

Symmetric and Dual PRFs from Standard Assumptions: A Generic Validation of a Prevailing Assumption

MIHIR BELLARE¹

ANNA LYSYANSKAYA²

February 2024

Abstract

A two-input function is a dual PRF if it is a PRF when keyed by either of its inputs. Dual PRFs are assumed in the design and analysis of numerous primitives and protocols including HMAC, AMAC, TLS 1.3 and MLS. But, not only do we not know whether particular functions on which the assumption is made really are dual PRFs; we do not know if dual PRFs even exist. What if the goal is impossible? This paper addresses this with a foundational treatment of dual PRFs, giving constructions based on standard assumptions. This provides what we call a generic validation of the dual PRF assumption. Our approach is to introduce and construct symmetric PRFs, which imply dual PRFs and may be of independent interest. We give a general construction of a symmetric PRF based on a function having a weak form of collision resistance coupled with a leakage hardcore function, a strengthening of the usual notion of hardcore functions we introduce. We instantiate this general construction in two ways to obtain two specific symmetric and dual PRFs, the first assuming any collision-resistant hash function, and the second assuming any one-way permutation. A construction based on any one-way function evades us and is left as an intriguing open problem.

¹ Department of Computer Science & Engineering, University of California San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. Email: mbellare@ucsd.edu. URL: <http://cseweb.ucsd.edu/~mihir/>. Supported in part by NSF grant CNS-2154272 and KACST. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467.

² Computer Science Department, Brown University, Providence, RI 02912, USA. Email: anna.lysyanskaya@brown.edu. URL: <https://cs.brown.edu/people/anna/>. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467.

Contents

1	Introduction	3
2	Basic definitions	6
3	Dual PRF security of existing PRF constructions	8
4	Leakage hardcore functions	9
5	The SPRF construction	11
6	Instantiations	14
6.1	Construction from (keyless) CR hash functions	15
6.2	Construction from any OWP	17
6.3	Ending remarks	18
	References	18

1 Introduction

A function family is a dual PRF [6] if it is a PRF and also remains so when its key and input are switched. This property was used as an assumption on the compression function in order to prove security of two hash-function based PRFs, namely the widely-used HMAC [8] and the newer AMAC [7]. Dual PRFs are also now being assumed in TLS 1.3 [22, 19] and other Internet security protocols [15, 32, 28, 1, 18].

We have, however, no constructions of dual PRFs under standard assumptions, and thus little idea how strong is the assumption, or if it is even valid. We address this with a foundational treatment of dual PRFs, giving constructions based on standard assumptions. This is the first theoretical evidence that dual PRFs exist, and provides what we call a generic validation of the dual PRF assumption. Tools that we introduce and use for our construction include leakage hardcore functions and symmetric PRFs.

PRFS. Let $F : F.Keys \times F.Inp \rightarrow F.Out$ be a function family taking a key $fk \in F.Keys$ and an input $x \in F.Inp$ to (deterministically) return the output $y = F(fk, x) \in F.Out$. We recall that F is a PRF [24] if an efficient adversary has negligible advantage in distinguishing whether its oracle is $F(fk, \cdot)$ or a random function, where fk is chosen at random from $F.Keys$. This well-known notion has seen an enormous number of applications in both theoretical and applied cryptography.

DUAL PRFS. Let $S : S_0 \times S_1 \rightarrow S.Out$ be a function family. Let $S^{swap} : S_1 \times S_0 \rightarrow S.Out$ be defined by $S^{swap}(a_0, a_1) = S(a_1, a_0)$. That is, the key for S^{swap} is the input for S and the input for S^{swap} is the key for S . Both S and S^{swap} are legitimate function families and we can ask if they are PRFs. We say that S is a dual PRF [6] if both S and S^{swap} are PRFs. That is (1) an oracle for $S(a_0, \cdot)$ is indistinguishable from an oracle for a random function when a_0 is chosen at random and, separately but also, (2) an oracle for $S(\cdot, a_1)$ is indistinguishable from an oracle for a random function when a_1 is chosen at random. The question we consider in this paper is, do dual PRFs exist, and, if so, under what assumptions?

CONTEXT. Dual PRFs were introduced by Bellare [6] in the context of HMAC. Recall that HMAC [8] is a cryptographic-hash-function-based PRF implemented in TLS and many other places. From the proof perspective, the underlying primitive is the compression function h of the hash function, and this is assumed in [6] to be a dual PRF in order to conclude PRF-security of HMAC. (In a little more detail, one starts with a related and simpler design, NMAC [8], that is PRF-secure assuming h is a PRF [6, 23, 3]. The dual PRF assumption on h arises in stepping from NMAC to HMAC [6].)

AMAC is a hash-function based PRF used in the widely deployed Ed25519 signature scheme [13], and its analysis also assumes the compression function is a dual PRF [7]. And since then, the use of dual PRFs has widened even further. Dual PRFs are now invoked in the design and analysis of many Internet security protocols, including TLS 1.3 [22, 19], hybrid key-exchange [15, 32], post-quantum versions of WireGuard [28] and Noise [1], and Message Layer Security (MLS) [18].

GENERIC VALIDATION. The assumption that a function h is a dual PRF could fail for two reasons. One is generic, namely that *nothing* can be a dual PRF. Dual PRFs may simply not exist. The second reason is specific, namely that, although some functions may be dual PRFs, the particular h used in some particular application isn't.

Generic failure can be ruled out by showing that the security goal is achievable under standard assumptions. We call this *generic validation*. It has value because generic failure is not an idle fear. It has happened for several (attractive) goals, for example virtual blackbox obfuscation [26, 5] and commitment secure against selective opening [10] to name just a few.

Generic validation won't show that a particular candidate practical construct satisfies the as-

sumption. This needs *dedicated validation*, meaning either a dedicated proof or cryptanalysis. But generic validation is the first step. In its absence, the goal may be just wishful thinking, and the candidate construct doomed. In its presence, the candidate is at least in principle plausible, and successful dedicated validation is a possibility. Generic validation is thus desirable for the security goal underlying any new assumption.

For (standard) PRFs, we have strong generic validation: classical foundational results say that PRFs exist assuming only that one-way functions exist. (OWFs imply PRGs [27] which imply PRFs [24].) We also have constructions from many particular assumptions [30, 29, 4]. Dual PRFs, in contrast, have at this point *no generic validation*. Despite their having been introduced ten years ago [6], and despite their use as an assumption in supporting the security of the widely-used HMAC [6], there has been no construction under any (standard or not) assumption. This is the gap we fill.

NEGATIVE RESULTS. One’s first thought may be that every PRF S is also a dual PRF. It is easy to see that this is not true. For example suppose $S : \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}^k$ is a PRF with the property that $S(0^k, a) = a$ for all a . This will not contradict PRF security of S because 0^k has negligible probability of being chosen as the key in the PRF game. However S^{swap} is clearly not a PRF because $S^{\text{swap}}(a, 0^k) = S(0^k, a) = a$ so an adversary can query its oracle at 0^k and it will get back the key a , using which it can easily violate PRF security.

Thus we need special constructions. The next natural question is whether known constructions of PRFs are dual PRFs. But they are not. For example, take the classic GGM construction [24] of a PRF from a PRG. We show in Section 3 that there is a choice of the PRG under which the constructed PRF is not a dual PRF. Or take the Naor-Reingold PRF. We give in Section 3 a direct attack violating dual PRF security. The Dodis-Yampolskiy PRF [21] is promising because the formula adds the key and input, thereby seeming to give them symmetric roles, but security requires that the input comes from a much smaller space than the key, and this precludes being a symmetric PRF as per our definition. See Section 3 for more information.

SYMMETRIC PRFS. Our approach to construct dual PRFs is based on the notion we introduce of a symmetric PRF. Let $S : S \times S \rightarrow S$.Out be a function family whose key space and input space are the same set, call it S . We say that S is *symmetric* if $S(a_0, a_1) = S(a_1, a_0)$ for all $a_0, a_1 \in S$. That is, S is unchanged if the order of its inputs is swapped. Then we make the following observation. Suppose S is (1) A PRF, and (2) is symmetric. Then it is a dual PRF. This is easy to see because the symmetry implies that $S^{\text{swap}} = S$, namely S^{swap} is in fact identical to S . So its PRF security follows directly from the fact that S is a PRF. We will construct symmetric PRFs.

SPRF. In Section 5 we give a general construction of a symmetric (hence dual) PRF $S : D \times D \rightarrow \{0, 1\}^k$. It is defined in terms of three other functions E, H, R as follows:

Function family $S(a_0, a_1)$

```

 $r_0 \leftarrow E(a_0) ; r_1 \leftarrow E(a_1)$ 
 $z_0 \leftarrow H(a_0) ; z_1 \leftarrow H(a_1)$ 
 $y_0 \leftarrow R(r_0, z_1) ; y_1 \leftarrow R(r_1, z_0)$ 
 $y \leftarrow y_0 \oplus y_1 ; \text{Return } y$ 

```

Here R is a PRF with range $\{0, 1\}^k$ and D is some appropriate domain. The functions E, H can be thought of roughly as “extract” and “hash,” and they will be instantiated in different ways. The idea is that r_0, z_0 depend on the input a_0 while r_1, z_1 depend on the input a_1 , and only in the application of R are the inputs “mixed.” Two applications of R are used, the key being an r -value

and the input the opposing z -value. Note that the use of this high-level structure with the xor already guarantees that S is symmetric, regardless of the choices of R, E, H .

