Black-Box Separations and Their Adaptability to the Non-uniform Model

Ahto Buldas[1,2,3] and Margus Niitsoo[4,⋆]

[1] Cybernetica AS, Mäealuse 2/1, 12618 Tallinn, Estonia
[2] Tallinn University of Technology, Raja 15, 12618 Tallinn, Estonia
[3] Guardtime AS, Tammsaare tee 60, 11316 Tallinn, Estonia
[4] University of Tartu, Liivi 2, 50409 Tartu, Estonia

Abstract. Oracle separation methods are used in cryptography to rule out black-box reductions between cryptographic primitives. It is sufficient to find an oracle relative to which the base primitive exists but there are no secure instances of the constructed primitive. It is often beyond our current reach to construct a fixed oracle with such properties because it is difficult to prove the existence of secure base primitives. To overcome this gap, randomized oracles are used to create random base primitives that are secure on average. After that, a fixed oracle is extracted from the probability distribution by using non-constructive probabilistic arguments and the countability of the set of adversaries. Such extraction only applies to uniform reductions because the set of non-uniform adversaries is not countable. We study how to adapt oracle separation results to the non-uniform model. The known separation techniques are capable of ruling out the so-called fully black-box reductions and a certain strong form of semi black-box reductions also in the non-uniform model. We study how to go beyond the barrier of strong semi black-box reductions and show that this is possible by using random oracles with auxiliary advice. For that end, we prove a conjecture of Unruh (2007) about pre-sampling being a sufficient substitute for advice for any oracle distribution.

1 Introduction

Complex cryptographic protocols are often built from simpler building blocks called primitives. Usually, the security of such protocols is proved based solely on the security guarantees of the original primitives independent of their actual implementation details. Such constructions are called *black-box reductions*. To date, almost all security proofs for efficient cryptographic constructions utilize black-box reductions.

Although back-box reductions are extremely useful cryptographic tools, there exist limits on where they can be applied. There are many known cases for which it is proved that such a reduction cannot exist. This usually means that a very clever proof construction is necessary if the reduction can be achieved at all. As very few of these clever constructions are known, the power and limits of black-box reductions are of a rather special interest to many people working in the field.

⋆ This research was supported by the European Regional Development Fund through the Estonian Center of Excellence in Computer Science, EXCS, and by Estonian Research Council's personal research grant PUT-2.

C. Boyd and L. Simpson (Eds.): ACISP 2013, LNCS 7959, pp. 152–167, 2013.

The first separation result involving black-box reductions was given in 1989 by Impagliazzo and Rudich [7]. They showed that there are no such reductions from key agreement protocols to one-way functions. Their seminal paper was followed by a long line of similar types of results [5,6,10]. The approach was even generalized by Kim, Simon and Tetali [8] to give bounds on reduction efficiency. Though all these results leave open the existence of more specific security proofs they are still valuable hardness indicators as they rule out the most obvious approaches.

Non-existence of black-box reductions is commonly shown by oracle separation techniques. In complexity theory, oracle separation has been widely applied to prove limits of the proof techniques capable of establishing set-theoretical inclusions between complexity classes. For example, with oracle separation one can easily show that diagonal arguments are insufficient to solve the famous $\mathbf{P}$ vs $\mathbf{NP}$ question. This is done by showing two oracles—one relative to which the conjecture holds and another for which it does not. In cryptography, oracle separation is used to argue that it is impossible to securely implement a primitive or a protocol $\mathcal{P}$ given only black-box access to a secure low-level primitive f as an instance of a class of primitives $\mathcal{Q}$. This is done by defining an oracle so that f remains secure even if the adversary can access the oracle but any instance $\mathcal{P}^f$ of the protocol $\mathcal{P}$ being considered is insecure relative to the oracle.

In classical separation results of complexity theory, oracles are defined as fixed functions that always behave the same way. In cryptographic separations, it is often hard to explicitly define a fixed separation oracle. For example, if one wishes that one way functions exist relative to the oracle, an explicit functional description of such function should then be given (as a part of the oracle). This is however still unreachable for the state of the art computer science—the existence of one way functions is conjectured but not yet proved. So, in cryptographic separations we often assume that oracles are chosen randomly from certain probability distributions. We then prove that the separation statements hold *on average* and then argue that there exists a particular choice of the oracle for which the statements hold. For example, one-way functions indeed exist in the *random oracle model* because random functions are *one-way on average* [7].

It would then seem natural that the oracle separation results would also be stated with respect to the random oracles. However, as the classical separation theorems are adopted from the classical model, the authors still try to make their oracle choice deterministic. Such an *oracle extraction* approach, though, has a big limitation—it usually requires that the number of adversaries is countable, and hence the whole approach is usable only in the *uniform* model, where the adversaries are ordinary Turing machines.

To avoid the countability argument, Buldas, Laur and Niitsoo [2] proposed an alternative oracle extraction approach where the oracle extraction step is unnecessary. Rather than trying to extract a suitable deterministic oracle from a probability distribution, they assumed that there exists a black-box reduction (that works for every f) and derived a contradiction by assuming the probabilistic separation condition and the average (over f) version of the reduction condition. We call the method they introduced the *averaging approach*. They proved that the averaging approach is capable of showing that there are no semi black-box reductions between two primitives. However, they were able to do this only for a strong version of semi black-box reduction where the simulator A does not depend on the instance f of the source primitive.

In this paper, we give an overview on both the traditional oracle extraction based separation and the averaging-based separation techniques. For each type of the reduction, we outline the main steps of the separation and point out the steps where the countability assumption is used. Thereby, we determine the exact reason why the separation fails in the non-uniform model. We achieved the following results:

- The traditional oracle extraction approach still works for the strong semi black-box reductions, because the separation oracle can be chosen for a fixed adversary A and a fixed simulator S. Similar to the averaging approach, this is not so for the weak semi black-box reductions, and from this viewpoint, the averaging approach has no advantage over the traditional methods.
- We derive sufficient average-based separation criterions for the weak semi black-box and the variable semi black-box reductions. It turns out that proving the security condition of the oracle separation for the weak semi black-box reduction is equivalent of proving the security of a cryptographic construction in a model where the adversary is given a function $\varphi(\mathcal{O})$ of the oracle.
- We generalize the results of Unruh [11] about oracles with auxiliary strings so that they would apply to arbitrary oracle distributions.

2 Notation

By $x \leftarrow \mathcal{D}$ we mean that x is chosen randomly according to a distribution $\mathcal{D}$. We use the Big Oh notation for describing asymptotic properties of functions. In particular, $f(k) = O(1)$ means that f is bounded and $f(k) = k^{-\omega(1)}$ means that $f(k)$ decreases faster than any polynomial, i.e., f is *negligible*. A Turing machine M is *poly-time* if it runs in time $k^{O(1)}$, where k denotes the input size that is mostly referred to as the *security parameter*.

By an *oracle Turing machine* we mean an incompletely specified Turing machine S that comprises calls to *oracles*. The description can be completed by defining the oracle as a function $\mathcal{O} : \{0,1\}^* \to \{0,1\}^*$. In this case, the machine is denoted by $S^{\mathcal{O}}$. The function $y \leftarrow \mathcal{O}(x)$ does not have to be computable but has a conditional running time $t(x)$, which does not necessarily reflect the actual amount of computations needed to produce y from x. The running time of $S^{\mathcal{O}}$ comprises the conditional running time of oracle calls—each call $\mathcal{O}(x)$ takes $t(x)$ steps. We assume that all the oracles $\mathcal{O}$ are *poly-time*, that is $t(x) = \{x\}^{O(1)}$, where $\{x\}$ denotes the bit-length of x. Note that though the classical complexity-theoretic oracles only require a single step, this more general notion is appropriate in cryptography where oracles often model abstract adversaries with running time t. We say that S is a *poly-time oracle machine* if $S^{\mathcal{O}}$ runs in poly-time, whenever $\mathcal{O}$ is poly-time. By a *non-uniform* poly-time oracle machine we mean an ordinary poly-time oracle machine S together with a family $\mathcal{A} = \{a_k\}_{k\in\mathbb{N}}$ of (advice) bit-strings a_k with length $k^{O(1)}$. For any oracle $\mathcal{O}$ and any input x, it is assumed that $S^{\mathcal{O}}(x)$ has access to the advice string $a_{\{x\}}$. Usually, the advice strings are omitted for simplicity, but their presence must always be assumed when S is non-uniform. One of the most important facts about non-uniform machines is that there are uncountably many of them, whereas the set of ordinary Turing machines is countable.

