

Hardness-Preserving Reductions via Cuckoo Hashing*

Itay Berman^{†§} Iftach Haitner^{†§} Ilan Komargodski[¶] Moni Naor^{||}

November 22, 2017

Abstract

The focus of this work is *hardness-preserving* transformations of somewhat limited pseudorandom functions families (PRFs) into ones with more versatile characteristics. Consider the problem of *domain extension* of pseudorandom functions: given a PRF that takes as input elements of some domain \mathcal{U} , we would like to come up with a PRF over a larger domain. Can we do it with little work and without significantly impacting the security of the system? One approach is to first hash the larger domain into the smaller one and then apply the original PRF. Such a reduction, however, is vulnerable to a “birthday attack”: after $\sqrt{|\mathcal{U}|}$ queries to the resulting PRF, a collision (i.e., two distinct inputs having the same hash value) is very likely to occur. As a consequence, the resulting PRF is *insecure* against an attacker making this number of queries.

In this work we show how to go beyond the aforementioned birthday attack barrier by replacing the above simple hashing approach with a variant of *cuckoo hashing*, a hashing paradigm that resolves collisions in a table by using two hash functions and two tables, cleverly assigning each element to one of the two tables. We use this approach to obtain: (i) a domain extension method that requires *just two calls* to the original PRF, can withstand as many queries as the original domain size, and has a distinguishing probability that is exponentially small in the amount of non-cryptographic work; and (ii) a *security-preserving* reduction from non-adaptive to adaptive PRFs.

*A preliminary version appeared in the proceedings of TCC 2013 [9].

[†]MIT Computer Science and Artificial Intelligence Laboratory (CSAIL). Email: itayberm@mit.edu. Most of this work was done while the author was in the School of Computer Science, Tel Aviv University. Supported in part by NSF Grants CNS-1413920 and CNS-1350619, and by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236.

[‡]School of Computer Science, Tel Aviv University. Email: iftachh@cs.tau.ac.il.

[§]Research supported in part by Check Point Institute for Information Security and the (I-CORE) program (Center No. 4/11) of the Planning and Budgeting Committee.

[¶]Cornell Tech, New York, USA. Email: komargodski@cornell.edu. Supported in part by a Packard Foundation Fellowship. Most of this work done while being a Ph.D student at the Weizmann Institute of Science, supported in part by a grant from the I-CORE Program (Center No. 4/11) of the Planning and Budgeting Committee and the Israel Science Foundation.

^{||}Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot 76100, Israel. Email: moni.naor@weizmann.ac.il. Research supported in part by a grant from the I-CORE Program (Center No. 4/11) of the Planning and Budgeting Committee and the Israel Science Foundation. Incumbent of the Judith Kleeman Professorial Chair.

1 Introduction

The focus of this work is *hardness-preserving* transformations of somewhat limited pseudorandom functions families (PRFs) into ones with more versatile characteristics. Examples of such somewhat limited families include those with *small* domain or those that can withstand only *non-adaptive* (also known as static) attacks, in which the attacker chooses its queries ahead of time, before seeing any of the answers. In contrast, less limited families might have *large* domain or be secure against *adaptive* (dynamic) attacks, in which the attacker’s queries might be chosen as a function of all previous answers.

A common paradigm, first suggested by Levin [30, §5.4], for increasing the usability and security of a PRF, is to “hash” the inputs into a smaller domain *before* applying the PRF. This approach was originally suggested in order to achieve “PRF domain extension” (using a short, e.g., fixed, input length PRF to get a variable-length PRF); more recently, it was used to transform non-adaptive PRFs into adaptive ones [8]. Such reductions, however, are vulnerable to the following “birthday attack”: after $\sqrt{|\mathcal{U}|}$ queries to the resulting PRF, where \mathcal{U} is the hash function range, a collision (i.e., two distinct inputs having the same hash value) is very likely to occur. Such collisions are an obstacle to the indistinguishability of the PRF, since in a random function we either do not expect to see a collision at all (if the range is large enough) or expect to see fewer collisions. Hence, the resulting PRF is *insecure* against an attacker making this number of queries.

In this work we study variants of the above hashing approach to go beyond the birthday attack barrier. At a high-level, our approach, which can be traced back to Siegel [53], is based on applying a dictionary data structure¹, in which the locations accessed in the search of an element are determined by its value and some fixed random string (i.e., the same string is used for all elements), and not on values seen during the search. Now to do the conversion to domain extension we assign random values to all locations (by the underlying PRF). We think of the large domain as the universe from which the elements of the dictionary are taken. The resulting value of the extended function at point x will be some (simple) function of the values assigned to the locations accessed during the search for x . For instance, one can view Levin’s construction above as an instance of this framework, where the fixed random string describes a hash function from large domain to a smaller-size set \mathcal{U} , and the PRF, whose domain is \mathcal{U} , assigns random values for $|\mathcal{U}|$ locations. The distinguishing probability of the resulting scheme is the distinguishing probability of the underlying PRF *plus* the probability of failure of the dictionary (which in Levin’s construction is determined by the “birthday paradox”). The cost of the extension is related to the worst-case search time of the dictionary (which in Levin’s construction is a single invocation of the hash function).

We focus on constructions based on *cuckoo hashing*: a hashing paradigm typically used for resolving hash collisions in a table by using two hash functions and two tables, assigning each element to one of the two tables, and enabling lookup using only two queries (see Section 1.2). We use this paradigm to present a new PRF domain extension method that requires *just two calls* to the original PRF, can withstand as many queries as the original domain size, and has a distinguishing probability that is exponentially small in the amount of non-cryptographic work. We also obtain a *security-preserving reduction* from non-adaptive to adaptive PRFs, an improvement upon the recent result of Berman and Haitner [8].

Before stating our results, we discuss in detail pseudorandom functions and cuckoo hashing.

¹In this context a dictionary is a data structure used for maintaining a set of elements while supporting membership queries.

1.1 Pseudorandom Functions

Pseudorandom function families (PRFs), introduced by Goldreich, Goldwasser, and Micali [22], are function families that cannot be distinguished from a family of *truly* random functions by an efficient distinguisher given an oracle access to a random member of the family. PRFs have an extremely important role in cryptography, allowing parties who share a common secret key to send secure messages, identify themselves, and authenticate messages [21, 31]. They have many other applications as well, and can be used in just about any setting that requires a random function provided as a black-box [5, 11, 14, 20, 32, 43]. Different PRF constructions, whose security is based on different hardness assumptions, are known in the literature. The construction most relevant to this work is the one of [22], hereafter the \mathcal{GGM} construction, which uses a length-doubling pseudorandom generator (and thus can be based on the existence of one-way functions [24]).

We use the following definitions: an efficiently computable function family ensemble $\mathcal{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ is a (q, t, ε) -PRF, if (for large enough n) a $q(n)$ -query oracle-aided algorithm (distinguisher) of running time $t(n)$, getting access to a random function from the family, distinguishes between \mathcal{F}_n and the family of all functions (with the same input/output domains), with probability at most $\varepsilon(n)$. \mathcal{F} is a *non-adaptive* (q, t, ε) -PRF if it is only required to be secure against non-adaptive distinguishers (i.e., ones that prepare all their queries in advance). Finally, \mathcal{F} is a t -PRF if q is only limited by t and $\varepsilon = 1/t$.

We also make use of the information-theoretic analog of a t -PRF, known as a t -wise independent family (see Definition 2.5).

1.2 Cuckoo Hashing and Many-wise Independent Hash Function

Cuckoo hashing, introduced by Pagh and Rodler [45], is an efficient technique for constructing dynamic dictionaries. Such data structures are used to maintain a set of elements, while supporting membership queries as well as insertions and deletions of elements. Cuckoo hashing maintains such a dynamic dictionary by keeping two tables of size only slightly larger than the number of elements to be inserted, and two hash functions mapping the elements into cells of those tables. It then applies a clever algorithm for placing at most a single element in each cell. Each membership query requires just two memory accesses (in the worst case) and they are determined by the hash functions. Many variants of cuckoo hashing have been proposed since its introduction, and extensive literature has been devoted to its analysis (cf., [18, 15, 29, 19, 2]).

Pagh and Pagh [44] used ideas in the spirit of cuckoo hashing to construct efficient many-wise independent hash functions. Let \mathcal{H} , \mathcal{G} and \mathcal{F} be function families from \mathcal{D} to \mathcal{S} , from \mathcal{D} to \mathcal{R} and from \mathcal{S} to \mathcal{R} respectively, where \mathcal{R} is a group with operation \oplus . Define the function family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ from \mathcal{D} to \mathcal{R} as

$$\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F}) = (\mathcal{F} \circ \mathcal{H}) \oplus (\mathcal{F} \circ \mathcal{H}) \oplus \mathcal{G},$$

where $\mathcal{F}_1 \circ \mathcal{F}_2$, for function families \mathcal{F}_1 and \mathcal{F}_2 , is the function family whose members are the elements of $\mathcal{F}_1 \times \mathcal{F}_2$ and $(f_1, f_2)(x)$ is defined by $f_1(f_2(x))$ ($\mathcal{F}_1 \oplus \mathcal{F}_2$ is analogously defined). In other words, given $f_1, f_2 \in \mathcal{F}, h_1, h_2 \in \mathcal{H}$ and $g \in \mathcal{G}$, design a function $\mathcal{PP}_{f_1, f_2, h_1, h_2, g}(x) = f_1(h_1(x)) \oplus f_2(h_2(x)) \oplus g(x)$. Pagh and Pagh [44] showed that when the families \mathcal{H} and \mathcal{G} are of “high enough” independence, that is, roughly $(c \cdot \log |\mathcal{S}|)$ -wise independent, then the family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \Pi)$ is $O(|\mathcal{S}|^{-c})$ -indistinguishable from random by a $|\mathcal{S}|$ -query, *non-adaptive* distinguisher, where Π is the set of all functions from \mathcal{S} to \mathcal{R} . Note that the security of the resulting family goes

well beyond the birthday attack barrier: it is indistinguishable from random by an attacker making $|\mathcal{S}| \gg \sqrt{|\mathcal{S}|}$ queries.

Aumüller et al. [4] (building on the work of Dietzfelbinger and Woelfel [16]) strengthen the result of [44] by using more sophisticated hash functions \mathcal{H} and \mathcal{G} (rather than the $O(\log |\mathcal{S}|)$ -wise independent that [44] used). Specifically, for a given $z \geq 0$, Aumüller et al. [4] constructed a function family $\mathcal{ADW}_z(\mathcal{H}, \mathcal{G}, \Pi)$ that is $O(|\mathcal{S}|^{-(z+1)})$ -indistinguishable from random by a $|\mathcal{S}|$ -query, *non-adaptive* distinguisher, where Π is the set of all functions from \mathcal{S} to \mathcal{R} .² The idea to use more sophisticated hash functions, in the sense that they require less combinatorial work, already appeared in previous works, e.g., the work of Arbritman et al. [2, §5.4].

In Section 3 we take the above results a step further, showing that they hold also for *adaptive* distinguishers.³ Our approach for this transformation has many predecessors. For instance, the work of Naor and Reingold [40], and of Jetchev et al. [27]. Furthermore, it turns out that by using the above function family with a pseudorandom function \mathcal{F} instead of a truly random function from Π , namely the family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ (or $\mathcal{ADW}_z(\mathcal{H}, \mathcal{G}, \mathcal{F})$), we get a pseudorandom function that is superior to \mathcal{F} (the actual properties of $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ are determined by the properties of \mathcal{F} and the choice of \mathcal{H} and \mathcal{G}). This understanding is the main conceptual contribution of this paper, and the basis for the results presented below.

We note that the works of Pagh and Pagh [44], Dietzfelbinger and Woelfel [16], and Aumüller et al. [4] have gone almost unnoticed in the cryptography literature so far.⁴ In this work we apply, in a black-box manner, the analysis of [44] and of [4] in cryptographic settings.

1.3 Our Results

We use a construction inspired by cuckoo hashing to improve upon two PRF reductions: PRF domain extension and non-adaptive to adaptive PRF.

1.3.1 PRF Domain Extension

PRF domain extensions use PRFs with “small” domain size to construct PRFs with larger (or even unlimited) domain size. These extensions reduce the cost of a single invocation of the PRF and increase its usability. Domain extension methods are typically measured by the security of the resulting PRFs, and by the number of calls the resulting PRF makes to the underlying PRF.

Among the known domain extension techniques are CBC-MAC and PMAC (a survey on their security can be found in [39]). The number of calls made by these constructions to the underlying (small domain) PRF can be as small as two. Assuming that the underlying PRF is a random function over $\{0, 1\}^n$, then the resulting family is $(q, \infty, O(q^2/2^n))$ -PRF (i.e., the ∞ in the second parameter means that the distinguisher’s running time is unlimited). A second technique is the Feistel or Beneš transformations (e.g., [1, 46, 47], a survey of which can be found in [48]). The Beneš based construction makes 8 calls to the underlying PRF and is $(q, \infty, O(q/2^n))$ -PRF, whereas a

²The \mathcal{ADW} ’s function family is in fact more complicated than the above simplified description. See Section 6 for the formal definition.

³Note that in some cases an adaptive adversary is a more powerful distinguisher than a non-adaptive one. For example, when trying to distinguish between a truly random function and a random involution (permutations where the cycle length is at most 2), there exists an adaptive distinguisher that will succeed with very high probability by asking two queries while any non-adaptive distinguisher will fail with very high probability (see [28, 42]).

⁴Pagh and Pagh [44] did notice this connection, and in particular mentioned the connection of their work to that of Bellare et al. [7].

5-round Feistel based construction (which makes 5 calls to the underlying PRF) is $(q, \infty, O(q/2^n))$ -PRF.⁵

Our cuckoo hashing based function family (see below) is $(q, \infty, O(q/2^n))$ -PRF and makes only two calls to the underlying PRF. Moreover, our construction can extend the domain size to any fixed $\text{poly}(n)$ length, unlike the aforementioned constructions, which only double the domain size.⁶

Theorem 1.1 (informal). *Let $k \leq n$, let \mathcal{H} and \mathcal{G} be efficient k -wise independent function families mapping strings of length $\ell(n)$ to strings of length n , and let Π be the family of all functions from $\{0, 1\}^n$ to $\{0, 1\}^n$. Then, the family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \Pi)$, mapping strings of length $\ell(n)$ to strings of length n , is a $(q, \infty, q/2^{\Omega(k)})$ -PRF, for $q \leq 2^{n-2}$.*

For $k = \Theta(n)$, Theorem 1.1 yields a domain extension that is $(q, \infty, q/2^n)$ -PRF, and makes only two calls to the underlying PRF. Replacing in the above construction the function family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \Pi)$ with the family $\mathcal{ADW}_z(\mathcal{H}, \mathcal{G}, \Pi)$ yields a more versatile domain extension that offers a tradeoff between the number of calls to the PRF and the independence required for the hash functions. For details, see Section 6.2.

PRG to PRF reductions. Theorem 1.1 can also be used to get a hardness-preserving construction of PRFs from pseudorandom generators (PRG) in settings where there is a non-trivial bound on the number of queries to the PRF. Jain et al. [26], who were the first to propose this goal, noted that one can realize it using a domain extension constructions. Thus, we apply Theorem 1.1 to get constructions of PRFs from PRGs which improves some of the parameters of Jain et al. [26], but require longer keys. See Appendix B for details.