Now we need to find choices of E, H under which S is a PRF. Intuitively, a difficulty in using the PRF security of R is that the construction does not use a key for R in a blackbox way. If we think of r_0 as the key, then z_0 is related information that is needed to simulate an attacker against S .

Very roughly, we want E to extract hardcore bits, and we want H to provide some kind of collision resistance (CR). In the proof that S is a PRF we would first use the security of E to move to a game in which r_0 is random. Then we would use the PRF security of R to replace $R(r_0, \cdot)$ with a random function R . Finally we would use the CR-security of H to say that the z_1 values do not repeat, which means in each xor the first component, and hence the whole, is random.

However getting this to work requires some care. We strive to make the conditions on E, H as weak and general as possible so as to allow the maximum flexibility in instantiation and the ability to instantiate under assumptions as weak as possible. In this spirit one choice we make is to allow both E and H to be keyed. Both the key and the input would be derived from the single input a_i above. Now the main difficulty is that no standard notion of hardcore function security suffices for E . Instead we introduce the notion of E being a leakage hardcore function for H . Roughly—the formal definition is in Section 4—this means that E with a target key applied to a hidden x_0 continues to look random even given an oracle that can get the results of H at x_0 under other, different keys of its choice. For H , we ask that it be computationally almost universal (CAU) [6]. This is a weak form of collision resistance in which the adversary must produce its collision without knowing the key. See Section 5 for the full construction and Theorem 5.2 for the formal claim and proof of PRF security.

INSTANTIATIONS. To obtain constructions of symmetric (and hence dual) PRFs under specific, standard assumptions, we instantiate the primitives in our general SPRF construction under the assumption in question. In Section 6 we give two corresponding results, one under one-way permutations (OWPs) and the other under collision-resistant (CR) hash functions, meaning either of these assumptions now yield symmetric and dual PRFs. The OWP instantiation uses the Blum-Micali-Yao (BMY) PRG [16, 33] to instantiate the leakage hardcore function E and an iterated OWP to instantiate H . The CR hash function instantiation uses CR hash to instantiate H and uses a strong randomness extractor to instantiate E .

DISCUSSION AND OPEN QUESTIONS. The main open question that evades us is a construction of a symmetric and dual PRF from any one-way function (OWF). The first question is whether one can instantiate our SPRF construction under a OWF. If not, the next question is whether there is some other, different construction.

We note that while our result about SPRF has striven to make as general and weak-as-possible assumptions on the component E, H functions, we have not, in our instantiations, found a way to take full advantage of this. The only way we have found to get a leakage hardcore function E for H is to make H a keyless CR function, in which case Lemma 4.1 says that E being a standard hardcore function for H suffices. But there may, potentially, exist choices of keyed, CAU functions H for which a leakage hardcore function E exists, and this may then be a direction towards a OWF-based dual PRF.

SUBSEQUENT WORK. The motivation for our new constructions of dual PRFs was primarily theoretical, namely to give a generic validation for the dual PRF assumption on the compression function used in the proof of PRF security of HMAC [6]. Following the posting of our paper on the Cryptology ePrint Archive [11], however, Aviram, Dowling, Komargodski, Paterson, Ronen and

Yogev (ADKPRY) [2] revisit the problem of constructing dual PRFs with a more practical motivation, namely the use of dual PRFs as key combiners in the TLS 1.3 key schedule. They extend our general construction above to apply, at the end, an output function G , meaning their dual PRF returns $G(S(a_0, a_1))$ where $S(a_0, a_1)$ is defined via R, E, H as above. They then instantiate R, E, H, G via HMAC to obtain an efficient dual PRF.

The assumption made in TLS 1.3 [22, 19] and the other above-mentioned Internet security protocols [15, 32, 28, 1, 18] is that HMAC itself is a dual PRF. This assumption has been validated by Backendal, Bellare, Günther and Scarlata (BBGS) [3] via a proof of dual PRF security of HMAC based on certain assumptions on the underlying compression function h . We note that these assumptions include that h is itself a dual PRF.

2 Basic definitions

Our treatment is concrete rather than asymptotic. For any security goal for a primitive, for example PRF security of a function family, we define an advantage metric, in this case the PRF advantage of an adversary against the function family, which is a number. There is no explicit security parameter; one way of thinking about it is to consider that the security parameter has been fixed. For a function family to be a PRF typically means, informally, that “efficient” adversaries have “negligible” PRF advantage; in the absence of a security parameter, this is defined in quantitative, rather than asymptotic, terms. Theorems are made formal by giving the concrete security of reductions. Discussion surrounding theorems will clarify what they mean qualitatively. The concrete treatment makes notation somewhat simpler, allows us to see the quantitative security of reductions, and is more in keeping with the motivating setting of HMAC, where there are no asymptotics.

NOTATION AND CONVENTIONS. We let ε denote the empty string. If y is a string then $|y|$ denotes its length and $y[i]$ denotes its i -th coordinate for $1 \leq i \leq |y|$. If X is a finite set, we let $x \leftarrow_s X$ denote picking an element of X uniformly at random and assigning it to x . Algorithms may be randomized unless otherwise indicated. Running time is worst case. If A is an algorithm, we let $y \leftarrow A(x_1, \dots; r)$ denote running A with random coins r on inputs x_1, \dots and assigning the output to y . We let $y \leftarrow_s A(x_1, \dots)$ be the result of picking r at random and letting $y \leftarrow A(x_1, \dots; r)$. We let $[A(x_1, \dots)]$ denote the set of all possible outputs of A when invoked with inputs x_1, \dots .

We use the code based game playing framework of [12]. (See Fig. 1 for an example.) By $\text{Pr}[G]$ we denote the event that the execution of game G results in the game returning true. We adopt the convention that the running time of an adversary refers to the worst-case execution time of the game with the adversary, so that the time for the execution of oracles to compute replies to oracle queries is included. This means that usually in reductions, adversary running time is roughly maintained. In writing a game, we assume boolean variables (e.g. `bad`) are automatically initialized to false.

FUNCTION FAMILIES. A function family $F : F.\text{Keys} \times F.\text{Inp} \rightarrow F.\text{Out}$ is a 2-argument function taking a key fk in the keyspace $F.\text{Keys}$ and an input x in the input space $F.\text{Inp}$ to return an output $F(fk, x)$ in the output space $F.\text{Out}$. For $fk \in F.\text{Keys}$ we let $F_{fk} : F.\text{Inp} \rightarrow F.\text{Out}$ be defined by $F_{fk}(x) = F(fk, x)$ for all $x \in F.\text{Inp}$. We say that F is a permutation family if $F.\text{Inp} = F.\text{Out}$ and F_{fk} is a permutation for every $fk \in F.\text{Keys}$. We say that F is keyless if $F.\text{Keys} = \{\varepsilon\}$ consists only of the empty string. (It is tempting in this case to just drop the key in the notation but it makes it harder to pattern-match with the definitions and so, somewhat pedantically, we tend to explicitly write ε as the key when dealing with keyless families.) The reason to consider such families is that some notions of security, such as one-wayness, hold just as well for them. (For others, like PRF-security, keying is crucial.)

<p><u>Game $\mathbf{G}_F^{\text{prf}}(\mathcal{A})$</u> $fk \leftarrow_s \mathbf{F}.\text{Keys}$ $c \leftarrow_s \{0, 1\}; c' \leftarrow_s \mathcal{A}^{\text{FN}}$ Return ($c = c'$)</p> <p><u>$\text{FN}(x)$</u> If $T[x] = \perp$ then If ($c = 1$) then $T[x] \leftarrow \mathbf{F}(fk, x)$ Else $T[x] \leftarrow_s \mathbf{F}.\text{Out}$ Return $T[x]$</p>	<p><u>Game $\mathbf{G}_F^{\text{ow}}(\mathcal{A})$</u> $fk \leftarrow_s \mathbf{F}.\text{Keys}$ $x \leftarrow_s \mathbf{F}.\text{Inp}; y \leftarrow \mathbf{F}(fk, x)$ $x' \leftarrow_s \mathcal{A}(fk, y)$ Return ($\mathbf{F}(fk, x') = y$)</p> <p><u>Game $\mathbf{G}_H^{\text{cau}}(\mathcal{A})$</u> $(x_0, x_1) \leftarrow_s \mathcal{A}$ $hk \leftarrow_s \mathbf{H}.\text{Keys}$ If ($x_0 = x_1$) then return false Return ($\mathbf{H}(hk, x_0) = \mathbf{H}(hk, x_1)$)</p> <p><u>Game $\mathbf{G}_H^{\text{cr}}(\mathcal{A})$</u> $hk \leftarrow_s \mathbf{H}.\text{Keys}$ $(x_0, x_1) \leftarrow_s \mathcal{A}(hk)$ If ($x_0 = x_1$) then return false Return ($\mathbf{H}(hk, x_0) = \mathbf{H}(hk, x_1)$)</p>
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 1: **Games for defining PRF and OWF security of a function family F, CAU-security of a function family H and HC being a hardcore function family for H.**

PSEUDO-RANDOM FUNCTIONS. The security of function family \mathbf{F} as a PRF is defined via game $\mathbf{G}_F^{\text{prf}}(\mathcal{A})$ of Fig. 1 associated to \mathbf{F} and adversary \mathcal{A} . Table T is assumed initially \perp everywhere. The PRF advantage of \mathcal{A} is

$$\begin{aligned} \mathbf{Adv}_F^{\text{prf}}(\mathcal{A}) &= 2 \Pr[\mathbf{G}_F^{\text{prf}}(\mathcal{A})] - 1 \\ &= \Pr[\mathbf{G}_F^{\text{prf}}(\mathcal{A}) | c = 1] - \left(1 - \Pr[\mathbf{G}_F^{\text{prf}}(\mathcal{A}) | c = 0]\right). \end{aligned} \quad (1)$$

The first equation is the definition, while the second is an alternative representation known to be equal by a standard conditioning argument.