3 Basic Lemmas

Lemma 1 (Probabilistic Argument). *Let $\mathcal{F}$ be a probability space and $\mathcal{P}$ be a predicate function. Then* $\Pr_{f\leftarrow\mathcal{F}}[\mathcal{P}(f)]>0 \;\Rightarrow\; \exists f\colon \mathcal{P}(f)$.

Lemma 2 (Countability Argument). *Let $\mathcal{F}$ be a probability space and $\mathcal{P}(f,A)$ be a predicate function where A varies over all poly-time Turing machines, then*

$$\forall_{pol} A\colon \Pr_{f\leftarrow\mathcal{F}}[\mathcal{P}(f,A)]=1 \quad\Rightarrow\quad \Pr_{f\leftarrow\mathcal{F}}\left[\forall_{pol} A\colon \mathcal{P}(f,A)\right]=1 \ .$$

Proof. Countable intersection of measure one sets is a measure one set. □

Lemma 3 (Borel-Cantelli). *Let $\{\mathsf{E}_i\}_{i\in\mathbb{N}}$ be a countable set of events and E_∞ be the event that infinitely many of these events happen. If $\sum_n \Pr[\mathsf{E}_n]<\infty$ then $\Pr[\mathsf{E}_\infty]=0$.*

Proof. Indeed, let $B_k=\bigcup_{n=k}^{\infty}\mathsf{E}_k$. If $x\in\mathsf{E}_\infty$ then $x\in\cap_k B_k$, because otherwise x only belongs to a finite sequence $\mathsf{E}_1,\ldots,\mathsf{E}_{k-1}$ of events. Hence, $\mathsf{E}_\infty\subseteq\cap_k\mathsf{B}_k$ and $\Pr[\mathsf{E}_\infty]\le\Pr[\cap_k\mathsf{B}_k]\le\Pr[\mathsf{B}_k]$. From $\sum_n\Pr[\mathsf{E}_n]<\infty$ it follows that for every $\epsilon>0$ there is k such that $\sum_{n=k}^{\infty}\Pr[\mathsf{E}_k]<\epsilon$. Thus, $\Pr[\mathsf{E}_\infty]\le\Pr[\mathsf{B}_k]=\Pr[\cup_{n=k}^{\infty}\mathsf{E}_k]\le\sum_{n=k}^{\infty}\Pr[\mathsf{E}_k]<\epsilon$, which implies $\Pr[\mathsf{E}_\infty]=0$. □

Lemma 4 (Negligible Average Argument). *Let $\mathcal{F}$ be a distribution so that for every $f\leftarrow\mathcal{F}$ there is a real-valued function $\delta_f\colon\mathbb{N}\to[0,1]$. If $\mathop{\mathbf{E}}_{f\leftarrow\mathcal{F}}[\delta_f(k)]=\varepsilon(k)=k^{-\omega(1)}$, then $\delta_f(k)=k^{-\omega(1)}$ for measure one of f's.*

Proof. As $\Pr_{f\leftarrow\mathcal{F}}\left[\delta_f(k)>k^2\cdot\varepsilon(k)\right]\le k^{-2}$ and $\Pr_{f\leftarrow\mathcal{F}}\left[\delta_f(k)\le k^2\cdot\varepsilon(k)\right]\ge 1-k^{-2}$ by Markov inequality, we define E_k as the event that $\delta_f(k)>k^2\cdot\varepsilon(k)$. Now we use the Borel-Cantelli lemma and $\sum_k\Pr[\mathsf{E}_k]\le\sum_k k^{-2}<\infty$ to imply

$$\Pr_{f\leftarrow\mathcal{F}}\left[\text{"}\delta_f(k)>k^2\cdot\varepsilon(k)\text{ for infinitely many }k\text{-s"}\right]=\Pr[\mathsf{E}_\infty]=0 \ .$$

Thus, for measure one of f's: $\exists k_0\forall k>k_0:\delta_f(k)\le k^2\cdot\varepsilon(k)=k^{-\omega(1)}$. □

Lemma 5 (Overwhelming Average Argument). *Let $\mathcal{F}$ be a distribution so that for every $f\leftarrow\mathcal{F}$ there is a function $\delta_f\colon\mathbb{N}\to[0,1]$. If $\mathop{\mathbf{E}}_{f\leftarrow\mathcal{F}}[\delta_f(k)]=1-k^{-\omega(1)}$, then $\delta_f(k)=1-k^{-\omega(1)}$ for measure one of f's.*

Proof. $\mathop{\mathbf{E}}_{f\leftarrow\mathcal{F}}[1-\delta_f(k)]=1-\mathop{\mathbf{E}}_{f\leftarrow\mathcal{F}}[\delta_f(k)]=1-(1-k^{-\omega(1)})=k^{-\omega(1)}$, which by Lemma 4 implies that $1-\delta_f(k)=k^{-\omega(1)}$ for measure one of f's. □

Lemma 6. *There exist quantities $\delta_i(k)=k^{-\omega(1)}$ for which $\mathop{\mathbf{E}}_{i}[\delta_i(k)]\ne k^{-\omega(1)}$.*

Proof. Let $I=\{1,2,\ldots\}$ and $p_i=\frac{6}{\pi^2 i^2}$ for all $i\in I$. Then $\sum_{i\in I}p_i=1$. For all $i\in I$ we define the function δ_i by $\delta_i(k)=\delta_{ik}$, where δ_{ik} is the Kronecker delta. Now we define a probability space on $\{\delta_i\}_{i\in I}$ such that $\mathsf{Pr}[\delta_i]=p_i$ for all $i\in I$. Note that $\delta_i(k)=k^{-\omega(1)}$ for all $i\in I$ but the average of all δ_i-s is non-negligible, because $\mathop{\mathbf{E}}_{i}[\delta_i(k)]=\frac{6}{\pi^2}\cdot k^{-2}=k^{-O(1)}\ne k^{-\omega(1)}$. □

4 Primitives

Complex cryptographic constructions can often be decomposed into a few fundamental building blocks that are called *primitives*. One is usually interested in proving the security of constructions based solely on the properties of the underlying primitives. Reingold et al. [9] give a formal definition by considering primitives as families of functions of type $f\colon \{0,1\}^* \to \{0,1\}^*$ along with a matching between these functions and Turing machines implementing them. Indeed, for many common primitives such as one-way permutations or collision-resistant hash functions this formalization coincides with our intuition—an instance of a primitive is indeed just one function.

In some more complicated cases that actually have more than one type of functionality, it may make more sense to define a primitive as a tuple of functions. However, we can usually concatenate these functions into one single function – we just use the first few input bits to tell the function which sub-function we want to use. This means that we can still formalize the primitive as one single function, although it may be a little counter-intuitive.

A primitive is usually defined in terms of functional requirements that the instances of the primitive must satisfy before it makes sense to talk about their security. These requirements, though, are just syntactic and have nothing to do with security. For example, every permutation is an instance of the one-way permutation primitive, however, it does not have to be a secure instance.

In cryptography we also have to define the *security* of primitives. Reingold et al. [9] define security as a relation between primitives and Turing machines that possibly break them. That is, a machine either breaks the primitive or not. In this work, we use a more specific (but still sufficiently general) definition of security given in [2], where the breakage advantage is a real-valued function that also depends on the security parameter k which is usually tied to the actual input (or output) lengths of the primitive:

Definition 1 (Primitives, Adversaries and Advantage). *A* primitive $\mathcal{P}$ *is a set of functions of type* $f\colon \{0,1\}^* \to \{0,1\}^*$. *Primitives have an* advantage function $\mathrm{ADV}_k^{\mathcal{P}}(\cdot,\cdot)$, *which given as input the security parameter* $k \in \mathbb{N}$, *an instance* f *of* $\mathcal{P}$, *and an oracle Turing machine* $\mathsf{A}^{\mathcal{O}}$ *(an* adversary*) returns a real number* $\mathrm{ADV}_k^{\mathcal{P}}(\mathsf{A}^{\mathcal{O}}, f) \in [0,1]$ *(the* advantage *of* $\mathsf{A}^{\mathcal{O}}$*). The function* $\mathrm{ADV}_k^{\mathcal{P}}(\cdot, f)$ *is extended to probabilistic Turing machines by taking the average over their randomness strings*[1]. *We say that* $\mathsf{A}^{\mathcal{O}}$ breaks *an instance* f *of* $\mathcal{P}$ *if* $\mathrm{ADV}_k^{\mathcal{P}}(\mathsf{A}^{\mathcal{O}}, f) \neq k^{-\omega(1)}$. *If for a fixed oracle* $\mathcal{O}$ *no probabilistic poly-time oracle Turing machine* $\mathsf{A}^{\mathcal{O}}$ *breaks* f *then* f *is said to be* secure relative to $\mathcal{O}$.