1.3.2 From Non-Adaptive to Adaptive PRF

Adaptive PRFs can be constructed from non-adaptive ones using general techniques such as using the PRG-based construction of Goldreich et al. [22] or the *synthesizer* based construction of Naor and Reingold [41]. These constructions, however, make (roughly) n calls to the underlying non-adaptive PRF (where n is the input length). Recently, Berman and Haitner [8] showed how to perform this security uplift at a much lower cost: the adaptive PRF makes only a *single* call to the non-adaptive PRF. Their construction, however, incurs a significant degradation in the security: assuming the underlying function is a non-adaptive t -PRF, then the resulting function is an (adaptive) $O(t^{1/3})$ -PRF. The reason for this degradation is the birthday attack we mentioned earlier.

We present a reduction from non-adaptive to adaptive PRFs that *preserves* the security of the non-adaptive PRF. The resulting adaptive PRF makes only two calls to the underlying non-adaptive one.

Theorem 1.2 (informal). *Let t be a polynomial-time computable integer function, let $\mathcal{H} = \{\mathcal{H}_n: \{0, 1\}^n \mapsto [4t(n)]_{\{0,1\}^n}\}_{n \in \mathbb{N}}$ (where $[4t(n)]_{\{0,1\}^n}$ are the first $4t(n)$ elements of $\{0, 1\}^n$) and $\mathcal{G} = \{\mathcal{G}_n: \{0, 1\}^n \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$ be efficient $O(\log t(n))$ -wise independent function families, and let \mathcal{F} be a length-preserving non-adaptive t -PRF. Then, $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ is a length-preserving $(t/4)$ -PRF.*

⁵This is by no means an exhaustive list of all domain extension constructions that achieve beyond-birthday security.

⁶Of course, one can use these constructions to extend the domain to any $\text{poly}(n)$ length by a recursive construction. This, however, will increase the number of calls to the underlying PRF by a polynomial factor.

Our construction, as well as the one of Berman and Haitner [8], depends on the number of queries made by the distinguisher. Thus, this does not give a single transformation for all poly-time adversaries. The PRG-based and synthesizer-based constructions are better in this sense as they are independent of the distinguisher.⁷

Finally, as it was the case in the domain extension, replacing the function family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ with the family $\mathcal{ADW}_z(\mathcal{H}, \mathcal{G}, \mathcal{F})$, yields a more versatile non-adaptive to adaptive transformation that offers a trade-off between the number of calls to the PRF and the requisite independence. See Section 6.3 for details.

1.4 More Related Work

Bellare et al. [7] introduced a paradigm for using PRFs in the symmetric-key settings that, in retrospect, is similar to cuckoo hashing. Assume that two parties, who share a secret function f would like to use it for (shared-key) encryption. The ‘textbook’ (stateless) solution calls for the sender to choose r at random and send $(r, f(r) \oplus M)$ to the receiver, where M is the message to be encrypted. This proposal breaks down if the sender chooses the same r twice (in two different sessions with different messages). Thus, the scheme is subject to the birthday attack and the length parameters should be chosen accordingly. This requires the underlying function to have a large domain. Instead, [7] suggested choosing $t > 1$ values at random, and sending $(r_1, \dots, r_t, f(r_1) \oplus \dots \oplus f(r_t) \oplus M)$. They were able to show much better security than the single r case. They also showed a similar result for message authentication. Our domain extension results (see Section 4) improve upon the results of [7].

The problem of transforming a scheme that is only resilient to non-adaptive attack into one that is resilient to adaptive attacks has received quite a lot of attention in the context of pseudo-random permutations (or block ciphers). Maurer and Pietrzak [34] showed that, given a family of permutations that are information-theoretic secure against non-adaptive attacks, if two members of this family are independently composed, then the resulting permutation is also secure against adaptive attacks (see [34] for the exact formulation). Pietrzak [49] showed, however, that this is not necessarily the case for permutations that are random-looking under a computational assumption (see also [38, 50]), reminding us that translating information-theoretic results to the computational realm is a tricky business.

Paper Organization

Basic notations and formal definitions are given in Section 2. In Section 3 we formally define the hashing paradigm of Pagh and Pagh [44] and show how to extend their result to hold against adaptive adversaries. Our domain extension reduction based on [44] is described in Section 4, and the improved non-adaptive to adaptive reduction, also based on [44], is described in Section 5. In Section 6 we present the more advanced (and more complex) hashing paradigm of Aumüller et al. [4], and use it to obtain a more versatile version of the above reductions. Some possible directions for future research are discussed in Section 7.

⁷Berman and Haitner [8] do give a single transformation for all poly-time adversaries with the additional cost of slightly increasing the number of calls to the underlying non-adaptive PRF. Their transformation can be carried to our setting and be applied with the \mathcal{PP} function family. However, our construction is superior to [8]’s only by reducing the security loss in a polynomial factor, a factor that is meaningless when considering all poly-time adversaries. Thus, we avoid stating our result in this all poly-time adversaries setting and refer to [8] for such a result.

2 Preliminaries

2.1 Notations

All logarithms considered here are in base two. We use calligraphic letters to denote sets, uppercase for random variables, and lowercase for values. Let ‘||’ denote string concatenation. For an integer $t \in \mathbb{N}$, let $[t] = \{1, \dots, t\}$. For a set \mathcal{S} , let \mathcal{S}^* be the power set of \mathcal{S} (i.e., the set of all subsets of \mathcal{S}). For a set \mathcal{S} and an integer $t \in \mathbb{N}$, let $\mathcal{S}^{\leq t} = \{\bar{s} \in \mathcal{S}^* : |\bar{s}| \leq t \wedge \bar{s}[i] \neq \bar{s}[j] \ \forall i \neq j \in [|\bar{s}|]\}$, where $\bar{s}[i]$ is the i th element of \bar{s} , and let $[t]_{\mathcal{S}}$ be the first t elements, in increasing lexicographic order, of \mathcal{S} (equal to \mathcal{S} in case $|\mathcal{S}| \leq t$). For sets \mathcal{U} and \mathcal{V} , let $\Pi_{\mathcal{U} \rightarrow \mathcal{V}}$ stands for the set of all functions from \mathcal{U} to \mathcal{V} , and for integers n and ℓ , let $\Pi_{n,\ell} = \Pi_{\{0,1\}^n \rightarrow \{0,1\}^\ell}$.

Let poly denote the set all polynomials, and let PPTM be abbreviation for probabilistic (strictly) polynomial-time Turing machine. For $s \in \mathbb{N}$ and $t, q: \mathbb{N} \mapsto \mathbb{N}$, we say that D is a t -time q -query s -oracle-aided algorithm if, when invoked on input of length n , D runs in time $t(n)$ and makes at most $q(n)$ queries to each of its s oracles. Given a random variable X , we write $X(x)$ to denote $\Pr[X = x]$, and write $x \leftarrow X$ to indicate that x is selected according to X . Similarly, given a finite set \mathcal{S} , we let $s \leftarrow \mathcal{S}$ denote that s is selected according to the uniform distribution on \mathcal{S} . The *statistical distance* of two distributions P and Q over a finite set \mathcal{U} , denoted as $\text{SD}(P, Q)$, is defined as $\max_{\mathcal{S} \subseteq \mathcal{U}} |P(\mathcal{S}) - Q(\mathcal{S})| = \frac{1}{2} \sum_{u \in \mathcal{U}} |P(u) - Q(u)|$.

2.2 Pseudorandom Generators

Definition 2.1 (Pseudorandom Generators). *A polynomial-time function $G: \{0, 1\}^n \mapsto \{0, 1\}^{\ell(n)}$ is (t, ε) -PRG, if $\ell(n) > n$ for every $n \in \mathbb{N}$ (G stretches the input), and*

$$\left| \Pr_{x \leftarrow \{0,1\}^n} [D(G(x)) = 1] - \Pr_{y \leftarrow \{0,1\}^{\ell(n)}} [D(y) = 1] \right| \leq \varepsilon(n)$$

for every algorithm (distinguisher) D of running time $t(n)$ and large enough n . A $(t, 1/t)$ -PRG is called a t -PRG. If $\ell(n) = 2n$, we say that G is length-doubling.

2.3 Function Families

2.3.1 Operating on Function Families

We consider two natural operations on function families.

Definition 2.2 (composition of function families). *Let $\mathcal{F}^1: \mathcal{D}^1 \mapsto \mathcal{R}^1$ and $\mathcal{F}^2: \mathcal{D}^2 \mapsto \mathcal{R}^2$ be two function families with $\mathcal{R}^1 \subseteq \mathcal{D}^2$. The composition of \mathcal{F}^1 with \mathcal{F}^2 , denoted $\mathcal{F}^2 \circ \mathcal{F}^1$, is the function family $\{(f_2, f_1) \in \mathcal{F}^2 \times \mathcal{F}^1\}$, where $(f_2, f_1)(x) := f_2(f_1(x))$.*

Definition 2.3 (group operation of function families). *Let $\mathcal{F}^1: \mathcal{D} \mapsto \mathcal{R}^1$ and $\mathcal{F}^2: \mathcal{D} \mapsto \mathcal{R}^2$ be two function families with $\mathcal{R}^1, \mathcal{R}^2 \subseteq \mathcal{R}$, where \mathcal{R} is a group with operation \oplus . The group operation of \mathcal{F}^1 with \mathcal{F}^2 , denoted $\mathcal{F}^2 \oplus \mathcal{F}^1$, is the function family $\{(f_2, f_1) \in \mathcal{F}^2 \times \mathcal{F}^1\}$, where $(f_2, f_1)(x) := f_2(x) \oplus f_1(x)$.*

In all of our applications, the group \mathcal{R} from the above definition will simply be $\{0, 1\}^n$ for some $n \in \mathbb{N}$, with XOR as the group operation.

2.3.2 Function Family Ensembles

A function family ensemble is an infinite set of function families, whose elements (families) are typically indexed by the set of integers. Let $\mathcal{F} = \{\mathcal{F}_n : \mathcal{D}_n \mapsto \mathcal{R}_n\}_{n \in \mathbb{N}}$ stands for an ensemble of function families, where each $f \in \mathcal{F}_n$ has domain \mathcal{D}_n and its range is a subset of \mathcal{R}_n . Such ensemble is *length preserving*, if $\mathcal{D}_n = \mathcal{R}_n = \{0, 1\}^n$ for every n . We naturally extend Definitions 2.2 and 2.3 to function family ensembles.

For a function family ensemble to be useful it should have an efficient sampling and evaluation algorithms.

Definition 2.4 (efficient function family ensembles). *A function family ensemble $\mathcal{F} = \{\mathcal{F}_n : \mathcal{D}_n \mapsto \mathcal{R}_n\}_{n \in \mathbb{N}}$ is efficient, if the following hold:*

Efficient sampling. *\mathcal{F} is samplable in polynomial-time: there exists a PPTM that given 1^n , outputs (the description of) a uniform element in \mathcal{F}_n .*

Efficient evaluation. *There exists a deterministic algorithm that given $x \in \mathcal{D}_n$ and (a description of) $f \in \mathcal{F}_n$, runs in time $\text{poly}(n, |x|)$ and outputs $f(x)$.*

2.3.3 Many-Wise Independent Hashing

Definition 2.5 (k -wise independent families). *A function family $\mathcal{H} = \{h : \mathcal{D} \mapsto \mathcal{R}\}$ is k -wise independent (with respect to \mathcal{D} and \mathcal{R}), if*

$$\Pr_{h \leftarrow \mathcal{H}}[h(x_1) = y_1 \wedge h(x_2) = y_2 \wedge \dots \wedge h(x_k) = y_k] = \frac{1}{|\mathcal{R}|^k},$$

for every distinct $x_1, x_2, \dots, x_k \in \mathcal{D}$ and every $y_1, y_2, \dots, y_k \in \mathcal{R}$.

For every $\ell, k \in \text{poly}$, the existence of efficient $k(n)$ -wise independent family ensembles mapping strings of length $\ell(n)$ to strings of length n is well known ([12, 54]). A simple and well known example of k -wise independent functions is the collection of all polynomials of degree $(k - 1)$ over a finite field. This construction has small size, and each evaluation of a function at a given point requires k operations in the field.

Fact 2.6. *For $\ell, n, k \in \mathbb{N}$, there exists a k -wise independent function family $\mathcal{H} = \{h : \{0, 1\}^\ell \mapsto \{0, 1\}^n\}$, such that sampling a random element in \mathcal{H} requires $k \cdot \max\{\ell, n\}$ random bits, and evaluating a function from \mathcal{H} is done in time $\text{poly}(\ell, n, k)$.*

We mention that a k -wise independent families (as defined in Definition 2.5) look random for k -query distinguishers, both non-adaptive and adaptive ones. On the other hand, *almost* k -wise independent families⁸ are only granted to be resistant against *non-adaptive* distinguishers.⁹ Yet, the result presented in Section 3 yields that, in some cases, the adaptive security of the latter families follows from their non-adaptive security.

⁸Formally, a function family $\mathcal{H} = \{h : \mathcal{D} \mapsto \mathcal{R}\}$ is (ε, k) -wise independent if for any $x_1, \dots, x_k \in \mathcal{D}$ and for any $y_1, \dots, y_k \in \mathcal{R}$ it holds that $\left| \Pr_{h \leftarrow \mathcal{H}}[h(x_1) = y_1 \wedge \dots \wedge h(x_k) = y_k] - |\mathcal{R}|^{-k} \right| \leq \varepsilon$. We call a family of functions an almost k -wise independent family, if it is (ε, k) -wise independent for some small $\varepsilon > 0$.

⁹See Footnote 3 and references therein.

2.4 Pseudorandom Functions

Definition 2.7 (Pseudorandom Functions). *An efficient function family ensemble $\mathcal{F} = \{\mathcal{F}_n: \{0, 1\}^{m(n)} \mapsto \{0, 1\}^{\ell(n)}\}_{n \in \mathbb{N}}$ is an (adaptive) (q, t, ε) -PRF, if for every t -time q -query oracle-aided algorithm (distinguisher) D , it holds that*

$$\left| \Pr_{f \leftarrow \mathcal{F}_n} [\mathsf{D}^f(1^n) = 1] - \Pr_{\pi \leftarrow \Pi_{m(n), \ell(n)}} [\mathsf{D}^\pi(1^n) = 1] \right| \leq \varepsilon(n),$$

for large enough n . If $q(n)$ is only bounded by $t(n)$ for every $n \in \mathbb{N}$, then \mathcal{F} is called (t, ε) -PRF. A $(t, 1/t)$ -PRF is called a t -PRF.

If D is limited to be non-adaptive (i.e., it has to write all his oracle calls before making the first call), then \mathcal{F} is called non-adaptive (q, t, ε) -PRF (and we apply the above notational conventions also for this case).

Some applications require the pseudorandom functions to be secure against distinguishers with access to many oracles (and not just a single oracle as in Definition 2.7).

Definition 2.8 (pseudorandom functions secure against many-oracle distinguishers). *An efficient function family ensemble $\mathcal{F} = \{\mathcal{F}_n: \{0, 1\}^{m(n)} \mapsto \{0, 1\}^{\ell(n)}\}_{n \in \mathbb{N}}$ is an s -oracle (q, t, ε) -PRF, if for every t -time q -query s -oracle-aided algorithm (distinguisher) D , it holds that*

$$\left| \Pr_{\bar{f} \leftarrow \mathcal{F}_n^s} [\mathsf{D}^{\bar{f}}(1^n) = 1] - \Pr_{\bar{\pi} \leftarrow \Pi_{m(n), \ell(n)}^s} [\mathsf{D}^{\bar{\pi}}(1^n) = 1] \right| \leq \varepsilon(n),$$

for large enough n .