ONE-WAY FUNCTIONS. The security of function family \mathbf{F} as a OWF is defined via game $\mathbf{G}_F^{\text{ow}}(\mathcal{A})$ of Fig. 1 associated to \mathbf{F} and adversary \mathcal{A} . The point x' returned by the latter is required to be in $\mathbf{F}.\text{Inp}$. The owf advantage of \mathcal{A} is defined as $\mathbf{Adv}_F^{\text{ow}}(\mathcal{A}) = \Pr[\mathbf{G}_F^{\text{ow}}(\mathcal{A})]$. In this case, \mathbf{F} may or may not be keyed. A one-way permutation (OWP) is simply a family of permutations that is a OWF.

UNIVERSAL AND CAU FUNCTIONS. Consider game $\mathbf{G}_H^{\text{cau}}(\mathcal{A})$ of Fig. 1 associated to \mathbf{H} and adversary \mathcal{A} . The points x_0, x_1 returned by the latter are required to be in $\mathbf{H}.\text{Inp}$. The CAU-advantage of \mathcal{A} is defined as $\mathbf{Adv}_H^{\text{cau}}(\mathcal{A}) = \Pr[\mathbf{G}_H^{\text{cau}}(\mathcal{A})]$. We say that \mathbf{H} is universal if $\mathbf{Adv}_H^{\text{cau}}(\mathcal{A}) = 1/|\mathbf{H}.\text{Out}|$ for all adversaries \mathcal{A} , regardless of their computing time. Computational almost universal functions, introduced by Bellare [6], are a relaxation of universal functions in which the advantage is treated as a computational metric in the usual way and adversaries may be computationally bounded.

CR FUNCTIONS. The security of function family \mathbf{H} as a collision-resistant (CR) function is defined via game $\mathbf{G}_H^{\text{cr}}(\mathcal{A})$ of Fig. 1 associated to \mathbf{H} and adversary \mathcal{A} . The points x_0, x_1 returned by the latter are required to be in $\mathbf{H}.\text{Inp}$. The cr advantage of \mathcal{A} is defined as $\mathbf{Adv}_H^{\text{cr}}(\mathcal{A}) = \Pr[\mathbf{G}_H^{\text{cr}}(\mathcal{A})]$.

Practical CR hash functions such as SHA-256 are keyless. A CR function family is CAU, giving an easy way to get the latter.

EXTRACTORS. Let X, Y be random variables. We define $\mathbf{SD}(X, Y)$, the statistical distance between X and Y ; $\mathbf{H}_\infty(X)$, the min-entropy of X ; and $\mathbf{H}_\infty(X|Y)$, the min-entropy of X given Y , via:

$$\begin{aligned}\mathbf{SD}(X, Y) &= \frac{1}{2} \sum_z |\Pr[X = z] - \Pr[Y = z]| \\ 2^{-\mathbf{H}_\infty(X)} &= \max_x \Pr[X = x] \\ 2^{-\mathbf{H}_\infty(X|Y)} &= \sum_y \Pr[Y = y] \cdot \max_x \Pr[X = x | Y = y].\end{aligned}$$

Recall, paraphrasing the definition above, that a function family $\text{Ext} : \{0, 1\}^s \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ is universal if for every distinct $x_1, x_2 \in \{0, 1\}^n$ we have $\Pr[\text{Ext}(sk, x_1) = \text{Ext}(sk, x_2)] = 2^{-m}$ where the probability is over $sk \leftarrow_s \{0, 1\}^s$. The following is a generalized version of the Leftover Hash Lemma (LHL) [27, 20].

Lemma 2.1 *Let $\text{Ext} : \{0, 1\}^s \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a function family that is universal. Let X be a random variable over $\{0, 1\}^n$. Let U_s, U_m be random variables distributed uniformly over $\{0, 1\}^s$ and $\{0, 1\}^m$, respectively, and let Y be a random variable. Assume the three random variables $(X, Y), U_s, U_m$ are independent. Then*

$$\mathbf{SD}((U_s, \text{Ext}(U_s, X), Y), (U_s, U_m, Y)) \leq \frac{1}{2} \sqrt{2^{m - \mathbf{H}_\infty(X|Y)}}. \quad (2)$$

SYMMETRIC PRFS. Let $\mathbf{S} : S_0 \times S_1 \rightarrow \mathbf{S.Out}$ be a function family. Let $\mathbf{S}^{\text{swap}} : S_1 \times S_0 \rightarrow \mathbf{S.Out}$ be defined by $\mathbf{S}^{\text{swap}}(a_0, a_1) = \mathbf{S}(a_1, a_0)$. We say that \mathbf{S} is a dual PRF if both \mathbf{S} and \mathbf{S}^{swap} are PRFs. We say that \mathbf{S} is *symmetric* if $S_0 = S_1$ and $\mathbf{S}(a_0, a_1) = \mathbf{S}(a_1, a_0)$ for every $a_0, a_1 \in S_1$. If \mathbf{S} is symmetric then $\mathbf{S}^{\text{swap}} = \mathbf{S}$. Thus if \mathbf{S} is symmetric and a PRF, it is automatically a dual PRF. We will accordingly target the stronger notion of a symmetric PRF and obtain a dual PRF as a consequence.

3 Dual PRF security of existing PRF constructions

If we seek dual PRFs, the first and natural question is whether existing constructions of PRFs might happen to already be dual. Here we look at a few popular ones and show this is not the case.

GGM. Let $F_1 : \{0, 1\}^k \times \{0, 1\} \rightarrow \{0, 1\}^k$ be a PRF with input space $\{0, 1\}$. The GGM construction [24] builds from it the PRF $\text{GGM} : \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}^k$ defined as follows.

Function family $\text{GGM}(x, y)$
 For $i = 1, \dots, k$ do $x \leftarrow F_1(x, y[i])$
 Return x

Suppose F_1 has the property that $F_1(0^k, 0) = F_1(0^k, 1) = 0^k$. It could still be a PRF and in particular if PRFs exist we can easily build a PRF F_1 with this property. But then $\text{GGM}^{\text{swap}}(y, 0^k) = \text{GGM}(0^k, y) = 0^k$ so GGM^{swap} is certainly not a PRF. Thus GGM is not a dual PRF. This shows that the GGM construction does not in general yield a dual PRF.

NAOR REINGOLD. Let \mathbb{G} be prime-order group in which the DDH problem is hard, and let $g \in \mathbb{G}$ be a generator of \mathbb{G} . Let $q = |\mathbb{G}|$. The Naor-Reingold PRF [30] $\text{NR} : \mathbb{Z}_q^{n+1} \times \{0, 1\}^n \rightarrow \mathbb{G}$ is defined by

Function family $\text{NR}(\mathbf{a}, x)$
 $b \leftarrow \mathbf{a}[0] \cdot \prod_{i=1}^n \mathbf{a}[i]^{x[i]} \pmod q$
 $y \leftarrow g^b$
 Return y

Here the key \mathbf{a} is a $(n + 1)$ -vector over \mathbb{G} and its i -th component is denoted $\mathbf{a}[i] \in \mathbb{G}$, with the components indexed from 0 to n . Let $1_{\mathbb{G}}$ denote the identity element of \mathbb{G} and let $\mathbf{0} = (0, \dots, 0) \in \mathbb{G}^{n+1}$ denote the $(n + 1)$ -vector all of whose components equal 0. Then $\text{NR}^{\text{swap}}(x, \mathbf{0}) = \text{NR}(\mathbf{0}, x) = g^0 = 1_{\mathbb{G}}$ for all $x \in \{0, 1\}^n$. Thus NR^{swap} cannot be a PRF and NR is not a dual PRF. This is true for all choices of \mathbb{G}, g .

Some variants of NR [9] restrict the keyspace to $(\mathbb{Z}_q^*)^{n+1}$, which would preclude the above attack on NR^{swap} . However, NR^{swap} is still subject to attack by setting \mathbf{a} to all 1s.

DODIS YAMPOLSKIY. Let $\mathbf{e} : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ be a non-degenerate bilinear map, where groups \mathbb{G}, \mathbb{G}_T have prime order p . Let g be a generator of \mathbb{G} and $S \subseteq \mathbb{Z}_p$ a set of size N . Then the Dodis Yampolskiy PRF [21] $\text{DY} : \mathbb{Z}_p \times S \rightarrow \mathbb{G}_T$ is defined by

Function family $\text{DY}(a, x)$
 If $(a + x) \pmod p = 0$ then $b \leftarrow 1$ else $b \leftarrow (a + x)^{-1} \pmod p$
 $y \leftarrow \mathbf{e}(g, g)^b$
 Return y

This construction is promising because the roles of a and x are symmetric, so we may think we can swap them and have a symmetric PRF. The difficulty is that for security the input x must come from a much smaller space than the key, meaning $N = |S|$ is much less than p . This is because security is based on the q -BDHI assumption, and as per [21, Theorem 2], security of the PRF requires $q = N$ and security of q -BDHI for adversaries with running time more than N . In particular, the construction is not shown secure when $S = \mathbb{Z}_p$. But to meet our definition of a symmetric PRF from Section 2, the key-space and domain must be the same set. This asymmetry in the key and input for DY, and how it precludes some applications, has been pointed out before in several contexts, including in BC [9] for security against related-key attack.

