Ef£cient Construction of (Distributed) Veri£able Random Functions

Yevgeniy Dodis*

October 16, 2002

Abstract

We give the £rst simple and ef£cient construction of *veri£able random functions* (VRFs). VRFs, introduced by Micali et al. [MRV99], combine the properties of regular pseudorandom functions (PRFs) [GGM86] (i.e., indistinguishability from a random function) and digital signatures [GMR88] (i.e., one can provide an unforgeable proof that the VRF value is correctly computed). The ef£ciency of our VRF construction is only slightly worse than that of a regular PRF construction of Naor and Reingold [NR97]. In contrast to ours, the previous VRF constructions [MRV99, Lys02] all involved an expensive generic transformation from veri£able unpredictable functions (VUFs), while our construction is simple and direct.

We also provide the £rst construction of *distributed* VRFs. Our construction is more ef£cient than the only known construction of distributed (non-veri£able) PRFs [Nie02], but has more applications than the latter. For example, it can be used to distributively implement the random oracle model in a *publicly veri£able* manner, which by itself has many applications (e.g., constructing threshold signature schemes).

Our main construction is based on a new variant of decisional Dif£e-Hellman (DDH) assumption on certain groups where the regular DDH assumption does *not* hold. We do not make any claims about the validity of our assumption (which we call *sum-free* DDH, or sf-DDH). However, this assumption seems to be plausible based on our *current* understanding of certain candidate elliptic and hyper-elliptic groups which were recently proposed for use in cryptography [JN01, Jou00]. We hope that the demonstrated power of our sf-DDH assumption will serve as a motivation for its closer study.

*Department of Computer Science, New York University, 251 Mercer Street, New York, NY 10012, USA. Email: dodis@cs.nyu.edu

1 Introduction

As a motivating example for our discussion, consider the problem of implementing the *random oracle model* [BR93]. Recall that in this model one assumes the existence of a publicly verifiable random function $\mathcal{O}$ (over some suitable domain and range, e.g. $\{0,1\}^n$). Each value $\mathcal{O}(x)$ is random and independent from the other values, and evaluating $\mathcal{O}$ on the same input twice yields the same (random) output. This model has found numerous applications in cryptography, which we do not even attempt to enumerate (for few examples, see [BR93, BR94, BR96, FS86, GQ88, Sch91, Oka92, Mic94, PS96, BF01]). It was shown by Canetti et al. [CGH98], though, that no fixed public function can generically replace the random oracle, so more elaborate solutions are needed.

PSEUDORANDOM FUNCTIONS. As the first attempt, we may assume the existence of a trusted (but computationally bounded) party T. Since a function is an exponential sized object, T cannot store it explicitly. While maintaining a dynamically growing look-up table is a possibility, it is very inefficient as it requires large storage and growing complexity. A slightly better option is to use a *pseudorandom function* (PRF) $F_{SK}(\cdot)$ [GGM86]. As indicated, this function is fully specified by a short secret key (or *seed*) SK, and yet, using F_{SK} (for randomly generated SK) is *computationally* indistinguishable from using exponential-sized $\mathcal{O}$. Put differently, F_{SK} is computationally indistinguishable from a truly random function to any polynomial time adversary *who does not know the secret key* SK.

Of course, PRFs have found numerous more practical applications (e.g., see [NR97] and the references therein), primarily in the area of symmetric-key cryptography (i.e., when the value SK can be shared between mutually trusted parties). For example, they gives very simple constructions of symmetric-key encryption and message authentication codes. In terms of constructing PRFs, there are several options. In principle, PRFs can be constructed from one-way functions [GGM86, HILL99], but this is quite inefficient. Another alternative is to assume that one already has a PRF of small or fixed size (e.g., a block cipher), and show how to extend its domain and range to get a fully functional PRF. For a simple example, if $H : \{0,1\}^* \to \{0,1\}^\ell$ is a collision-resistant hash function [Dam87] and $F_{SK} : \{0,1\}^\ell \to R$ is our fixed-sized PRF, then $F_{SK} \circ H : \{0,1\}^* \to R$ is also a PRF (many other constructions are possible too; see [BKR00, BCK96] and the references therein). Of course, we are still left with the question of constructing the needed small-sized PRF.

The last alternative is to construct PRFs from some well studied number-theoretic assumption. The most popular such construction is due to Naor and Reingold [NR97] and is based on the decisional Diffie-Hellman (DDH) assumption (for related construction based on factoring, see [NRR00]). This assumption in some group $\mathbb{G}$ of prime order q states that given elements g, g^a and g^b of (where g is the generator of $\mathbb{G}$), it is hard to distinguish the value g^{ab} from a truly random value g^c (where a, b, c are random in $\mathbb{Z}_q$). The PRF of [NR97] is a tree-based construction similar to the PRF construction of [GGM86] from a pseudorandom generator. Namely, the secret key $SK = (g, a_1, \ldots a_\ell)$ consists of a random generator g of $\mathbb{G}$ and ℓ random exponents in $\mathbb{Z}_q$ (where ℓ is the length of the input to our PRF $F_{SK} : \{0,1\}^\ell \to \mathbb{G}$). Given $x = x_1 \ldots x_\ell \in \{0,1\}^n$, the PRF is defined by:

$$F_{g,a_1,\ldots,a_\ell}(x_1 \ldots x_\ell) \stackrel{\text{def}}{=} g^{\prod_{\{i|x_i=1\}} a_i \bmod q} \tag{1}$$

VERIFIABLE RANDOM FUNCTIONS. Coming back to our motivating application, replacing random oracle with a PRF has several problems. The first one is the question of verifiability and transferability. Even if everybody trusts T (which we will revisit soon), T has to be contacted not only when the value of F has to be computed for the first time, but even if one party needs to verify that another party has used the correct value of F. Thus, it would be much nicer if each value of $F_{SK}(x)$ would come with a proof $\pi_{SK}(x)$ of correctness, so that the recipient and everybody else can use this proof without the need to contact T again. As a side product, the ability to give such proof will also insure that T himself cannot "cheat" by giving inconsistent values of F, or denying a correctly computed value of the function. This leads to the notion of *verifiable (pseudo)random functions*, or VRFs [MRV99].

Slightly more formally, the key generation outputs a public veri£cation key PK in addition to the secret evaluation key SK, and the function family $\{F_{SK}\}$ has the following properties: (1) Given SK, it is easy to compute the value of the function $y = F_{SK}(x)$ and the corresponding proof $\pi_{SK}(x)$; (2) given PK, x, y, π, one can ef£ciently verify if $y = F_{SK}(x)$ (and only one such value of y can be proven for any x and PK); (3) given only PK and oracle access to both $F_{SK}(\cdot)$ and $\pi_{SK}(\cdot)$, no adversary can distinguish the value $F_{SK}(x)$ from a truly random value without explicitly asking one of the oracles on input x (the last property is sometimes called *residual pseudorandomness*). Put differently, the function remains (pseudo)random when restricted to all inputs whose function values were not previously revealed (and proved).

VRFs already found several applications. For example, using VRFs one can reduce the number of rounds for resettable zero-knowledge proofs to 3 in the bare model [MR01]. As another interesting application, they can be used in a non-interactive lottery system used in micropayments [MR02]. The lottery organizer commits to a VRF by publishing the public key PK. Any participant is allowed to choose his lottery ticket x by himself and send it to the organizer (with the only requirement that x was not used before). The value $y = F_{SK}(x)$ somehow determines whether the user has won the lottery. The organizer sends y to the user together with the proof $\pi_{SK}(x)$, which ensures that the organizer cannot cheat. On the other hand, the unpredictability of y ensures that the participant cannot bias the lottery in his favor. Another set of applications, given by Naor et al. [NPR99] for the case of regular PRFs (or their distributed variants; see below), can be also enhanced by the veri£ability property of VRFs. For example, VRFs could be used to implement a trusted key distribution center (KDC). For a group of users S with "descriptor" x_S (which could be the name of the group or a common password), the value $k_S = F_{SK}(x_S)$ can be used as a common random key used by the members of S. When a party proves his right to get this key (which is done by some application dependent mechanism), KDC would provide this party with k_S together with the proof of its correctness. Another application in similar spirit is that of long-term encryption of information [NPR99]. Finally, the pseudorandomness and veri£ability of a VRF immediately imply that VRF by itself is an unforgeable signature scheme secure against adaptive chosen message attack [GMR88].

CONSTRUCTIONS OF VRFS. Unfortunately, VRFs are not very well studied yet. Currently, we have two constructions of VRFs: based on RSA [MRV99], and based on a separation between computational and decisional Dif£e-Hellman problems in certain groups [Lys02]. Both of these constructions roughly proceed as follows. First, they construct a relatively simple and ef£cient veri£able *unpredictable* function (VUF) based on the corresponding assumption. Roughly, a VUF is the same veri£able object as a VRF, except each "new" value $F_{SK}(x)$ is only unpredictable (i.e., hard to compute) rather than pseudorandom. From VUFs, a generic construction to VRFs is given, as introduced by [MRV99]. Unfortunately, this construction is very ineff£cient and also looses a very large factor in its exact security. Essentially, £rst one uses the Goldreich-Levin theorem [GL89] to construct a VRF with very small (slightly super-logarithmic) input size and output size 1 (and pretty dramatic security loss too).[1] Then one makes enough such computations to amplify the output size to roughly match that of the input. Then one follows another rather ineff£cient tree-based construction on the resulting VRF to get a VRF with arbitrary input size and small output size. Finally, one evaluates the resulting convoluted VRF several times to increase the output size to the desired level. In some sense, the ineff£ciency of the above construction is expected given its generality and the fact that it has to convert pure unpredictability into a much stronger property of pseudorandomness. Still, this means that the resulting VRF constructions are very bulky and inelegant. In this work we present the £rst simple, ef£cient and "direct" VRF construction.