We emphasize that our definition says nothing about the efficiency of f. The function may even be non-computable, as long as the advantage that can be gained by any adversary is negligible. In practice, one needs an instantiation of a primitive that is both efficient and secure. Commonly, it is required that we can compute the function f with a (uniform) poly-time Turing machine for f to be called *efficient*.

[1] Each fixed randomness string gives a deterministic poly-time Turing machine for which $\mathrm{ADV}_k^{\mathcal{P}}()$ is already defined.

5 Black-Box Reductions

Reductions capture the relations between different security notions. A primitive $\mathcal{P}$ can be reduced to a primitive $\mathcal{Q}$ if there exists a construction that given a secure instance of the primitive $\mathcal{Q}$ yields a secure instance of $\mathcal{P}$. Most common cryptographic reductions are *black-box reductions*, where an instance f of a primitive $\mathcal{Q}$ is treated as an atomic object in the construction and in the security proof. In this work, we consider four sub-notions of black-box reductions: fully-black box reductions, strong and weak semi black box reductions, and variable semi black box reductions.

The first step towards classification of black-box reductions was made by Reingold, Trevisan and Vadhan [9]. They showed a complex hierarchy of 7 types of different reductions. Our classification is based on their work but leaves out some of the more general reductions and introduces one that actually arises quite often in practice. Also, we assume that the reduction construction is deterministic whereas the original hierarchy uses probabilistic Turing machines everywhere. This is necessary for reductions between deterministic primitives as the reduction cannot be allowed any randomization in that case. If we consider randomized primitives, G can usually be made deterministic even when it is randomized in essence – the key idea is to use the oracle as a source of randomness. This approach was already used by Impagliazzo and Rudich [7] in the first paper about oracle separations in cryptology.

In the first three definitions, the construction of a derived primitive is independent of the actual implementation of $\mathcal{P}$, whereas the construction itself may depend on the implementation in the last definition. The reduction in question is *uniform* if all oracle machines in the corresponding definition are uniform, otherwise the reduction is *non-uniform*. We assume that the construction G is always deterministic but the adversaries are allowed to be randomized.

Definition 2 (Fully black-box reduction). *A fully black-box reduction $\mathcal{P} \overset{\mathrm{f}}{\Longrightarrow} \mathcal{Q}$ is determined by two poly-time oracle machines G and S, satisfying the next two conditions:*
(C) *If f implements $\mathcal{Q}$ then G^f implements $\mathcal{P}$.*
(S) *For every instance $f \in \mathcal{Q}$, if A breaks G^f (as $\mathcal{P}$) then $S^{\mathsf{A},f}$ breaks f (as $\mathcal{Q}$).*

In brief, we must to provide a universal oracle algorithm S that can handle all successful adversaries to establish a fully black-box reduction. So-called *semi-black-box* reductions weaken this restriction by allowing for some specialization in the security proofs. The degree to which the specialization can go is different for different authors. We give two possible definitions, one just slightly stronger than the other.

Definition 3 (Strong semi black-box reduction). *A strong semi-black-box reduction $\mathcal{P} \overset{ss}{\Longrightarrow} \mathcal{Q}$ is a poly-time oracle machine G, satisfying the next two conditions:*
(C) *If f correctly implements $\mathcal{Q}$ then G^f correctly implements $\mathcal{P}$.*
(S) *For all poly-time oracle machines A there exists a poly-time oracle machine B such that for every instance $f \in \mathcal{Q}$ if A^f breaks G^f then B^f breaks f.*

Definition 4 (Weak semi black-box reduction). *By weak semi-black-box reduction $\mathcal{P} \overset{ws}{\Longrightarrow} \mathcal{Q}$ we mean a poly-time oracle machine G, satisfying the next two conditions:*[2]

[2] This was the reduction given by Reingold et al. [9] as the semi-black-box reduction.

(C) *If f correctly implements $\mathcal{Q}$ then G^f correctly implements $\mathcal{P}$.*
(S) *For every instance $f \in \mathcal{Q}$ and a poly-time oracle machine* A*, there exists a poly-time oracle machine* B *such that if* A^f *breaks G^f then* B^f *breaks f.*

The difference between the two reduction types is very subtle. In the strong case, the construction of B may non-constructively depend on A and G but it has to be universal for all $f \in \mathcal{Q}$. In the weak case, such a universal B might not exist – the construction of B may depend on f as well as on A and G. This subtle difference is however extremely important, as it seems to create a theoretical boundary for at least one general separation method we show later.

In both semi-black-box reductions G must be universal for all valid instances $f \in \mathcal{Q}$ and as such, specific properties of an instance f cannot be used in the construction. Variable semi-black-box reductions [3] weaken even this restriction so that the constructions of both G and B may depend on the instance f. However, such constructions must exist for all instances of $\mathcal{Q}$ regardless of the actual efficiency of f. If we restrict f in the following definition to efficiently implementable instances of $\mathcal{Q}$, then we get the definition of *white-box reductions*, which is the most general type of reductions.

Definition 5 (Variable semi black-box reduction). *We say that there is a* variable semi-black-box reduction $\mathcal{P} \overset{v}{\Longrightarrow} \mathcal{Q}$ *iff for any correct implementation f of $\mathcal{Q}$:*
(C) *there exists a poly-time oracle machine G^f that correctly implements $\mathcal{P}$;*
(S) *for every instance $f \in \mathcal{Q}$ and for any poly-time oracle machine* A*, there exists a poly-time oracle machine* B *such that if* A^f *breaks G^f, then* B^f *breaks f.*

These reduction types form a linear hierarchy with fully black-box reductions being the strongest and white-box reductions being the weakest. Existence of a reduction of one type also trivially implies the existence of reductions of all weaker types. This is important as it means that non-existence of a weaker reduction also implies non-existence of all stronger reductions.

6 Oracle-Extraction Based Separation

Showing the non-existence of black-box reductions of a primitive $\mathcal{P}$ to a primitive $\mathcal{Q}$ by oracle separation involves two major steps:

- (Breakage argument) Define an oracle $\mathcal{O}$ relative to which there is no secure $\mathcal{P}$;
- (Security argument) Show that there is secure $\mathcal{Q}$ relative to $\mathcal{O}$.

In cryptography, oracle separation is almost never done with an explicitly defined oracle. Instead of that, the existence of a suitable oracle is proved by using probabilistic arguments, i.e. it is shown that an oracle with the desired properties can be extracted from a probability space of oracles. So, the first step of an oracle separation is to define a probability distribution $\mathcal{O} \leftarrow \mathcal{F}$ of oracles and show that there is secure instance $f^{\mathcal{O}}$ of $\mathcal{Q}$ but no instance $G^{\mathcal{O}}$ of $\mathcal{P}$ is secure relative to $\mathcal{O}$. By using the so-called oracle embedding techniques first introduced by Simon [10], the secure instance f of $\mathcal{Q}$ can be identified with the oracle $\mathcal{O}$, i.e. the oracle distribution is $f \leftarrow \mathcal{F}$. To show the security argument in the oracle-extraction based separation techniques we show that:

[3] They were called $\forall\exists$-semi-black-box by Reingold et al.

s_1: Every fixed poly-time adversary S that uses f as an oracle can break f only with negligible success, on average, i.e. $\forall\limits_{\text{pol}} S\colon \mathop{\mathbf{E}}\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(S^f, f)\right] = k^{-\omega(1)}$.

s_2: $\Pr\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(S^f, f) = k^{-\omega(1)}\right] = 1$, i.e. for measure one of f's, no poly time S can break f (by s_1 and Lemma 4).

s_3: For measure one of oracles f, no poly-time S^f can break f better than with negligible success, i.e. $\Pr\limits_{f\leftarrow\mathcal{F}}\left[\forall\limits_{\text{pol}} S\colon\ \text{ADV}_k(S^f, f) = k^{-\omega(1)}\right] = 1$ (by s_2, Lemma 2).

