. This follows from the security of $\mathsf{i}\mathcal{O}$. A simple counting argument implies that with high probability (at least $1 - 2^{-\lambda}$), there is no circuit C of size $q(\lambda)$ such that $C(i) = y_i$ for all $i \in \{0, \ldots, q(\lambda) + \lambda\}$. Thus the circuit $A^1_{b,f}$ with truly random y_i's is functionally equivalent to the constant zero circuit with high probability. □

This completes the proof of Claim 2. □

3 Extensions

3.1 Variants of the Superfluous Padding Assumption

One possible restriction on Sam_0 and Sam_1, proposed by [BM15] as a weaker assumption, requires the marginal distribution of aux_0 to be the same as the marginal distribution of aux_1. While this does not hold for our counterexample, it can be easily modified to have this property. Rather than having aux_b output the bit b, aux_b outputs a random string r. On input r, C_b outputs b.

The proof techniques above, when applied to this modified construction, show how to move to a hybrid where aux_b is independent of r. We can then apply a standard injective PRG trick to make C_b independent of r and of b.

3.2 Implication About Double Obfuscation

The necessity of superfluous padding implies a surprising result. If $\mathsf{i}\mathcal{O}$ increases the size of circuits, then there are efficiently sampleable distributions on (aux_0, C_0) and (aux_1, C_1) such that $(\mathsf{aux}_0, \mathsf{i}\mathcal{O}^k(C_0))$ is indistinguishable from $(\mathsf{aux}_1, \mathsf{i}\mathcal{O}^k(C_1))$ for some integer $k > 1$, but $(\mathsf{aux}_0, \mathsf{i}\mathcal{O}(C_0))$ is perfectly distinguishable from $(\mathsf{aux}_1, \mathsf{i}\mathcal{O}(C_1))$. This follows from our construction of Sam_0 and Sam_1 because the inner $k-1$ obfuscations are functionally equivalent to padding.

References

[BCC+14] Nir Bitansky, Ran Canetti, Henry Cohn, Shafi Goldwasser, Yael Tauman Kalai, Omer Paneth, and Alon Rosen. The impossibility of obfuscation with auxiliary input or a universal simulator. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part II*, volume 8617 of *Lecture Notes in Computer Science*, pages 71–89. Springer, 2014.

[BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *Public-Key Cryptography PKC 2014*, volume 8383 of *Lecture Notes in Computer Science*, pages 501–519. Springer Berlin Heidelberg, 2014.

[BM15] Christina Brzuska and Arno Mittelbach. Universal computational extractors and the superfluous padding assumption for indistinguishability obfuscation. Cryptology ePrint Archive, Report 2015/581, 2015. `http://eprint.iacr.org/`.

[BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors,

Advances in Cryptology - ASIACRYPT 2013, volume 8270 of *Lecture Notes in Computer Science*, pages 280–300. Springer Berlin Heidelberg, 2013.

[GGH+13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 40–49, 2013.

[GK05] Shafi Goldwasser and Yael Tauman Kalai. On the impossibility of obfuscation with auxiliary input. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 553–562. IEEE, 2005.