DISTRIBUTED PRFS. Coming back again to our target application of implementing the random oracle, the biggest problem of both PRF/VRF-based solutions is the necessity to fully trust the honest party T holding the secret key for F. Of course, VRFs slightly reduced this trust level, but T still singlehandedly knows all the values of F. Clearly, this approach (1) puts to much trust into T, (2) makes T is bottleneck of all the computations; (3)

[1] For example, one needs to assume a super-polynomial hardness for the given VUF to make sure that the resulting VRF is polynomially secure. Is it an interesting open question to improve this reduction.

makes T is "single point of failure": compromising T will break the security of any application which depends on the random oracle assumption.

The natural solution to this problem is to distribute the role of T among n servers. This leads to the notion of *distributed* PRF*s* (DPRFs) and *distributed* VRF*s* (DVRFs). Since the latter concept was not studied prior to our work, we start with DPRFs, thus ignoring the issue of veri£ability for now. Intuitively, DPRFs with threshold $1 \leq t < n$ allow any $(t+1)$ out of n servers to jointly compute the function using their shares, while no coalition of up to t servers to be in a better situation that any outside party. Namely, the function remains pseudorandom to any such coalition. In the most ambitious form, the computation of DPRF should be *non-interactive* and single-round. The £rst requirement means that the servers do not need to interact with each other in order to help the user compute the value of the function. Instead, the only communication goes between the user and the servers. The second requirement means that the entire computation should proceed in one round: the user gives to (at least) $t+1$ honest servers the needed input x, each server i computes the share y_i of the output $y = F_{SK}(x)$ by using its secret key share SK_i, and £nally the user combines the shares y_i and recovers y.

Not surprisingly, DPRFs have a variety of applications, including distributed KDCs, threshold evaluation of the Cramer-Shoup cryptosystem [CG99], ef£cient metering of the web [NP98], asynchronous Byzantine agreement [Nie02] and several others (see [NPR99, Nie02]). DPRFs £rst originate in the work of Micali and Sidney [MS95]. However, their construction (later improved by [NPR99]) can tolerate only a moderate number of servers or a small threshold, since its complexity is proportional to n^t. The next in¤uential work is that of Naor et al. [NPR99], who give several ef£cient constructions of certain weak variants of DPRFs. Ironically, one of the constructions (namely, that of distributed *weak* PRF) can be turned into an ef£cient DPRF by utilizing random oracles. Even though this is non-trivial (since nobody should compute the value of a DPRF without the cooperation of $t+1$ servers), we would certainly prefer a solution in the plain model, since elimination of the random oracle was one of the main motivation for DPRFs!

The £rst regular DPRF was recently constructed by Nielsen [Nie02] by distributing a slightly modi£ed variant of the Naor-Reingold PRF [NR97], given in Equation (1) (in the £nal version of their work, [NR97] also give essentially the same construction). Unfortunately, the resulting DPRF in highly *interactive* and requires a lot of rounds (proportional to the length of the input). Thus, the question of non-interactive (and, hopefully, round-ef£cient) DPRF construction remained open.

DISTRIBUTED VRFS. Even though DVRFs were not explicitly studied prior to this work, they seem to provide the most satisfactory solution to our original problem of implementing the random oracle. Indeed, distributing the secret key ensures that no coalition of up to t servers can compromise the security (i.e., pseudorandomness) of the resulting random oracle. On the other hand, veri£ability ensures that one does not need to contact the servers again once the random oracle was computed once: the proof can convince any other party of the correctness of the VRF value. For example, DVRFs by themselves provide an ordinary threshold signature scheme, which can be veri£ed without further involvement of the servers. And, of course, using DVRFs is likely to enhance the security, robustness or functionality of many applications originally designed for plain PRFs, VRFs and DPRFs.

OUR CONTRIBUTIONS. We give the £rst simple and direct construction of VRFs, based on a new "DDH-like" assumption which seems to be plausible on certain recently proposed elliptic and hyper-elliptic groups (e.g., [JN01]). We call this assumption *sum-free decisional Dif£e-Hellman* (sf-DDH) assumption. We will comment more on this assumption below. We mention, however, that in the proposed groups the regular regular DDH assumption is *false* (in fact, this is what gives us veri£ability!), and yet the sf-DDH or some similar assumption seems plausible. Our construction is similar to the Naor-Reingold (NR) construction given by Equation (1), except we utilize some carefully chosen encoding C before applying the NR-construction. Speci£cally, if $C : \{0,1\}^\ell \to \{0,1\}^L$ is some injective encoding, we consider the function of the form

$$F_{g,a_1,\ldots,a_L}(x_1 \ldots x_\ell) \stackrel{\text{def}}{=} g^{\prod_{\{i|C(x)_i=1\}} a_i \bmod q} \tag{2}$$

Identifying the properties of the encoding C and constructing C satisfying these properties will be one of the main technical challenges we will have to face. At the end we will achieve $L = O(\ell)$ (speci£cally, $L = 2\ell$ to get a regular PRF, and $L = 3\ell + 2$ to get a VRF), making our ef£ciency very close to the NR-construction. We also mention that our construction is very similar "in syntax" to the VUF construction of Lysyanskaya [Lys02]. In fact, the "only" differences are as follows: (1) we build a VRF while [Lys02] builds a VUF (which is a weaker notion); (2) we use different (seemingly orthogonal to each other) assumptions, even though suggest the same groups were these assumptions hold; (3) we use different encoding functions C. Speci£cally, [Lys02] uses any error-correcting code, but only for the purposes of making a slightly weaker assumption (i.e., identity with appended 1 would yield a VUF under a slightly stronger, but reasonable assumption). On the other hand, we use a very different kind of encoding, because the fact that DDH is easy in our group implies that the identity *cannot* be used, irrespective of what stronger assumption we make. Accidentally, by using a second (now error-correcting) code on top of our encoding C, we can weaken our underling assumption too, similarly to [Lys02].

Our second main contribution is the £rst construction of a distributed VRF (DVRF). Namely, we show that our VRF construction can be made distributed and *non-interactive* (although multi-round). This is the £rst non-interactive construction of a distributed PRF (let alone VRF), since the only previous DPRF construction of [Nie02, NR97] is highly interactive among the servers. In fact, our DVRF construction is more ef£cient than the above mention DPRF construction, despite achieving the extra veri£ability. We already mentioned the big saving in communication complexity (roughly, from $n^2\ell k$ to $n\ell k$, where k is the security parameter), since the servers do not have to interact in our construction. Another important advantage, though, is that we dispense with the need to perform somewhat expensive (concurrently composable) zero knowledge proofs for the equality of discrete logs. This is because in our groups the DDH problem is easy, so it can be locally checked by each party without the need for the proof. In particular, even though we need to apply the encoding C to the message, while the construction of [Nie02, NR97] does not, the lack of ZK-proofs makes our round complexity again slightly better. We also remark that our construction can be easily made *proactive* using standard techniques (see [OY91, HJJ+97]).

Finally, we remark that the same distributed construction can be applied to distribute the VUF of Lysyanskaya [Lys02] (which results in a threshold "unique signature" scheme under a different assumption than the one we propose).

OUR NEW sf-DDH ASSUMPTION. Finally, we elaborate on the sf-DDH assumption, putting it in comparison to the other related assumptions. Recently, Joux and Nguyen [JN01] demonstrated the existence of groups where the DDH assumption is certainly false, but its computational version CDH (i.e., compute g^{ab} from g, g^a, g^b) still seems to be hard. These groups with various ¤avors of the above CDH/DDH separation have already found numerous applications, e.g. [Jou00, BF01, Lys02, BLS01]. Intuitively, the fact that DDH is easy gives many useful properties (like veri£ability), while the hardness of some appropriate CDH-like assumption can be still put to use in security. As already observed by [BF01], our current techniques in such groups only allow us to test DDH relations by means of a certain bilinear mapping (details are not important), for which we do not know a multi-linear variant. In fact, Boneh and Silverberg [BS02] observe that a multi-linear variant of such mapping seems unlikely to exist in the currently proposed groups, and pose as a major open problem to exhibit groups where such mappings exist. This suggests that many natural, but more restrictive ¤avors of DDH seem to hold in the currently proposed groups (where regular DDH is easy). And this is exactly the approach we take. We assume a reasonable "DDH-like" assumption which seems quite possible even when regular DDH is false. Intuitively (see formal de£nition in Section 2.2), it states that given elements of the form $G(I) = g^{\prod_{i \in I} a_i}$ for various subsets $I \subseteq \{1 \dots L\}$, any other element $G(J)$ is pseudorandom unless once can explicitly £nd a "DDH-tuple" $(G(I_1), G(I_2), G(I_3), G(J))$ which would allow to trivially verify $G(J)$. We notice, however, that we do not make a strong claim that this so called sf-DDH assumption is true. Rather, that it seems plausible given what we know. Thus, one can view our work as a strong motivation to study this and related *decisional* assumptions in the above mentioned "gap groups".

2 De£nitions

2.1 Veri£able Random Functions and Friends

In this section we give the de£nition of veri£able random functions (VRFs). For putting our results in perspective with prior works, we also give de£nitions of veri£able unpredictable functions (VUFs) and a new de£nition of regular pseudorandom functions (PRFs). Here and everywhere, $\mathsf{negl}(k)$ refers to some function negligible in the security parameter k.