To show the breakage argument we have to show that:

b_1: Every instance G^f of $\mathcal{P}$ can be broken by a poly-time A with overwhelming probability, i.e. $\forall\limits_{\text{pol}} G \exists\limits_{\text{pol}} A\colon \mathop{\mathbf{E}}\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(A^f, G^f)\right] = 1 - k^{-\omega(1)}$.

b_2: For every instance G^f of $\mathcal{P}$ there is a poly-time A, so that A^f breaks G^f for measure one of f's, i.e. $\forall\limits_{\text{pol}} G \exists\limits_{\text{pol}} A\colon \Pr\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(A^f, G^f) = 1 - k^{-\omega(1)}\right] = 1$ (b_1, Lemma 5).

b_3: For measure one of oracles f, every G^f can be broken by a poly-time machine A, i.e. $\Pr\limits_{f\leftarrow\mathcal{F}}\left[\forall\limits_{\text{pol}} G \exists\limits_{\text{pol}} A\colon\ \text{ADV}_k(A^f, G^f) = 1 - k^{-\omega(1)}\right] = 1$ (b_2, Lemma 2).

Finally, by combining these two sets of measure one, we have that measure one of oracles satisfy both the breakage and the security conditions, which means that by the probabilistic argument (Lemma 1) there exists a fixed separation oracle. Note that (b_1) cannot be replaced with the weaker statement $\mathop{\mathbf{E}}\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(A^f, G^f)\right] = k^{-O(1)}$, because it does not imply that there is an f for which $\text{ADV}_k(A^f, G^f) = k^{-O(1)}$. There exists a counterexample (Lemma 6) for which $\text{ADV}_k(A^f, G^f) = k^{-\omega(1)}$ for all f.

Table 1. Reduction types and separation conditions for oracle extraction based separations. The quantifier $\forall\limits_{\text{of}}$ means that the quantified variable varies over all oracle functions.

Type	Reduction Condition	Separation Condition
Fully bb	$\exists\limits_{\text{of}} p \exists\limits_{\text{pol}} S \forall f \forall A$: $A\,\mathbf{br}\,p(f) \Rightarrow S^{A,f}\,\mathbf{br}\,f$	$\forall\limits_{\text{of}} p \forall\limits_{\text{pol}} S \exists\mathcal{F}$: $\mathop{\mathbf{E}}\limits_{f,A\leftarrow\mathcal{F}}[\text{ADV}_k(A, p(f))] = 1 - k^{-\omega(1)}$ $\mathop{\mathbf{E}}\limits_{f,A\leftarrow\mathcal{F}}\left[\text{ADV}_k(S^{f,A}, f)\right] = k^{-\omega(1)}$
Strong Semi bb	$\exists\limits_{\text{of}} p \forall\limits_{\text{pol}} A \exists\limits_{\text{pol}} S \forall f$: $A^f\,\mathbf{br}\,p(f) \Rightarrow S^f\,\mathbf{br}\,f$	$\forall\limits_{\text{of}} p \exists\limits_{\text{pol}} A \forall\limits_{\text{pol}} S \exists\mathcal{F}$: $\mathop{\mathbf{E}}\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(A^f, p(f))\right] = 1 - k^{-\omega(1)}$ $\mathop{\mathbf{E}}\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(S^f, f)\right] = k^{-\omega(1)}$
Weak Semi bb	$\exists\limits_{\text{of}} p \forall\limits_{\text{pol}} A \forall f \exists\limits_{\text{pol}} S$: $A^f\,\mathbf{br}\,p(f) \Rightarrow S^f\,\mathbf{br}\,f$	$\forall\limits_{\text{of}} p \exists\limits_{\text{pol}} A \exists\mathcal{F}$: $\mathop{\mathbf{E}}\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(A^f, p(f))\right] = 1 - k^{-\omega(1)}$ $\forall\limits_{\text{pol}} S \mathop{\mathbf{E}}\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(S^f, f)\right] = k^{-\omega(1)}$ Countability argument for S
Variable Semi bb	$\forall f \exists\limits_{\text{pol}} G \forall\limits_{\text{pol}} A \exists\limits_{\text{pol}} S$: $A^f\,\mathbf{br}\,G^f \Rightarrow S^f\,\mathbf{br}\,f$	$\exists\mathcal{F}$: $\forall\limits_{\text{pol}} G \exists\limits_{\text{pol}} A \mathop{\mathbf{E}}\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(A^f, G^f)\right] = 1 - k^{-\omega(1)}$ $\forall\limits_{\text{pol}} S \mathop{\mathbf{E}}\limits_{f\leftarrow\mathcal{F}}\left[\text{ADV}_k(S^f, f)\right] = k^{-\omega(1)}$ Countability arguments for G and S

The steps s_1, s_2, b_1 and b_2 also apply to non-uniform reductions, whereas the steps s_3 and b_3 do not, because there are uncountably many non-uniform machines S and G.

In order to still apply the separation technique in the non-uniform model, there are several ways to go on. The first way is to use the fact that stronger types of black-box reductions may need weaker separation arguments and perhaps there is no need to use the countability argument, i.e. the existence of the separation oracle may be deducible without extraction. Indeed, if we examine the negations of the reduction statements for all four types of black-box reductions (Table 1), we observe that the non-existence of fully black-box and the strong semi black-box reductions can be proven without the countability argument. The main reason is that the oracle distribution $\mathcal{F}$ must be defined for a particular choice of A and S. Formally, this means that in the separation condition, the quantifier $\exists \mathcal{F}$ stands after the quantifiers $\underset{\text{pol}}{\exists}\, A$ and $\underset{\text{pol}}{\forall}\, S$ (Table 1).

It also turns out in a natural way that in order to show non-existence of fully black-box reductions, we may use two separate oracles A and f, where the secure instance (f) of $\mathcal{Q}$ only has access to f, while the adversary S has access to both oracles. This fact was first pointed out by Hsiao and Reyzin [6]. In Table 1, we list the separation conditions for all four types of reductions. All proofs are given in Appendix A.

To conclude, the oracle extraction based separation techniques are applicable to the *strong semi black-box reductions* (and hence also for the *fully black-box reductions*) but not for the *weak semi black-box* and *variable semi black-box reductions*, because the latter would require the countability argument.

7 Averaging-Based Separation

In practical separations, both the security- and the breakage assumption are probabilistic, i.e. involve an average success over the oracle. The reduction condition (for fully black-box reduction) is deterministic and has the form:

$$\textsc{Adv}_k^{\mathcal{P}}(A, G^f) \neq k^{-\omega(1)} \quad \Rightarrow \quad \textsc{Adv}_k^{\mathcal{Q}}(S^{f,A}, f) \neq k^{-\omega(1)} \ . \tag{1}$$

For showing that there are no fully black-box reductions of $\mathcal{P}$ to $\mathcal{Q}$, we have to derive a contradiction based on the reduction condition (1) and the separation conditions:

$$\text{(S)} \ \mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}} \left[\textsc{Adv}_k^{\mathcal{P}}(A, G^f)\right] = 1 - k^{-\omega(1)}$$

$$\text{(B)} \ \underset{\text{pol}}{\forall} S\colon \ \mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}} \left[\textsc{Adv}_k^{\mathcal{Q}}(S^{f,A}, f)\right] = k^{-\omega(1)} \ .$$

To derive a contradiction from (S), (B), and (1), the traditional approach focuses on conditions (S) and (B) and tries to derive a negation of (1) from these two conditions. This is done by using oracle extraction, i.e. by extracting a fixed oracles f and A from $\mathcal{F}$ so that (1) is not satisfied. The average-based separation technique [2] does the opposite: it first focuses on (1) and tries to derive the following averaged version:

$$\mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}} \left[\textsc{Adv}_k^{\mathcal{P}}(A, G^f)\right] \neq k^{-\omega(1)} \quad \Rightarrow \quad \mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}} \left[\textsc{Adv}_k^{\mathcal{Q}}(S^{f,A}, f)\right] \neq k^{-\omega(1)} \ , \tag{2}$$

and then derive a contradiction based on (S), (B) and (2). Indeed, from (S) it follows that $\mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}}\left[\mathrm{ADV}_k^{\mathcal{P}}(A,G^f)\right] = 1 - k^{-\omega(1)} \neq k^{-\omega(1)}$. By (2) we imply that $\mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}}\left[\mathrm{ADV}_k^{\mathcal{Q}}(S^{f,A},f)\right] \neq k^{-\omega(1)}$ which contradicts (B).