The following lemma shows that a standard (single oracle) PRF is also a many-oracle one, with a loss that depends multiplicatively on the number of oracles. This lemma has been used several times in the past, for example, in [6, Lemma 3.3] and [10, Theorem 1]. The lemma is proven using a standard hybrid argument, see [6, 10].¹⁰

Lemma 2.9. *Let $\mathcal{F} = \{\mathcal{F}_n: \{0, 1\}^{m(n)} \mapsto \{0, 1\}^{\ell(n)}\}_{n \in \mathbb{N}}$ be a function family ensemble. Then for every t -time q -query s -oracle adaptive [resp., non-adaptive] distinguisher D , there exists a $(t + s \cdot q \cdot e_{\mathcal{F}})$ -time q -query single-oracle adaptive [resp., non-adaptive] distinguisher $\widehat{\mathsf{D}}$, where $e_{\mathcal{F}}$ stands for the evaluation time of \mathcal{F} ,¹¹ with*

$$\begin{aligned} & \left| \Pr_{f \leftarrow \mathcal{F}_n} [\widehat{\mathsf{D}}^f(1^n) = 1] - \Pr_{\pi \leftarrow \Pi_{m(n), \ell(n)}} [\widehat{\mathsf{D}}^\pi(1^n) = 1] \right| \\ & \geq \frac{1}{s} \cdot \left| \Pr_{\bar{f} \leftarrow \mathcal{F}_n^s} [\mathsf{D}^{\bar{f}}(1^n) = 1] - \Pr_{\bar{\pi} \leftarrow \Pi_{m(n), \ell(n)}^s} [\mathsf{D}^{\bar{\pi}}(1^n) = 1] \right|, \end{aligned}$$

for every $n \in \mathbb{N}$.

3 From Non-Adaptive to Adaptive Hashing

In this section we show that non-adaptively secure function families with a certain combinatorial property are also adaptively secure, yielding that the function families of Pagh and Pagh [44]

¹⁰[10] only states and proves the adaptive case, but the very same argument also yields the non-adaptive case.

¹¹That is, $\mathsf{D}(1^n)$ runs in time $t(n) + s \cdot q(n) \cdot e_{\mathcal{F}}(n)$. Moreover, we implicitly assume that the evaluation time of \mathcal{F} is greater than the time needed to sample $\ell(n)$ random bits.

and Aumüller et al. [4] are adaptively secure. In the next sections we take advantage of the latter implication to derive our hardness-preserving PRF reductions. We note that the above approach is not useful for (arbitrary) non-adaptive PRFs, since a PRF might not possess the required combinatorial property (indeed, not every non-adaptive PRF is an adaptive one).

To define the aforementioned combinatorial property, we use the notion of left-monotone sets.

Definition 3.1 (left-monotone sets). *Let \mathcal{S} and \mathcal{T} be sets. A set $\mathcal{M} \subseteq \mathcal{S}^* \times \mathcal{T}$ is left-monotone, if for every $(\overline{s}_1, t) \in \mathcal{M}$ and every $\overline{s}_2 \in \mathcal{S}^*$ that has \overline{s}_1 as a prefix, it holds that $(\overline{s}_2, t) \in \mathcal{M}$.*

Namely, a product set is left monotone, if it is monotone with respect to its left-hand-side part, where all sequences having a prefix in a monotone set, are also in the set.

The main result of this section states roughly that for a function family $\mathcal{F} = \{f_{u,v}: \mathcal{D} \mapsto \mathcal{R}\}_{(u,v) \in \mathcal{U} \times \mathcal{V}}$ and a left-monotone set $\text{BAD} \subseteq \mathcal{D}^* \times \mathcal{U}$, if for every set of queries $\overline{q} = (q_1, \dots, q_{|\overline{q}|})$ it holds that (1) for every $u \in \mathcal{U}$ such that $(\overline{q}, u) \notin \text{BAD}$, the outputs of f on $q_1, \dots, q_{|\overline{q}|}$ is uniform over $\mathcal{R}^{|\overline{q}|}$, and (2) the probability over the choice of $u \leftarrow \mathcal{U}$ that $(\overline{q}, u) \in \text{BAD}$ is small, then \mathcal{F} is adaptively secure.

Lemma 3.2. *Let \mathcal{U} and \mathcal{V} be non-empty sets, let $\mathcal{F} = \mathcal{F}(\mathcal{U}, \mathcal{V}) = \{f_{u,v}: \mathcal{D} \mapsto \mathcal{R}\}_{(u,v) \in \mathcal{U} \times \mathcal{V}}$ be a function family and let $\text{BAD} \subseteq \mathcal{D}^* \times \mathcal{U}$ be left-monotone. Let $t \in \mathbb{N}$, and assume that for every $\overline{q} = (q_1, \dots, q_{|\overline{q}|}) \in \mathcal{D}^{\leq t}$ it holds that¹²*

1. $(f(q_1), \dots, f(q_{|\overline{q}|}))_{f \leftarrow \{f_{u,v}: v \in \mathcal{V}\}}$ is uniform over $\mathcal{R}^{|\overline{q}|}$, for every $u \in \mathcal{U}$ with $(\overline{q}, u) \notin \text{BAD}$, and
2. $\Pr_{u \leftarrow \mathcal{U}}[(\overline{q}, u) \in \text{BAD}] \leq \varepsilon$.

Then,

$$\left| \Pr_{\substack{u \leftarrow \mathcal{U} \\ v \leftarrow \mathcal{V}}}[\mathsf{D}^{f_{u,v}} = 1] - \Pr_{\pi \leftarrow \Pi}[\mathsf{D}^\pi = 1] \right| \leq \varepsilon$$

for every t -query oracle-aided adaptive algorithm D , letting Π be the set of all functions from \mathcal{D} to \mathcal{R} .

Note that the above properties of the family \mathcal{F} mean that \mathcal{F} is non-adaptively secure. The proof that \mathcal{F} is also *adaptively* secure, which critically uses the above structure of the set BAD , can be found in Appendix A.¹³

3.1 The Pagh and Pagh [44] Function Family

We show that Lemma 3.2 can be applied to the function family of Pagh and Pagh [44].

Definition 3.3 (The Pagh and Pagh [44] function family). *Let \mathcal{H} be a function family from \mathcal{D} to \mathcal{U} , let \mathcal{G} be a function family from \mathcal{D} to \mathcal{R} and let \mathcal{F} be a function family from \mathcal{S} to \mathcal{R} , with $\mathcal{U} \subseteq \mathcal{S}$ and \mathcal{R} being a group with respect to the operation \oplus . The function family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ from \mathcal{D} to \mathcal{R} , is defined by*

$$\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F}) := (\mathcal{F} \circ \mathcal{H}) \oplus (\mathcal{F} \circ \mathcal{H}) \oplus \mathcal{G}.$$

For $h_1, h_2 \in \mathcal{H}$, let $\mathcal{PP}_{h_1, h_2}(\mathcal{G}, \mathcal{F}) := (\mathcal{F} \circ h_1) \oplus (\mathcal{F} \circ h_2) \oplus \mathcal{G}$.




Figure 1: The function family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$. \mathcal{H} hashes down a domain \mathcal{D} to a domain \mathcal{S} . Then \mathcal{F} maps \mathcal{S} to \mathcal{R} . We do this twice and xor it with \mathcal{G} , that hashes the domain \mathcal{D} directly to \mathcal{R} .

Graphically, this function family is given in Figure 1. Pagh and Pagh [44] showed that when instantiated with the proper function families, the above function family has the following properties:

Theorem 3.4 ([44]). *Let $t \in \mathbb{N}$, let $\mathcal{H} = \{h: \mathcal{D} \mapsto \mathcal{U}\}$ and $\mathcal{G} = \{g: \mathcal{D} \mapsto \mathcal{R}\}$ be function families with \mathcal{R} being a group with respect to the operation \oplus , and let $\Pi = \Pi_{\mathcal{S} \mapsto \mathcal{R}}$.*

If $\mathcal{U} \subseteq \mathcal{S}$ and $|\mathcal{U}| \geq 4t$, then for every $k \in \mathbb{N}$ there exists a left-monotone set $\text{BAD} \subseteq \mathcal{D}^{\leq t} \times \mathcal{H}^2$ such that the following holds for every $\bar{q} = (q_1, \dots, q_{|\bar{q}|}) \in \mathcal{D}^{\leq t}$:

1. *Assuming that \mathcal{G} is k -wise independent over the elements of \bar{q} , then $(f(q_1), \dots, f(q_{|\bar{q}|}))_{f \leftarrow \mathcal{PP}_{h_1, h_2}(\mathcal{G}, \Pi)}$ is uniform over $\mathcal{R}^{|\bar{q}|}$ for every $u \in \mathcal{U}$ such that $(\bar{q}, u) \notin \text{BAD}$.*
2. *Assuming that \mathcal{H} is k -wise independent over the elements of \bar{q} , then $\Pr_{u \leftarrow \mathcal{H}^2}[(\bar{q}, u) \in \text{BAD}] \leq t/2^{\Omega(k)}$.¹⁴*

Pagh and Pagh [44] concluded that for (the many) applications where the analysis is applied with respect to a static set it is safe to use this family instead. However, as we can see, the function family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \Pi)$ is not only close to being uniform in the eyes of a non-adaptive distinguisher, but also allows us to apply Lemma 3.2 to deduce its security in the eyes of adaptive distinguishers.

¹²Recall that for a set \mathcal{S} and an integer t , $\mathcal{S}^{\leq t}$ denotes the set $\{\bar{s} \in \mathcal{S}^*: |\bar{s}| \leq t \wedge \bar{s}[i] \neq \bar{s}[j] \ \forall i \neq j \in [|\bar{s}|]\}$.

¹³Lemma 3.2 can be derived as a special case of a result given in [27, Theorem 12] (closing a gap in the proof appearing in [33]). Yet, for the sake of completeness, we include an independent proof of this lemma here.

¹⁴The function family we consider above (i.e., \mathcal{PP}) is slightly different than the one given in [44]. Their construction maps element $x \in \mathcal{D}$ to $F_1[h_1(x)] \oplus F_2[h_2(x)] \oplus g(x)$, where F_1 and F_2 are uniformly chosen vectors from \mathcal{R}^t , $h_1, h_2: \mathcal{D} \mapsto [t]$ are uniformly chosen from a function family \mathcal{H} and $g: \mathcal{D} \mapsto \mathcal{R}$ is chosen uniformly from a function family \mathcal{G} . Yet, the correctness of Theorem 3.4 follows in a straightforward manner from [44] original proof (specifically from Lemma 3.3 and Lemma 3.4).

By plugging in Theorem 3.4 into the general framework lemma (Lemma 3.2), we get the following result:

Lemma 3.5. *Let $t \in \mathbb{N}$, let \mathcal{H} , \mathcal{G} and Π be as in Theorem 3.4, and let D be an adaptive, t -query oracle-aided algorithm. Assuming that \mathcal{H} and \mathcal{G} are k -wise independent, then*

$$\left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}, \mathcal{G}, \Pi)}[D^f = 1] - \Pr_{\pi \leftarrow \Pi}[D^\pi = 1] \right| \leq t/2^{\Omega(k)}.$$

Proof. Let $\mathcal{U} = \mathcal{H} \times \mathcal{H}$ and $\mathcal{V} = \Pi \times \Pi \times \mathcal{G}$. For $(h_1, h_2) \in \mathcal{U}$ and $(\pi_1, \pi_2, g) \in \mathcal{V}$, let $F_{(h_1, h_2), (\pi_1, \pi_2, g)} = \pi_1 \circ h_1 \oplus \pi_2 \circ h_2 \oplus g$, and let $\mathcal{F} = \{F_{u,v} : \mathcal{D} \mapsto \mathcal{R}\}_{(u,v) \in \mathcal{U} \times \mathcal{V}}$. Finally, let BAD be the set BAD of Theorem 3.4. We prove the lemma showing that the above sets meet the requirements stated in Lemma 3.2.

Item 1 of Theorem 3.4 and the assumed independence of \mathcal{G} and \mathcal{H} , yield that the first requirement of Lemma 3.2 is satisfied. Item 2 of Theorem 3.4 yields that the second requirement of Lemma 3.2 is satisfied for $\varepsilon = t/2^{\Omega(k)}$. Hence, the proof of the lemma follows by Lemma 3.2. \square

Remark 3.6. *For some of our applications, see Sections 4 to 6, we need to apply Lemma 3.5 with efficient k -wise independent function family ensembles mapping strings of length n to the set $[t(n)]_{\{0,1\}^n}$, where t is an efficiently computable function. It is easy to see (cf., [8]) that such ensembles exist for any efficiently computable t that is a power of two. By considering $t'(n) = 2^{\lfloor \log(t(n)) \rfloor}$, we use these ensembles for our applications, while only causing factor of two loss in the resulting security.*

We use the above function family of Pagh and Pagh [44] to extend the domain of pseudorandom functions (see Section 4), and to transform a non-adaptive pseudorandom function into an adaptive one (see Section 5). In Section 6 we instantiate the above framework with the more advanced function family of Aumüller et al. [4], to get more versatile variants of the above applications.

4 PRF Domain Extension

In this section we use the function family \mathcal{PP} of Pagh and Pagh [44] (see Section 3) to extend a domain of a given PRF.

Theorem 4.1 (Restating Theorem 1.1). *Let $\mathcal{H} = \{\mathcal{H}_n : \{0,1\}^{d(n)} \mapsto \{0,1\}^{s(n)}\}_{n \in \mathbb{N}}$ and $\mathcal{G} = \{\mathcal{G}_n : \{0,1\}^{d(n)} \mapsto \{0,1\}^{r(n)}\}_{n \in \mathbb{N}}$ be efficient $k(n)$ -wise independent function family ensembles, and let $\mathcal{F} = \{\mathcal{F}_n : \{0,1\}^{s(n)} \mapsto \{0,1\}^{r(n)}\}_{n \in \mathbb{N}}$ be a (q, t, ε) -PRF. Then,*

$$\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F}) = \{\mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n) : \{0,1\}^{d(n)} \mapsto \{0,1\}^{r(n)}\}_{n \in \mathbb{N}}$$

is a $(q, t - p \cdot q, 2\varepsilon + q/2^{\Omega(k)})$ -PRF, where p is a polynomial determined by the evaluation and sampling time of \mathcal{H} , \mathcal{G} and \mathcal{F} and $q(n) \leq 2^{s(n)-2}$ for every $n \in \mathbb{N}$.

To prove Theorem 4.1 we first show that $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ is computationally close to $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \Pi)$, where \mathcal{F} is a PRF family and Π is the set of all functions. Then, we invoke Lemma 3.5 that shows that $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \Pi)$ is statistically close to the set of all functions with appropriate domain and range.

Proof of Theorem 4.1. Let \mathcal{H}, \mathcal{G} and \mathcal{F} be as in the statement.

Claim 4.2. For every $(t - 2q \cdot (e_{\mathcal{H}} + e_{\mathcal{G}} + e_{\mathcal{F}}))$ -time q -query distinguisher D , where $e_{\mathcal{H}}, e_{\mathcal{G}}, e_{\mathcal{F}}: \mathbb{N} \mapsto \mathbb{N}$ are the evaluation and sampling times of \mathcal{H} , \mathcal{G} and \mathcal{F} , respectively, and for all large enough n :

$$\left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)}[D^f(1^n) = 1] - \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_{s(n), r(n)})}[D^f(1^n) = 1] \right| \leq 2\varepsilon(n).$$

Proof. Assume that there is a distinguisher D as in the statement for which

$$\left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)}[D^f(1^n) = 1] - \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_{s(n), r(n)})}[D^f(1^n) = 1] \right| > 2\varepsilon(n).$$

We show how it can be used to break the PRF. We define a $(t - 2q(n) \cdot e_{\mathcal{F}}(n))$ -time q -query two-oracle distinguisher \tilde{D} :

Algorithm 4.3 (\tilde{D}).

