

Reusable Designated-Verifier NIZKs for all NP from CDH

Willy Quach¹, Ron D. Rothblum², and Daniel Wichs¹

¹ Northeastern University

quach.w@husky.neu.edu

wichs@ccs.neu.edu

² Technion

rothblum@cs.technion.ac.il

Abstract. Non-interactive zero-knowledge proofs (NIZKs) are a fundamental cryptographic primitive. Despite a long history of research, we only know how to construct NIZKs under a few select assumptions, such as the hardness of factoring or using bilinear maps. Notably, there are no known constructions based on either the computational or decisional Diffie-Hellman (CDH/DDH) assumption without relying on a bilinear map.

In this paper, we study a relaxation of NIZKs in the *designated verifier* setting (DV-NIZK), in which the public common-reference string is generated together with a secret key that is given to the verifier in order to verify proofs. In this setting, we distinguish between *one-time* and *reusable* schemes, depending on whether they can be used to prove only a single statement or arbitrarily many statements. For reusable schemes, the main difficulty is to ensure that soundness continues to hold even when the malicious prover learns whether various proofs are accepted or rejected by the verifier. One-time DV-NIZKs are known to exist for general NP statements assuming only public-key encryption. However, prior to this work, we did not have any construction of reusable DV-NIZKs for general NP statements from any assumption under which we didn't already also have standard NIZKs.

In this work, we construct reusable DV-NIZKs for general NP statements under the CDH assumption, without requiring a bilinear map. Our construction is based on the *hidden-bits paradigm*, which was previously used to construct standard NIZKs. We define a cryptographic primitive called a *hidden-bits generator (HBG)*, along with a designated-verifier variant (DV-HBG), which modularly abstract out how to use this paradigm to get both standard NIZKs and reusable DV-NIZKs. We construct a DV-HBG scheme under the CDH assumption by relying on techniques from the Cramer-Shoup hash-proof system, and this yields our reusable DV-NIZK for general NP statements under CDH.

We also consider a strengthening of DV-NIZKs to the *malicious designated-verifier* setting (MDV-NIZK) where the setup consists of an honestly generated common random string and the verifier then gets to choose his own (potentially malicious) public/secret key pair to generate/verify proofs. We construct MDV-NIZKs under the “one-more CDH” assumption without relying on bilinear maps.

1 Introduction

(Non-Interactive) Zero-Knowledge. Zero-knowledge proofs, introduced in the seminal work of Goldwasser, Micali, and Rackoff [GMR85,GMR89], allow a prover to convince a verifier that a statement is valid without revealing anything beyond its validity. Standard zero-knowledge proof systems are interactive. Blum, Feldman, and Micali [BFM88] introduced the concept of non-interactive zero-knowledge (NIZK) proofs, which consist of a single message from the prover to the verifier. Such NIZKs cannot exist in the plain model, and are therefore considered in the *common reference string (CRS)* model, where a trusted third party chooses some common string (either uniformly at random or from some designated distribution) which is given to both the prover and the verifier. Such NIZKs for general **NP** statements have been constructed from a few select assumptions such as: (doubly-enhanced) trapdoor permutations which can be instantiated from factoring [BFM88,DMP88,FLS99,Go11], the Diffie-Hellman assumption over bilinear groups [CHK03,GOS06] indistinguishability obfuscation [SW14] or fully exponential KDM hardness [CCRR18]. We also have such NIZKs in the random-oracle model [FS87].³ However, despite a long history of research, we don't have any constructions based on several common standard assumptions: most notably the computational or decisional Diffie-Hellman assumptions (CDH, DDH) without requiring a bilinear map, or the learning-with-errors (LWE) assumption.

Designated-Verifier NIZK. In this work, we focus on a relaxed notion of NIZKs in the *designated-verifier* setting (DV-NIZK). In this model a trusted-third party generates a CRS together with secret key which is given to the verifier and is used to verify whether proofs are accepting or rejecting. We distinguish between schemes having *one-time* (a.k.a. single-theorem) security versus *reusable* (a.k.a. multi-theorem) security. One-time secure schemes only guarantee soundness for a single proof of a single statement. However, since the verifier's decision whether to accept or reject a proof depends on the secret key, a malicious prover may be able to learn something about the secret key over time by producing many proofs and seeing whether they are accepted or rejected by the verifier. Reusable DV-NIZKs ensure that soundness continues to hold even in such settings, where a prover can test whether the verifier accepts or rejects various proofs. In terms of constructions, there appears to be a huge gap between these notions. One-time secure DV-NIZKs were constructed for general **NP** statements assuming only the existence of public-key encryption [PsV06]. On the other hand, prior to this work, we did not have any constructions of reusable DV-NIZKs for general **NP** statements based on any assumptions under which we don't already also have standard NIZKs.