Finally we note that if $S = \mathbb{Z}_p$ then DY is symmetric. Hence, if it is a PRF then it is also a dual PRF. So is it a PRF when $S = \mathbb{Z}_p$? To the best of our knowledge, this is an open question; we are aware of neither a proof nor an attack.

DISCUSSION. Although this should be obvious, we should nonetheless clarify that the above attacks do not represent any bugs or critiques. These constructions were not designed or claimed to be dual PRFs. But the first question one should ask in seeking dual PRFs is whether existing constructions of PRFs happen to be dual PRFs. The above indicates that this is not the case and one must seek new constructions.

4 Leakage hardcore functions

For our construction we will introduce an extension of the standard notion of a hardcore function. We call it a leakage hardcore function. To understand it, it is useful to begin by recalling the usual notion.

HARDCORE FUNCTIONS. Suppose H is a function family. A *hardcore function* for H is a function family $\text{HC} : \text{HC.Keys} \times (\text{H.Keys} \times \text{H.Inp}) \rightarrow \text{HC.Out}$, so that an input is a pair (hk, x) consisting of a key for H and an input for H . We say that HC is a hardcore predicate for H if $\text{HC.Out} = \{0, 1\}$.

<p>Game $\mathbf{G}_{H,HC}^{\text{hc}}(\mathcal{A})$</p> <p>$hk_0 \leftarrow_{\\$} \text{H.Keys}$</p> <p>$hck_0 \leftarrow_{\\$} \text{HC.Keys}$</p> <p>$x_0 \leftarrow_{\\$} \text{H.Inp}$</p> <p>$w_0 \leftarrow \text{H}(hk_0, x_0)$</p> <p>$s_1 \leftarrow \text{HC}(hck_0, (hk_0, x_0))$</p> <p>$s_0 \leftarrow_{\\$} \text{HC.Out}$</p> <p>$c \leftarrow_{\\$} \{0, 1\}$</p> <p>$c' \leftarrow_{\\$} \mathcal{A}(hk_0, hck_0, w_0, s_c)$</p> <p>Return $(c = c')$</p>	<p>Game $\mathbf{G}_{H,HC}^{\text{lhc}}(\mathcal{A})$</p> <p>$hk_0 \leftarrow_{\\$} \text{H.Keys}$</p> <p>$hck_0 \leftarrow_{\\$} \text{HC.Keys}$</p> <p>$x_0 \leftarrow_{\\$} \text{H.Inp}$</p> <p>$s_1 \leftarrow \text{HC}(hck_0, (hk_0, x_0))$</p> <p>$s_0 \leftarrow_{\\$} \text{HC.Out}$</p> <p>$c \leftarrow_{\\$} \{0, 1\}$</p> <p>$c' \leftarrow_{\\$} \mathcal{A}^{\text{LK}}(hk_0, hck_0, s_c)$</p> <p>Return $(c = c')$</p> <hr/> <p>$\text{LK}(hk)$</p> <p>$w_0 \leftarrow \text{H}(hk, x_0)$</p> <p>Return w_0</p>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 2: **Games for defining security of HC as a standard and leakage hardcore function for H.**

(Some hardcore functions are unkeyed; in fact both the RSA and the DL function families have unkeyed hardcore functions. On the other hand, the Goldreich-Levin hardcore predicate has a key that is a randomly chosen string.) Recall that security considers an adversary given a key hk_0 defining the function $\text{H}(hk_0, \cdot)$, a key hck_0 for the hardcore function, and the result $w \leftarrow \text{H}(hk_0, x_0)$ of evaluating the function at $x_0 \leftarrow_{\$} \text{H.Inp}$. Now the adversary gets s_c for a challenge bit c where $s_1 = \text{HC}(hck_0, (hk_0, x_0))$ is the output of the hardcore function on x_0 and s_0 is a random string of the same length. The adversary should have a hard time figuring out c . Formally the security of HC as a hardcore function for H is defined via game $\mathbf{G}_{H,HC}^{\text{hc}}(\mathcal{A})$ of Fig. 2 associated to H, HC and adversary \mathcal{A} . The hcf advantage of \mathcal{A} is defined as $\text{Adv}_{H,HC}^{\text{hc}}(\mathcal{A}) = 2 \Pr[\mathbf{G}_{H,HC}^{\text{hc}}(\mathcal{A})] - 1$.

LEAKAGE HARDCORE FUNCTIONS. A *leakage hardcore (LHC) function* for H is again a function family $\text{HC} : \text{HC.Keys} \times (\text{H.Keys} \times \text{H.Inp}) \rightarrow \text{HC.Out}$, so that an input is a pair (hk, x) consisting of a key for H and an input for H. Again we say that HC is a leakage hardcore predicate for H if $\text{HC.Out} = \{0, 1\}$. The new element in a leakage hardcore function is that the adversary has an oracle LK via which it can obtain “leakage” about x_0 . This leakage has a very particular form (although one could define LHC functions more generally, allowing other leakage as well), namely the adversary can obtain the value of the same function family H on x_0 under any key $hk \in \text{H.Keys}$ of its choice. Thus LK takes input hk and returns $\text{H}(hk, x_0)$, the result of evaluating H on the given key under the hidden input x_0 . The requirement is that figuring out the challenge bit remains hard. The formalization uses game $\mathbf{G}_{H,HC}^{\text{lhc}}(\mathcal{A})$ of Fig. 2 associated to H, HC and adversary \mathcal{A} . The lhc advantage of \mathcal{A} is defined as $\text{Adv}_{H,HC}^{\text{lhc}}(\mathcal{A}) = 2 \Pr[\mathbf{G}_{H,HC}^{\text{lhc}}(\mathcal{A})] - 1$. Since \mathcal{A} could in particular call its oracle on hk_0 , we omit giving it $\text{H}(hk_0, x_0)$ as input as in the standard game.

BUILDING LEAKAGE HARDCORE FUNCTIONS. Towards getting a leakage hardcore function for a given function family H, one simple observation is that if H is keyless then a standard hardcore function is leakage hardcore. This is captured by the following lemma.

Lemma 4.1 *Suppose H is a keyless function family and $\text{HC} : \text{HC.Inp} \times (\{\varepsilon\} \times \text{H.Inp}) \rightarrow \text{HC.Out}$ is a function family. Let \mathcal{A} be a lhc-adversary. Then the proof constructs a hc-adversary \mathcal{A}_0 such that*

$$\text{Adv}_{H,HC}^{\text{lhc}}(\mathcal{A}) \leq \text{Adv}_{H,HC}^{\text{hc}}(\mathcal{A}_0).$$

<p>Function family $S(a_0, a_1)$</p> <p>$(x_0, hk_0, hck_0) \leftarrow a_0; (x_1, hk_1, hck_1) \leftarrow a_1$</p> <p>$r_0 \leftarrow \text{HC}(hck_0, (hk_0, x_0)); r_1 \leftarrow \text{HC}(hck_1, (hk_1, x_1))$</p> <p>$w_0 \leftarrow \text{H}(hk_1, x_0); w_1 \leftarrow \text{H}(hk_0, x_1)$</p> <p>$z_0 \leftarrow (w_0, hk_0, hck_0); z_1 \leftarrow (w_1, hk_1, hck_1)$</p> <p>$y_0 \leftarrow \text{R}(r_0, z_1); y_1 \leftarrow \text{R}(r_1, z_0)$</p> <p>$y \leftarrow y_0 * y_1$</p> <p>Return y</p>

Figure 3: **Our SPRF construction.**

Adversary \mathcal{A}_0 has about the same running time as adversary \mathcal{A} .

Proof of of Lemma 4.1: Adversary \mathcal{A}_0 gets inputs hk_0, hck_0, w_0, s_c and runs \mathcal{A} on inputs hk_0, hck_0, s_c . Since $\text{H.Keys} = \{\varepsilon\}$, the LK oracle is intuitively useless to \mathcal{A} . Formally, if a query hk is made by \mathcal{A} to LK then it must be that $hk = \varepsilon$, and thus \mathcal{A}_1 can simulate the oracle, returning w_0 as the response. Eventually \mathcal{A} outputs a bit c' , and \mathcal{A}_1 outputs the same bit. ■

Our construction of a symmetric PRF will need a CAU function family that has a leakage hardcore function which outputs lots of bits. In Section 5 we will assume it. Later we will give various constructions from various assumptions.

5 The SPRF construction

We provide our general SPRF construction of a symmetric, and hence dual, PRF.