Input: 1^n .

Oracle: functions ϕ_1, ϕ_2 from \mathcal{D}_n to \mathcal{R}_n .

1. Set $h_1, h_2 \leftarrow \mathcal{H}_n, g \leftarrow \mathcal{G}_n$.
2. Set $f = (\phi_1 \circ h_1) \oplus (\phi_2 \circ h_2) \oplus g$.
3. Emulate $D^f(1^n)$.

Note that $\tilde{D}(1^n)$ makes $q(n)$ queries to its oracle where each query consists of a call to ϕ_1 and ϕ_2 and an evaluation of h_1, h_2 and g , so it can be implemented to run in time as D plus at most $2q(n) \cdot (e_{\mathcal{H}}(n) + e_{\mathcal{G}}(n))$ which is exactly $(t - 2q(n) \cdot e_{\mathcal{F}}(n))$. Observe that in case ϕ_1 and ϕ_2 are uniformly drawn from \mathcal{F}_n , then the emulation of $D^f(1^n)$ done in $\tilde{D}^{\phi_1, \phi_2}(1^n)$ is identical to a random execution of $D^f(1^n)$ with $f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)$. Similarly, in case ϕ_1 and ϕ_2 are uniformly drawn from $\Pi_{s(n), r(n)}$, then the emulation is identical to a random execution of $D^f(1^n)$ with $f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_{s(n), r(n)})$. Thus,

$$\begin{aligned} & \left| \Pr_{(f_1, f_2) \leftarrow \mathcal{F}_n \times \mathcal{F}_n}[\tilde{D}^{f_1, f_2}(1^n) = 1] - \Pr_{(\pi_1, \pi_2) \leftarrow \Pi_{s(n), r(n)} \times \Pi_{s(n), r(n)}}[\tilde{D}^{\pi_1, \pi_2}(1^n) = 1] \right| \\ &= \left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)}[D^f(1^n) = 1] - \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_{s(n), r(n)})}[D^f(1^n) = 1] \right|. \end{aligned}$$

Hence, Lemma 2.9 yields that there exists a single-oracle distinguisher \hat{D} that when invoked on input of length n makes $q(n)$ queries to its oracle and runs in time $t(n) + 2q(n) \cdot (e_{\mathcal{H}}(n) + e_{\mathcal{G}}(n)) + 2q(n) \cdot e_{\mathcal{F}}(n) \leq t(n) + 2q(n) \cdot (e_{\mathcal{H}}(n) + e_{\mathcal{G}}(n) + e_{\mathcal{F}}(n))$, such that

$$\begin{aligned} & \left| \Pr_{f \leftarrow \mathcal{F}_n}[\hat{D}^f(1^n) = 1] - \Pr_{\pi \leftarrow \Pi_{s(n), r(n)}}[\hat{D}^\pi(1^n) = 1] \right| \\ & \geq \frac{1}{2} \cdot \left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)}[D^f(1^n) = 1] - \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_{s(n), r(n)})}[D^f(1^n) = 1] \right| > \varepsilon(n), \end{aligned}$$

which is a contradiction. \square

Let D be a q -query $(t - p \cdot q)$ -time distinguisher, for $p := 2(e_{\mathcal{H}} + e_{\mathcal{G}} + e_{\mathcal{F}})$. Since \mathcal{H} and \mathcal{G} are k -wise independent function families, we invoke Lemma 3.5 (instantiated with the domain of bit strings by setting $t = q$, $\mathcal{D} = \{0, 1\}^d$, $\mathcal{U} = \mathcal{S} = \{0, 1\}^s$ and $\mathcal{R} = \{0, 1\}^r$; note that $|U| = 2^s \geq 4q = 4t$) and using Claim 4.2 and the triangle inequality to get that

$$\left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)}[D^f(1^n) = 1] - \Pr_{\pi \leftarrow \Pi_{d(n), r(n)}}[D^\pi(1^n) = 1] \right| \leq 2\varepsilon(n) + q(n)/2^{\Omega(k(n))}$$

for large enough n . \square

Note that in order for Theorem 4.1 to be useful, we have to set $k(n) = \Omega(\log q(n))$. In Section 6 we show how to achieve domain extension using functions with less independence, but with the cost of additional calls to the PRF.

5 From Non-Adaptive to Adaptive PRF

In this section we use the function family \mathcal{PP} of Pagh and Pagh [44] (see Section 3), to transform an non-adaptive PRF into an adaptive one in a security preserving manner. To ease notations, we assume that the given non-adaptive PRF is length preserving.

Theorem 5.1. *Let q be a polynomial-time computable integer function with $q(n) \leq 2^{n-2}$ for every $n \in \mathbb{N}$, let $\mathcal{H} = \{\mathcal{H}_n: \{0, 1\}^n \mapsto [4q(n)]_{\{0,1\}^n}\}_{n \in \mathbb{N}}$ and $\mathcal{G} = \{\mathcal{G}_n: \{0, 1\}^n \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$ be efficient $(c \cdot \log q)$ -wise independent function family ensembles, where $c > 0$ is universal.*

If $\mathcal{F} = \{\mathcal{F}_n: \{0, 1\}^n \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$ is a non-adaptive $(4q, p \cdot t, \varepsilon)$ -PRF for some $p \in \text{poly}$ determined by the evaluation time of $q, \mathcal{H}, \mathcal{G}$ and \mathcal{F} , then $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ is an adaptive $(q, t, 2\varepsilon + 1/q)$ -PRF.

Theorem 1.2 is a special case of Theorem 5.1, starting with a $(p \cdot q)$ -PRF, where $p \in \text{poly}$ is determined by q and the evaluation time of \mathcal{H}, \mathcal{G} and \mathcal{F} (assuming that $p(n) \geq 4$), and ending with $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ which is an adaptive q -PRF.

To prove Theorem 5.1 we begin by showing that if $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ is distinguishable from $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \Pi)$ by an adaptive distinguisher, where \mathcal{F} is non-adaptive pseudorandom function and Π is a truly random function with the same domain and range as \mathcal{F} , then there is a non-adaptive distinguisher that can succeed in the same task almost as well. Then, we invoke Lemma 3.5 to get the theorem.

Proof of Theorem 5.1. Let \mathcal{H}, \mathcal{G} and \mathcal{F} be as in the statement and let $p = e_q + 8q(e_{\mathcal{H}} + e_{\mathcal{G}} + e_{\mathcal{F}})$, where e_q is the evaluation time of q , and $e_{\mathcal{H}}, e_{\mathcal{G}}$ and $e_{\mathcal{F}}$ are the sampling and evaluation time of \mathcal{H}, \mathcal{G} and \mathcal{F} , respectively.

Claim 5.2. *For every t -time q -query oracle-aided adaptive distinguisher \mathbb{D} , there exists a $(p \cdot t)$ -time $(4q)$ -query, non-adaptive, oracle-aided distinguisher $\tilde{\mathbb{D}}$ such that for every $n \in \mathbb{N}$ and $q(n) \leq 2^{n-2}$:*

$$\begin{aligned} & \left| \Pr_{f \leftarrow \mathcal{F}_n} [\widehat{\mathbb{D}}^f(1^n) = 1] - \Pr_{\pi \leftarrow \Pi_n} [\widehat{\mathbb{D}}^\pi(1^n) = 1] \right| \\ & \geq \frac{1}{2} \cdot \left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)} [\mathbb{D}^f(1^n) = 1] - \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_n)} [\mathbb{D}^f(1^n) = 1] \right|. \end{aligned}$$

Proof. The proof follows along similar lines to the proof of [8, Lemma 3.3]. Let $\tilde{\mathbb{D}}$ be the following two-oracle distinguisher:

Algorithm 5.3 ($\tilde{\mathbb{D}}$).

Input: 1^n .

Oracles: Functions ϕ_1 and ϕ_2 from $\{0, 1\}^n$ to $\{0, 1\}^n$.

1. Compute $\phi_1(x)$ and $\phi_2(x)$ for every $x \in [4q(n)]_{\{0,1\}^n}$.
2. Set $f = (\phi_1 \circ h_1) \oplus (\phi_2 \circ h_2) \oplus g$, where $h_1, h_2 \leftarrow \mathcal{H}_n$ and $g \leftarrow \mathcal{G}_n$.

3. Emulate $D^f(1^n)$: answer a query x to ϕ_1 and ϕ_2 made by D with $f(x)$, using the information obtained in Step 1.

Note that $\tilde{D}(1^n)$ makes $4q(n)$ non-adaptive queries to ϕ_1 and ϕ_2 , and it can be implemented to run in time $e_q(n) + 8q(n) + t(n) + q(n) \cdot (e_{\mathcal{H}}(n) + e_{\mathcal{G}}(n))$. Observe that the definition of \mathcal{PP} guarantees that ϕ_1 and ϕ_2 are only queried on the first $4q(n)$ elements of their domain. Hence, in case ϕ_1 and ϕ_2 are uniformly drawn from \mathcal{F}_n , then the emulation of $D^f(1^n)$ done in $\tilde{D}^{\phi_1, \phi_2}$ is identical to a random execution of $D^f(1^n)$ with $f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)$. Similarly, in case ϕ_1 and ϕ_2 are uniformly drawn from Π_n , then the emulation is identical to a random execution of $D^f(1^n)$ with $f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_n)$. Thus,

$$\begin{aligned} & \left| \Pr_{(f_1, f_2) \leftarrow \mathcal{F}_n \times \mathcal{F}_n} [\tilde{D}^{f_1, f_2}(1^n) = 1] - \Pr_{(\pi_1, \pi_2) \leftarrow \Pi_n \times \Pi_n} [\tilde{D}^{\pi_1, \pi_2}(1^n) = 1] \right| \\ &= \left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)} [D^f(1^n) = 1] - \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_n)} [D^f(1^n) = 1] \right|. \end{aligned}$$

Hence, Lemma 2.9 yields that there exists a non-adaptive, single-oracle distinguisher \hat{D} that when invoked on input of length n makes $4q(n)$ queries and runs in time $e_q(n) + 8q(n) + t(n) + q(n) \cdot (e_{\mathcal{H}}(n) + e_{\mathcal{G}}(n)) + 2q(n) \cdot e_{\mathcal{F}}(n) \leq p(n) \cdot t(n)$, such that

$$\begin{aligned} & \left| \Pr_{f \leftarrow \mathcal{F}_n} [\hat{D}^f(1^n) = 1] - \Pr_{\pi \leftarrow \Pi_n} [\hat{D}^\pi(1^n) = 1] \right| \\ & \geq \frac{1}{2} \cdot \left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)} [D^f(1^n) = 1] - \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_n)} [D^f(1^n) = 1] \right|, \end{aligned}$$

for every $n \in \mathbb{N}$. □

Since \mathcal{F} is a non-adaptive $(4q, p \cdot t, \varepsilon)$ -PRF (for a large enough polynomial p as defined above), Claim 5.2 implies that for every adaptive t -time q -query oracle-aided algorithm D , it holds that

$$\left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)} [D^f(1^n) = 1] - \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \Pi_n)} [D^f(1^n) = 1] \right| \leq 2\varepsilon(n),$$

for large enough n . Since \mathcal{H} and \mathcal{G} are $k(n)$ -wise independent, we can use Lemma 3.5 (instantiated for our purpose by setting $t = q$, $\mathcal{D} = \{0, 1\}^d$, $\mathcal{S} = \{0, 1\}^n$, $\mathcal{U} = [4q]_{\{0, 1\}^n}$ and $\mathcal{R} = \{0, 1\}^n$; note that $|\mathcal{U}| = 4q = 4t$) and the triangle inequality to get that

$$\left| \Pr_{f \leftarrow \mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{F}_n)} [D^f(1^n) = 1] - \Pr_{\pi \leftarrow \Pi_n} [D^\pi(1^n) = 1] \right| \leq 2\varepsilon(n) + q(n)/2^{\Omega(k(n))}$$

for large enough $n \in \mathbb{N}$. Setting $k(n) = \Theta(\log q(n))$ finishes the proof. □

6 Hardness-Preserving Reductions via Advanced Cuckoo Hashing

In this section we apply the reductions given in Sections 4 and 5 using the function family of Pagh and Pagh [44], with the function family of Aumüller et al. [4], to get a more versatile reduction (for comparison see Section 6.2.1). Roughly speaking, the function family of Aumüller et al. [4] requires less combinatorial work (i.e., smaller independence) than the Pagh and Pagh [44] family. On the other hand, the function family of Aumüller et al. [4] requires more “randomness” (i.e., has a longer description) and is harder to describe. In Section 6.1 we formally define the hash function family of Aumüller et al. [4], state their (non-adaptive) result, and apply Lemma 3.2 to get an adaptive variant of this result. In Section 6.2 we use the function family of Aumüller et al. [4] to obtain a PRF domain extension, and in Section 6.3 we use it to get a non-adaptive to adaptive transformation of PRFs.

6.1 The Aumüller et al. [4] Function Family

The function family of Aumüller et al. [4] (building upon Dietzfelbinger and Woelfel [16]) follows the same basic outline as the Pagh and Pagh [44] function family, but uses more complex hash functions. Recall that the members of the Pagh and Pagh [44] function family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{F})$ are of the form $(f_1 \circ h_1) \oplus (f_2 \circ h_2) \oplus g$, for $f_1, f_2 \in \mathcal{F}$, $h_1, h_2 \in \mathcal{H}$ and $g \in \mathcal{G}$. In the function family $\mathcal{ADW}(\mathcal{H}, \mathcal{L}, \mathcal{G}, \mathcal{F}, \mathcal{M}, \mathcal{Y})$ described below, the role of $h_1, h_2 \in \mathcal{H}$ is taken by some variant of *tabulation hashing* (and not simply from a relatively high k -wise independent family as in [44]). At the heart of these functions lies a function $a_{h, \bar{g}, \bar{m}} : \mathcal{D} \mapsto \mathcal{S}$ of the form:

$$a_{h, \bar{g}, \bar{m}}(x) := h(x) \oplus \bigoplus_{1 \leq i \leq z} m_i(g_i(x)),$$

for $h : \mathcal{D} \mapsto \mathcal{S}$, $\bar{g} = (g_1, \dots, g_z)$ and $\bar{m} = (m_1, \dots, m_z)$, where $g_i : \mathcal{D} \mapsto \mathcal{U}$ and $m_i : \mathcal{U} \mapsto \mathcal{S}$. Jumping ahead, the m_i 's will be chosen to be random functions (or pseudorandom functions) and the g_i 's and h will be chosen from a relatively low independence family. The Aumüller et al. [4] construction uses several functions of the above form that, unlike [44], are chosen in a *correlated* manner (in particular, sharing the *same* function vector \bar{g}). The precise definition is:

Definition 6.1 (The Aumüller et al. [4] function family). *For $z \in \mathbb{N}$ let*

1. \mathcal{D}, \mathcal{U} be sets and \mathcal{S}, \mathcal{R} be commutative groups defined with respect to an operation $\oplus_{\mathcal{S}}$ and $\oplus_{\mathcal{R}}$, respectively (we will omit the subscript when it is clear);
2. function families $\mathcal{H} = \{h : \mathcal{D} \mapsto \mathcal{S}\}$, $\mathcal{L} = \{\ell : \mathcal{D} \mapsto \mathcal{R}\}$, $\mathcal{F} = \{f : \mathcal{S} \mapsto \mathcal{R}\}$, $\mathcal{G} = \{g : \mathcal{D} \mapsto \mathcal{U}\}$, $\mathcal{M} = \{m : \mathcal{U} \mapsto \mathcal{S}\}$ and $\mathcal{Y} = \{y : \mathcal{U} \mapsto \mathcal{R}\}$;
3. functions $h_1, h_2 \in \mathcal{H}$, $\ell \in \mathcal{L}$, $f_1, f_2 \in \mathcal{F}$;
4. function vector $\bar{g} = (g_1, \dots, g_z)$, where $g_i \in \mathcal{G}$ for every $1 \leq i \leq z$;
5. function vectors $\bar{m}^1 = (m_1^1, \dots, m_z^1)$, $\bar{m}^2 = (m_1^2, \dots, m_z^2)$, where $m_i^j \in \mathcal{M}$ for each $j \in \{1, 2\}$ and $1 \leq i \leq z$; and
6. function vector $\bar{y} = (y_1, \dots, y_z)$, where $y_i \in \mathcal{Y}$ for $1 \leq i \leq z$.