De£nition 1 *A function family $F_{(\cdot)}(\cdot) : \{0,1\}^{\ell(k)} \to \{0,1\}^{m(k)}$ is a family of* VRF*s, if there exists a probabilistic polynomial time algorithm* Gen *and deterministic algorithms* Prove *and* Verify *such that:* $\mathsf{Gen}(1^k)$ *outputs a pair of keys (PK, SK);* $\mathsf{Prove}_{SK}(x)$ *outputs a pair $\langle F_{SK}(x), \pi_{SK}(x)\rangle$, where $\pi_{SK}(x)$ is the proof of correctness; and* $\mathsf{Verify}_{PK}(x, y, \pi)$ *veri£es that $y = F_{SK}(x)$ using the proof π. More formally, we require the following:*

1. *Uniqueness: no values $(PK, x, y_1, y_2, \pi_1, \pi_2)$ can satisfy $\mathsf{Verify}_{PK}(x, y_1, \pi_1) = \mathsf{Verify}_{PK}(x, y_2, \pi_2)$ when $y_1 \neq y_2$.*
2. *Provability: if $(y, \pi) = \mathsf{Prove}_{SK}(x)$, then $\mathsf{Verify}_{PK}(x, y, \pi) = 1$.*
3. *Pseudorandomness: for any* PPT *$A = (A_1, A_1)$ who did not call its oracle on x (see below)*

$$\Pr\left[\, b = b' \;\middle|\; \begin{array}{c} (PK, SK) \leftarrow \mathsf{Gen}(1^k);\ (x, st) \leftarrow A_1^{\mathsf{Prove}(\cdot)}(PK);\ y_0 = F_{SK}(x); \\ y_1 \leftarrow \{0,1\}^{m(k)};\ b \leftarrow \{0,1\};\ b' \leftarrow A_2^{\mathsf{Prove}(\cdot)}(y_b, st) \end{array} \,\right] \leq \frac{1}{2} + \mathsf{negl}(k)$$

Intuitively, the de£nition states that no value of the function can be distinguished from a random string, even after seeing any other function values together with the corresponding proofs.

De£nition 2 *A function family $F_{(\cdot)}(\cdot) : \{0,1\}^{\ell(k)} \to \{0,1\}^{m(k)}$ is a family of* VUF*s, if satis£es the same syntax, uniqueness and provability properties as the family of* VRF*s, except pseudorandomness is replaced by the following weaker property:*

3'. *Unpredictability: for any* PPT *A_1 who did not call its oracle on x (see below)*

$$\Pr\left[\, y = F_{SK}(x) \;\middle|\; (PK, SK) \leftarrow \mathsf{Gen}(1^k);\ (x, y) \leftarrow A_1^{\mathsf{Prove}(\cdot)}(PK) \,\right] \leq \mathsf{negl}(k)$$

Regular PRFs form the non-veri£able analogs of VRFs. Namely, $PK = \emptyset$, $\pi_{SK}(\cdot) = \emptyset$, there is no algortihm Verify, no uniqueness and provability properties, and pseudorandomness is the only remaining property. Speci£cally, it states

$$\Pr\left[\, b = b' \;\middle|\; \begin{array}{c} SK \leftarrow \mathsf{Gen}(1^k);\ (x, st) \leftarrow A_1^{F(\cdot)}(1^k);\ y_0 = F_{SK}(x); \\ y_1 \leftarrow \{0,1\}^{m(k)};\ b \leftarrow \{0,1\};\ b' \leftarrow A_2^{F(\cdot)}(y_b, st) \end{array} \,\right] \leq \frac{1}{2} + \mathsf{negl}(k)$$

(provided A did not call its oracle on x). We notice that the above de£nition is not the typical de£nition for PRFs as given by [GGM86]: namely, that no adversary can distinguish having oracle access to a truly random function from having oracle access to a pseudorandom function. However, it is easy to see that our de£nition is equivalent to that usual one, so will we use it as the more convenient in the context of VRFs.

2.2 Dif£e-Hellman Assumptions

In what follows, it should be understood that all the objects below will be parametrized by the security parameter k, but we will not explicitly mention this unless needed. Assume $\mathsf{Setup}(1^k)$ outputs the description of some cyclic group $\mathbb{G}$ of prime order q together with its random generator g. Let $L = L(k)$ be some integer and $a_1 \dots a_L$ be random elements of $\mathbb{Z}_q$. Let $[L]$ denote $\{1 \dots L\}$, and given a subset $I \subseteq [L]$, we denote $a_I = \prod_{i \in I} a_i \bmod q$ (where $a_\emptyset = 1$), $G(I) = G_I = g^{a_I}$. Also, DLog_g stands for the discrete logarithm base g (where we often omit g when clear). For example, $\mathsf{DLog}(G_I) = a_I$. Finally, we will often view an element $z \in \{0,1\}^L$ as either a subset $\{i \mid z_i = 1\}$, or an L-dimensional vector over $GF(2)$ (and vice versa).

GENERALIZED DIFFIE-HELLMAN ASSUMPTIONS. The security of ours, as well as the previous related constructions [NR97, Lys02], will rely on various assumptions of the following common ¤avor. The adversary A has oracle access to $G(\cdot)$, and tries to "obtain information" about some value $G(J)$. The meaning of obtaining information depends on whether we are making a computational or a decisional assumption. In the former case, A has to compute $G(J)$, while in the latter case A has to distinguish $G(J)$ from a random element of $\mathbb{G}$. While the decisional assumption is stronger, it has a potential of yielding a (veri£able) *pseudorandom* function, while the computational assumption can yield at best[2] a (veri£able) *unpredictable* function.

In either case, we require that it should be hard to any polynomial time adversary to succeed. Of course, to make non-trivial sense of "success", one has to make some non-trivial restrictions on when the adversary is considered suceessful. Formally, given that the adversary called its oracle on subsets $I_1, \dots, I_t$ and "obtained information" about $G(J)$, we can de£ne a predicate $\mathcal{R}(J, I_1, \dots I_t)$ which indicates whether the adversary's actions are "legal". For example, at the very least the predicate should be false if $J \in \{I_1 \dots I_t\}$. We call any such predicate *non-trivial*. We will certainly restrict ourselves to non-trivial predicates, but may sometimes place some more restrictions on $\mathcal{R}$ in order to make a more plausible and weaker assumption (see below).

De£nition 3 *Given $L = L(k)$, we say that the group $\mathbb{G}$ satis£es the* generalized decisional Dif£e-Hellman *(*gDDH*) assumption of order L* relative to a non-trivial predicate $\mathcal{R}$*, if for any* legal PPT *adversary $A = (A_1, A_1)$ we have*

$$\Pr\left[\, b = b' \;\middle|\; \begin{array}{c} (\mathbb{G}, q, g) \leftarrow \mathsf{Setup}(1^k);\ (a_1 \dots a_L) \leftarrow \mathbb{Z}_q,\ (J, st) \leftarrow A_1^{G(\cdot)}(\mathbb{G}, q); \\ y_0 = G(J);\ y_1 \leftarrow \mathbb{G};\ b \leftarrow \{0,1\};\ b' \leftarrow A_2^{G(\cdot)}(y_b, st) \end{array} \,\right] \le \frac{1}{2} + \mathsf{negl}(k)$$

where A is legal if it called its oracle on subsets $I_1 \dots I_t$ satisfying $\mathcal{R}(J, I_1, \dots, I_t) = 1$.
Very similarly, the group $\mathbb{G}$ satis£es the generalized computational Dif£e-Hellman *(*gCDH*) assumption of order L* relative to $\mathcal{R}$*, if for any legal (see above)* PPT *adversary A_1 we have*

$$\Pr\left[\, G(J) = y \;\middle|\; (\mathbb{G}, q, g) \leftarrow \mathsf{Setup}(1^k);\ (a_1 \dots a_L) \leftarrow \mathbb{Z}_q,\ (J, y) \leftarrow A_1^{G(\cdot)}(\mathbb{G}, q) \,\right] \le \mathsf{negl}(k)$$

We notice that the more restrictions $\mathcal{R}$ places on the I_i's and the "target" set J, the harder it is for the adversary to succeed, so the assumption becomes weaker (and more preferable). Thus, the strongest possible assumption of the above type is to put no further restrictions on $\mathcal{R}$ other than non-triviality (i.e., $J \notin \{I_1, \dots I_t\}$). We call the two resulting assumptions simply gDDH and gCDH (without specifying $\mathcal{R}$). A slightly weaker assumption results when we require that the target set is equal to the full set $J = [L]$, i.e. the adversary has to obtain information about $g^{a_1 \cdots a_L}$. We call the resulting assumptions *full target* $\mathsf{gDDH/gCDH}$ (where $L = 2$ yields regular $\mathsf{DDH/CDH}$). We remark that these full target assumptions are the "standard" way to de£ne generalized (aka group) Dif£e-Hellman assumptions (e.g., in [STW96, BCP01, BCPQ01, Lys02]), but we will £nd our distinction (and, therefore, terminology) more convenient. Finally, making L larger generally makes the assumption stronger (e.g.,

[2] Unless a generic ineff£cient conversion is used, or one assumes the existence of a random oracle, in which case applying the random oracle to a computationally hard object trivially gives a pseudorandom object.

for unrestricted or full target gCDH/gCDH), since the adversary can always choose to concentrate on some subset of L. Thus, it is preferable to base the security of some contsruction on as small L and as restrictive $\mathcal{R}$ as possible.