³ Additionally, we have constructions of NIZKs with an inefficient prover based on one-way permutations [FLS99]. In this work, we restrict ourselves to NIZKs where the prover can generate proofs efficiently given an **NP** witness.

Malicious-Designated-Verifier NIZK. We also consider a strengthening of (reusable) DV-NIZKs to the malicious-designated-verifier setting (MDV-NIZKs). In this setting, the trusted party only generates a common uniformly random string. The verifier then gets to choose a public/secret key pair where the public key is used by the prover to generate proofs and the secret key is used by the verifier to verify proofs. The main difference between DV-NIZKs and MDV-NIZKs is that, in the latter, we require zero-knowledge to hold even if the public key is chosen maliciously by the verifier. Therefore, an MDV-NIZK is similar to standard NIZKs in that the only *trusted setup* consists of a common random string, but an MDV-NIZK also requires additional *potentially untrusted setup* where the verifier publishes a public-key for which it keeps the corresponding secret key.

The notion of (reusable) MDV-NIZKs is equivalent to 2-round malicious-verifier ZK protocols in the common random string model (where the verifier’s first-round message is reusable) by thinking of the verifier’s public key as the first-round message. It is easy to see that the construction of *non-reusable* DV-NIZKs of [PsV06] extends naturally to yield *non-reusable* MDV-NIZKs assuming 2-round maliciously secure oblivious transfer in the common random string model. However, prior to this work, we did not have any constructions of *reusable* MDV-NIZKs for general NP statements based on any assumptions under which we don’t already also have standard NIZKs.

Prior Work on DV-NIZKs and NIZKs with Pre-Processing. In prior work, the notion of DV-NIZKs was mainly studied in the context of non-malleable and CCA secure encryption. It is known that one-time DV-NIZKs allow us to compile any CPA secure (public-key) encryption scheme into a non-malleable one [PsV06] and reusable DV-NIZKs can compile it into a CCA secure one (by adapting the [NY90,DDN91] paradigm to the designated-verifier case). In this context, the work of Cramer and Shoup [CS98,CS02] constructed “hash-proof systems” which are unconditionally secure reusable DV-NIZKs for specific “algebraic” languages (e.g., the equality of two discrete logarithms) and used them to get practical CCA secure encryption. However, reusable DV-NIZKs have received surprisingly little attention as a general primitive. We believe that this notion is naturally interesting beyond its applications to non-malleable and CCA encryption. For example, it can take the place of standard NIZKs in the context of multiparty computation in scenarios where there is some (reusable) trusted setup.

DV-NIZKs can be thought of as a special case of a more general notion of “NIZKs with preprocessing” in which a trusted-third party creates a CRS together with two secret keys: td_V given to the verifier and td_P given to the prover. We can consider two special cases of such NIZKs with preprocessing: if td_P is empty then this corresponds to the “designated-verifier” (DV-NIZK) setting that we study in this work, and if td_V is empty then we can think of this as a “designated-prover” (DP-NIZK). Several prior works study NIZKs with preprocessing [DMP90,KMO90,LS91,Dam93,DFN06,CC18] but all either

(1) only consider specific “algebraic” languages rather than general **NP**, (2) are not reusable or (3) require assumptions such as factoring from which we already have standard NIZKs. The one exception is a very recent work of Kim and Wu [KW18] (CRYPTO 2018), which gave a novel construction of reusable DP-NIZKs for general **NP** languages under the LWE assumption. In that work, they explicitly asked the question whether one can construct reusable NIZKs in the preprocessing model under the CDH/DDH assumption. We answer their open question positively in this paper by constructing reusable DV-NIZKs under CDH. It remains a fascinating open question whether one can construct reusable DV-NIZKs under LWE, and conversely, whether one can construct reusable DP-NIZKs under CDH/DDH.

1.1 Our Results.

In this work, we construct reusable DV-NIZKs for general **NP** languages under the computational Diffie-Hellman (CDH) assumption without requiring a bilinear map.

Theorem 1.1. *Under the CDH assumption, there exists an (adaptively secure, statistically sound) reusable DV-NIZK proof system for all NP.*

We also construct reusable MDV-NIZKs for general **NP** languages under the one-more CDH (OM-CDH) assumption without requiring a bilinear map.

Theorem 1.2. *Under the One-More CDH assumption (Definition 6.3), there exists an (adaptively secure, statistically sound) reusable MDV-NIZK proof system for all NP.*

Our construction goes through the *hidden-bits* paradigm introduced by Feige, Lapidot and Shamir [FLS99] (see also [Gol01,Gol11]) to construct standard NIZKs. This paradigm consists of two steps. First, construct a NIZK for general **NP** statements in an idealized model called the “hidden-bits model” where the prover is given a long string of uniformly random bits and can choose to reveal some subset of them to the verifier. Such NIZKs in the hidden-bits model were constructed unconditionally with statistical soundness and zero knowledge. Second, use a cryptographic tool to compile NIZKs in the hidden-bits model to NIZKs in the CRS model. Such a compiler was constructed concretely using (doubly enhanced) trapdoor permutations, which can be instantiated based on factoring.