Define $\text{adw}_{\bar{m}^1, \bar{m}^2, \bar{y}, h_1, h_2, \ell, \bar{g}, f_1, f_2} : \mathcal{D} \mapsto \mathcal{R}$ by

$$\text{adw}_{\bar{m}^1, \bar{m}^2, \bar{y}, h_1, h_2, \ell, \bar{g}, f_1, f_2} := (f_1 \circ a_{h_1, \bar{g}, \bar{m}^1}) \oplus_{\mathcal{R}} (f_2 \circ a_{h_2, \bar{g}, \bar{m}^2}) \oplus_{\mathcal{R}} a_{\ell, \bar{g}, \bar{y}},$$

where $a_{h, \bar{g}, \bar{m}}(x) := h(x) \oplus \bigoplus_{1 \leq i \leq z} m_i(g_i(x))$.¹⁵

For $\bar{m}^1, \bar{m}^2 \in \mathcal{M}^z$, $h_1, h_2 \in \mathcal{H}$ and $\bar{g} \in \mathcal{G}^z$, function family \mathcal{L}, \mathcal{F} and \mathcal{Y} as above, let

$$\mathcal{ADW}_{z, (\bar{m}^1, \bar{m}^2, h_1, h_2, \bar{g})}(\mathcal{L}, \mathcal{F}, \mathcal{Y}) := \{\text{adw}_{\bar{m}^1, \bar{m}^2, \bar{y}, h_1, h_2, \ell, \bar{g}, f_1, f_2} : \bar{y} \in \mathcal{Y}^z, \ell \in \mathcal{L}, f_1, f_2 \in \mathcal{F}\}.$$

Finally, let

$$\begin{aligned} \mathcal{ADW}_z(\mathcal{H}, \mathcal{L}, \mathcal{G}, \mathcal{F}, \mathcal{M}, \mathcal{Y}) := & \{\text{adw}_{\bar{m}^1, \bar{m}^2, \bar{y}, h_1, h_2, \ell, \bar{g}, f_1, f_2} : \\ & \bar{m}^1, \bar{m}^2 \in \mathcal{M}^z, \bar{y} \in \mathcal{Y}^z, h_1, h_2 \in \mathcal{H}, \\ & \ell \in \mathcal{L}, \bar{g} \in \mathcal{G}^z, f_1, f_2 \in \mathcal{F}\}. \end{aligned}$$

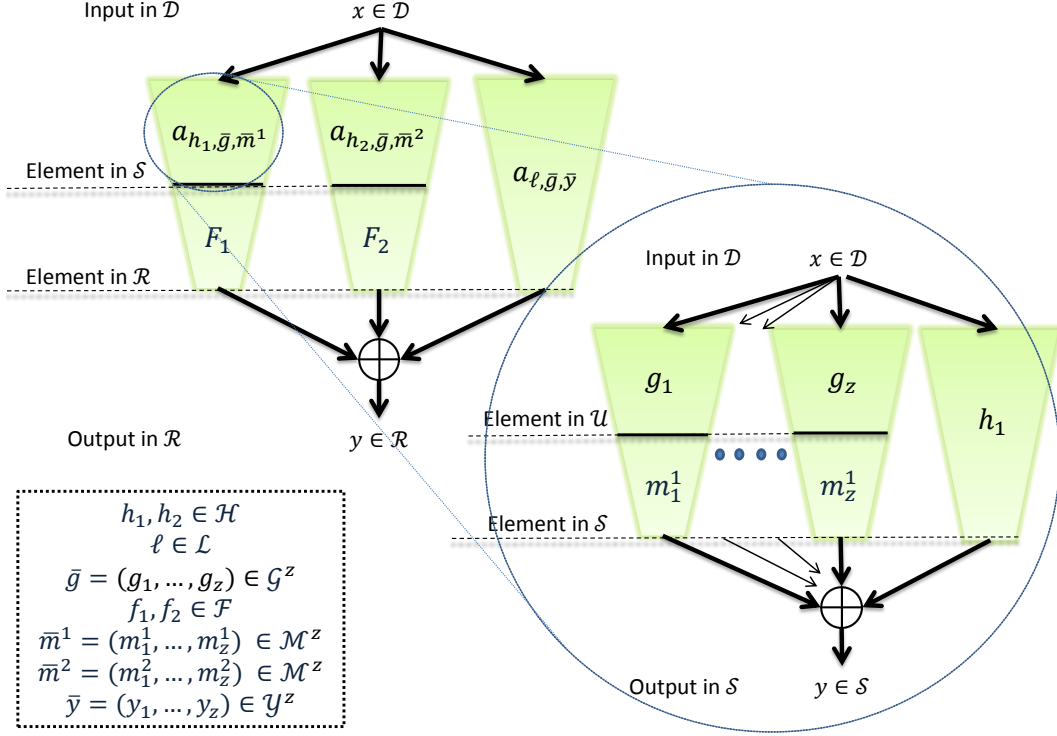


Figure 2: The function family $\mathcal{ADW}_z(\mathcal{H}, \mathcal{L}, \mathcal{G}, \mathcal{F}, \mathcal{M}, \mathcal{Y})$ is in the top left corner. On the bottom right corner the function $a_{h_1, \bar{g}, \bar{m}^1}$ is depicted. The function $a_{h_2, \bar{g}, \bar{m}^2}$ is similar. In $a_{\ell, \bar{g}, \bar{y}}$ the range of the functions ℓ and $y_1 \dots, y_z$ is \mathcal{R} (rather than \mathcal{S}).

Graphically, this function family is described in Figure 2.

Aumüller et al. [4] proved the following result with respect to the above function family.

Theorem 6.2 ([4]). *Let $t, k, z \in \mathbb{N}$ and let $\mathcal{D}, \mathcal{R}, \mathcal{S}, \mathcal{U}$ be commutative groups defined with respect to an operation \oplus such that $|\mathcal{U}| \in [t]$ and $|\mathcal{S}| \geq 4t$.¹⁶ Assume that $z \cdot k \in O(\log t)$. Let $\mathcal{H} = \{h: \mathcal{D} \mapsto \mathcal{S}\}$, $\mathcal{L} = \{\ell: \mathcal{D} \mapsto \mathcal{R}\}$ and $\mathcal{G} = \{g: \mathcal{D} \mapsto \mathcal{U}\}$ be $2k$ -wise independent hash families.*

Then, there exist a universal constant $\text{const} > 0$ and a left-monotone set $\text{BAD} \subseteq \mathcal{D}^{\leq t} \times ((\Pi_{\mathcal{U} \mapsto \mathcal{S}})^z)^2 \times \mathcal{H}^2 \times \mathcal{G}^z$, such that the following holds for every $\bar{q} = (q_1, \dots, q_{|\bar{q}|}) \in \mathcal{D}^{\leq t}$:

1. $(f(q_1), \dots, f(q_{|\bar{q}|}))_{f \leftarrow \mathcal{ADW}_{z,u}(\mathcal{L}, \Pi_{\mathcal{S} \mapsto \mathcal{R}}, \Pi_{\mathcal{U} \mapsto \mathcal{R}})}$ is uniform over $\mathcal{R}^{|\bar{q}|}$ for every $u = (\bar{m}^1, \bar{m}^2, h_1, h_2, \bar{g}) \in ((\Pi_{\mathcal{U} \mapsto \mathcal{S}})^z)^2 \times \mathcal{H}^2 \times \mathcal{G}^z$ such that $(\bar{q}, u) \notin \text{BAD}$, and
2. $\Pr_{u \leftarrow ((\Pi_{\mathcal{U} \mapsto \mathcal{S}})^z)^2 \times \mathcal{H}^2 \times \mathcal{G}^z} [(\bar{q}, u) \in \text{BAD}] = \text{const} \cdot t / |\mathcal{U}|^{z \cdot k / 2}$.

Remark 6.3. *The construction from Definition 6.1 and Theorem 6.2 are taken from [3, Section 6] which is the conference version of [4]. In [3] the authors only considered the case where $z \cdot k$ is*

¹⁵Note that $a_{h_1, \bar{g}, \bar{m}^1}, a_{h_2, \bar{g}, \bar{m}^2}: \mathcal{D} \mapsto \mathcal{S}$ uses $\oplus_{\mathcal{S}}$ and $a_{\ell, \bar{g}, \bar{y}}: \mathcal{D} \mapsto \mathcal{R}$ uses $\oplus_{\mathcal{R}}$.

¹⁶The actual setting in [4] is more general. Specifically, an additional parameter $\varepsilon > 0$ is used to set the size of \mathcal{S} as $(1 + \varepsilon)t$. For simplicity of presentation, comparison with the statement of Theorem 3.4, and since we use this theorem only when m is an integer, we set $\varepsilon = 3$.

constant, independent of t . The proof for the case where $z \cdot k \in O(\log t)$ follows from the proof of [4, Theorem 2].

Applying Lemma 3.2, the above yields that, for the right choice of parameters, the function family $\mathcal{ADW}_z(\mathcal{H}, \mathcal{L}, \mathcal{F}, \mathcal{G}, \mathcal{M}, \mathcal{Y})$ is not only close to being uniform in the eyes of a *non-adaptive* distinguisher, but also in the eyes of an *adaptive* one. Specifically, combining Lemma 3.2 and theorem 6.2 yields the following result.

Lemma 6.4. *Let $t, k, z, \text{const}, \mathcal{D}, \mathcal{R}, \mathcal{S}, \mathcal{U}, \mathcal{H}, \mathcal{L}, \mathcal{G}$ be as in Theorem 6.2 and let D be an adaptive, t -query oracle-aided algorithm. Then,*

$$\left| \Pr_{f \leftarrow \mathcal{ADW}_z(\mathcal{H}, \mathcal{L}, \mathcal{G}, \Pi_{\mathcal{S} \rightarrow \mathcal{R}}, \Pi_{\mathcal{U} \rightarrow \mathcal{S}}, \Pi_{\mathcal{U} \rightarrow \mathcal{R}})}[D^f = 1] - \Pr_{\pi \leftarrow \Pi}[D^\pi = 1] \right| \leq \text{const} \cdot t / |\mathcal{U}|^{z \cdot k / 2}.$$

Proof. Let $\mathcal{U}' = ((\Pi_{\mathcal{U} \rightarrow \mathcal{S}})^z)^2 \times \mathcal{H}^2 \times \mathcal{G}^z$, $\mathcal{V} = (\Pi_{\mathcal{U} \rightarrow \mathcal{R}})^z \times (\Pi_{\mathcal{S} \rightarrow \mathcal{R}})^2 \times \mathcal{L}$. For $(\bar{m}^1, \bar{m}^2, h_1, h_2, \bar{g}) \in \mathcal{U}'$ and $(\bar{y}, \pi_1, \pi_2, \ell) \in \mathcal{V}$, let

$$F_{(\bar{m}^1, \bar{m}^2, h_1, h_2, \bar{g}), (\bar{y}, \pi_1, \pi_2, \ell)} = (\pi_1 \circ a_{h_1, \bar{g}, \bar{m}^1}) \oplus (\pi_2 \circ a_{h_2, \bar{g}, \bar{m}^2}) \oplus a_{\ell, \bar{g}, \bar{y}},$$

and let

$$\mathcal{F} = \{F_{u', v} : \mathcal{D} \mapsto \mathcal{R}\}_{(u', v) \in \mathcal{U}' \times \mathcal{V}}.$$

Finally, let BAD be the set BAD of Theorem 6.2.

The above sets meet the requirements stated in Lemma 3.2: Item 1 of Theorem 6.2 assures that the first property of Lemma 3.2 is satisfied, and according to Item 2 of Theorem 6.2 we set ε of Lemma 3.2 to be $\text{const} \cdot t / |\mathcal{U}|^{k \cdot z}$, and thus the second property is also satisfied. Hence, applying Lemma 3.2 concludes the proof of the lemma. \square

We note that for large enough z , and in contrast to the Pagh and Pagh [44] family, using \mathcal{ADW} we get meaningful results even when using an underlying $k = o(\log t)$ -wise independent family. (For a thorough comparison between \mathcal{PP} and \mathcal{ADW} see Section 6.2.1.) In particular, for specific settings of parameters we get the following corollary.

Corollary 6.5. *Let $t, \text{const}, \mathcal{D}, \mathcal{R}, \mathcal{S}, \mathcal{U}, \mathcal{H}, \mathcal{L}, \mathcal{G}$ be as in Theorem 6.2. Let D be an adaptive, t -query oracle-aided algorithm and $c \in \mathbb{N}$. If either*

1. $z = 2(c + 2)$, $|\mathcal{U}| = t$ and $k = 1$, or
2. $z = 2(c + 2) \cdot \log t$, $|\mathcal{U}| = 2$ and $k = 1$,

then

$$\left| \Pr_{f \leftarrow \mathcal{ADW}_z(\mathcal{H}, \mathcal{L}, \mathcal{G}, \Pi_{\mathcal{S} \rightarrow \mathcal{R}}, \Pi_{\mathcal{U} \rightarrow \mathcal{S}}, \Pi_{\mathcal{U} \rightarrow \mathcal{R}})}[D^f = 1] - \Pr_{\pi \leftarrow \Pi}[D^\pi = 1] \right| \leq \text{const} / t^{c+1}.$$

Proof. Since \mathcal{H} and \mathcal{L} are pairwise independent, we can apply Lemma 6.4. Using either setting of parameters the advantage of D is $\text{const} \cdot t / |\mathcal{U}|^{z \cdot k / 2} \leq \text{const} / t^{c+1}$. \square

6.2 PRF Domain Extension via the ADW Family

In this subsection we present a PRF a domain extension using the Aumüller et al. [4] family, ADW . This allows us to avoid the large independence required for using the Pagh and Pagh [44] family \mathcal{PP} .

In what follows we state two domain extension results (whose proofs are similar to that of Theorem 4.1). In both of the results we show how to extend the domain of a (q, t, ε) -PRF $\mathcal{F} = \{\mathcal{F}_n: \{0, 1\}^{s(n)} \mapsto \{0, 1\}^{r(n)}\}_{n \in \mathbb{N}}$. We let $s, r, q, t, \varepsilon, d, u$ be integer functions such that $q(n) \leq 2^{n-2}$, and let $\mathcal{H} = \{\mathcal{H}_n: \{0, 1\}^{d(n)} \mapsto \{0, 1\}^{s(n)}\}_{n \in \mathbb{N}}$, $\mathcal{L} = \{\mathcal{L}_n: \{0, 1\}^{d(n)} \mapsto \{0, 1\}^{r(n)}\}$ and $\mathcal{G} = \{\mathcal{G}_n: \{0, 1\}^{d(n)} \mapsto \{0, 1\}^{u(n)}\}$ be efficient pairwise independent function family ensembles. Let $p \in \text{poly}$ be a large enough polynomial determined by the evaluation and sampling time of \mathcal{H} , \mathcal{L} , \mathcal{G} , and \mathcal{F} . Furthermore, both of the results will rely on Corollary 6.5 to argue that the function family $ADW_z(\mathcal{H}, \mathcal{L}, \mathcal{G}, \Pi_{s,r}, \Pi_{u,s}, \Pi_{u,r})$ is indistinguishable from random (i.e., instantiate Corollary 6.5 with $\mathcal{S} = \{0, 1\}^s$, $\mathcal{R} = \{0, 1\}^r$ and $\mathcal{U} = \{0, 1\}^u$).