INGREDIENTS. Our construction of a symmetric PRF has the following ingredients:

- A CAU function family $\text{H} : \text{H.Keys} \times \text{H.Inp} \rightarrow \text{H.Out}$
- A leakage hardcore function family $\text{HC} : \text{HC.Keys} \times (\text{H.Keys} \times \text{H.Inp}) \rightarrow \text{HC.Out}$ for H .
- A PRF $\text{R} : \text{HC.Out} \times \text{R.Inp} \rightarrow \text{R.Out}$ such that $\text{H.Out} \times \text{H.Keys} \times \text{HC.Keys} \subseteq \text{R.Inp}$ and the range R.Out is a commutative group whose operation we denote $*$. Thus a key for R is an output of HC while a triple consisting of an output of H , a key for H and a key for HC is a valid input for R .

We refer to a triple $(\text{H}, \text{HC}, \text{R})$ of function families satisfying the above conditions as a *suite*. The simplest case for the group is that $\text{R.Out} = \{0, 1\}^{\text{R.ol}}$ is the set of all strings of some length R.ol , and $y_1 * y_2 = y_1 \oplus y_2$, but the existence of efficient PRFs with algebraic ranges [30] motivates being more general.

SPRF CONSTRUCTION. Our construction associates to any suite $(\text{H}, \text{HC}, \text{R})$ as above the function family $\text{S} = \text{SPRF}[\text{H}, \text{HC}, \text{R}]$ defined as follows. It has $\text{S.Keys} = \text{S.Inp} = \text{H.Inp} \times \text{H.Keys} \times \text{HC.Keys}$, meaning a key or input is a triple $a = (x, hk, hck)$ consisting of a point $x \in \text{H.Inp}$, a key hk for the CAU family H and a key hck for the hardcore function family HC . It has range the group $\text{S.Out} = \text{R.Out}$. The function family is then defined as shown in Fig. 3.

Proposition 5.1 *Let $(\text{H}, \text{HC}, \text{R})$ be a suite of function families. Let $\text{S} = \text{SPRF}[\text{H}, \text{HC}, \text{R}]$ be the function family associated to them as above. Then S is symmetric.*

Games G_0, G_1	Games $\boxed{G_2}, G_3$	Game G_4
$hk_0 \leftarrow_{\$} \text{H.Keys}$	$hk_0 \leftarrow_{\$} \text{H.Keys}$	$c' \leftarrow_{\$} \mathcal{A}^{\text{FN}}$
$hck_0 \leftarrow_{\$} \text{HC.Keys}$	$hck_0 \leftarrow_{\$} \text{HC.Keys}$	Return ($c' = 1$)
$x_0 \leftarrow_{\$} \text{H.Inp}$	$x_0 \leftarrow_{\$} \text{H.Inp}$	$\text{FN}((x_1, hk_1, hck_1))$
$r_0 \leftarrow \text{HC}(hck_0, (hk_0, x_0)) \ // \ G_0$	$c' \leftarrow_{\$} \mathcal{A}^{\text{FN}}$	$y \leftarrow_{\$} \text{R.Out}$
$r_0 \leftarrow_{\$} \text{HC.Out} \ // \ G_1$	Return ($c' = 1$)	Return y
$c' \leftarrow_{\$} \mathcal{A}^{\text{FN}}$	$\text{FN}((x_1, hk_1, hck_1))$	
Return ($c' = 1$)	$r_1 \leftarrow \text{HC}(hck_1, (hk_1, x_1))$	
$\text{FN}((x_1, hk_1, hck_1))$	$w_0 \leftarrow \text{H}(hk_1, x_0)$	
$r_1 \leftarrow \text{HC}(hck_1, (hk_1, x_1))$	$w_1 \leftarrow \text{H}(hk_0, x_1)$	
$w_0 \leftarrow \text{H}(hk_1, x_0)$	$z_0 \leftarrow (w_0, hk_0, hck_0)$	
$w_1 \leftarrow \text{H}(hk_0, x_1)$	$z_1 \leftarrow (w_1, hk_1, hck_1)$	
$z_0 \leftarrow (w_0, hk_0, hck_0)$	$y_0 \leftarrow_{\$} \text{R.Out}$	
$z_1 \leftarrow (w_1, hk_1, hck_1)$	If ($R[z_1] \neq \perp$) then	
$y_0 \leftarrow \text{R}(r_0, z_1)$	$\text{bad} \leftarrow \text{true}; \boxed{y_0 \leftarrow R[z_1]}$	
$y_1 \leftarrow \text{R}(r_1, z_0)$	$R[z_1] \leftarrow y_0$	
$y \leftarrow y_0 * y_1$	$y_1 \leftarrow \text{R}(r_1, z_0)$	
Return y	$y \leftarrow y_0 * y_1$	

Figure 4: **Games for proof of Theorem 5.2.**

Proof of Proposition 5.1: The first condition, that the keyspace and input space of S are the same set, is met by definition. For a_0, a_1 in this common set we now need to show that $\text{S}(a_0, a_1) = \text{S}(a_1, a_0)$. This follows from the symmetry in the description of S and the assumption that the group R.Out is commutative. ■

PRF SECURITY OF SPRF. To show S is a dual PRF, it suffices by Proposition 5.1 to show that S is a PRF. This is the claim of the following theorem.

Theorem 5.2 *Let $(\text{H}, \text{HC}, \text{R})$ be a suite of function families. Let $\text{S} = \text{SPRF}[\text{H}, \text{HC}, \text{R}]$ be the (symmetric) function family associated to them as above. Let \mathcal{A} be an adversary making at most q queries to its FN oracle. Then the proof constructs adversaries $\mathcal{A}_{\text{H}}, \mathcal{A}_{\text{HC}}, \mathcal{A}_{\text{R}}$ such that*

$$\text{Adv}_{\text{S}}^{\text{prf}}(\mathcal{A}) \leq \text{Adv}_{\text{H}, \text{HC}}^{\text{hlc}}(\mathcal{A}_{\text{HC}}) + \text{Adv}_{\text{R}}^{\text{prf}}(\mathcal{A}_{\text{R}}) + \frac{q(q-1)}{2} \cdot \text{Adv}_{\text{H}}^{\text{cau}}(\mathcal{A}_{\text{H}}). \quad (3)$$

The running times of the constructed adversaries are about the same as that of the original.

Proof of Theorem 5.2: Consider games G_0 – G_4 of Fig. 4. In the code for games G_0, G_1 , if a line is followed by the name of a game, then that line is included *only* in the named game. Unmarked lines are included in both games. Game G_2 includes the boxed code while game G_3 does not.

We assume wlog that the oracle queries of \mathcal{A} are always all distinct. This means the “If $T[x] = \perp$ ” test in game $\mathbf{G}_{\text{S}}^{\text{prf}}(\mathcal{A})$ of Fig. 1 will always return `true` and so we can drop it. The $c = 1$ case of $\mathbf{G}_{\text{S}}^{\text{prf}}(\mathcal{A})$ is thus captured by game G_0 . On the other hand, game G_4 captures the $c = 0$ case of game $\mathbf{G}_{\text{S}}^{\text{prf}}(\mathcal{A})$ except that it returns `true` iff the latter returns `false`. From Equation (1) we thus have

$$\text{Adv}_{\text{S}}^{\text{prf}}(\mathcal{A}) = \Pr[\mathbf{G}_{\text{S}}^{\text{prf}}(\mathcal{A}) | c = 1] - \left(1 - \Pr[\mathbf{G}_{\text{S}}^{\text{prf}}(\mathcal{A}) | c = 0]\right)$$

$$\begin{aligned}
&= \Pr[G_0] - \Pr[G_4] \\
&= p_0 + p_1 + p_2 + p_3,
\end{aligned} \tag{4}$$

where for $i \in \{0, 1, 2, 3\}$ we have let

$$p_i = \Pr[G_i] - \Pr[G_{i+1}].$$

We will build adversaries $\mathcal{A}_H, \mathcal{A}_{HC}, \mathcal{A}_R$ such that

$$p_0 \leq \mathbf{Adv}_{H,HC}^{\text{lhC}}(\mathcal{A}_{HC}) \tag{5}$$

$$p_1 \leq \mathbf{Adv}_R^{\text{prf}}(\mathcal{A}_R) \tag{6}$$

$$p_2 \leq \frac{q(q-1)}{2} \cdot \mathbf{Adv}_H^{\text{cau}}(\mathcal{A}_H). \tag{7}$$

We will also observe that

$$p_3 = 0. \tag{8}$$

Putting together Equations (4), (5), (6), (7) and (8) we get Equation (3). We now justify the above claims.

In game G_1 , the key r_0 for the first application of R is chosen at random rather than obtained as $HC(hck_0, (hk_0, x_0))$. Consider adversary \mathcal{A}_{HC} shown in Fig. 5. It is playing game $\mathbf{G}_{H,HC}^{\text{lhC}}(\mathcal{A}_{HC})$, so it has input hk_0, hck_0, s . It runs \mathcal{A} , simulating the latter's FN oracle via a procedure FNSIM that is shown in the code. The key point is that \mathcal{A}_{HC} invokes its LK oracle to compute w_0 . Letting c be the challenge bit in game $\mathbf{G}_{H,HC}^{\text{lhC}}(\mathcal{A}_{HC})$ we have

$$\begin{aligned}
&\mathbf{Adv}_{H,HC}^{\text{lhC}}(\mathcal{A}_{HC}) \\
&= \Pr[\mathbf{Adv}_{H,HC}^{\text{lhC}}(\mathcal{A}_{HC}) | c = 1] - \left(1 - \Pr[\mathbf{Adv}_{H,HC}^{\text{lhC}}(\mathcal{A}_{HC}) | c = 0]\right) \\
&= \Pr[G_0] - \Pr[G_1] = p_0
\end{aligned}$$

which establishes Equation (5).