7.1 Poly-Preserving Reductions

The main problem with the the averaging approach described above is that the averaged condition (2) cannot be derived from the general reduction condition (1). Indeed, let $a_f(k) = \mathrm{ADV}_k^{\mathcal{P}}(A,G^f)$ and $b_f(k) = \mathrm{ADV}_k^{\mathcal{Q}}(S^{f,A},f)$. We would like to prove that if (for all f) $b_f(k) = k^{-\omega(1)}$ implies $a_f(k) = k^{-\omega(1)}$, then

$$\mathop{\mathbf{E}}_{f}[b_f(k)] = k^{-\omega(1)} \Rightarrow \mathop{\mathbf{E}}_{f}[a_f(k)] = k^{-\omega(1)} \ .$$

The negligible average argument (Lemma 4) implies that $a_f(k) = k^{-\omega(1)}$ for measure one of f's, but this does not mean that $\mathop{\mathbf{E}}_{f}[a_f(k)]$ is negligible (Lemma 6).

So, for average-based separation, the guarantee condition (S) is too weak—much weaker than what is usually expected when constructing practical reductions. For these reasons, the guarantee condition was strengthened [2] and the class of reductions restricted in the following reasonable way:

Definition 6 (Poly-preserving reductions). *A reduction of $\mathcal{P}$ to $\mathcal{Q}$ is* poly-preserving *if the security guarantee (S) decreases the advantage by at most a polynomial amount, i.e. there exists $c \geq 1$ (independent of f, A and k) such that*

$$\mathrm{ADV}_k^{\mathcal{Q}}(S^f,f) \geq \left[\mathrm{ADV}_k^{\mathcal{P}}(A^f,p(f))\right]^c \ . \tag{3}$$

7.2 Averaging-Based Separation for Poly-Preserving Reductions

For fully-black box reductions, (3) is in the form $\mathrm{ADV}_k^{\mathcal{Q}}(S^{f,A},f) \geq \left[\mathrm{ADV}_k^{\mathcal{P}}(A,G^f)\right]^c$. For poly-preserving reductions, the averaged reduction condition (2) easily follows:

$$\mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}}\left[\mathrm{ADV}_k^{\mathcal{Q}}(S^{f,A},f)\right] \geq \mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}}\left[\left(\mathrm{ADV}_k^{\mathcal{P}}(A,G^f)\right)^c\right] \geq \left(\mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}}\left[\mathrm{ADV}_k^{\mathcal{P}}(A,G^f)\right]\right)^c \ ,$$

where the second inequality is an application of the Jensen inequality. This implies that if $\mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}}\left[\mathrm{ADV}_k^{\mathcal{P}}(A,G^f)\right]$ is non-negligible, then so is $\mathop{\mathbf{E}}_{f,A\leftarrow\mathcal{F}}\left[\mathrm{ADV}_k^{\mathcal{Q}}(S^{f,A},f)\right]$.

Table 2 lists the separation conditions for all four types of reductions. Note that the breakage condition for averaging-based separation can be somewhat weaker than in the traditional extraction-based approach. We only have to assume that the success of A is non-negligible. All proofs are given in Appendix B.

Table 2. Reduction types and separation conditions for averaging-based separation in the case of poly-preserving black-box reductions

Type	Reduction Condition	Separation Condition
Full bb	$\underset{\text{of}}{\exists}\, p \underset{\text{pol}}{\exists}\, S \forall f \forall A$: $\text{ADV}_k(S^{f,A}, f) \geq \left[\text{ADV}_k(A^f, p(f))\right]^c$	$\underset{\text{of}}{\forall}\, p \underset{\text{pol}}{\forall}\, S \exists \mathcal{F}$: $\underset{f,A\leftarrow\mathcal{F}}{\mathbf{E}} \left[\text{ADV}_k(A, p(f))\right] \neq k^{-\omega(1)}$ $\underset{f,A\leftarrow\mathcal{F}}{\mathbf{E}} \left[\text{ADV}_k(S^{f,A}, f)\right] = k^{-\omega(1)}$
Str s-bb	$\underset{\text{of}}{\exists}\, p \underset{\text{pol}}{\forall}\, A \underset{\text{pol}}{\exists}\, S \forall f$: $\text{ADV}_k(S^f, f) \geq \left[\text{ADV}_k(A^f, p(f))\right]^c$	$\underset{\text{of}}{\forall}\, p \underset{\text{pol}}{\exists}\, A \underset{\text{pol}}{\forall}\, S \exists \mathcal{F}$: $\underset{f\leftarrow\mathcal{F}}{\mathbf{E}} \left[\text{ADV}_k(A^f, p(f))\right] \neq k^{-\omega(1)}$ $\underset{f\leftarrow\mathcal{F}}{\mathbf{E}} \left[\text{ADV}_k(S^f, f)\right] = k^{-\omega(1)}$
Weak s-bb	$\underset{\text{of}}{\exists}\, p \underset{\text{pol}}{\forall}\, A \forall f \underset{\text{pol}}{\exists}\, S_f$: $\text{ADV}_k(S_f^f, f) \geq \left[\text{ADV}_k(A^f, p(f))\right]^c$	$\underset{\text{of}}{\forall}\, p \underset{\text{pol}}{\exists}\, A \exists \mathcal{F}$: $\underset{f\leftarrow\mathcal{F}}{\mathbf{E}} \left[\text{ADV}_k(A^f, p(f))\right] \neq k^{-\omega(1)}$ $\underset{\text{pol}}{\forall}\, S \underset{\text{of}}{\forall}\, \varphi \underset{f\leftarrow\mathcal{F}}{\mathbf{E}} \left[\text{ADV}_k(S^f_{\varphi(f)}, f)\right] = k^{-\omega(1)}$
Vari s-bb	$\forall f \underset{\text{pol}}{\exists}\, G_f \underset{\text{pol}}{\forall}\, A \underset{\text{pol}}{\exists}\, S_f$: $\text{ADV}_k(S_f^f, f) \geq \left[\text{ADV}_k(A^f, G_f^f)\right]^c$	$\underset{\text{of}}{\forall}\, \psi \exists \mathcal{F}$: $\forall G \underset{\text{pol}}{\exists}\, A \underset{\text{pol}}{\exists}\, \underset{f\leftarrow\mathcal{F}}{\mathbf{E}} \left[\text{ADV}_k(A^f, G^f_{\psi(f)})\right] \neq k^{-\omega(1)}$ $\underset{\text{pol}}{\forall}\, S \underset{\text{of}}{\forall}\, \varphi \underset{f\leftarrow\mathcal{F}}{\mathbf{E}} \left[\text{ADV}_k(S^f_{\varphi(f)}, f)\right] = k^{-\omega(1)}$

8 Going beyond the Strong Semi Black-Box Boundary

The strong Semi Black-box separations are a clear boundary for oracle extraction based approaches, as anything below that requires the use of countability arguments, which fail in the non-uniform model. It therefore seems that the best hope of proving stronger reductions would rest with the averaging-based approach. For this to succeed, however, one would have to be able to account for oracle-dependent advice strings $\varphi(\mathcal{F})$ being available to the adversary construction. This is by not a trivial obstacle to overcome.

The most promising approach for doing so stems from the work of Unruh [11], where they showed that under reasonable assumptions, the oracle could be switched out with one fairly independent from the original without the adversary having a significant chance of noticing the switch. A problem with his approach was that it only applies to standard random oracles, which separation oracles rarely are. If this idea was to be used, the result first needed to be generalized to work for other, less standard oracles as well. This turned out to be possible as we managed to prove one of the conjectures presented in the original paper [11] pertaining to the fully general choice of oracle distributions.