Before moving to our new sum-free gDDH assumption, let us brie¤y state some simple facts about gDDH/gCDH. It was already observed by [STW96] that gDDH assumption of any polynomial order $L(k)$ (with or without full target) follows from the regular DDH assumption (which corresponds to $L = 2$). Unfortunately, we do not know of the same result for the gCDH problem. The best analog of this result was implicitly obtained by [Lys02], who more or less showed that the regular gCDH assumption of logarithmic order $O(\log k)$ (even with full target) implies the gCDH assumption of any polynomial order $L(k)$, *provided* in the latter we restrict the adversary to operate on the codewords of any good error-correcting code (i.e., $J, I_1 \dots I_t \subseteq [L]$ must be all "far" from each other in order to satisfy $\mathcal{R}$).

SUM-FREE gDDH. We already saw that the regular DDH assumption is a very strong security assumption in that it implies the gDDH assumption. This useful fact almost immediately implies, for example, that the Naor-Reingold construction in Equation (1) is a PRF under DDH, illustrating the power of DDH for proving pseudorandomness. Unfortunately, groups were DDH is true are not convenient for making *veri£able* random functions, since nobody can verify the equality of discrete logs. On the other hand, we will see shortly that it is very easy to obtain veri£ability in groups where DDH is solvable in polynomial time (such as the group suggested by [JN01]). Unfortunately, such groups certainly do not satisfy the gDDH assumption too, which seems to imply that we have to settle for the computational assumption (like gCDH) in such groups, which in turn implies that we settle only for the VUF construction rather than the desired VRF. Indeed, obtaining such a VUF is exactly what was recently done by Lysyanskaya [Lys02] in groups where DDH is easy but gCDH is hard.

However, we make the curcial observation that the fact that DDH is easy does *not* mean that no version of gDDH assumption can be true: it only means *we might have to put more restrictions on the predicate* $\mathcal{R}$ in order to make it hard for the adversary to break the gDDH assumption relative to $\mathcal{R}$. Indeed, for the current elliptic groups for which we believe in a separation between DDH and CDH, we only know how to test if (h, u, v, w) is of the form $u = h^a, v = h^b, w = h^{ab}$ (this is called a DDH-tuple). For example, as was mentioned by Boneh and Franklin [BF01], it seems reasonable to assume that it is hard to distinguidh a tuple $(h, h^a, h^b, h^c, h^{abc})$ from a random tuple (h, h^a, h^b, h^c, h^d). Put differently, when $a_1 \dots a_L$ are chosen at random and given a sample $g = G(\emptyset), G(I_1) \dots G(I_t)$, the only way we know how to distinguish $G(J)$ from a random element of such groups is by exhibiting three sets I_m, I_p, I_s (where $0 \le m, p, s \le t$, and I_0 denotes the empty set) such that $a_J \cdot a_{I_m} = a_{I_p} \cdot a_{I_s} \bmod q$.[3] The last equation implies that "$J + I_m = I_p + I_s$", where we view the sets as L-bit $0/1$-vectors, and the addition is bitwise over the integers. In other words, one has to explicitly £nd a DDH-tuple among the samples $G(I_i)$'s and the target $G(J)$.

We formalize this intuition into the following predicate $\mathcal{R}(J, I_1, \dots, I_t)$. Let us denote $I_0 = \emptyset$. We say that J is DDH-*dependent* on $I_1 \dots I_t$ if there are indices $0 \le m, p, s \le t$ satisfying $J + I_m = I_p + I_s$ (see explanation above). For example, 10101 is DDH-dependent on $01010, 00001$ and 11111, since $10101 + 01011 = 11111 + 00001 = 11112$. Then we de£ne the DDH-*free* relation $\mathcal{R}$ to be true if and only if J is DDH-independent from $I_1 \dots I_t$.

De£nition 4 *Given $L = L(k)$, we say that the group $\mathbb{G}$ (where regular* DDH *is easy) satis£es the* sum-free decisional Dif£e-Hellman *(*sf-DDH*) assumption of order L if if satis£es the* gDDH *assumption of order L relative to the* DDH*-free relation $\mathcal{R}$ above. $\mathbb{G}$ satis£es the* full target sf-DDH *assumption if we additionally require $J = [L]$.*

For our purposes we notice that DDH-dependence also implies that $J \oplus I_m = I_p \oplus I_s$, where $\oplus$ indicates the bitwise addition moduo 2 (i.e., we make "$2 = 0$"), or $J \oplus I_m \oplus I_p \oplus I_s = 0$. Let us call J 4-*wise independent* from $I_1 \dots I_t$ if no three sets I_m, I_j, I_s yield $J \oplus I_m \oplus I_p \oplus I_s = 0$. Hence, if we let $\mathcal{R}'(J, I_1, \dots, I_t) = 1$ if and only

[3] One can also try to £nd the additive relations, but since the a_i's are all random, it seems that the only such relations one can £nd would trivially follow from some multiplicative relations.

if J is 4-wise independent from the I_i's, we get that $\mathcal{R}'$ is a stricter relation that our DDH-free $\mathcal{R}$. But this means that gDDH assumption relative to $\mathcal{R}'$ is a *weaker* assumption than sf-DDH, so we call it *weak* sf-DDH. Our actual construction will in fact be based on weak sf-DDH.

To summarize, sf-DDH is the strongest possible assumption which is conceivable in the groups were regular DDH is false. We chose this assumption to get the simplest and most ef£cient VRF construction possible when DDH is false (in fact, we only need weak sf-DDH in our case). However, even if the ambitious sf-DDH assumption we propose turns out to be false in the current groups where DDH is easy — which we currently have no indication of — it seems plausible that some reasonable weaker gDDH assumptions (relative to more restrictive $\mathcal{R}$) might still hold. And our approach seems to be general enough to allow some easy modi£cation to our construction (at slight ef£ciency loss) meet many such weaker gDDH assumptions.

3 Constructions

Assume $\mathbb{G}$ is the group where DDH is easy while some version of sf-DDH holds (we will be more speci£c soon). We consider the natural the type of functions given by Equation (2); in our new notation, $F_{g,a_1,\ldots,a_L}(x_1 \ldots x_\ell) = G(C(x))$,[4] where C is some currently unspeci£ed (but ef£ciently computable) injective mapping from $\{0,1\}^\ell$ to $\{0,1\}^L$. As we will see, the properties of encoding C will be crucial in showing the properties of the resulting function. To emphasize this dependence on C, we will also call the above function $NR_C(\cdot)$ when other parameters are clear form the context.

3.1 Building PRFs

We notice, that the de£nition above already suf£ces to give a candidate for a regular PRF. As a warm-up towards VRFs, we £rst determine the conditions on C and the kind of gDDH assumption we need in order to get a regular PRF.

Lemma 1 *Given encoding $C : \{0,1\}^\ell \to \{0,1\}^L$, assume predicate $\mathcal{R}$ satis£es $\mathcal{R}(C(w), C(x_1), \ldots, C(x_t)) = 1$ for any $w \notin \{x_1, \ldots, x_t\}$. Then $NR_C(\cdot)$ is a* PRF *under the* gDDH *assumption of order L relative to $\mathcal{R}$.*

Proof: The proof follows almost immediately by comparing the deginition of gDDH relative to $\mathcal{R}$ (De£nition 3) and the de£nition of PRF given in Section 2.1. Indeed, the adversary can query $NR_C(\cdot)$ at any points $x_1, \ldots x_t$, which corresponds to querying $G(\cdot)$ on $C(x_1) \ldots C(x_t)$, and has to distinguish $NR_C(w) = G(C(w))$ for some $w \notin \{x_1 \ldots x_t\}$. Since our assumption implies that $\mathcal{R}(C(w), C(x_1), \ldots, C(x_t)) = 1$, this adversary is legal for breaking gDDH (of order L) relative to $\mathcal{R}$, which is a contradiction. □

As an immediate corollary, usual gDDH assumption implies that $NR_C(\cdot)$ is a PRF for any (injective) C, including the identity. This in turn gives the result of [NR97], since we mentioned that regular DDH implies gDDH [STW96].

More interestingly, we will now determine the properties of C which suf£ce to show that NR_C is a PRF under the much weaker sf-DDH assumption (for now, of the same large order L; we will reduce the order later). In the following, view every subset of $[L]$ (or element of $\{0,1\}^L$) as an L-dimensional vector over $GF(2)$. Recall our de£nition of a vector J being 4-wise independent from the collection $I_1 \ldots I_t$. To generalize this notion, we say that the collection of vectors $I_1 \ldots I_t$ is 4-wise independent, if no 4 or fewer vectors are linearly dependent.

Theorem 1 *Assume $C : \{0,1\}^\ell \to \{0,1\}^L$ is such that the collection $\{C(x) \mid x \in \{0,1\}^\ell\}$ is 4-wise independent. Then $NR_C(\cdot)$ is a* PRF *under the weak (and thus regular)* sf-DDH *assumption of order L.*

[4] Notice, we output a (pseudo)random element of $\mathbb{G}$ instead of a (pseudo)random m-bit string. However, standard hashing techniques imply we can extract an almost uniform string of length close to $\log |\mathbb{G}|$ from such an output. See [NR97].

Proof: Obvious from Lemma 1 and the de£nition of weak sf-DDH. □

CONSTRUCTING 4-WISE INDEPENDENT ENCODINGS. To get our PRF under the sf-DDH assumption (i.e., in groups were regular DDH might be false), it suf£cing to construct a 4-wise independent encoding C. Naturally, the goal is to make L as close to ℓ as possible. Such encodings come up quite often in the theory of derandomization (see [ABI86, AS00]), and are closely related to coding theory.[5] In our case, the well known construction is very simple and ef£cient, so we present it in a self-contained manner.