Game G_2 maintains a table $R[\cdot]$ that is initially everywhere \perp . It optimistically picks y_0 at random and sets $R[z_1]$ to this value. However, in between these two steps, it first checks whether $R[z_1]$ was already defined, and if so, sets the flag `bad` to `true`. This means that the setting of $R[z_1]$ to the newly-chosen y_0 was wrong. Accordingly (via the boxed code which is included in game G_2) a correction is made, resetting y_0 back to $R[z_1]$, so that in this game, $R[z_1]$ is the result of a random function on z_1 . Now consider adversary \mathcal{A}_R shown in Fig. 5. It has an FN oracle, and runs \mathcal{A} . In the simulation of \mathcal{A} 's oracle, it applies FN to z_1 to get y_0 . With c the challenge bit in game $\mathbf{G}_R^{\text{prf}}(\mathcal{A}_R)$, we have

$$\begin{aligned}
\mathbf{Adv}_R^{\text{prf}}(\mathcal{A}_R) &= \Pr[\mathbf{G}_R^{\text{prf}}(\mathcal{A}_R) | c = 1] - \left(1 - \Pr[\mathbf{G}_R^{\text{prf}}(\mathcal{A}_R) | c = 0]\right) \\
&= \Pr[G_1] - \Pr[G_2] = p_1
\end{aligned}$$

which establishes Equation (6).

In game G_3 , we may set `bad`, but, since the boxed code is absent, y_0 is always a fresh, random value. Games G_2, G_3 are identical until `bad` (differ only in code following the setting of `bad` to `true`) so by the Fundamental Lemma of Game Playing [12],

$$p_2 = \Pr[G_2] - \Pr[G_3] \leq \Pr[G_3 \text{ sets bad}]. \tag{9}$$

Adversary $\mathcal{A}_{\text{HC}}^{\text{LK}}(hk_0, hck_0, s)$	Adversary $\mathcal{A}_{\text{R}}^{\text{FN}}$	Adversary \mathcal{A}_{H}
$r_0 \leftarrow s$	$hk_0 \leftarrow_s \text{H.Keys}$	$i \leftarrow 0$
$c' \leftarrow_s \mathcal{A}^{\text{FNSIM}}$	$hck_0 \leftarrow_s \text{HC.Keys}$	$c' \leftarrow_s \mathcal{A}^{\text{FNSIM}}$
Return c'	$x_0 \leftarrow_s \text{H.Inp}$	$u_1, u_2 \leftarrow_s \text{H.Inp}$
$\text{FNSIM}((x_1, hk_1, hck_1))$	$c' \leftarrow_s \mathcal{A}^{\text{FNSIM}}$	If $(i \leq 1)$ then
$r_1 \leftarrow \text{HC}(hck_1, (hk_1, x_1))$	Return c'	Return (u_1, u_2)
$w_0 \leftarrow \text{LK}(hk_1)$	$\text{FNSIM}((x_1, hk_1, hck_1))$	$j_1 \leftarrow_s \{2, \dots, i\}$
$w_1 \leftarrow \text{H}(hk_0, x_1)$	$r_1 \leftarrow \text{HC}(hck_1, (hk_1, x_1))$	$j_2 \leftarrow_s \{1, \dots, j_1 - 1\}$
$z_0 \leftarrow (w_0, hk_0, hck_0)$	$w_0 \leftarrow \text{H}(hk_1, x_0)$	Return (v_{j_1}, v_{j_2})
$z_1 \leftarrow (w_1, hk_1, hck_1)$	$w_1 \leftarrow \text{H}(hk_0, x_1)$	$\text{FNSIM}((x_1, hk_1, hck_1))$
$y_0 \leftarrow \text{R}(r_0, z_1)$	$z_0 \leftarrow (w_0, hk_0, hck_0)$	$i \leftarrow i + 1$
$y_1 \leftarrow \text{R}(r_1, z_0)$	$z_1 \leftarrow (w_1, hk_1, hck_1)$	$v_i \leftarrow x_1$
$y \leftarrow y_0 * y_1$	$y_0 \leftarrow \text{FN}(z_1)$	$y \leftarrow_s \text{R.Out}$
Return y	$y_1 \leftarrow \text{R}(r_1, z_0)$	Return y
	$y \leftarrow y_0 * y_1$	
	Return y	

Figure 5: Adversaries for proof of Theorem 5.2.

We now design \mathcal{A}_{H} so that

$$\Pr[\text{G}_3 \text{ sets bad}] \leq \frac{q(q-1)}{2} \cdot \text{Adv}_{\text{H}}^{\text{cau}}(\mathcal{A}_{\text{H}}). \quad (10)$$

Adversary \mathcal{A}_{H} is shown in Fig. 5. The integer i is the number of FN queries made by \mathcal{A} , and we consider two cases. The first is if $i \leq 1$. Then the probability that **bad** is set in G_3 is zero, so Equation (10) is true no matter what \mathcal{A}_{H} returns. So, as a default, we just have \mathcal{A}_{H} return a pair (u_1, u_2) of random inputs. Now assume $i \geq 2$. This permits the choices of j_1, j_2 as shown. Now we note that for game G_3 to set **bad**, a z_1 value must repeat across queries. By assumption the queries are distinct, so the only way this could happen is if there were queries $j_1 < j_2$ such that the w_1, hk_1, hck_1 values in these queries were the same but the x_1 values were different. This would be a collision for $\text{H}(hk_0, \cdot)$. Now we have to argue that such a collision can be found by a CAU-adversary \mathcal{A}_{H} . This adversary does not know hk_0 , so how can it simulate \mathcal{A} ? In game G_3 , the point y_0 is always random. Since R.Out is a group, y is also random. So \mathcal{A}_{H} can simulate \mathcal{A} 's oracle by just returning random values. It does this, collecting all the x_1 values in the queries. In the end it picks at random two of these values and returns them. This justifies Equation (10), which, combined with Equation (9), justifies Equation (7).

As we have just said, in game G_3 , the point y_0 is always random and independent of anything else. Since R.Out is a group, y is also random. This justifies Equation (8) and completes the proof. ■

6 Instantiations

We instantiate our SPRF construction to get symmetric and dual PRFs under specific assumptions.

6.1 Construction from (keyless) CR hash functions

We give a construction from any keyless collision-resistant hash function. It itself will play the role of H . The following lemma says that for suitable choices of parameters, an extractor (see Section 2 for background) will provide a leakage hardcore function.

Lemma 6.1 *Let $H : \{\varepsilon\} \times \{0, 1\}^n \rightarrow \{0, 1\}^r$ be a keyless function family. Let $\text{Ext} : \{0, 1\}^s \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a function family that is universal. Let $\text{HC} : \{0, 1\}^s \times (\{\varepsilon\} \times \{0, 1\}^n) \rightarrow \{0, 1\}^m$ be defined by $\text{HC}(hck, (\varepsilon, x)) = \text{Ext}(hck, x)$. Let \mathcal{A} be a LHC-adversary. Then*

$$\text{Adv}_{H, \text{HC}}^{\text{lhc}}(\mathcal{A}) \leq 2^{-(n+2-m-r)/2}. \quad (11)$$

The result is information-theoretic, meaning it is true regardless of the running time of \mathcal{A} .

Proof of Lemma 6.1: Let random variable X be uniformly distributed over $\{0, 1\}^n$. Let U_s, U_m be random variables distributed uniformly over $\{0, 1\}^s$ and $\{0, 1\}^m$, respectively, and let $Y = H(\varepsilon, X)$. The following chain of inequalities, which establishes the lemma, is justified below:

$$\text{Adv}_{H, \text{HC}}^{\text{lhc}}(\mathcal{A}) \leq \text{SD}((U_s, \text{Ext}(U_s, X), Y), (U_s, U_m, Y)) \quad (12)$$

$$\leq \frac{1}{2} \sqrt{2^{m-\mathbf{H}_\infty(X|Y)}} \quad (13)$$

$$\leq 2^{-(n+2-m-r)/2}. \quad (14)$$

Let X and U_s represent, respectively, the randomly chosen x_0 and hck in game $\mathbf{G}_{H, \text{HC}}^{\text{lhc}}(\mathcal{A})$ of Fig. 2. Then $\text{Ext}(U_s, X)$ represents s_1 while U_m represents s_0 . Since H is keyless, the only information \mathcal{A} can get from its LK oracle is $Y = H(\varepsilon, X)$. The statistical distance of Equation (12) then represents the maximum possible advantage that \mathcal{A} can obtain. The three random variables $(X, Y), U_s, U_m$ are independent so we can apply Lemma 2.1 to get Equation (13). Since $|Y| = r$ we have $\mathbf{H}_\infty(X|Y) \geq n - r$, which, together with some simplification, yields Equation (14). ■

Our symmetric and dual PRF $\mathbf{S}_{m,r}$ is parameterized by integers m, r . Given these, we proceed as follows:

- We select n so that $2^{-(n+2-m-r)/2}$ is negligible. Specifically, set $n = 3(m + r)$, so that $2^{-(n+2-m-r)/2} = 2^{-(m+r+1)}$.
- Then we select a function family $\text{Ext} : \{0, 1\}^s \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ that is universal.
- Next we select a keyless, collision-resistant function family $H : \{\varepsilon\} \times \{0, 1\}^n \rightarrow \{0, 1\}^r$. Since it is collision resistant, it is certainly CAU.
- We let $\text{HC} : \{0, 1\}^s \times (\{\varepsilon\} \times \{0, 1\}^n) \rightarrow \{0, 1\}^m$ be defined as in Lemma 6.1 based on H, Ext , namely $\text{HC}(hck, (\varepsilon, x)) = \text{Ext}(hck, x)$.
- Finally we select a PRF $R : \{0, 1\}^m \times \text{R.Inp} \rightarrow \text{R.Out}$ such that $\{0, 1\}^r \times \{\varepsilon\} \times \{0, 1\}^s \subseteq \text{R.Inp}$, and also R.Out is a commutative group, for simplicity $\{0, 1\}^l$ for some l with the group operation being bitwise xor. As we explain below, this can ultimately be built from a CR hash function, making the latter the only assumption.