Theorem 1. *Let $\mathcal{F}$ be any distribution of Oracles and let $f \leftarrow \mathcal{F}$. We say that f is consistent with a matching $M = \{x_1 \rightarrow y_1, \ldots, x_m \rightarrow y_m\}$ if $f(x_i) = y_i$ for all $i \in \{1, 2, \ldots, m\}$. Let $\varphi(f)$ be an oracle function with an output of length p. Then there is an oracle function S such that S^f is a matching of length m and the following holds: For any probabilistic oracle Turing machine B that makes at most q queries to its oracle, $\Delta(B^{\mathcal{F}}_{\varphi(\mathcal{F})}; B^{\mathcal{F}/S}_{\varphi(\mathcal{F})}) \leq \sqrt{\frac{pq}{2m}}$, where $\mathcal{F}/S$ is an oracle sampled according to $\mathcal{F}$ conditioned only on being consistent with $S^{\mathcal{F}}$ (which is also a random variable).*

We will generalize the proof for Theorem 2 of Unruh [11] to work for arbitrary oracle distributions. The proof is quite similar to the original one, with only the notion of

information $J(\mathsf{M})$ gathered by an adaptive list M about the advice string z given a more general definition. Since most of the proofs are completely analogous, we will only give a detailed description of the parts that have to be changed. We will use the notion of adapative lists from the original paper. To reiterate, a adaptive list is a Turing Machine that takes as input a finite advice string z and then proceeds to make a number of oracle queries, terminating in finite time, outputting all the oracle queries that it made along with their results. It is assumed that the query responses are cached so that oracle is never queried twice with the same input. It is further assumed that F is deterministic, although this does not seem to be an essential assumption in our treatment. We will define a TM G so that when given such an input, it will methodically query the oracle on all the yet unqueried inputs.

For an adaptive list F, we will define the variable $F_k^{\mathbf{o},z}$ as the response to the k-th oracle query made by $\mathsf{G} \circ \mathsf{F}^{\mathbf{o}}(z)$ when run with the oracle $\mathbf{o}$ and input z. According to the preceeding assumptions, $F_k^{\mathbf{o},z}$ is well defined for all $k \leq |\mathrm{Range}|$, $\mathbf{o} \in \mathcal{O}$ and $z \in \mathcal{Z}$. Let $\mathbf{F}_k^{\mathcal{O},\mathbf{Z}}$ be the variable induced by F_k by choosing $\mathbf{o} \leftarrow \mathcal{O}$ and $z \leftarrow \mathbf{Z}$. For convenience, we note $\mathbf{F}_{k\to l} = \mathbf{F}_{k+1}, \ldots, \mathbf{F}_l$ and $\mathbf{F}_{*l} = \mathbf{F}_{0\to l}$.

Let the distribution of advice strings $\mathbf{Z}$ be dependent on the distribution of oracles $\mathcal{O}$. Let $\mathcal{O}'$ stand for the distribution of oracles that is distributed identically to $\mathcal{O}$ but that is independent from $\mathbf{Z}$. Let $\mathcal{O}'/S_k$ denote the distribution $\mathcal{O}'$ conditioned on agreeing with $\mathcal{O}$ on all the queries $F_{*k}^{\mathcal{O},\mathbf{Z}}$. The goal is then to show a bound on the statistical distance $\Delta(\mathbf{F}_{*k}^{\mathcal{O},\mathbf{Z}}, \mathbf{F}_{k\to k+q}^{\mathcal{O}'/S_k,\mathbf{Z}}; \mathbf{F}_{*(k+q)}^{\mathcal{O},\mathbf{Z}})$. We will use the Kullback-Leibler distance:

$$D(\mathbf{X}||\mathbf{Y}|\mathbf{Z}) = \sum_{z\in\mathcal{Z}} \Pr[\mathbf{Z}=z] \sum_{x\in\mathcal{X}} \Pr[\mathbf{X}=x|\mathbf{Z}=z] \lg\left(\frac{\Pr[\mathbf{X}=x|\mathbf{Z}=z]}{\Pr[\mathbf{Y}=x|\mathbf{Z}=z]}\right) \ .$$

The following useful properties (Gibbs inequality, chain rule, and funcion applications only decreasing the distance) are well known and easy to verify.

$$\begin{aligned}
&D(\mathbf{X}||\mathbf{Y}) \geq 0 \ , \\
&D(\mathbf{X}_1, \ldots, \mathbf{X}_k||\mathbf{Y}_1, \ldots, \mathbf{Y}_k) = \\
&\qquad D(\mathbf{X}_1||\mathbf{X}_2) + D(\mathbf{X}_2||\mathbf{Y}_2|\mathbf{X}_1) + \ldots D(\mathbf{X}_k||\mathbf{Y}_k|\mathbf{X}_1, \ldots, \mathbf{X}_{k-1}) \ , \\
&D(f(\mathbf{X})||f(\mathbf{Y})) \leq D(\mathbf{X}||\mathbf{Y}) \ .
\end{aligned}$$

It is worth noting that for the chain rule,

$$D(\mathbf{X}_2||\mathbf{Y}_2|\mathbf{X}_1) = \sum_{x_1\in\mathcal{X}_1} \Pr[\mathbf{X}_1=x_1] \sum_{x_2\in\mathcal{X}_2} \Pr[\mathbf{X}=x_2|\mathbf{X}=x_1] \lg\left(\frac{\Pr[\mathbf{X}_2=x_2|\mathbf{X}_1=x_1]}{\Pr[\mathbf{Y}_2=x_2|\mathbf{Y}_1=x_1]}\right) \ ,$$

the conditioning is actually over both the values of $\mathbf{X}_1$ and $\mathbf{Y}_1$. We define[4] $J_k(\mathsf{F}) = D(\mathbf{F}_{*k}^{\mathcal{O},\mathbf{Z}}||\mathbf{F}_{*k}^{\mathcal{O}',\mathbf{Z}}|\mathbf{Z})$ and $J_{k\to l}(\mathsf{F}) = J_l(\mathsf{F}) - J_k(\mathsf{F}) = D(\mathbf{F}_{k\to l}^{\mathcal{O},\mathbf{Z}}||\mathbf{F}_{k\to l}^{\mathcal{O}'/S_k,\mathbf{Z}}|\mathbf{F}_{*k}^{\mathcal{O},\mathbf{Z}}, \mathbf{Z})$. We note that although we use a slightly different notation for J that makes the length

[4] Introduction of S_k is due to the chain rule conditioning over both distibutions.

of the list k explicit, this is purely for syntactic convenience. As in the original, let $J_k = \max_{\mathsf{F}} J_k(\mathsf{F})$.

The proof in the original paper requires three properties from $J(\mathsf{F})$. Two of them ($J_k(\mathsf{F}) \geq 0$ and $J_l(\mathsf{F}) \leq J_{k\to l}(\mathsf{F}) + J_k(\mathsf{F})$) follow directly from the properties of Kullback-Leibler distance. The third property $J_k(\mathsf{F}) \leq H(\mathbf{Z})$ is just slightly trickier, but follows trivially from the following lemma.

Lemma 7. *Let* $\mathbf{X}$ *and* $\mathbf{Y}$ *be identically distributed. Additionally, let* $\mathbf{Z}$ *be variable independent from* $\mathbf{Y}$ *(but possibly related to* $\mathbf{X}$*). In such a case,* $D(\mathbf{X}||\mathbf{Y}|\mathbf{Z}) \leq H(\mathbf{Z})$.