In the first result we rely on the first setting of parameters in Corollary 6.5 and implement the function families $\Pi_{s,r}, \Pi_{u,s}$ and $\Pi_{u,r}$ using a single pseudorandom function family $\mathcal{F} = \{\mathcal{F}_n: \{0, 1\}^{s(n)} \mapsto \{0, 1\}^{r(n)}\}_{n \in \mathbb{N}}$. Assuming $u(n) \leq s(n)$ and $s(n) \leq r(n)$, the implementation is done in the natural way by padding with leading zeroes. Since $ADW_z(\mathcal{H}, \mathcal{L}, \mathcal{F}, \mathcal{G}, \mathcal{M}, \mathcal{Y})$ makes two calls to \mathcal{F} , $2z$ calls to \mathcal{M} , and z calls to \mathcal{Y} (in total $3z + 2$ calls), this results with a PRF that makes $3z + 2$ calls to the underlying PRF.

Theorem 6.6. *Let \mathcal{H}, \mathcal{L} and \mathcal{G} be defined as above such that $u(n) \leq s(n) \leq r(n)$. Let $c > 1$ be a constant and let $z = 2(c + 2)$.*

Then, $ADW_z(\mathcal{H}, \mathcal{L}, \mathcal{G}, \mathcal{F}, \mathcal{F}, \mathcal{F}) = \{ADW_z(\mathcal{H}_n, \mathcal{L}_n, \mathcal{G}_n, \mathcal{F}_n, \mathcal{F}_n, \mathcal{F}_n): \{0, 1\}^{d(n)} \mapsto \{0, 1\}^{r(n)}\}_{n \in \mathbb{N}}$ is a $(q, t - p \cdot q, (3z + 2) \cdot \varepsilon + 1/q^{c+1})$ -PRF

In the second result we rely on the first setting of parameters in Corollary 6.5 and take advantage of the fact that $u = 2$. This enables us to implement the function families \mathcal{M} and \mathcal{Y} as small tables of random values which are embedded into the PRF key. This results with a PRF that makes just two calls to the underlying PRF (but has a longer key).

Theorem 6.7. *Let \mathcal{H}, \mathcal{L} and \mathcal{G} be defined as above. Let $c > 1$ be a constant and let $z = z(n) = 2(c + 2) \cdot \log q(n)$. Let $\mathcal{M} = \{\mathcal{M}_n\}_{n \in \mathbb{N}}$ (resp., $\mathcal{Y} = \{\mathcal{Y}_n\}_{n \in \mathbb{N}}$) be family of tables, such that \mathcal{M}_n (resp., \mathcal{Y}_n) is a table of two (resp., $z(n)$) random elements from $\{0, 1\}^{s(n)}$ (resp., $\{0, 1\}^{r(n)}$).*

Then, $ADW_z(\mathcal{H}, \mathcal{L}, \mathcal{G}, \mathcal{F}, \mathcal{M}, \mathcal{Y}) = \{ADW_{z(n)}(\mathcal{H}_n, \mathcal{L}_n, \mathcal{G}_n, \mathcal{F}_n, \mathcal{M}_n, \mathcal{Y}_n): \{0, 1\}^{d(n)} \mapsto \{0, 1\}^{r(n)}\}_{n \in \mathbb{N}}$ is a $(q, t - p \cdot q, 2\varepsilon + 1/q^{c+1})$ -PRF.

A notable difference between Theorems 6.6 and 6.7 is that in Theorem 6.6 the resulting PRF makes a large constant (i.e., $6c + 12$) number of queries to the underlying PRF, whereas in Theorem 6.7 it only makes two calls to the underlying PRF but the PRF key is longer (i.e., it has $(6c + 12) \cdot \log q$ random values embedded into it).

6.2.1 Comparing the \mathcal{PP} and ADW Based Constructions

Theorems 4.1, 6.6 and 6.7 present different tradeoffs between two types of resources: “cryptographic work” — the total evaluation time of the calls to the underlying short-domain PRF, and “combinatorial work” — the independence needed from the function families used.¹⁷ In the \mathcal{PP} -based

¹⁷The independence affects the amount of random bits and evaluation time needed for these families.

construction (Theorem 4.1), we minimize the number of calls to the PRF, thus keeping the cryptographic work small. But on the other hand, we require relatively high independence, which results in much combinatorial work. In the first \mathcal{ADW} -based construction (Theorem 6.6), the situation is somewhat reversed: we minimize the independence needed, but make more calls to the PRF. In the second \mathcal{ADW} -based construction (Theorem 6.7), we minimize both the number of calls to the PRF and the independence needed, but we require much more hash functions, thus increasing again the combinatorial work.

In the following we instantiate Theorems 4.1, 6.6 and 6.7 in a specific setting, and compare their performance in terms of required randomness complexity, and evaluation time.¹⁸ Our starting point is a length-preserving (q, t, ε) -PRF $\mathcal{F} = \{\mathcal{F}_n: \{0, 1\}^n \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$. Our goal is to construct a function family with a larger domain, $\{0, 1\}^d$ for $d = d(n)$, whose security only deteriorates, comparing to that of \mathcal{F} , by an additive factor of $1/q^c$, for $c := c(n) > 1$. We consider the following function families:

1. **Family 1:** $\mathcal{PP}(\mathcal{H}, \mathcal{F}) = \{\mathcal{PP}(\mathcal{H}_n, \mathcal{F}_n): \{0, 1\}^d \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$, where $\mathcal{H} = \{\mathcal{H}_n: \{0, 1\}^d \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$ is $\Omega(c \cdot \log q)$ -wise independent function family ensemble.

This instantiation, based on Theorem 4.1, makes only two calls to \mathcal{F} , but requires hash functions of high independence.

2. **Family 2:** $\mathcal{ADW}_z(\mathcal{H}, \mathcal{H}, \mathcal{H}, \mathcal{F}, \mathcal{F}, \mathcal{F}) = \{\mathcal{ADW}_z(\mathcal{H}_n, \mathcal{H}_n, \mathcal{H}_n, \mathcal{F}_n, \mathcal{F}_n, \mathcal{F}_n): \{0, 1\}^d \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$, where $z = 2(c + 2)$ and $\mathcal{H} = \{\mathcal{H}_n: \{0, 1\}^d \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$ is a pairwise independent function family ensemble.

This instantiation, based on Theorem 6.6, reduce the needed independence, in the price of increasing the number of calls to \mathcal{F} .

3. **Family 3:** $\mathcal{ADW}_z(\mathcal{H}, \mathcal{H}, \mathcal{H}, \mathcal{F}, \mathcal{M}, \mathcal{Y}) = \{\mathcal{ADW}_{z(n)}(\mathcal{H}_n, \mathcal{H}_n, \mathcal{H}_n, \mathcal{F}_n, \mathcal{M}_n, \mathcal{Y}_n): \{0, 1\}^d \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$, where $z(n) = 2(c + 2) \cdot \log q$, $\mathcal{H} = \{\mathcal{H}_n: \{0, 1\}^d \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$ is a pairwise independent function family ensemble, $\mathcal{M} = \{\mathcal{M}_n\}_{n \in \mathbb{N}}$ (resp., $\mathcal{Y} = \{\mathcal{Y}_n\}_{n \in \mathbb{N}}$) is a family of tables, and \mathcal{M}_n (resp., \mathcal{Y}_n) is a random table of two (resp., z) elements from $\{0, 1\}^n$.

This instantiation, based on Theorem 6.7, makes only two calls to \mathcal{F} and keeps the independence low, in the price of needing many more hash functions.

To ease the comparison between the above families, we introduce relevant notation:

1. $r_{\mathcal{F}}$ – the amount of random bits required to sample a random element in \mathcal{F}_n
2. $e_{\mathcal{F}}$ – the evaluation time of a single call to an element in \mathcal{F}_n
3. $e^{(k)}$ – the evaluation time of a k -wise independent function family from $\{0, 1\}^d$ to $\{0, 1\}^n$. Recall that (Fact 2.6), sampling a random element in the latter function family requires $k \cdot \max\{d, n\}$ random bits.

The randomness complexity and evaluation time of each instantiation are summarized in Table 1.

¹⁸Another criterion of comparison is the sampling time. In our settings it is analogous to the evaluation time, so we omit it.

Instantiation	Randomness Complexity (Key Size)	Evaluation Time
\mathcal{PP} (Family 1)	$O(c \cdot d \cdot \log q) + 2 \cdot r_{\mathcal{F}}$	$O(e^{(c \cdot \log q)}) + 2 \cdot e_{\mathcal{F}}$
\mathcal{ADW} (Family 2)	$O(c \cdot d) + (6c + 14) \cdot r_{\mathcal{F}}$	$O(c \cdot e^{(2)}) + (6c + 14) \cdot e_{\mathcal{F}}$
\mathcal{ADW} (Family 3)	$O(c \cdot d \cdot \log q) + 2 \cdot r_{\mathcal{F}}$	$O(c \cdot \log q \cdot e^{(2)}) + 2 \cdot e_{\mathcal{F}}$

Table 1: Comparison between the domain extension results based on the function family \mathcal{PP} from Theorem 4.1 and \mathcal{ADW} from Theorems 6.6 and 6.7 instantiated in Items 1 to 3 above to achieve security deterioration $1/q^c$ for $c \in \mathbb{N}$.

6.3 From Non-Adaptive to Adaptive PRF Via the \mathcal{ADW} Family

In this subsection we state the result for the non-adaptive to adaptive transformation using the Aumüller et al. [4] family, \mathcal{ADW} . As in the previous section, this allows us to avoid the large independence required for using the function family \mathcal{PP} of Pagh and Pagh [44]. For simplicity we state only the reduction that follows from the second settings of parameters of Corollary 6.5 (analogous to Theorem 5.1). The first settings of parameters of Corollary 6.5 also yields non-adaptive to adaptive reduction, with different tradeoff between the randomness complexity (key size) of the PRF to its evaluation time (see Section 6.2.1).

Recall that our reduction from Section 5 requires that all calls to the PRF \mathcal{F} are within the first $4q(n)$ first elements of $\{0,1\}^n$. We make sure the above holds by the setting this to be the range of the hash function we use.

Theorem 6.8. *Let q be integer functions, let $c > 1$ be a constant and let $z = z(n) = 2(c + 2) \cdot \log q(n)$. Fix three efficient pairwise independent function family ensembles: $\mathcal{H} = \{\mathcal{H}_n: \{0,1\}^n \mapsto [4q(n)]_{\{0,1\}^n}\}_{n \in \mathbb{N}}$, $\mathcal{L} = \{\mathcal{L}_n: \{0,1\}^n \mapsto \{0,1\}^n\}_{n \in \mathbb{N}}$ and $\mathcal{G} = \{\mathcal{G}_n: \{0,1\}^n \mapsto \{0,1\}^n\}_{n \in \mathbb{N}}$. Let $\mathcal{M} = \{\mathcal{M}_n\}_{n \in \mathbb{N}}$ (resp., $\mathcal{Y} = \{\mathcal{Y}_n\}_{n \in \mathbb{N}}$) be family of tables, such that \mathcal{M}_n (resp., \mathcal{Y}_n) is a random table of two (resp., $z(n)$) elements from $[4q(n)]_{\{0,1\}^n}$ (resp., $\{0,1\}^n$).*

If $\mathcal{F} = \{\mathcal{F}_n: \{0,1\}^n \mapsto \{0,1\}^n\}_{n \in \mathbb{N}}$ is a non-adaptive $(4q, p \cdot t, \varepsilon)$ -PRF for some $p \in \text{poly}$ determined by the evaluation time of $q, \mathcal{H}, \mathcal{G}, \mathcal{L}$ and \mathcal{F} , then $\mathcal{ADW}_z(\mathcal{H}, \mathcal{L}, \mathcal{G}, \mathcal{F}, \mathcal{M}, \mathcal{Y}) = \{\mathcal{ADW}_{z(n)}(\mathcal{H}_n, \mathcal{L}_n, \mathcal{G}_n, \mathcal{F}_n, \mathcal{M}_n, \mathcal{Y}_n): \{0,1\}^n \mapsto \{0,1\}^n\}_{n \in \mathbb{N}}$ is a $(q, t, 2\varepsilon + 1/q^{c+1})$ -PRF.

7 Further Research

The focus of this paper is on PRFs. Specifically, in Sections 4 to 6 we have shown domain extension techniques and non-adaptive to adaptive transformations for PRFs that provide a nice tradeoff between combinatorial work, cryptographic work and error. In general, hardness-preserving reductions between pseudorandom objects have led to fruitful research with many results (some of which we review next). It is an interesting question whether our technique has any bearing on other models.

Perhaps the most interesting model for this kind of reductions is pseudorandom permutations (PRPs) (without going through a PRP-to-PRF reduction). Given a family of (q, t, ε) -PRPs from n -bits to n bits, how can we construct a family of PRPs with *larger* domain while preserving its security? How about constructing a family of PRPs with *smaller* domain while preserving its security? Finally, it is also interesting how to transform a family of (q, t, ε) -PRPs that is secure against non-adaptive adversaries to a family of PRPs that are also secure against adaptive adversaries (see discussion in Section 1.4). One related paper is that of Hoang et al. [25], that gives

a method to convert a PRF into a PRP with beyond-birthday security (see also [52, 36]). Other related works are of Håstad [23] and Morris et al. [37] showed how to extend the domain of a PRP.

A different model of interest is message authentication codes (MACs). In this model, we are interested in designing domain extension techniques that given an n -bit to n -bit MAC with MAC security ε against q queries provide variable-length MAC with some (good enough) promise on the MAC security in terms of q and ε . The best answer to-date for this question was given by Dodis and Steinberger [17] that showed that given an n -bit to n -bit MAC with MAC security ε against q queries, it is possible to get a variable-length MAC achieving MAC security $O(\varepsilon \cdot q \cdot \text{poly}(n))$ against queries of total length qn .

Another interesting model is public random functions. A public random function $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is a system with a public and private interface which behaves as the same random function at both interfaces. In other words, a public random function can be interpreted as a random oracle. In this model, again, the domain extension problem is very interesting. To date, the best construction is of Maurer and Tessaro [35] that presented a construction $\mathbf{C}_{\varepsilon, m, \ell}$ that extends public random functions $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to a function $\mathbf{C}_{\varepsilon, m, \ell}(f) : \{0, 1\}^{m(n)} \rightarrow \{0, 1\}^{\ell(n)}$ with time complexity $\text{poly}(n, 1/\varepsilon)$ and which is secure against adversaries which make up to $\Theta(2^{(1-\varepsilon)n})$ queries.

On a different note, as we mentioned in the text following Theorem 1.2, our non-adaptive to adaptive transformation depends on the number of queries made by the distinguisher which does not give a single transformation for all poly-time adversaries. Whether a single reduction with a constant number of calls to the non-adaptive PRF that works for all poly-time adversaries exists is left as an open problem.

Pătraşcu and Thorup [51] have shown that for many data structure problems it is possible to use tabulation hashing even though it is ‘merely’ 3-wise independent. The question is whether this has any bearing on cryptographic constructions.