We now have a suite (H, HC, R) and can apply our **SPRF** transform. The resulting symmetric and dual PRF is $\mathbf{S}_{m,r} : (\{0, 1\}^n \times \{\varepsilon\} \times \{0, 1\}^s) \times (\{0, 1\}^n \times \{\varepsilon\} \times \{0, 1\}^s) \rightarrow \{0, 1\}^l$, defined as follows:

Function family $\mathcal{S}_{m,r}((x_0, \varepsilon, sk_0), (x_1, \varepsilon, sk_1))$

$$\begin{aligned}
r_0 &\leftarrow \text{Ext}(sk_0, x_0) ; r_1 \leftarrow \text{Ext}(sk_1, x_1) \\
w_0 &\leftarrow \text{H}(\varepsilon, x_0) ; w_1 \leftarrow \text{H}(\varepsilon, x_1) \\
z_0 &\leftarrow (w_0, \varepsilon, sk_0) ; z_1 \leftarrow (w_1, \varepsilon, sk_1) \\
y_0 &\leftarrow \text{R}(r_0, z_1) ; y_1 \leftarrow \text{R}(r_1, z_0) \\
y &\leftarrow y_0 \oplus y_1 \\
&\text{Return } y
\end{aligned}$$

The following shows that $\mathcal{S}_{m,r}$ is a PRF. Since it is symmetric, it is thus also a dual PRF.

Theorem 6.2 *Let $m, r \geq 1$ be integers, and select $n, \text{Ext}, \text{H}, \text{HC}, \text{R}$ as above to define the (symmetric) function family $\mathcal{S}_{m,r}$ also as above. Let \mathcal{A} be an adversary. Then the proof constructs adversaries $\mathcal{A}'_{\text{H}}, \mathcal{A}_{\text{R}}$ such that*

$$\text{Adv}_{\mathcal{S}_{m,r}}^{\text{prf}}(\mathcal{A}) \leq 2^{-(m+r+1)} + \text{Adv}_{\text{R}}^{\text{prf}}(\mathcal{A}_{\text{R}}) + \text{Adv}_{\text{H}}^{\text{cr}}(\mathcal{A}'_{\text{H}}). \quad (15)$$

The running times of the constructed adversaries are about the same as that of the original.

The above Theorem assumes that H is CR and R is a PRF. Our ultimate claim is to rely only on the CR assumption. This is possible because (compressing) CR functions imply OWFs, which in turn imply PRGs [27] which in turn imply PRFs [24]. (A direct construction of a PRG from a CR function is also possible [17] but assumes regularity and exponential hardness of the CR function, which we do not want to assume.) We do not give a formal result encompassing the final claim of a dual PRF from just a CR function because, in our concrete-security framework, the statement would need concrete bounds, and we do not know these bounds for the chain of just-mentioned reductions from prior work. Instead we leave this final theoretical result (CR hash functions imply dual PRFs) as understood asymptotically.

Proof of Theorem 6.2: Theorem 5.2 yields adversaries $\mathcal{A}_{\text{H}}, \mathcal{A}_{\text{HC}}, \mathcal{A}_{\text{R}}$ such that

$$\text{Adv}_{\mathcal{S}_{m,r}}^{\text{prf}}(\mathcal{A}) \leq \text{Adv}_{\text{H},\text{HC}}^{\text{lhc}}(\mathcal{A}_{\text{HC}}) + \text{Adv}_{\text{R}}^{\text{prf}}(\mathcal{A}_{\text{R}}) + \frac{q(q-1)}{2} \cdot \text{Adv}_{\text{H}}^{\text{cau}}(\mathcal{A}_{\text{H}}),$$

where q is the number of queries \mathcal{A} makes to its FN oracle. Lemma 6.1 together with the choice of n made above imply that

$$\text{Adv}_{\text{H},\text{HC}}^{\text{lhc}}(\mathcal{A}_{\text{HC}}) \leq 2^{-(n+2-m-r)/2} = 2^{-(m+r+1)},$$

explaining the first term in Equation (15). Now we perform a small optimization. Cau-Adversary \mathcal{A}_{H} in the proof of Theorem 5.2 guessed a colliding pair of inputs for H , but our H is keyless and we assume CR. A CR-adversary \mathcal{A}'_{H} can instead try all candidate pairs and return one (if any) that works. So we can replace $q(q-1)/2 \cdot \text{Adv}_{\text{H}}^{\text{cau}}(\mathcal{A}_{\text{H}})$ by $\text{Adv}_{\text{H}}^{\text{cr}}(\mathcal{A}'_{\text{H}})$. This justifies Equation (15). ■

Remark 6.3 While unkeyed hash functions assumed to be CR are a practical reality (SHA-256 is an example), their formal treatment involves some subtleties. In the asymptotic setting, they cannot exist if we allow non-uniform adversaries. (Such an adversary could hardwire a collision for each choice of the security parameter.) If adversaries are assumed uniform, however, this anomaly goes away, and the assumption of the existence of such a family is meaningful. The concrete setting is inherently non uniform [14] but results (like ours) are still meaningful because they give explicit reductions (adversary constructions). Further elaboration can be found in [31].

6.2 Construction from any OWP

We show that the existence of one-way permutations (OWPs) implies the existence of dual PRFs. We do this by instantiating our SPRF construction using an iterated OWP for H and a leakage hardcore function obtained via the BMY PRG [16, 33].

Let $F : \{\varepsilon\} \times X \rightarrow X$ be a keyless one-way family of permutations with domain and range a set X . (The standard definition of a OWP is indeed keyless.) For $i \geq 1$ let $F^{(i)} : \{\varepsilon\} \times X \rightarrow X$ be the i -th iterate of F , defined inductively by

$$F^{(0)}(\varepsilon, x) = x \quad \text{and} \quad F^{(i)}(\varepsilon, x) = F(\varepsilon, F^{(i-1)}(\varepsilon, x)) \text{ for } i \geq 1.$$

Our symmetric and dual PRF S_m is parameterized by an integer m . Let $R : \{0, 1\}^m \times R.\text{Inp} \rightarrow R.\text{Out}$ be a PRF such that $X \times \{\varepsilon\} \times \{\varepsilon\} \subseteq R.\text{Inp}$, and also $R.\text{Out}$ is a commutative group, for simplicity $\{0, 1\}^l$ for some l with the group operation being bitwise xor. This is not an extra assumption because OWPs imply PRGs [16, 33, 25] which in turn imply PRFs [24]. Let $H = F^{(m)}$ be the m -fold iterate of F . We assume a hardcore predicate $\text{HC}_1 : \{\varepsilon\} \times (\{\varepsilon\} \times H.\text{Inp}) \rightarrow \{0, 1\}$ for F . (Any OWP can be modified to one that has a keyless hardcore predicate, making this assumption wlog.) Let $\text{HC} : \{\varepsilon\} \times (\{\varepsilon\} \times H.\text{Inp}) \rightarrow \{0, 1\}^m$ be defined by

Function family $\text{HC}(\varepsilon, (\varepsilon, x))$
 For $i = 0, \dots, m$ do
 $b_i \leftarrow \text{HC}_1(\varepsilon, (\varepsilon, x)); x \leftarrow F(\varepsilon, x)$
 Return $b_1 b_2 \dots b_m$

Then HC is a hardcore function for $H = F^{(m)}$ assuming only one-wayness of F . Now we have two observations. First, since F , and hence H , is keyless, and we know that HC is a hardcore function for H , Lemma 4.1 implies that it is also a *leakage* hardcore function for H . Second, H is trivially CAU, because it is a permutation family, so there simply do not exist collisions. We can thus apply our **SPRF** transform to the suite (H, HC, R) to get a symmetric function family S_m that, by Theorem 5.2, is a PRF.

The following says that S_m is a PRF. Since it is symmetric, it is also a dual PRF.

Theorem 6.4 *Let $m \geq 1$ be an integer, and select F, H, HC, R as above to define the (symmetric) function family S_m also as above. Let \mathcal{A} be an adversary. Then the proof constructs adversaries $\mathcal{A}_{\text{HC}}, \mathcal{A}_R$ such that*

$$\text{Adv}_{S_m}^{\text{prf}}(\mathcal{A}) \leq \text{Adv}_{H, \text{HC}}^{\text{hc}}(\mathcal{A}_{\text{HC}}) + \text{Adv}_R^{\text{prf}}(\mathcal{A}_R). \quad (16)$$

The running times of the constructed adversaries are about the same as that of the original.