We generalize the second step of the hidden bits paradigm by defining a cryptographic primitive called a “hidden-bits generator” (HBG) which can be used to compile NIZKs in the hidden-bits model into ones in the CRS model.⁴ This primitive modularizes the “hidden-bits paradigm” and simplifies the task of constructing NIZKs by reducing it to the task of constructing a HBG. We

⁴ A similar primitive called a “verifiable pseudorandom generator” was defined by [DN00] for the purpose of constructing ZAPs, which also lead to a construction of NIZKs.

also clarify how to use HBG to get adaptive ZK security via the “hidden bits paradigm”, which turns out to be surprisingly subtle and was not very clear from prior presentations of this paradigm. To get our main result, we generalize the hidden bits paradigm even further by extending the notion of HBG to the designated-verifier setting (DV-HBG) and the malicious-designated-verifier setting (MDV-HBG) and showing that the same compiler allows us to go from DV-HBG (resp. MDV-HBG) to reusable DV-NIZKs (resp. MDV-NIZKs). We then show how to construct DV-HBG from the computational Diffie-Hellman (CDH) assumption without bilinear maps. The last step uses the Cramer-Shoup hash-proof system, which can be thought of as a reusable DV-NIZK for equality of two discrete logarithms. Therefore we are in some sense bootstrapping a reusable DV-NIZK for this specific language to get a reusable DV-NIZK for all of NP. Finally, we show how to construct MDV-HBG from the one-more CDH (OM-CDH) assumption. This essentially starts with our construction of DV-HBG, which is clearly insecure in the malicious-designated-verifier setting, and shows how to immunize it against malicious attacks. While the high level idea is simple, the proof of security is quite involved and uses techniques which may be of independent interest.

1.2 Technical Overview

NIZKs via the Hidden-Bits Paradigm. We first review the “hidden-bits paradigm” proposed by [FLS99]; see [Gol01,Gol11] for a modern presentation which we follow here.

The starting point of this paradigm is a construction of NIZKs in an idealized model called the “hidden-bits model”. In this model, there is a trusted third party that generates uniformly random bits r_1, \dots, r_k and gives them to the prover. The prover outputs a proof π along with a subset $I \subseteq [k]$ of the bits to open. The verifier gets (I, π) from the prover together with the bits $\{r_i\}_{i \in I}$ from the trusted third party. Note that the verifier does not learn anything about the unopened bits $\{r_i\}_{i \notin I}$ and the prover cannot modify the values of the opened bits $\{r_i\}_{i \in I}$. Such NIZKs in the hidden-bits model can be constructed unconditionally with security against an unbounded prover/verifier where the soundness error can be made exponentially small.

The second step compiles NIZKs in the hidden-bits model into NIZKs in the CRS model. Such a compiler was presented by [FLS99,Gol01,Gol11] using doubly-enhanced trapdoor permutations (TDPs) (see also [BY93,GR13,CL17]). On a high level, the CRS consists of random values y_1, \dots, y_k in the range of the TDP. The prover chooses a random permutation f_{com} along with an inversion trapdoor sk and inverts all of the values in the CRS to get preimages x_1, \dots, x_k . Define r_1, \dots, r_k to be hardcore bits of x_1, \dots, x_k . The prover then runs the hidden-bits prover with r_1, \dots, r_k to generate some proof (π, I) to which it appends the values $\text{com}, \{x_i\}_{i \in I}$. The verifier checks $y_i = f_{\text{com}}(x_i)$, computes $\{r_i\}_{i \in I}$ to be the hardcore bit of x_i and then runs the hidden bits verifier on (π, I) . Intuitively, a malicious prover has a extremely limited ability to control the randomness r_1, \dots, r_k by choosing com ; by relying on an exponentially small

soundness error of the hidden-bits proofs which survives a union-bound over all such com 's, this flexibility is insufficient to break soundness. On the other hand, the verifier does not learn anything about the values $\{r_i\}_{i \notin I}$ by the security of the TDP.⁵ While this is the high level approach, there are some subtleties involved; see [BY93, Gol01, Gol04, Gol11, GR13, CL17].

Hidden-Bits Generator (HBG). We begin by defining an abstract cryptographic primitive, which we call a hidden-bits generator (HBG), that can be used to compile NIZKs in the hidden-bits model to NIZKs in the CRS model. An HBG that generates k bits consists of three algorithms:

- **Setup** creates a crs .
- **GenBits**(crs) outputs a *short* commitment com whose size is much smaller than k , along with hidden-bits $\{r_i\}_{i \in [k]}$, and certificates $\{\pi_i\}_{i \in [k]}$.
- **Verify**($\text{crs}, \text{com}, i, r_i, \pi_i$) checks the certificate π_i to verify that r_i is indeed the i 'th hidden bit.