Proof.

$$\begin{aligned} D(\mathbf{X}||\mathbf{Y}|\mathbf{Z}) &= \sum_{z\in\mathcal{Z}} \Pr[\mathbf{Z}=z] \sum_{x\in\mathcal{X}} \Pr[\mathbf{X}=x|\mathbf{Z}=z] \lg\left(\frac{\Pr[\mathbf{X}=x|\mathbf{Z}=z]}{\Pr[\mathbf{Y}=x|\mathbf{Z}=z]}\right) \\ &= \sum_{z\in\mathcal{Z}} \Pr[\mathbf{Z}=z] \sum_{x\in\mathcal{X}} \Pr[\mathbf{X}=x|\mathbf{Z}=z] \lg\left(\frac{\Pr[\mathbf{X}=x,\mathbf{Z}=z]}{\Pr[\mathbf{Y}=x]\Pr[\mathbf{Z}=z]}\right) \\ &\leq \sum_{z\in\mathcal{Z}} \Pr[\mathbf{Z}=z] \sum_{x\in\mathcal{X}} \Pr[\mathbf{X}=x|\mathbf{Z}=z] \lg\left(\frac{1}{\Pr[\mathbf{Z}=z]}\right) \\ &= \sum_{z\in\mathcal{Z}} \Pr[\mathbf{Z}=z] \lg\left(\frac{1}{\Pr[\mathbf{Z}=z]}\right) = H(\mathbf{Z}) \ , \end{aligned}$$

where the inequality is due to $\frac{\Pr[\mathbf{X}=x,\mathbf{Z}=z]}{\Pr[\mathbf{X}=x]} \leq 1$ and $\mathbf{X}, \mathbf{Y}$ are identically distributed. □

Corollary 1. $J_k(\mathsf{F}) \leq H(\mathbf{Z})$.

Proof.

$$\begin{aligned} J_k(\mathsf{F}) &= D(\mathbf{F}_{*k}^{\mathcal{O},\mathbf{Z}}||\mathbf{F}_{*k}^{\mathcal{O}',\mathbf{Z}}|\mathbf{Z}) = D(f(\mathcal{O},\mathbf{Z})||f(\mathcal{O}',\mathbf{Z})|\mathbf{Z}) \leq D(\mathcal{O},\mathbf{Z}||\mathcal{O}',\mathbf{Z}|\mathbf{Z}) \\ &= D(\mathcal{O}||\mathcal{O}'|\mathbf{Z}) \leq H(\mathbf{Z}) \ . \end{aligned}$$

□

The only piece missing is to use $D(\mathbf{X}||\mathbf{Y}|\mathbf{Z})$ for bounding $\Delta(X,Z;Y,Z)$. This is also completely analogous to the proof in Unruh, as

$$\begin{aligned} \Delta(\mathbf{X},\mathbf{Z};\mathbf{Y},\mathbf{Z}) &= \sum_{z\in\mathcal{Z}} \Pr[\mathbf{Z}=z]\Delta(\mathbf{X};\mathbf{Y}|\mathbf{Z}=z) \leq \sum_{z\in\mathcal{Z}} \Pr[\mathbf{Z}=z]\sqrt{\frac{D(\mathbf{X}||\mathbf{Y}|\mathbf{Z}=z)}{2}} \\ &\leq \sqrt{\frac{1}{2}\sum_{z\in\mathcal{Z}} \Pr[\mathbf{Z}=z] D(\mathbf{X}||\mathbf{Y}|\mathbf{Z}=z)} = \sqrt{\frac{1}{2}D(\mathbf{X}||\mathbf{Y}|\mathbf{Z})} \ , \end{aligned}$$

where the first inequality is due to Kullback-Leibler and the second is an application of Jensen's inequality. All the other parts of the proof remain fairly unaltered, with a few pieces (such as replacing G with ∇G) becoming obsolete due to independence requirements being relaxed.

This theorem basically allows one to formally replace a polynomial-length oracle-dependent advice string with just fixing a super-polynomial number of responses to oracle queries, with only negligible chance of the adversary behaving differently. This fits in well with many of the already known proofs for separation results, which will still work even when the number of queries is slightly super-polynomial just as long as it is still negligible when compared with the full domain and range of the oracle function. In such cases, one can then replace the original usually oracle extraction based argumentation with the averaging-based argumentation to yield a stronger result that holds also in the non-uniform model. For instance, this seems to be the case with the work of Simon [10] where it was shown that collision-resistant hash functions cannot be constructed based purely on one-way functions. As his argumentation still remains valid when the adversary makes a super-polynomial number of queries, the result can be generalized to the non-uniform model.

References

1. Buldas, A., Jürgenson, A., Niitsoo, M.: Efficiency bounds for adversary constructions in black-box reductions. In: Boyd, C., González Nieto, J. (eds.) ACISP 2009. LNCS, vol. 5594, pp. 264–275. Springer, Heidelberg (2009)
2. Buldas, A., Laur, S., Niitsoo, M.: Oracle separation in the non-uniform model. In: Pieprzyk, J., Zhang, F. (eds.) ProvSec 2009. LNCS, vol. 5848, pp. 230–244. Springer, Heidelberg (2009)
3. Gennaro, R., Gertner, Y., Katz, J.: Lower bounds on the efficiency of encryption and digital signature schemes. In: STOC 2003, pp. 417–425 (2003)
4. Gennaro, R., Gertner, Y., Katz, J., Trevisan, L.: Bounds on the efficiency of generic cryptographic constructions. SIAM Journal on Computing 35, 217–246 (2006)
5. Gertner, Y., Kannan, S., Malkin, T., Reingold, O., Viswanathan, M.: The relationship between public key encryption and oblivious transfer. In: FOCS 2000, pp. 325–335 (2000)
6. Hsiao, C.-Y., Reyzin, L.: Finding collisions on a public road, or do secure hash functions need secret coins? In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 92–105. Springer, Heidelberg (2004)
7. Impagliazzo, R., Rudich, S.: Limits on the provable consequences of one-way permutations. In: STOC 1989, pp. 44–61 (1989)
8. Kim, J.H., Simon, D.R., Tetali, P.: Limits on the efficiency of one-way permutation-based hash functions. In: FOCS 1999, pp. 535–542 (1999)
9. Reingold, O., Trevisan, L., Vadhan, S.: Notions of reducibility between cryptographic primitives. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 1–20. Springer, Heidelberg (2004)
10. Simon, D.R.: Findings collisions on a one-way street: Can secure hash functions be based on general assumptions? In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 334–345. Springer, Heidelberg (1998)
11. Unruh, D.: Random oracles and auxiliary input. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 205–223. Springer, Heidelberg (2007)

A Oracle Extraction Examples

Theorem 2. *If* $\underset{pol}{\forall} G \underset{pol}{\exists} A \underset{pol}{\forall} S \exists \mathcal{F}$: *so that (I)* $\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(A^f, G^f) \right] = 1 - k^{-\omega(1)}$ *and (II)* $\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(S^f, f) \right] = k^{-\omega(1)}$*, there exist no strong semi black-box reductions.*

Proof. Using the overwhelming average argument for (I) and the negligible average argument for (II), we imply that $\Pr_{f\leftarrow\mathcal{F}}\left[A^f \mathbf{br} G^f\right] = 1$ and $\Pr_{f\leftarrow\mathcal{F}}\left[S^f \not\mathbf{br} f\right] = 1$, which implies $\forall_{\text{pol}} G \exists_{\text{pol}} A \forall_{\text{pol}} S \exists \mathcal{F}\colon \Pr_{f\leftarrow\mathcal{F}}\left[A^f \mathbf{br} G^f \wedge S^f \not\mathbf{br} f\right] = 1$ and hence from the probabilistic argument: $\forall_{\text{pol}} G \exists_{\text{pol}} A \forall_{\text{pol}} S \exists f\colon \left[A^f \mathbf{br} G^f \wedge S^f \not\mathbf{br} f\right]$, which is the negation of the strong semi black-box reduction condition.[5] □

Theorem 3. *If* $\forall_{pol} G \exists_{pol} A \exists \mathcal{F}\colon$ *so that (I)* $\mathop{\mathbf{E}}_{f\leftarrow\mathcal{F}}\left[\text{Adv}_k(A^f, G^f)\right] = 1 - k^{-\omega(1)}$ *and (II)* $\forall_{pol} S\colon \mathop{\mathbf{E}}_{f\leftarrow\mathcal{F}}\left[\text{Adv}_k(S^f, f)\right] = k^{-\omega(1)}$*, there exist no uniform weak semi b-b reductions.*