Let us now view any element of $x \in \{0,1\}^\ell$ as an element of the £eld $GF(2^\ell)$, which can be represented as an ℓ-dimensional vector over $GF(2)$. This gives us the same bitwise addition operation $\oplus$, but now we also have a multiplication operation. Then we set $L = 2\ell$ and de£ne $C(x) = (x^3 \| x)$, which is interpreted as follows. We £rst cube x, which gives us another ℓ-dimentional vector x^3, and then we append x to it. Notice, the code C is explicit and extremely ef£cient to evaluate. It is also very easy to see that any 4-wise independent encoding we can come up with must have $L \geq 2\ell$,[6] so our encoding is optimal. Now, assume there are some *non-zero* distinct $x_1, x_2, x_3, x_4 \in GF(2^\ell)$ and constants $\alpha_1, \alpha_2, \alpha_3, \alpha_4 \in \{0,1\}$ such that $\sum_{i=1}^4 \alpha_i C(x_i) = 0$. We will show that $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 0$, which yields 4-wise independence.

Since our bitwise addition is the same as in the £eld, we get $\sum_{i=1}^4 \alpha_i x_i = 0$ and $\sum_{i=1}^4 \alpha_i x_i^3 = 0$ over $GF(2^\ell)$. Next, we square the £rst equation. Since $GF(2^\ell)$ has characteristic 2 and $\alpha_i^2 = \alpha_i$, the only surviving terms are $\alpha_i x_i^2$, which gives us $\sum_{i=1}^4 \alpha_i x_i^2 = 0$. Similarly, raising the £rst equation to the power 4 gives $\sum_{i=1}^4 \alpha_i x_i^4 = 0$. Thus, we have a linear system (with unknowns $\alpha_1, \alpha_2, \alpha_3, \alpha_4$) saying that $\sum_{i=1}^4 \alpha_i x_i^j = 0$ for $j = 1, 2, 3, 4$. The system corresponds to the famous Vandermonde matrix whose determinant is $x_1 x_2 x_3 x_4 \cdot \prod_{i<j}(x_i - x_j) \neq 0$, since all the x_i's are distinct and non-zero. Thus, the only solution to the system is the trivial all-zero solution, completing the proof.

As a small technicality, we get the 4-wise independent encoding $C : \{0,1\}^\ell \setminus \{0^\ell\} \to \{0,1\}^{2\ell}$, i.e. we explude the all-zero vector. This implies that we get the PRF whose input domain excludes the all-zero vector too. This is typically not a problem since we are "loosing" only one out of 2^ℓ points. Of course, one can always increase L by 1 and add a "dedicated" random $a_{L+1} \in \mathbb{Z}_q$ to point 0^ℓ, but this seems to be going through too much trouble for such a small technicality. To summarize,

Theorem 2 *The encoding C above de£nes a* PRF *mapping ℓ bits (except 0^ℓ) to an element of $\mathbb{G}$, which is secure under the (weak)* sf-DDH *assumption of order 2ℓ.*

REDUCING THE ORDER. While Theorem 2 gives a simple PRF construction, it is based on the sf-DDH assumption of high polynomial order $2\ell(k)$. While this assumption is reasonable, we now show how to reduce the order to $O(\log k)$ at only a marginal ef£ciency loss. So let $C : \{0,1\}^\ell \to \{0,1\}^L$ be any 4-wise independet encoding satisfying Theorem 1 (like the one we constructied above). The idea, similar to that of [Lys02], is to use an error-correcting code $E : \{0,1\}^L \to \{0,1\}^N$ on top of our encoding C. However, since we are dealing with linear dependence, we will have to restrict ourselves to *linear* codes (which was not needed in [Lys02]), and the analysis will be slightly more involved. Thus, let E be a linear error correcting code of distance δN (where $\delta > 0$ and $N = O(L)$), and de£ne $\tilde{C} = E \circ C : \{0,1\}^\ell \to \{0,1\}^N$.

Theorem 3 *Assume (weak)* sf-DDH *assumption holds for any order $p = O(\log k)$. Then $NR_{\tilde{C}}(\cdot)$ is a* PRF.

Proof: Assume some adversary $A = (A_1, A_2)$ breaks the pseudorandomness of $NR_{\tilde{C}}(\cdot)$. We constuct an adversary $B = (B_1, B_2)$ which breaks sf-DDH assumption of some order p (to be speci£ed later). Assume B has oracle

[5] In particular, obtaining the 4-wise independent encoding C we need is equivalent to designing a parity check matrix of any linear code of distance 5. Our speci£c code gives such matrix for the famous (and optimal) BCH code of designed distance 5. See [MS77].

[6] Since all pairwise sums $C(x_1) \oplus C(x_2)$ have to be distinct non-zero elements of $\{0,1\}^L$.

access to $F_{g,b_1,\ldots,b_p}(H)$ for any subset $H \subseteq [p]$. B chooses a random subset $I \subset [N]$ of cardinality exactly p, and implicitly sets $a_i = b_i$ for $i \in I$. It also picks on its own random a_i for all $i \notin I$. B (actually B_1) now runs A_1 with these implicit assignment in mind. Speci£cally, when A_1 asks for the value $NR_{\tilde{C}}(x)$, B_1 computes $z = \tilde{C}(x)$, the restriction $z_I \in \{0,1\}^p$ of z to the positions in I, and the restriction $z_{\bar{I}}$ of z to the complement of I. It asks its oracle for the value $y_I = F_{g,b_1,\ldots,b_p}(z_I)$, and returns to A_1 the value $y = (y_I)^{\prod_{i \notin I} a_i}$. When A_1 output the input challenge value x', B_1 outputs the challenge input value $z'_I = \tilde{C}(x')_I$. Next, when B_2 gets back the challenge output value, which we call y'_I (the reason will be clear), B_2 sets $y' = (y'_I)^{\prod_{i \notin I} a_i}$ and passes it as the challenge output value to A_2. Then B_2 simulates oracle queries of A_2 in the same way as B_1 did for A_1. Finally, B_2 outputs the same decision as A_2.

We notice that B simulated A is a completely perfect way. Indeed, when the challenge for B was the correct output value, B translated it to the correct output value for A. Otherwise, B raised a random element to some (wlog, non-zero) power, which left it a radnom group element. Thus, B distinguishes with the same advantage as A modulo the problem that B could be illegal. Namely, assume A asked $t = \mathsf{poly}(k)$ queries $x^1, \ldots, x^t$ altogether. Then B asked t queries $z_I^1, z_I^2, \ldots z_I^t$ and outputted the challenge z'_I. Now it could be that z'_I is 4-wise dependent on $z_I^1, z_I^2, \ldots z_I^t$. If this happens (say with probability α), then we modify the behavior of B to stop simulating A and let B output a random bit. We will show that it suf£ces to set $p = O(\log k)$ in order to make $\alpha \leq 1/2$, which would complete the proof.

And here is the reason. Since B simulates A in a perfect way, A gets no information about I during its run. Thus, A cannot choose its actions based on I, so we may assume that A chose some values $z^1 = \tilde{C}(x^1), \ldots, z^t = \tilde{C}(x^t), z' = \tilde{C}(x')$, where $x' \notin \{x^1, \ldots, x^t\}$, and only then (when the above values are *£xed*) the subset I was actually chosen. Now, α exactly measures the probability that the random "p-projection" z'_I is 4-wise dependent on the other projections $z_I^1, \ldots, z_I^1$. We will apply the union bound. There are at most $T = O(t^3)$ choices for indices s, m, r (including 0 for the empty set) which can cause the dependence of the projections. Thus, to show $\alpha \leq 1/2$ it suf£ce to show that for any three (or less, but this will only be easier) £xed indices s, m, r the probability of linear dependence in the projection is at most $1/2T$. De£ne $\tilde{z} \stackrel{\text{def}}{=} z' \oplus z^s \oplus z^m \oplus z^r = \tilde{C}(x') \oplus \tilde{C}(x^m) \oplus \tilde{C}(x^s) \oplus \tilde{C}(x^r)$ and $v \stackrel{\text{def}}{=} C(x') \oplus C(x^m) \oplus C(x^s) \oplus C(x^r)$. Notice that since our error-correcting code E is linear, we have $\tilde{z} = E(v)$. Also, since C is 4-wise independent, we have that $v \neq 0$. Thus, it remains to estimate the probability that $\tilde{z}_I = 0$, or that $E(v)_I = 0$ when $v \neq 0$. But since $v \neq 0$, $E(v)$ has at least δN entries which are 1. Since I picks p out of N entries of $E(v)$ at random, the probability of picking all 0's (without replacement) is at most $(1-\delta)^p$. It remains to pick p so that $(1-\delta)^p = O(1/t^3)$, and we see that setting $p = O(\log k)$ suf£ces indeed (recall that k is our security parameter, t is polynomial in k and δ is a constant). This completes the proof. □

We remark that since error-correcting code can in principle approach a rate of 1, using Theorem 2 we can get a PRF construction with £nal expansion $N = (2+\varepsilon)\ell$ based of the sf-DDH assumption of order $O(\log k)$.

3.2 Building VRFs

So far we saw how to construct plain PRFs based on sf-DDH assumption. We now show how extend the above techniques to get a VRF. As before the construction is parameterized by some encoding $C : \{0,1\}^\ell \to \{0,1\}^L$.