8 Acknowledgments

We thank Eylon Yogev and the anonymous referees for their helpful comments. The third author would like to thank his M.Sc advisor Ran Raz for his support.

References

- [1] W. Aiello and R. Venkatesan. Foiling birthday attacks in length-doubling transformations - benes: A non-reversible alternative to feistel. In *Advances in Cryptology – EUROCRYPT ’96*, pages 307–320, 1996.
- [2] Y. Arbitman, M. Naor and G. Segev. Backyard cuckoo hashing: Constant worst-case operations with a succinct representation. In *Proceedings of the 51th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 787–796, 2010.
- [3] M. Aumüller, M. Dietzfelbinger and P. Woelfel. Explicit and efficient hash families suffice for cuckoo hashing with a stash. In L. Epstein and P. Ferragina, editors, *ESA*, volume 7501 of *Lecture Notes in Computer Science*, pages 108–120. Springer, 2012.

- [4] M. Aumüller, M. Dietzfelbinger and P. Woelfel. Explicit and efficient hash families suffice for cuckoo hashing with a stash. *Algorithmica*, 70(3):428–456, 2014.
- [5] M. Bellare and S. Goldwasser. New paradigms for digital signatures and message authentication based on non-interactive zero knowledge proofs. In *Advances in Cryptology – CRYPTO ’89*, pages 194–211, 1989.
- [6] M. Bellare, R. Canetti and H. Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *FOCS*, pages 514–523, 1996.
- [7] M. Bellare, O. Goldreich and H. Krawczyk. Stateless evaluation of pseudorandom functions: Security beyond the birthday barrier. In *Advances in Cryptology – CRYPTO ’99*, pages 270–287, 1999.
- [8] I. Berman and I. Haitner. From non-adaptive to adaptive pseudorandom functions. *J. Cryptology*, 28(2):297–311, 2015.
- [9] I. Berman, I. Haitner, I. Komargodski and M. Naor. Hardness preserving reductions via cuckoo hashing. In *Theory of Cryptography - 10th Theory of Cryptography Conference, TCC 2013*, pages 40–59, 2013.
- [10] O. Billet, J. Etrog and H. Gilbert. Lightweight privacy preserving authentication for rfid using a stream cipher. In *18th International Symposium on the Foundations of Software Engineering (FSE)*, pages 55–74, 2010.
- [11] M. Blum, W. S. Evans, P. Gemmell, S. Kannan and M. Naor. Checking the correctness of memories. *Algorithmica*, 12(2/3):225–244, 1994.
- [12] L. J. Carter and M. N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, pages 143–154, 1979.
- [13] N. Chandran and S. Garg. Balancing output length and query bound in hardness preserving constructions of pseudorandom functions. In *Progress in Cryptology - INDOCRYPT 2014*, pages 89–103, 2014.
- [14] B. Chor, A. Fiat, M. Naor and B. Pinkas. Tracing traitors. *IEEE Transactions on Information Theory*, 46(3):893–910, 2000.
- [15] M. Dietzfelbinger and C. Weidling. Balanced allocation and dictionaries with tightly packed constant size bins. *Theor. Comput. Sci.*, 380(1-2):47–68, 2007.
- [16] M. Dietzfelbinger and P. Woelfel. Almost random graphs with simple hash functions. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 629–638, 2003.
- [17] Y. Dodis and J. P. Steinberger. Domain extension for macs beyond the birthday barrier. In *Advances in Cryptology – EUROCRYPT 2011*, pages 323–342, 2011.
- [18] D. Fotakis, R. Pagh, P. Sanders and P. G. Spirakis. Space efficient hash tables with worst case constant access time. *Theory Comput. Syst.*, 38(2):229–248, 2005.

- [19] A. M. Frieze, P. Melsted and M. Mitzenmacher. An analysis of random-walk cuckoo hashing. *SIAM J. Comput.*, 40(2):291–308, 2011.
- [20] O. Goldreich. Towards a theory of software protection. In *Advances in Cryptology – CRYPTO ’86*, pages 426–439, 1986.
- [21] O. Goldreich, S. Goldwasser and S. Micali. On the cryptographic applications of random functions. In *Advances in Cryptology – CRYPTO ’84*, pages 276–288, 1984.
- [22] O. Goldreich, S. Goldwasser and S. Micali. How to construct random functions. *Journal of the ACM*, pages 792–807, 1986.
- [23] J. Håstad. The square lattice shuffle. *Random Structures & Algorithms*, 29(4):466–474, 2006.
- [24] J. Håstad, R. Impagliazzo, L. A. Levin and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, pages 1364–1396, 1999.
- [25] V. T. Hoang, B. Morris and P. Rogaway. An enciphering scheme based on a card shuffle. In *Advances in Cryptology – CRYPTO 2012*, pages 1–13, 2012.
- [26] A. Jain, K. Pietrzak and A. Tentes. Hardness preserving constructions of pseudorandom functions. In *Theory of Cryptography, 9th Theory of Cryptography Conference, TCC 2012*, pages 369–382, 2012.
- [27] D. Jetchev, O. Özen and M. Stam. Understanding adaptivity: Random systems revisited. In *Advances in Cryptology – ASIACRYPT 2012*, pages 313–330, 2012.
- [28] E. Kaplan, M. Naor and O. Reingold. Derandomized constructions of k -wise (almost) independent permutations. *Algorithmica*, 55(1):113–133, 2009.
- [29] A. Kirsch, M. Mitzenmacher and U. Wieder. More robust hashing: Cuckoo hashing with a stash. *SIAM J. Comput.*, 39(4):1543–1561, 2009.
- [30] L. A. Levin. One-way functions and pseudorandom generators. *Combinatorica*, 7(4):357–363, 1987.
- [31] M. Luby. *Pseudorandomness and cryptographic applications*. Princeton computer science notes. Princeton University Press, 1996.
- [32] M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal on Computing*, 17(2):373–386, 1988.
- [33] U. M. Maurer. Indistinguishability of random systems. In *Advances in Cryptology – EURO-CRYPT 2002*, pages 110–132, 2002.
- [34] U. M. Maurer and K. Pietrzak. Composition of random systems: When two weak make one strong. In *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004*, pages 410–427, 2004.
- [35] U. M. Maurer and S. Tessaro. Domain extension of public random functions: Beyond the birthday barrier. In *Advances in Cryptology – CRYPTO 2007*, pages 187–204, 2007.

- [36] B. Morris and P. Rogaway. Sometimes-recurse shuffle - almost-random permutations in logarithmic expected time. In *Advances in Cryptology – EUROCRYPT 2014*, pages 311–326, 2014.
- [37] B. Morris, P. Rogaway and T. Stegers. How to encipher messages on a small domain. In *Advances in Cryptology – CRYPTO 2009*, pages 286–302, 2009.
- [38] S. Myers. Black-box composition does not imply adaptive security. In *Advances in Cryptology – EUROCRYPT 2004*, pages 189–206, 2004.
- [39] M. Nandi. A unified method for improving prf bounds for a class of blockcipher based macs. In *Fast Software Encryption, 17th International Workshop, FSE 2010, Seoul, Korea*, pages 212–229, 2010.
- [40] M. Naor and O. Reingold. On the construction of pseudorandom permutations: Luby-rackoff revisited. *Journal of Cryptology*, 12(1):29–66, 1999.
- [41] M. Naor and O. Reingold. Synthesizers and their application to the parallel construction of pseudo-random functions. *Journal of Computer and System Sciences*, 58(2):336–375, 1999.
- [42] M. Naor and O. Reingold. Constructing pseudo-random permutations with a prescribed structure. *Journal of Cryptology*, 15(2):97–102, 2002.
- [43] R. Ostrovsky. An efficient software protection scheme. In *Advances in Cryptology – CRYPTO ’89*, 1989.
- [44] A. Pagh and R. Pagh. Uniform hashing in constant time and optimal space. *SIAM Journal on Computing*, 38(1):85–96, 2008.
- [45] R. Pagh and F. F. Rodler. Cuckoo hashing. *J. Algorithms*, 51(2):122–144, 2004.
- [46] J. Patarin. Security of random feistel schemes with 5 or more rounds. In *Advances in Cryptology – CRYPTO 2004*, pages 106–122, 2004.
- [47] J. Patarin. A proof of security in $o(2^n)$ for the benes scheme. In *Progress in Cryptology - AFRICACRYPT 2008*, pages 209–220, 2008.
- [48] J. Patarin. Security of balanced and unbalanced feistel schemes with linear non equalities. *IACR Cryptology ePrint Archive*, 2010:293, 2010.
- [49] K. Pietrzak. Composition does not imply adaptive security. In *Advances in Cryptology – CRYPTO 2005*, pages 55–65, 2005.
- [50] K. Pietrzak. Composition implies adaptive security in minicrypt. In *Advances in Cryptology – EUROCRYPT 2006*, pages 328–338, 2006.
- [51] M. Pătraşcu and M. Thorup. The power of simple tabulation hashing. *J. ACM*, 59(3):14, 2012.
- [52] T. Ristenpart and S. Yilek. The mix-and-cut shuffle: Small-domain encryption secure against N queries. In *Advances in Cryptology – CRYPTO 2013*, pages 392–409, 2013.

- [53] A. Siegel. On universal classes of extremely random constant-time hash functions. *SIAM Journal on Computing*, 33(3):505–543, 2004.
- [54] M. N. Wegman and L. Carter. New hash functions and their use in authentication and set equality. *J. Comput. Syst. Sci.*, 22(3):265–279, 1981.

A Proof of Lemma 3.2

In this section we prove Lemma 3.2 from Section 3. We mentioned again that Lemma 3.2 can be derived as a special case of a result given in [27, Theorem 12] (closing a gap in the proof appearing in [33]). Yet, for the sake of completeness, we include an independent proof of this lemma below.

Definition A.1 (Restating Definition 3.1). *Let \mathcal{S} and \mathcal{T} be sets. A set $\mathcal{M} \subseteq \mathcal{S}^* \times \mathcal{T}$ is left-monotone, if for every $(\overline{s}_1, t) \in \mathcal{M}$ and every $\overline{s}_2 \in \mathcal{S}^*$ that has \overline{s}_1 as a prefix, it holds that $(\overline{s}_2, t) \in \mathcal{M}$.*

Lemma A.2 (Restating Lemma 3.2). *Let \mathcal{U} and \mathcal{V} be non-empty sets, let $\mathcal{F} = \mathcal{F}(\mathcal{U}, \mathcal{V}) = \{f_{u,v} : \mathcal{D} \mapsto \mathcal{R}\}_{(u,v) \in \mathcal{U} \times \mathcal{V}}$ be a function family and let $\text{BAD} \subseteq \mathcal{D}^* \times \mathcal{U}$ be left-monotone. Let $t \in \mathbb{N}$, and assume that for every $\overline{q} = (q_1, \dots, q_{|\overline{q}|}) \in \mathcal{D}^{\leq t}$ it holds that¹⁹*

1. $(f(q_1), \dots, f(q_{|\overline{q}|}))_{f \leftarrow \{f_{u,v} : v \in \mathcal{V}\}}$ is uniform over $\mathcal{R}^{|\overline{q}|}$, for every $u \in \mathcal{U}$ with $(\overline{q}, u) \notin \text{BAD}$, and
2. $\Pr_{u \leftarrow \mathcal{U}}[(\overline{q}, u) \in \text{BAD}] \leq \varepsilon$.

Then,

$$\left| \Pr_{\substack{u \leftarrow \mathcal{U} \\ v \leftarrow \mathcal{V}}}[\mathsf{D}^{f_{u,v}} = 1] - \Pr_{\pi \leftarrow \Pi}[\mathsf{D}^\pi = 1] \right| \leq \varepsilon$$

for every t -query oracle-aided adaptive algorithm D , letting Π be the set of all functions from \mathcal{D} to \mathcal{R} .

Proof. Let D be a t -query distinguisher. We assume for simplicity that D is deterministic (the reduction to the randomized case is standard) and makes exactly t valid (i.e., inside \mathcal{D}) distinct queries. To prove the lemma we consider a process (Algorithm A.3) that runs D twice: one giving it completely random answers and the second time, choosing $u \leftarrow \mathcal{U}$ and continuing answering with the same answers as the first round until we hit a BAD event according to the queries and the chosen u . We then choose a random v that is consistent with the answers given so far and continue answering with it.

Intuitively, the answers provided to D in the first round are distributed like the answers D expects to get from a *truly random function*, while the answers provided to D in the second round are distributed like the answers D expects to get from a random function in \mathcal{F} . But, since these answers are the same until a BAD event occurs, the distinguishing ability of D is bounded by the probability of such an event to occur. Since, u is chosen *after* D has already “committed” to the queries it is going to make, this probability is bounded by the non-adaptive property of \mathcal{F} .

For a vector $\overline{v} = (v_1, \dots, v_t)$, let $\overline{v}_{1,\dots,i}$ be the first i element in \overline{v} (i.e., $\overline{v}_{1,\dots,i} = (v_1, \dots, v_i)$) and let $\overline{v}_{1,\dots,0} = \lambda$, where λ is the empty vector. Consider the following random process:

¹⁹Recall that for a set \mathcal{S} and an integer t , $\mathcal{S}^{\leq t}$ denotes the set $\{\overline{s} \in \mathcal{S}^* : |\overline{s}| \leq t \wedge \overline{s}[i] \neq \overline{s}[j] \ \forall i \neq j \in [|\overline{s}|]\}$.

Algorithm A.3.

1. Emulate D , while answering the i^{th} query q_i with $a_i \leftarrow \mathcal{R}$.
Set $\bar{q} = (q_1, \dots, q_t)$ and $\bar{a} = (a_1, \dots, a_t)$.
2. Choose $u \leftarrow \mathcal{U}$ and set $v = \perp$.
3. If $(\lambda, u) \in \text{BAD}$, set $v \leftarrow \mathcal{V}$.
4. Emulate D again, while answering the i^{th} query q'_i according to the following procedure:
 - (a) If $(\bar{q}'_{1,\dots,i} = (q'_1, \dots, q'_i), u) \notin \text{BAD}$, answer with $a'_i = a_i$ (the same a_i from Step 1).
 - (b) Otherwise $((\bar{q}'_{1,\dots,i}, u) \in \text{BAD})$:
 - i. If $v = \perp$, set $v \leftarrow \{v' \in \mathcal{V} : \forall j \in [i-1]: f_{u,v'}(q'_j) = a'_j\}$.
 - ii. Answer with $a'_i = f_{u,v}(q'_i)$.
5. Set $\bar{q}' = (q'_1, \dots, q'_t)$ and $\bar{a}' = (a'_1, \dots, a'_t)$. In case $v = \perp$, set $v \leftarrow \{v' \in \mathcal{V} : \forall j \in [t]: f_{u,v'}(q'_j) = a'_j\}$.

Let $\bar{A}, \bar{Q}, \bar{A}', \bar{Q}', U$ and V be the (jointly distributed) random variables induced by the values of $\bar{a}, \bar{q}, \bar{a}', \bar{q}', u$ and v respectively, in a random execution of Algorithm A.3. By definition \bar{A} has the same distribution as the oracle answers in a random execution of D^π with $\pi \leftarrow \Pi$. In Claim A.4 we show that \bar{A}' is distributed the same as the oracle answers in a random execution of $D^{f_{u,v}}$ with $(u, v) \leftarrow \mathcal{U} \times \mathcal{V}$. Using it, we now conclude the proof by bounding the statistical distance between \bar{A} and \bar{A}' .