As with Theorem 6.2, we stop short of a formal statement encompassing the final theoretical claim that OWPs alone imply dual PRFs, due to the challenges of casting this in a concrete framework. We have however already discussed above how it is obtained asymptotically. Briefly, OWPs imply PRFs and, if OWPs exist, so do OWPs with keyless hardcore predicates as assumed above.

Proof of Theorem 6.4: Theorem 5.2 yields adversaries $\mathcal{A}_H, \mathcal{A}_{\text{HC}}, \mathcal{A}_R$ such that

$$\text{Adv}_{S_m, r}^{\text{prf}}(\mathcal{A}) \leq \text{Adv}_{H, \text{HC}}^{\text{hc}}(\mathcal{A}_{\text{HC}}) + \text{Adv}_R^{\text{prf}}(\mathcal{A}_R) + \frac{q(q-1)}{2} \cdot \text{Adv}_H^{\text{cau}}(\mathcal{A}_H),$$

where q is the number of queries \mathcal{A} makes to its FN oracle. However $\text{Adv}_H^{\text{cau}}(\mathcal{A}_H) = 0$ since H is a permutation, so this term disappears. Also since H is keyless, the hc-advantage is the same as the hc-advantage. This justifies Equation (16). ■

6.3 Ending remarks

A construction of a dual PRF from any OWF eludes us, and we see this as an interesting open question. Since PRFs are known to exist given a OWF [27, 24], Theorem 5.2 reduces the task of building a dual PRF from a OWF to the task of building, from a OWF, a CAU function family H with a leakage hardcore function HC with long-enough output. However at present we do not know a way to do this.

One may ask what is the conclusion for HMAC. As discussed in Section 1, our intent was to give a generic validation of the dual PRF assumption made in various places including on HMAC’s compression function h in [6]. We have successfully done this through constructions of dual PRFs under standard assumptions. We could, in principle, plug one of our dual PRFs in as the choice of h for HMAC. Then the results of [6] combined with ours would imply PRF security of this alternative HMAC, the assumptions being (only) the ones in our results. However, we are not aware of any practical utility, or value, of this alternative HMAC.

Acknowledgments

We thank Stefano Tessaro for helpful comments on a previous draft. We thank the reviewers of the *Journal of Cryptology* for their careful reading and their constructive comments and corrections. We thank Kirthivaasan Puniamurthy for a typo correction.

References

- [1] Y. Angel, B. Dowling, A. Hülsing, P. Schwabe, and F. J. Weber. Post quantum noise. In H. Yin, A. Stavrou, C. Cremers, and E. Shi, editors, *ACM CCS 2022*, pages 97–109. ACM Press, Nov. 2022. 3, 6
- [2] N. Aviram, B. Dowling, I. Komargodski, K. G. Paterson, E. Ronen, and E. Yogev. Practical (post-quantum) key combiners from one-wayness and applications to TLS. *Cryptology ePrint Archive*, Report 2022/065, 2022. <https://eprint.iacr.org/2022/065>. 6
- [3] M. Backendal, M. Bellare, F. Günther, and M. Scarlata. When messages are keys: Is HMAC a dual-PRF? In H. Handschuh and A. Lysyanskaya, editors, *CRYPTO 2023, Part III*, volume 14083 of *LNCS*, pages 661–693. Springer, Heidelberg, Aug. 2023. 3, 6
- [4] A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 719–737. Springer, Heidelberg, Apr. 2012. 4
- [5] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In J. Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, Aug. 2001. 3
- [6] M. Bellare. New proofs for NMAC and HMAC: Security without collision resistance. *Journal of Cryptology*, 28(4):844–878, Oct. 2015. Preliminary version in C. Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 602–619, Springer, Heidelberg, Aug. 2006. 3, 4, 5, 7, 18
- [7] M. Bellare, D. J. Bernstein, and S. Tessaro. Hash-function based PRFs: AMAC and its multi-user security. In M. Fischlin and J.-S. Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 566–595. Springer, Heidelberg, May 2016. 3
- [8] M. Bellare, R. Canetti, and H. Krawczyk. Keying hash functions for message authentication. In N. Kobitz, editor, *CRYPTO’96*, volume 1109 of *LNCS*, pages 1–15. Springer, Heidelberg, Aug. 1996. 3

- [9] M. Bellare and D. Cash. Pseudorandom functions and permutations provably secure against related-key attacks. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 666–684. Springer, Heidelberg, Aug. 2010. 9
- [10] M. Bellare, R. Dowsley, B. Waters, and S. Yilek. Standard security does not imply security against selective-opening. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 645–662. Springer, Heidelberg, Apr. 2012. 3
- [11] M. Bellare and A. Lysyanskaya. Symmetric and dual PRFs from standard assumptions: A generic validation of an HMAC assumption. Cryptology ePrint Archive, Report 2015/1198, 2015. <https://eprint.iacr.org/2015/1198>. 5
- [12] M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006. 6, 13
- [13] D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang. High-speed high-security signatures. In B. Preneel and T. Takagi, editors, *CHES 2011*, volume 6917 of *LNCS*, pages 124–142. Springer, Heidelberg, Sept. / Oct. 2011. 3
- [14] D. J. Bernstein and T. Lange. Non-uniform cracks in the concrete: The power of free precomputation. In K. Sako and P. Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 321–340. Springer, Heidelberg, Dec. 2013. 16
- [15] N. Bindel, J. Brendel, M. Fischlin, B. Goncalves, and D. Stebila. Hybrid key encapsulation mechanisms and authenticated key exchange. In J. Ding and R. Steinwandt, editors, *Post-Quantum Cryptography - 10th International Conference, PQCrypto 2019*, pages 206–226. Springer, Heidelberg, 2019. 3, 6
- [16] M. Blum and S. Micali. How to generate cryptographically strong sequences of pseudorandom bits. *SIAM Journal on Computing*, 13(4):850–864, 1984. 5, 17
- [17] A. Boldyreva and V. Kumar. A new pseudorandom generator from collision-resistant hash functions. In O. Dunkelman, editor, *CT-RSA 2012*, volume 7178 of *LNCS*, pages 187–202. Springer, Heidelberg, Feb. / Mar. 2012. 16
- [18] C. Brzuska, E. Cornelissen, and K. Kohbrok. Security analysis of the MLS key derivation. In *2022 IEEE Symposium on Security and Privacy*, pages 2535–2553. IEEE Computer Society Press, May 2022. 3, 6
- [19] C. Brzuska, A. Delignat-Lavaud, C. Egger, C. Fournet, K. Kohbrok, and M. Kohlweiss. Key-schedule security for the TLS 1.3 standard. In S. Agrawal and D. Lin, editors, *ASIACRYPT 2022, Part I*, volume 13791 of *LNCS*, pages 621–650. Springer, Heidelberg, Dec. 2022. 3, 6
- [20] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. Cryptology ePrint Archive, Report 2003/235, 2003. <https://eprint.iacr.org/2003/235>. 8
- [21] Y. Dodis and A. Yampolskiy. A verifiable random function with short proofs and keys. In S. Vaudenay, editor, *PKC 2005*, volume 3386 of *LNCS*, pages 416–431. Springer, Heidelberg, Jan. 2005. 4, 9
- [22] B. Dowling, M. Fischlin, F. Günther, and D. Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol. *Journal of Cryptology*, 34(4):37, Oct. 2021. 3, 6
- [23] P. Gaži, K. Pietrzak, and M. Rybár. The exact PRF-security of NMAC and HMAC. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 113–130. Springer, Heidelberg, Aug. 2014. 3
- [24] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, Oct. 1986. 3, 4, 8, 16, 17, 18
- [25] O. Goldreich and L. A. Levin. A hard-core predicate for all one-way functions. In *21st ACM STOC*, pages 25–32. ACM Press, May 1989. 17

- [26] S. Hada. Zero-knowledge and code obfuscation. In T. Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 443–457. Springer, Heidelberg, Dec. 2000. 3
- [27] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. 4, 8, 16, 18
- [28] A. Hülsing, K.-C. Ning, P. Schwabe, F. J. Weber, and P. R. Zimmermann. Post-quantum WireGuard. In *2021 IEEE Symposium on Security and Privacy*, pages 304–321. IEEE Computer Society Press, May 2021. 3, 6
- [29] A. B. Lewko and B. Waters. Efficient pseudorandom functions from the decisional linear assumption and weaker variants. In E. Al-Shaer, S. Jha, and A. D. Keromytis, editors, *ACM CCS 2009*, pages 112–120. ACM Press, Nov. 2009. 4
- [30] M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. *Journal of the ACM*, 51(2):231–262, Mar. 2004. 4, 8, 11
- [31] P. Rogaway. Formalizing human ignorance. In P. Q. Nguyen, editor, *Progress in Cryptology - VI-ETCRYPT 06*, volume 4341 of *LNCS*, pages 211–228. Springer, Heidelberg, Sept. 2006. 16
- [32] D. Stebila, S. Fluhrer, and S. Gueron. Hybrid key exchange in TLS 1.3 – draft-ietf-tls-hybrid-design-05. <https://datatracker.ietf.org/doc/html/draft-ietf-tls-hybrid-design-05>, Aug. 2022. 3, 6
- [33] A. C.-C. Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd FOCS*, pages 80–91. IEEE Computer Society Press, Nov. 1982. 5, 17