- $\mathsf{Gen}(1^k)$: runs $(G, q, g) \leftarrow \mathsf{Setup}(1^k)$, picks random $a_1, \ldots, a_{L+1} \in \mathbb{Z}_q$, sets $h = g^{a_{L+1}}$, $y_1 = h^{a_1}$,..., $y_L = h^{a_L}$. Outputs public key $PK = (\mathbb{G}, q, g, h, y_1 = h^{a_1}, \ldots, y_L = h^{a_L})$, secret key $SK = (g, a_1, \ldots, a_L)$.
- $\mathsf{Prove}_{SK}(x)$: outputs $(\sigma_1, \ldots \sigma_L)$, where $\sigma_0 = g$ and $\sigma_j = g^{\prod_{\{i \leq j | C(x)=1\}} a_i}$ for $j = 1 \ldots L$. In particular, the value σ_L is $F_{SK}(x)$, while $(\sigma_1, \ldots, \sigma_{L-1})$ is the proof $\pi_{SK}(x)$.
- $\mathsf{Verify}_{PK}(\sigma_1, \ldots, \sigma_L)$: sets $\sigma_0 = g$ and checks, for every $1 \leq i \leq L$, that $(\sigma_{i-1}, \sigma_i, h, y_i)$ form a DDH-tuple (recall, DDH is easy!) when $C(x) = 1$, ot that $\sigma_{i-1} = \sigma_i$ is $C(x)_i = 0$. Accept if all the tests pass.

To satisfy the de£nition of VRFs (De£nition 1), we need to examine uniqueness, provability and pseudorandomness. The £rst two properties are very easy. Uniqueness follows from the fact that discrete logs are unique in $\mathbb{G}$ (and that our assumed algorithm for DDH will never accept an invalid tuple), while provability is obvious by construction.

Thus, we only need to examine the pseudorandomness property. Luckily, a lot of machinery has been already developed in Section 3.1. Essentially, the main difference we have is that when the adversary asks $\mathsf{Prove}(x)$, not only does he get $F(x) = G(C(x))$, but he also gets "for free" the proof values $G(I)$ for all $I \in \mathsf{Prefixes}(C(x))$, where for a set $J \subseteq [L]$ we de£ne $\mathsf{Prefixes}(J) \stackrel{\text{def}}{=} \{\emptyset, J \cap [1], J \cap [2], \ldots, J \cap [L-1], J\}$. Additionally, the public key gives the adversary the values $G(\{L+1\}), G(\{L+1,1\}), \ldots, G(\{L+1,L\})$. We denote this collection of $L+1$ subsets of $[L+1]$ involving element $L+1$ by $Pub(L+1)$. With these in mind, we easily get the following analog of Lemma 1.

Lemma 2 *Given encoding $C : \{0,1\}^\ell \to \{0,1\}^L$, assume that for any $w \notin \{x_1, \ldots, x_t\}$ the predicate $\mathcal{R}$ satis£es $\mathcal{R}(C(w), \mathsf{Prefixes}(C(x_1)), \ldots, \mathsf{Prefixes}(C(x_t)), Pub(L+1)) = 1$. Then our construction is a* VRF*, under the* gDDH *assumption of order $L+1$ relative to $\mathcal{R}$.*

Next, we can appropriately generalize the notion of 4-wise independence to that of 4-wise *pre£x-independence*. Namely, a vector J is 4-wise pre£x independent from vectors $I_1 \ldots I_t$ if there exist no $1 \leq p, r, s, \leq t$ and $I'_p \in \mathsf{Prefixes}(I_p)$, $I'_r \in \mathsf{Prefixes}(I_r)$, $I'_s \in \mathsf{Prefixes}(I_s)$ such that $J \oplus I'_p \oplus I'_r \oplus I'_s = 0$. A collection $\{I_1 \ldots I_t\}$ is said to be 4-wise pre£x independent if every vector I_i is 4-wise pre£x independent from the remaining vectors. Finally, we will say that the above collection has *pre£x-distance* at least 3, if for any $i \neq j$ and $I'_j \in \mathsf{Prefixes}(I_j)$, we have that I_i and I'_j differ in at least 3 positions when viewed as binary vectors of length L (in particular, every I_i has weight at least 3). Then, we get the following analog of Theorem 1.

Theorem 4 *Assume $C : \{0,1\}^\ell \to \{0,1\}^L$ is such that the collection $\left\{C(x) \mid x \in \{0,1\}^\ell\right\}$ is 4-wise pre£x-independent and has pre£x-distance at least 3. Then our construction is a* VRF *under the weak (and thus regular)* sf-DDH *assumption of order $L+1$.*

Proof: By Lemma 2, we only need to show that no vector $C(w)$ is linearly dependent on 3 (or fewer) vectors z_1, z_2, z_3 inside the sets $\mathsf{Prefixes}(C(x_1)), \ldots, \mathsf{Prefixes}(C(x_t)), Pub(L+1)$. Assuming the contrary, if none of z_1, z_2, z_3 comes from $Pub(L+1)$, we would exactly get that the collection $\left\{C(x) \mid x \in \{0,1\}^\ell\right\}$ is 4-wise pre£x-dependent, which is a contradiction. Otherwise, some z_i's (say, z_1) is one of $\{\{L+1\}, \{L+1,1\}, \ldots, \{L+1,L\}\}$. Since these are the only sets containing element $(L+1)$, in order to "cancel" $(L+1)$ one other z_i (say, z_2) also comes from this collection, which means that $z_1 \oplus z_2$ is some subset of I of $[L]$ or cardinality *at most* 2. The only way we can now have $C(w) \oplus I \oplus z_3 = 0$, is if some z_3 was a pre£x of some $C(x_j)$ (where $x_j \neq w$) which differs from $C(w)$ in at most 2 coordinates. But this is exactly what is ruled out by the fact the collection $\left\{C(x) \mid x \in \{0,1\}^\ell\right\}$ has pre£x-distance at least 3. □

CONSTRUCTING THE ENCODING. It remains again to construct a 4-wise pre£x-independent encoding of pre£x distance at least 3. We do it by giving a simple generic transformation from any regular 4-wise independent encoding $C : \{0,1\}^\ell \to \{0,1\}^L$, such as the encoding $(x^3 \| x)$ considered in the previous section. We will assume without loss of generality that every two distinct elements $C(x)$ and $C(w)$ differ in at least two positions. For example, this is true with the 4-wise independent encoding $(x^3 \| x)$ constructed in the previous section. However, even if originally false in C, one can always increase L by 1 by adding a "parity" bit to C (i.e., the XOR of all the bits of $C(x)$) and get the required distance at least 2 between distinct codewords. Also, for a technical reason we will exclude the zero vector 0^ℓ from the domain of our new encoding.

Lemma 3 *If C is 4-wise independent (and has distance at least 2), then $C'(x) = (C(x) \| 1 \| x \| 1)$ is 4-wise pre£x-independent and has pre£x-distance at least 3.*

Proof: Below we will refer to the two 1's in the de£nition of C' as "middle" and "last". We start with showing the pre£x distance. Take any $x \neq w$ and consider any pre£x I of $C'(w)$. This pre£x either "crosses" both the middle and the last 1, only the middle 1, or none of them. In the £rst case (i.e., we look at $C'(w)$ itself), we get distance three between $C'(x)$ and $C'(w)$ since $C(x)$ differs from $C(w)$ in at least two locations, and x differs from w in at least one more location. In the second case, $C(x)$ still differs from $C(w)$ in at least two locations, and now also I does not have the last 1 which $C'(x)$ has. Finally, in the last case (no 1's are crossed), I does not have both 1's that $C'(x)$ has, and also in between the 1's x is non-zero (this is where we exclude 0^ℓ) while the pre£x I is zero, giving distance at least 3 again.

Next, we show the 4-wise pre£x independence. Take any x, w_1, w_2, w_3 where $x \notin \{w_1, w_2, w_3\}$, and let z_1, z_2, z_3 be some pre£xes of $C'(w_1), C'(w_2), C'(w_2)$ such that $(C(x)\|1\|x\|1) \oplus z_1 \oplus z_2 \oplus z_3 = 0$. Notice, in order to cancel the last 1 of $C'(x)$, at least one of the pre£xes, say z_1 has to be full; i.e., $z_1 = C'(w_1) = C(w_1)\|1\|w_1\|1$. Since the middle 1's cancel out in $C'(x) \oplus C'(w_1)$, we have two possibilities for them to cancel in the full sum $C'(x) \oplus C'(w_1) \oplus z_2 \oplus z_3$. Either both pre£xes z_2 and z_3 cross the middle 1, or none does. In the £rst case, taking the "C-pre£xes" we get that $C(x) \oplus C(w_1) \oplus C(w_2) \oplus C(w_3) = 0$, which contradicts the fact that C is 4-wise independent. In the second case, we get that the identity parts between the 1's yield $x \oplus w_1 = 0$, i.e. $x = w_1$, which is again a contradiction. □

Applying the above Lemma to the 4-wise independent code $C(x) = (x^3\|x)$ used in Theorem 2, we get:

Theorem 5 *The encoding $C'(x) = (x^3\|x\|1\|x\|1)$ de£nes a* VRF *mapping ℓ bits (except 0^ℓ) to an element of $\mathbb{G}$, which is secure under the (weak)* sf-DDH *assumption of order $3\ell + 3$.*

REDUCING THE ORDER. Similarly to Theorem 3, we apply an "outer" error-correcting code to reduce the order of the sf-DDH assumption we need for Theorem 5. However, we need to be sure that our construction preserves pre£x-independence. Here is one direct way of doing it if we start — as in Lemma 3 — from any regular 4-wise independent (but perhaps not pre£x-independent) $C : \{0,1\}^\ell \to \{0,1\}^L$ with minimum distance 2. Let $E_1 : \{0,1\}^L \to \{0,1\}^{N_1}$ and $E_2 : \{0,1\}^\ell \to \{0,1\}^{N_2}$ be two linear error correcting codes, both correcting some constant fraction of errors. We de£ne the £nal encoding $\tilde{C}(x) = (E_1(C(x))\|1\|E_2(x)\|1)$ which maps ℓ non-zero bits to $N_1 + N_2 + 2 = O(\ell)$ bits. By carefully combining the arguments in Theorem 3 with the technique in Lemma 3, we get the following corollary whose proof we omit to avoid repetition.