Since the queries and answers in both emulations of D at Algorithm A.3 are the same until $(\bar{Q}_{1,\dots,i}, U) \in \text{BAD}$ for some $i \in [t]$, and since BAD is monotone, it holds that

$$\Pr[\bar{A} \neq \bar{A}'] \leq \Pr[(\bar{Q}, U) \in \text{BAD}] \tag{1}$$

In addition, since U is chosen *after* \bar{Q} , the second condition of Lemma A.2 yields that

$$\Pr[(\bar{Q}, U) \in \text{BAD}] \leq \varepsilon \tag{2}$$

It follows that $\Pr[\bar{A} \neq \bar{A}'] \leq \varepsilon$ and therefore $\text{SD}(\bar{A}, \bar{A}') \leq \varepsilon$.

We conclude that

$$\left| \Pr_{\substack{u \leftarrow \mathcal{U} \\ v \leftarrow \mathcal{V}}} [D^{f_{u,v}} = 1] - \Pr_{\pi \leftarrow \Pi} [D^\pi = 1] \right| \leq \text{SD}(\bar{A}, \bar{A}') \leq \varepsilon.$$

□

Claim A.4. \bar{A}' has the same distribution as the oracle answers in a random execution of $D^{f_{u,v}}$ with $(u, v) \leftarrow \mathcal{U} \times \mathcal{V}$.

Proof. It is easy to verify that \bar{A}' is the oracle answers in $D^{f_{u,v}}$. Hence, to obtain the claim we need to show that (U, V) is uniformly distributed over $\mathcal{U} \times \mathcal{V}$. The definition of Algorithm A.3 assures that U is uniformly distributed over \mathcal{U} , so it is left to show that conditioned on any fixing u of U , the value of V is uniformly distributed over \mathcal{V} .

In the following we condition on $U = u \in \mathcal{U}$. For an answers vector $\bar{w} \in \mathcal{R}^k$, let $\bar{q}_{\bar{w}}$ [resp., $\bar{q}_{\bar{w}}^+$] be the first k [resp., $k + 1$] queries asked by \mathcal{D} , assuming that it gets \bar{w} as the first k answers (since \mathcal{D} is deterministic these values are well defined). Let $\mathcal{S}_{\bar{w}} = \{v \in \mathcal{V} : f_{u,v}(\bar{q}_{\bar{w}}) = \bar{w}\}$ and let $W = \{\bar{w} \in \mathcal{R}^* : |\bar{w}| \leq t \wedge (\bar{q}_{\bar{w}}, u) \notin \text{BAD}\}$. If $\lambda \notin W$, it follows that $(\lambda, u) \in \text{BAD}$, and thus Algorithm [A.3](#) chooses v at Step [3](#). Hence V is uniformly distributed over \mathcal{V} . In case $\lambda \in W$, we conclude the proof by applying the following claim (proven below) with $\bar{w} = \lambda$ (note that $\mathcal{S}_{\lambda} = \mathcal{V}$).

Claim A.5. *Conditioned on $\bar{A}'_{1,\dots,i} = \bar{w} \in W$ for some $i \in \{0, \dots, t\}$, the value of V is uniformly distributed over $\mathcal{S}_{\bar{w}}$.* □

Proof of Claim A.5. We prove by reverse induction on $i = |\bar{w}|$. For the base case $i = t$, we note that (by definition) Algorithm [A.3](#) chooses v at Step [5](#), and thus V is uniformly distributed over $\mathcal{S}_{\bar{w}}$. In the following we assume the hypothesis holds for $i + 1$, and condition on $\bar{A}'_{1,\dots,i} = \bar{w} \in W$. In case $(\bar{q}_{\bar{w}}^+, u) \in \text{BAD}$, Algorithm [A.3](#) chooses v at Step [4\(b\)i](#), and thus V is uniformly distributed over $\mathcal{S}_{\bar{w}}$. So it is left to handle the case $(\bar{q}_{\bar{w}}^+, u) \notin \text{BAD}$.

Fix $v' \in \mathcal{S}_{\bar{w}}$ and let $a \in \mathcal{R}$ be such that $v' \in \mathcal{S}_{\bar{w} \circ a}$, where ‘ \circ ’ denotes vector concatenation (i.e., for $\bar{w} = (w_1, \dots, w_i)$, $\bar{w} \circ a = (w_1, \dots, w_i, a)$). Conditioning on $A'_{i+1} = a$, we can apply the induction hypothesis on $\bar{w} \circ a$ (since $\bar{w} \circ a \in W$) to get that V is uniformly distributed over $\mathcal{S}_{\bar{w} \circ a}$. It follows that

$$\begin{aligned} \Pr[V = v' \mid \bar{A}'_{1,\dots,i} = \bar{w}] &= \Pr[A'_{i+1} = a \mid \bar{A}'_{1,\dots,i} = \bar{w}] \cdot \Pr[v = v' \mid \bar{A}'_{1,\dots,i+1} = \bar{w} \circ a] \\ &= \frac{1}{|\mathcal{R}|} \cdot \frac{1}{|\mathcal{S}_{\bar{w} \circ a}|} \\ &= \frac{1}{|\mathcal{R}|} \cdot \frac{|\mathcal{R}|^{|\bar{w}|+1}}{|\mathcal{V}|} \\ &= \frac{|\mathcal{R}|^{|\bar{w}|}}{|\mathcal{V}|} = \frac{1}{|\mathcal{S}_{\bar{w}}|}, \end{aligned}$$

concluding the induction step. The second equality holds by the induction hypothesis, and for the third one we note that

$$\frac{|\mathcal{S}_{\bar{w}'}|}{|\mathcal{V}|} = \Pr_{v \leftarrow \mathcal{V}}[v \in \mathcal{S}_{\bar{w}'}] = \Pr_{v \leftarrow \mathcal{V}}[f_{u,v}(\bar{q}_{\bar{w}'}) = \bar{w}'] = \frac{1}{|\mathcal{R}|^{|\bar{w}'|}}, \quad (3)$$

for every $\bar{w}' \in W$, where the third equality of Equation [\(3\)](#) holds by the first property of $\mathcal{F}(\mathcal{U}, \mathcal{V})$ (as stated in Lemma [A.2](#)). □

B Hardness-Preserving PRG to PRF Reductions

Another application of our technique is a hardness-preserving construction of PRFs from pseudorandom generators (PRGs). For instance, constructing $2^{c'n}$ -PRF for some $0 < c' < c$, from a 2^{cn} -PRG. The efficiency of such constructions is measured by the number of calls made to the underlying PRG as well as other parameters such as representation size.

The construction of Goldreich et al. [\[22\]](#) (i.e., \mathcal{GGM}) is in fact hardness preserving according to the above criterion. Their construction, however, makes n calls to the underlying PRG, which might be too expensive in some settings.

Proposition B.1 ([22]). *Let G be a length-doubling (t, ε) -PRG whose evaluation time is e_G . For any efficiently-computable integer functions m and ℓ , there exists an efficient oracle-aided function family ensemble whose n 'th function family, denoted $\mathcal{GGM}_{m(n) \rightarrow \ell(n)}^G$, maps strings of length $m(n)$ to strings of length $\ell(n)$, makes $m(n)$ calls to G and is a $(q, t - m \cdot q \cdot e_G(\ell), m \cdot q \cdot \varepsilon(\ell))$ -PRF for any integer function q .²⁰*

As already mentioned in the introduction, in order to reduce the number of calls to the underlying PRG, Levin [30] suggested to first hash the input to a smaller domain, and only then apply \mathcal{GGM} (this is known as ‘‘Levin’s trick’’). The resulting construction, however, is not hardness preserving due to the ‘‘birthday attack’’ described in Section 1.

While the \mathcal{GGM} construction seems optimal for the security it achieves (as shown in [26]), in some settings the number of queries the distinguisher can make is *strictly less* than its running time. Consider a distinguisher of running time 2^{cn} that can only make $2^{\sqrt{n}} \ll 2^{cn}$ queries. In such settings the security of the \mathcal{GGM} construction seems like an overkill and raises the question of whether there exist more efficient reductions. Jain et al. [26] (who raised the above question) gave the following partial answer, by designing a domain extension method tailored to PRG to PRF reductions for a specific range of parameters.

Theorem B.2 ([26]). *Let G be a length-doubling 2^{cn} -PRG. Let $c > 0$, $1/2 \leq \alpha < 1$ and $q(n) = 2^{n^\alpha}$. There exists a length-preserving function family \mathcal{JPT}^G that on input of length n makes $O(\log(q(n))) = O(n^\alpha)$ calls to G and is a $(q(n), 2^{c'n}, 2^{c'n})$ -PRF for every $0 < c' < c$.*

A restriction of Theorem B.2 is that it dictates that the resulting PRF family makes *at least* $\Omega(\sqrt{n})$ calls to the underlying PRG (since $1/2 \leq \alpha < 1$). We note that the restriction that $\alpha > 1/2$ (and hence $q(n) > 2^{\sqrt{n}}$) in the construction of [26] is inherent due to their hashing technique (and is not a mere side-effect of the parameters above).

Using better hashing constructions (based on cuckoo hashing) yields a more versatile version of the above theorem, that in particular allows α to be arbitrary. Specifically, combining Proposition B.1 with Theorem 4.1 yields the following result.

Corollary B.3. *Let G be a length-doubling (t, ε) -PRG. Let $\mathcal{H} = \{\mathcal{H}_n: \{0, 1\}^n \mapsto \{0, 1\}^{m(n)}\}_{n \in \mathbb{N}}$ and $\mathcal{G} = \{\mathcal{G}_n: \{0, 1\}^n \mapsto \{0, 1\}^n\}_{n \in \mathbb{N}}$ be efficient $k(n)$ -wise independent function family ensembles. Let $q(n) \leq 2^{m(n)-2}$.*

Then, the length-preserving function family ensemble $\{\mathcal{PP}(\mathcal{H}_n, \mathcal{G}_n, \mathcal{GGM}_{m(n) \rightarrow n}^G)\}_{n \in \mathbb{N}}$, when invoked on input of length n , makes $m(n)$ calls to G and is a $(q, t - p \cdot m \cdot q, 2m \cdot q \cdot \varepsilon + q/2^{\Omega(k)})$ -PRF, where $p(\cdot)$ is a polynomial determined by the evaluation and sampling time of \mathcal{H} , \mathcal{G} and G .

In particular, for $c > 0$, $0 < \alpha < 1$, $t(n) = 2^{cn}$, $\varepsilon(n) = 1/t(n)$, $q(n) = 2^{n^\alpha}$, $m(n) = \Theta(\log(q(n)))$ and $k(n) = \Theta(n^\alpha + cn)$, the function family $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{GGM}_{m(n) \rightarrow n}^G)$ makes $m(n) = O(n^\alpha)$ calls to G and is a $(q(n), 2^{c'n}, 2^{c'n})$ -PRF, for every $0 < c' < c$.

Proof. We prove the ‘‘In particular’’ part of the corollary. Set $q(n) = 2^{n^\alpha}$, $m(n) = \lceil n^\alpha \rceil + 2$, and \mathcal{H} and \mathcal{G} to be $k(n)$ -wise independent for $k(n) = \Theta(n^\alpha + cn)$, with an appropriate constant, such that

²⁰ $\mathcal{GGM}_{m(n) \rightarrow \ell(n)}^G$ is a variant of the standard \mathcal{GGM} function family, that on input of length $m(n)$ uses seed of length $\ell(n)$ for the underlying generator, rather than seed of length $m(n)$. Formally, $\mathcal{GGM}_{m \rightarrow \ell}^G$ is the function family ensemble $\{\mathcal{GGM}_{m(n) \rightarrow \ell(n)}^G\}_{n \in \mathbb{N}}$, where $\mathcal{GGM}_{m(n) \rightarrow \ell(n)}^G = \{f_r\}_{r \in \{0, 1\}^{\ell(n)}}$, and for $r \in \{0, 1\}^{\ell(n)}$, the oracle-aided function $f_r: \{0, 1\}^{m(n)} \mapsto \{0, 1\}^{\ell(n)}$ is defined as follows: given oracle access to a length-doubling function G and input $x \in \{0, 1\}^{m(n)}$, $f_r^G(x) = r_x$, where r_x is recursively defined by $r_\varepsilon = r$, and, for a string w , $r_{w||0} || r_{w||1} = G(r_w)$. The original \mathcal{GGM} construction was length-preserving, i.e., $m(n) = \ell(n) = n$.

Family	#queries limitation	description (key) size	evaluation time
\mathcal{JPT} [26]	$2^{n^{1/2}} < q < 2^n$	$\Theta(\log q \cdot n)$	$\Theta(\log q \cdot e_G + e_{(\log q)})$
$\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{GGM}_m^G)$	$0 < q < 2^n$	$\Theta(n^2)$	$\Theta(\log q \cdot e_G + e_{(n)})$

Table 2: Comparison between the family \mathcal{JPT} in Theorem B.2 to $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{GGM}_m^G)$ in Corollary B.3. The notation $e_{(k)}$ in the table refers to the evaluation time of a k -wise independent function.

$\frac{q(n)}{2^{\Omega(k(n))}} < 2^{-cn}$. Let $t' = t - p \cdot m \cdot q$ and $\varepsilon' = 2m \cdot q \cdot \varepsilon + q/2^{\Omega(k)}$. By the first part of the corollary we get that $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{GGM}_{m(n) \rightarrow n}^G)$ makes $m(n)$ calls to G and is a (q, t', ε') -PRF. Next, we show that $t(n) > 2^{c'n}$ and $\varepsilon(n) < 2^{-c'n}$ for large enough n .

Let $c'' \in \mathbb{N}$ such that $n^{c''} > p(n)$ for large enough n . It follows that $t(n) > 2^{cn} - n^{c''} 2^{n^\alpha} (n^\alpha + 2)$ and $\varepsilon(n) < 2^{1+\log(n^\alpha+2)+n^\alpha-cn} + 2^{-cn}$. Hence, for every $c' < c$, we have $\varepsilon(n) < 2^{-c'n}$ and $t(n) > 2^{c'n}$ for large enough n , as required. \square

Comparison with the Jain et al. reduction The advantage of Corollary B.3 is that when the adversaries are allowed to make less than $2^{\sqrt{n}}$ queries, the number of calls to the PRG is reduced accordingly, and below $O(\sqrt{n})$ calls. This improves upon the function family \mathcal{JPT} , that for such adversaries must make at least $O(\sqrt{n})$ calls to the PRG.

The function family \mathcal{JPT} , however, might have shorter description (key) and evaluation time. Specifically, let q denote the number of queries the adversaries are allowed to make (i.e., $q = 2^{n^\alpha}$). According to Corollary B.3, the parameter k (the independence required) needs to be set to $\Theta(n)$. Hence, by Fact 2.6 it takes $\Theta(n^2)$ bits to describe a function in $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{GGM}_m^G)$. The evaluation time of a single call to $\mathcal{PP}(\mathcal{H}, \mathcal{G}, \mathcal{GGM}_m^G)$ is $\Theta(\log q \cdot e_G + e_{(n)})$, where e_G is the evaluation time of G and $e_{(k)}$ is the evaluation time of a k -wise independent function. In comparison, it takes $\Theta(\log q \cdot n)$ bits to describe a member in \mathcal{JPT} , and its evaluation time is $\Theta(\log q \cdot e_G + e_{(\log q)})$. This is summarize in Table 2.

Independent work. Independently and concurrently with this work, Chandran and Garg [13] showed that a variant of the construction of [26] achieves similar security parameters to [26] and also works for 2^{n^α} queries for any $0 < \alpha < 1/2$. The construction of [13], however, outputs only $n^{2\alpha}$ bits, as opposed to n bits in the construction of [26] and in our construction.