Proof. The overwhelming average argument for (I) and the negligible average argument for (II) imply $\Pr_{f\leftarrow\mathcal{F}}\left[A^f \mathbf{br} G^f\right] = 1$ and $\forall_{\text{pol}} S\colon \Pr_{f\leftarrow\mathcal{F}}\left[S^f \not\mathbf{br} f\right] = 1$. By the countability argument for S, we obtain $\Pr_{f\leftarrow\mathcal{F}}\left[\forall_{\text{pol}} S\colon S^f \not\mathbf{br} f\right] = 1$, which implies $\forall_{\text{pol}} G \exists_{\text{pol}} A \exists \mathcal{F}\colon \Pr_{f\leftarrow\mathcal{F}}\left[A^f \mathbf{br} G^f \wedge \forall_{\text{pol}} S\colon S^f \not\mathbf{br} f\right] = 1$, and we have the negation of the weak semi black-box reduction: $\forall_{\text{pol}} G \exists_{\text{pol}} A \exists f \forall_{\text{pol}} S\colon \left[A^f \mathbf{br} G^f \wedge S^f \not\mathbf{br} f\right]$.[6] □

Theorem 4. *If* $\exists \mathcal{F}\colon$ *so that (I)* $\forall_{pol} G \exists_{pol} A\colon \mathop{\mathbf{E}}_{f\leftarrow\mathcal{F}}\left[\text{Adv}_k(A^f, G^f)\right] = 1 - k^{-\omega(1)}$ *and (II)* $\forall_{pol} S\colon \mathop{\mathbf{E}}_{f\leftarrow\mathcal{F}}\left[\text{Adv}_k(S^f, f)\right] = k^{-\omega(1)}$*, there are no uniform variable semi b-b reductions.*

A proof was already presented by the steps s_1-b_3 in Section 6.

B Averaging Examples

Lemma 8. *The existence of weak semi black-box reductions is equivalent to:*

$$\exists_{of} p \forall_{pol} A \exists_{pol} \mathcal{S} \exists_{of} \varphi \forall f\colon \quad \text{Adv}_k(\mathcal{S}^f_{\varphi(f)}, f) \geq \left[\text{Adv}_k(A^f, p(f))\right]^c \quad \textit{where } c \geq 1\ . \tag{4}$$

Proof. Assume first that $\exists_{\text{of}} p \forall_{\text{pol}} A \forall f \exists_{\text{pol}} S_f\colon \text{Adv}_k(S^f_f, f) \geq \left[\text{Adv}_k(A^f, p(f))\right]^c$, i.e. there exists a weak semi black-box reduction and prove (4). Let φ be an oracle function so that $\varphi(f)$ is a bit-representation of S_f. Let $\mathcal{S}$ be the universal f-oracle machine, which when given as input a bit-representation $\varphi(f)$ behaves exactly like S_f. This means that $\text{Adv}_k(\mathcal{S}^f_{\varphi(f)}, f) = \text{Adv}_k(S^f_f, f)$. Moreover, as such simulation is possible with logarithmic overhead, it follows that $\mathcal{S}_{\varphi(f)}$ is poly-time. As $\mathcal{S}$ and φ are the same for all instances of f, the statement (4) follows.

From (4) by defining $S_f := \mathcal{S}_{\varphi(f)}$, there exists p such that for all poly-time A and for all f there is S_f, so that $\text{Adv}_k(S^f_f, f) = \text{Adv}_k(\mathcal{S}^f_{\varphi(f)}, f) \geq \left[\text{Adv}_k(A^f, p(f))\right]^c$, which proves the existence of weak semi black-box reduction. □

[5] No countability arguments were used.

[6] As we used the countability argument for S, the result does not apply to the models where the adversaries are allowed to be non-uniform.

Theorem 5. *If* $\underset{of}{\forall} p \underset{pol}{\exists} A \exists \mathcal{F}$: *so that (I)* $\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(A^f, p(f)) \right] \neq k^{-\omega(1)}$ *and (II)* $\underset{pol}{\forall} S \underset{of}{\forall} \varphi \underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(S^f_{\varphi(f)}, f) \right] = k^{-\omega(1)}$, *there exist no weak semi b-b reductions.*

Proof. By using (4), (I) and (II), we will derive a contradiction. Let p be as in (4). By applying the assumption of the theorem to this p, we conclude that there exist a poly-time oracle machine A and a distribution $\mathcal{F}$ with the properties (I) and (II). Now, from (4) it follows that for this A, there exist a poly-time oracle machine $\mathcal{S}$ and an oracle function φ such that (*) $\mathrm{ADV}_k(\mathcal{S}^f_{\varphi(f)}, f) \geq \left[\mathrm{ADV}_k(A^f, p(f)) \right]^c$ holds for all f. By (II), we have (**) $\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(\mathcal{S}^f_{\varphi(f)}, f) \right] = k^{-\omega(1)}$. Finally, by averaging (*) and using the Jensen's inequality we have:

$$\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(\mathcal{S}^f_{\varphi(f)}, f) \right] \geq \underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(A^f, p(f))^c \right] \geq \left[\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(A^f, p(f)) \right] \right]^c ,$$

which is a contradiction between (I) and (**). □

Lemma 9. *The existence of variable semi black-box reductions is equivalent to:*

$$\underset{of}{\exists} \psi \underset{pol}{\exists} \mathcal{P} \underset{pol}{\forall} A \underset{pol}{\exists} \mathcal{S} \underset{of}{\exists} \varphi \forall f : \quad \mathrm{ADV}_k(\mathcal{S}^f_{\varphi(f)}, f) \geq \left[\mathrm{ADV}_k(A^f, \mathcal{P}^f_{\psi(f)}) \right]^{O(1)} . \tag{5}$$

Proof. Assume first that $\forall f \underset{\text{pol}}{\exists} G_f \underset{\text{pol}}{\forall} A \underset{\text{pol}}{\exists} S_f$: $\mathrm{ADV}_k(S^f_f, f) \geq \left[\mathrm{ADV}_k(A^f, G^f_f) \right]^c$, i.e. there exists a variable semi black-box reduction, and prove (5). Let ψ be a mapping so that $\psi(f)$ is the bit-string representation of G_f. Let $\mathcal{P}$ be the universal f-oracle machine so that $\mathcal{P}_{\psi(f)}$ behaves identical to G_f. Hence, $\mathrm{ADV}_k(A^f, \mathcal{P}^f_{\psi(f)}) = \mathrm{ADV}_k(A^f, G^f_f)$, and due to the efficiency of simulation, $\mathcal{P}_{\psi(f)}$ is poly-time. For every A we define $\mathcal{S}$ and φ like in Lemma 8. The statement (5) follows. □

Theorem 6. *If* $\forall \psi \exists \mathcal{F}$: *so that (I)* $\underset{pol}{\forall} G \underset{pol}{\exists} A \ \underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(A^f, G_{\psi(f)}) \right] \neq k^{-\omega(1)}$ *and (II)* $\underset{pol}{\forall} S \underset{of}{\forall} \varphi \ \underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(S^f_{\varphi(f)}, f) \right] = k^{-\omega(1)}$, *there exist no weak semi b-b reductions.*

Proof. By using (5), (I) and (II), we derive a contradiction. Let ψ and $\mathcal{P}$ be as in (4). By applying the assumption of the theorem to ψ, we conclude that a distribution $\mathcal{F}$ with the properties (I) and (II). By applying (I) to $\mathcal{P}$, we conclude that there exists A such that (*) $\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(A^f, \mathcal{P}_{\psi(f)}) \right] \neq k^{-\omega(1)}$. From (5) it follows that there exist a poly-time $\mathcal{S}$ and an oracle function φ so that (**) $\mathrm{ADV}_k(\mathcal{S}^f_{\varphi(f)}, f) \geq \left[\mathrm{ADV}_k(A^f, \mathcal{P}^f_{\psi(f)}) \right]^c$ for all f. By (II), we have (***) $\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(\mathcal{S}^f_{\varphi(f)}, f) \right] = k^{-\omega(1)}$. Finally, by averaging (**) and using the Jensen's inequality, we have

$$\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(\mathcal{S}^f_{\varphi(f)}, f) \right] \geq \underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(A^f, \mathcal{P}^f_{\psi(f)})^c \right] \geq \left[\underset{f \leftarrow \mathcal{F}}{\mathbf{E}} \left[\mathrm{ADV}_k(A^f, \mathcal{P}^f_{\psi(f)}) \right] \right]^c .$$

A contradiction between (*) and (***). □