Theorem 6 *Assume (weak)* sf-DDH *assumption holds for any order $p = O(\log k)$. Then the code $\tilde{C}$ above de£nes a* VRF.

As earlier, we remark that since error-correcting codes can in principle approach a rate of 1, using Theorem 5 we can get a VRF construction with £nal expansion $N = (3+\varepsilon)\ell$ based of the sf-DDH assumption of order $O(\log k)$.

Finally, we remark that with an extra overhead of 2 in the expansion of $\tilde{C}$ (and a large polynomial loss in exact security), we can reduce our PRF and VRF constructions in both Theorem 3 and Theorem 6 to using the *full target* sf-DDH assumption of order $O(\log k)$. Since we have no evidence that full-target sf-DDH is a signi£cantly better assumption than regular sf-DDH, it is not clear if losing these overheads is worthwhile. Thus, we leave the details of this extension to the full version.

4 Distributed VRF

In this section we show that our VRF construction can be easily made distributed, which results in the £rst DVRF construction. Our construction is extremely simple and reminds DPRF construction of Nielsen [Nie02] based on regular DDH. However, the fact that DDH is easy implies we can make our construction non-interactive (i.e.,

servers do not need to know about each other) and more ef£cient than that of Nielsen. We start by presenting our model, and then show our simple construction.

THE MODEL. We assume there are n servers $S_1, \ldots, S_n$ and that we have a regular VRF $V = (\mathsf{Gen}, \mathsf{Prove}, \mathsf{Setup})$ which we want to distribute. First, we de£ne the syntax of the new generation algorithm $\mathsf{Gen}'(\cdot)$ run by the trusted party. $\mathsf{Gen}'(1^k)$ not only outputs the public/secret keys PK and SK for V, but also a pair of public/secret key (PK_i, SK_i) for each server S_i. The global secret key SK is then erased, each server S_i gets SK_i, and the values $(PK, PK_1, \ldots, PK_n)$ are published. When a user U approaches the server S_i with input x, the server determines if the user is quali£ed to learn the value/proof of $F(x)$. How this is done is specifed by the application at hand and is unimportant to us. If U is successful, though, we say that S_i was *initiated* on input x, and U and S_i engage in a possibly interactive protocol. To successfully complete this protocol, the user might have to simultaneously interact with several servers in some possibly prede£ned order (see below), but the servers do not need to interact to each other or know each other's state. Given a threshold t of the systems, the robustness property states that if U contacts s servers on input x, and at least at least $(t+1)$ of these servers are honest (plus, of course, each honest server accepts the user's request), then at the end of the protocol the user learns the unique correct output of $\mathsf{Prove}(x)$; i.e., the value $F(x)$ and the proof $\pi(x)$. This should hold even if the remaining $(s-t-1)$ of the contacted servers are malicious. We notice also that while the user U needs to know the "local' public key PK_i of server i in order to interact with server S_i, any outside party only needs to know the "global" public key PK in order to verify the consistency of $F(x)$ and $\pi(x)$. In other words, the veri£cation algortihm Verify does not have to be changed from the non-distributed setting.

The *security* property of the DVRF protocol states that for any t indices $i_1, \ldots, i_t$ and for any adversary $A = (A_1, A_2)$ who "breaks" the security of DVRF by "corrupting" servers $S_{i_1}, \ldots, S_{i_t}$ (see below), there exists an adversary $B = (B_1, B_2)$ which breaks the pseudorandomness property of our original VRF, as given by De£nition 1. We now de£ne what it means to break the security of DVRF. In addition to the public key $(PK, PK_1, \ldots, PK_n)$, A learns the values $SK_{i_1}, \ldots, SK_{i_t}$ of the corrupted servers. Then, A_1 runs in the £rst stage, in which it is given the ability to interact with any honest servers S_j on arbitrary inputs and in any manner that A_1 desires. However, we keep track of the set of inputs I which were initiated by A_1. At the end of the phase, A_1 outputs the challenge input x (and the state information for A_2). Then A_2 is given back a challenge y_b (for random b), which is either the value $y_0 = F(x)$ or a random element y_1 in the range of F. A_2 can then again interect with honest servers, just like A_1 did. At the end, A_2 outputs the guess $\tilde{b}$ and succeeds if $\tilde{b} = b$ and neither A_1 nor A_2 initiated the input x with any of the servers. A breaks the scheme if it succeeds with non-negligible advantage over $1/2$.

CONSTRUCTION. In Section 3.2 we de£ned a general candidate for VRF parametrized by any encoding C. We now show how to make such construction distributed for any C for which the basic construction is a VRF. The construction is quite simple, but it shows how convenient it is to have veri£ability (given by the easiness of DDH) "for free". Recall that we had $SK = (g, a_1, \ldots, a_L)$; $PK = (\mathbb{G}, q, g, h, y_1 = h^{a_1}, \ldots, y_L = h^{a_L})$; and $\mathsf{Prove}_{SK}(x) = (\sigma_1, \ldots \sigma_L)$, where $\sigma_0 = g$, $\sigma_j = \sigma_{j-1}^{a_j}$ if $C(x)_j = 1$ and $\sigma_j = \sigma_{j-1}$ otherwise.

To distribute this process, for every $j = 1 \ldots L$ we use Shamir's $(t+1, n)$-secret sharing [Sha79] over $\mathbb{Z}_q$ to split each a_j into n shares $(a_{j,1}, \ldots, a_{j,n})$, so that any $t+1$ of these shares suf£ce to recover a_j, while t or fewer shares give no information about a_j. We set the secret key SK_i of server i to $(a_{1,i}, \ldots, a_{L,i})$, and its public key PK_i to $(y_{1,i} = h^{a_{1,i}}, \ldots, y_{L,i} = h^{a_{L,i}})$. To compute $\mathsf{Prove}(x)$, the user U needs to contact at least $(t+1)$ honest servers. The protocol with the contacted S_i's proceeds in rounds. Assuming inductively that the value σ_{j-1} is known to both the user and the servers (with the base being $\sigma_0 = g$ which is known to everybody), we show how to compute σ_j. If $C(x)_j = 0$, $\sigma_j = \sigma_{j-1}$, so we are done. Otherwise, each server S_i sends the value $\sigma_{j,i} = \sigma_{j-1}^{a_{j,i}}$ to the user. The user locally checks that $(\sigma_{j-1}, \sigma_{j,i}, h, y_{j,i})$ form a proper DDH-tuple. If they do not, U discards the share and stops interacting with S_i. Upon receiving at least $(t+1)$ correct shares, U uses the corresponding Lagrange interpolation in the exponent to compute the (necessarily correct) value σ_j, and sends σ_j to all the servers

it is communicating with. Each server S_i, upon receiving σ_j, checks if $(\sigma_{j-1}, \sigma_j, h, y_j)$ form a valid DDH-tuple. If they do not, the server stops the interaction with U. Then the protocol proceeds to the next round until the entire output is computed.

SECURITY. The security of the above scheme is quite straightforward. Robustness is immediate since every share is checked for consistency. As for pseudorandomness, consider any successful distributed adversary $A = (A_1, A_2)$ who corrupts servers $i_1 \dots i_t$. We build $B = (B_1, B_2)$ for our original VRF as follows. B picks random values $a_{j,i_s} \in \mathbb{Z}_q$ for every $j \in [L]$ and $s \in [t]$, and gives the resulting secret keys $SK_{i_1}, \dots, SK_{i_t}$ to A. It then computes the induced public keys $PK_{i_1}, \dots, PK_{i_t}$ and uses its own public key $h^{a_1}, \dots, h^{a_L}$ to compute the remaining public keys PK_i for all non-corrupted users. This is done by performing the appropriate Lagrange interpolation in the exponent which computes the value $y_{j,i}$ from $y_j, y_{j,i_1}, \dots, y_{j,i_t}$. It hands all these public keys to A, after which B_1 starts running A_1. When A_1 initiates any server on input x, B_1 asks for the value $\mathsf{Prove}(x)$, and uses the response $(\sigma_1, \dots, \sigma_L)$, together with the knowledge of $SK_{i_1}, \dots, SK_{i_t}$, to compute all the relevant shares $\sigma_{j,i}$ (by again doing straightforward Lagrange interpolation in the exponent; details are obvious and omitted). This allows B_1 to simulate all the responses to A_1. After B_1 outputs the same challenge x' as A_1, B_2 gets the output challenge y', which it forwards to A_2 as well. Then B_2 simulates A_2's interaction with the servers in exactly the same way B_1 did it for A_1. Finally, B_2 outputs the same guess $\tilde{b}$ as A_2, which completes the reduction and the proof of security.

EFFICIENCY. The above protocol is quite ef£cient. The communication complexity is $O(t^2 \ell k)$, and the round complexity is $L = O(\ell)$. This is more ef£cient than the complexity of the (non-veri£able) DPRF construction of [Nie02] since no server interaction or expensive interactive zero-knowledge proofs are needed.

Finally, we remark that we can achieve proactive security as well (i.e., periodically refresh the sharing of the secret key to withstand "mobile" attacks [OY91]) by using standard share renewal techniques (see [HJJ+97]). Essentially, each server (veri£ably) distributes 0's to other servers, and all servers locally add these shares to their old secret shares (also correspondingly updating the public shares).

Acknowledgments

I would like to thank Alexander Barg and Venkatesan Guruswami for useful discussions about BCH codes. I would also like to thank Dan Boneh and Don Coppersmith for their preliminary (positive) evaluation of the sf-DDH assumption. Finally, I would like to thank Anna Lysyanskaya for inspiring this work.

References

[ABI86] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7:567–583, 1986.

[AS00] Noga Alon and Joel Spencer. *Probabilistic Method*. Wiley, John and Sons, 2000.

[BCK96] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Keyed hash functions and message authentication. In Neal Koblitz, editor, *Advances in Cryptology—CRYPTO '96*, volume 1109 of *Lecture Notes in Computer Science*, pages 1–15. Springer-Verlag, 18–22 August 1996.

[BCP01] Emmanuel Bresson, Olivier Chevassut, and David Pointcheval. Provably authenticated group dif£e-hellman key exchange — the dynamic case. In Boyd [Boy01], pages 290–309.

[BCPQ01] Emmanuel Bresson, Olivier Chevassut, David Pointcheval, and Jean-Jacques Quisquater. Provably authenticated group dif£e-hellman key exchange. In *Eighth ACM Conference on Computer and Communication Security*, pages 255–264. ACM, November 5–8 2001.

[BF01] Dan Boneh and Matthew Franklin. Identity based encryption from the weil pairing. In Kilian [Kil01], pages 213–229.

[BKR00] Mihir Bellare, Joe Kilian, and Phillip Rogaway. The security of the cipher block chaining message authentication code. *Journal of Computer and System Sciences*, 61:3:362–399, 2000.

[BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In Boyd [Boy01], pages 514–532.

[Boy01] Colin Boyd, editor. *Advances in Cryptology—ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, Gold Coast, Australia, 9–13 December 2001. Springer-Verlag.

[BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing ef£cient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communication Security*, pages 62–73, November 1993. Revised version appears in `http://www-cse.ucsd.edu/users/mihir/papers/crypto-papers.html`.

[BR94] Mihir Bellare and Phillip Rogaway. Optimal asymmetric encryption. In Alfredo De Santis, editor, *Advances in Cryptology—EUROCRYPT 94*, volume 950 of *Lecture Notes in Computer Science*, pages 92–111. Springer-Verlag, 1995, 9–12 May 1994. Revised version available from `http://www-cse.ucsd.edu/users/mihir/`.

[BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In Maurer [Mau96], pages 399–416. Revised version appears in `http://www-cse.ucsd.edu/users/mihir/papers/crypto-papers.html`.

[BS02] Dan Boneh and Alice Silverberg. Applications of multilinear forms to cryptography. IACR E-print Archive. Available from `http://eprint.iacr.org/2002/080/`, 2002.

[CG99] Ran Canetti and Sha£ Goldwasser. An ef£cient threshold public key cryptosystem secure against adaptive chosen ciphertext attack. In Stern [Ste99], pages 90–106.

[CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, pages 209–218, Dallas, Texas, 23–26 May 1998.

[Dam87] Ivan Damgård. Collision-free hash functions and public-key signature schemes. In David Chaum and Wyn L. Price, editors, *Advances in Cryptology—EUROCRYPT 87*, volume 304 of *Lecture Notes in Computer Science*. Springer-Verlag, 1988, 13–15 April 1987.

[FS86] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identi£cation and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology—CRYPTO '86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer-Verlag, 1987, 11–15 August 1986.

[GGM86] Oded Goldreich, Sha£ Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, October 1986.

[GL89] O. Goldreich and L. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, pages 25–32, Seattle, Washington, 15–17 May 1989.

[GMR88] Sha£ Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988.

[GQ88] Louis Claude Guillou and Jean-Jacques Quisquater. A "paradoxical" indentity-based signature scheme resulting from zero-knowledge. In Sha£ Goldwasser, editor, *Advances in Cryptology—CRYPTO '88*, volume 403 of *Lecture Notes in Computer Science*, pages 216–231. Springer-Verlag, 1990, 21–25 August 1988.

[HILL99] J. Håstad, R. Impagliazzo, L.A. Levin, and M. Luby. Construction of pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

[HJJ+97] Amir Herzberg, Markus Jakobsson, Stanislaw Jarecki, Hugo Krawczyk, and Moti Yung. Proactive public key and signature systems. In *Fourth ACM Conference on Computer and Communication Security*, pages 100–110. ACM, April 1–4 1997.

[JN01] Antoine Joux and Kim Nguyen. Separating decision Dif£e-Hellman from Dif£e-Hellman in cryptographic groups. IACR E-print Archive. Available from `http://eprint.iacr.org/2001/003/`, 2001.

[Jou00] Antoine Joux. A one-round protocol for tripartite dif£e-hellman. In *ANTS-IV Conference*, volume 1838 of *Lecture Notes in Computer Science*, pages 385–394. Spring-Verlag, 2000.

[Kil01] Joe Kilian, editor. *Advances in Cryptology—CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*. Springer-Verlag, 19–23 August 2001.

[Lys02] Anna Lysyanskaya. Unique signatures and veri£able random functions from the dh-ddh separation. In Yung [Yun02]. Available from `http://theory.lcs.mit.edu/ anna/papers/lys02.ps`.

[Mau96] Ueli Maurer, editor. *Advances in Cryptology—EUROCRYPT 96*, volume 1070 of *Lecture Notes in Computer Science*. Springer-Verlag, 12–16 May 1996.

[Mic94] Silvio Micali. A secure and ef£cient digital signature algorithm. Technical Report MIT/LCS/TM-501, Massachusetts Institute of Technology, Cambridge, MA, March 1994.

[MR01] Silvio Micali and Leonid Reyzin. Soundness in the public-key model. In Kilian [Kil01].

[MR02] Silvio Micali and Ronald Rivest. Micropayments revisited. In Bart Preneel, editor, *Progress in Cryptology — CT-RSA 2002*, volume 2271 of *Lecture Notes in Computer Science*. Springer-Verlag, February 18-22 2002.

[MRV99] Silvio Micali, Michael Rabin, and Salil Vadhan. Veri£able random functions. In *40th Annual Symposium on Foundations of Computer Science*, pages 120–130, New York, October 1999. IEEE.

[MS77] Jessie MacWilliams and Neil Sloane. *The theory of error-correcting codes*. North-Holland Publicing Company, Amsterdam, 1977.

[MS95] Silvio Micali and Ray Sidney. A simple method for generating and sharing pseudo-random functions. In Don Coppersmith, editor, *Advances in Cryptology—CRYPTO '95*, volume 963 of *Lecture Notes in Computer Science*, pages 185–196. Springer-Verlag, 27–31 August 1995.

[Nie02] Jesper Nielsen. Threshold pseudorandom function construction and its applications. In Yung [Yun02].

[NP98] Moni Naor and Benny Pinkas. Secure and ef£cient metering. In Kaisa Nyberg, editor, *Advances in Cryptology—EUROCRYPT 98*, volume 1403 of *Lecture Notes in Computer Science*, pages 576–590. Springer-Verlag, May 31–June 4 1998.

[NPR99] Moni Naor, Benny Pinkas, and Omer Reingold. Distributed pseudo-random functions and KDCs. In Stern [Ste99], pages 327–346.

[NR97] Moni Naor and Omer Reingold. Number-theoretic constructions of ef£cient pseudo-random functions. In *38th Annual Symposium on Foundations of Computer Science*, pages 458–467, Miami Beach, Florida, 20–22 October 1997. IEEE.

[NRR00] Moni Naor, Omer Reingold, and Alom Rosen. Pseudo-random functions and factoring. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 11–20, Portland, Oregon, 21–23 May 2000.

[Oka92] Tatsuaki Okamoto. Provably secure and practical identi£cation schemes and corresponding signature schemes. In Ernest F. Brickell, editor, *Advances in Cryptology—CRYPTO '92*, volume 740 of *Lecture Notes in Computer Science*, pages 31–53. Springer-Verlag, 1993, 16–20 August 1992.

[OY91] Rafail Ostrovsky and Moti Yung. How to withstand mobile virus attacks. In *10-th Annual ACM Symp. on Principles of Distributed Computing*, pages 51–59, 1991.

[PS96] David Pointcheval and Jacques Stern. Security proofs for signature schemes. In Maurer [Mau96], pages 387–398.

[Sch91] Claus-Peter Schnorr. Ef£cient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.

[Sha79] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

[Ste99] Jacques Stern, editor. *Advances in Cryptology—EUROCRYPT '99*, volume 1592 of *Lecture Notes in Computer Science*. Springer-Verlag, 2–6 May 1999.

[STW96] Michael Steiner, Gene Tsudik, and Michael Waidner. Dif£e-hellman key distribution extended to group communicatio. In *Third ACM Conference on Computer and Communication Security*, pages 31–37. ACM, March 14–16 1996.

[Yun02] Moti Yung, editor. *Advances in Cryptology—CRYPTO 2002*, Lecture Notes in Computer Science. Springer-Verlag, 18–22 August 2002.