An HBG should satisfy two simple properties. Firstly, we require the scheme to be *statistically binding*, meaning that (crs, com) together completely determine some sequence of bits r_1, \dots, r_k and no (even inefficient) prover can come up with a valid certificate π'_i for the wrong bit $r'_i \neq r_i$. Intuitively, by combining the above property together with the requirement that com is short, we ensure that the prover does not have much control over the bits r_i that he can open and the limited control that he does have is insufficient to break the soundness of the hidden-bits NIZK (by amplifying its soundness sufficiently so that it survives a union bound over all the com 's that the prover can choose). Secondly, we require the scheme to be *computationally hiding*, meaning that for any set $I \subseteq [k]$, if we are given honestly generated $\text{crs}, \text{com}, \{r_i, \pi_i\}_{i \in I}$ then the “unopened” hidden bits $\{r_j\}_{j \notin I}$ are computationally indistinguishable from uniform.

Compiling from Hidden-Bits Model to CRS Model. Intuitively, we would like to use HBG to compile NIZKs from the hidden-bits model to the CRS model by letting the prover generate the hidden-bits via the HBG **GenBits** algorithm. There are two issues with this basic approach:

- For soundness, if the malicious prover chooses a “bad” (not uniformly random) com then the HBG abstraction does not provide any guarantees that the bits r_i to which he is committed are random and hence we cannot rely on the soundness of the hidden-bits NIZK.
- For zero-knowledge, we notice that the honest hidden-bits prover may choose the set I adaptively depending on all of the bits $\{r_i\}_{i \in [k]}$ (and indeed this is the case for the hidden-bits NIZK constructed in [FLS99]) and we still need to argue that the unopened bits $\{r_j\}_{j \notin I}$ are hidden. The hiding property of

⁵ The basic compiler only achieves zero-knowledge for a single theorem and [FLS99] then relies on another generic compiler via the “or trick” to go from single-theorem to multi-theorem zero-knowledge.

HBG only guarantees that the unopened bits are hidden when I is chosen ahead of time.

To fix both of the above issues we add additional uniformly random bits s_1, \dots, s_k to the CRS of the NIZK and define the hidden-bits to be $r_i \oplus s_i$ where r_i comes from the HBG. This ensures that for any fixed com chosen by a malicious prover the hidden-bits that he can open are uniform over the choice of s_i .⁶ It also ensures that the choice of the set I chosen by the honest hidden-bits prover is independent of the outputs r_i of the HBG and therefore allows us to rely on HBG security.

We uncover an additional complication when proving *adaptive* ZK, where the malicious verifier can choose the statement to be proven adaptively depending on the CRS. The work of [FLS99] showed adaptive ZK for their particular protocol (using particular hidden-bits NIZK) but it did not give a modular proof. Indeed, our attempts to prove that the compiler can generically start with any hidden-bits NIZK and achieve adaptive ZK failed for subtle reasons involving “selective opening” failures. Instead, we were able to abstract out a special property of the hidden-bits NIZK of [FLS99] which we call “special ZK”, which we show to be sufficient to get adaptive ZK in the CRS model via the above compiler.

Using the compiler, we reduce the task of constructing NIZKs to that of constructing an HBG, which is a conceptually much simpler primitive.

Designated Verifier Setting: (M)DV-HBG to (M)DV-NIZK. We generalize the notion of HBG to the designated-verifier setting (DV-HBG). The only differences are that: (1) the **Setup** algorithm generates a crs together with a trapdoor td which is given to the verifier and the **Verify** algorithm takes the trapdoor td as an input, (2) we modify the statistically binding security property to hold even if a computationally unbounded prover can make polynomially many queries to the $\text{Verify}(\text{crs}, \text{td}, \dots)$ oracle which allows it to check whether various certificates are valid or invalid, and (3) we modify the computationally hiding property to hold even given td . To get our main result, we naturally extend our compiler to show that DV-HBG allows us to compile NIZKs in the hidden-bits model into reusable DV-NIZKs. Therefore, we reduce the task of constructing reusable DV-NIZKs to that of constructing DV-HBG.

We further generalize the notion of HBG to the malicious-designated-verifier setting (MDV-HBG). Now, in addition to a **Setup** algorithm that generates the crs there is a **KeyGen** algorithm that generates a public key pk along an associated secret key sk . Essentially, we think of crs, pk as together corresponding to the crs in the previous definition, and of sk as the trapdoor. The binding property is essentially the same as before. However, we require that hiding holds even if pk is generated maliciously (and adaptively depending on crs). We show that MDV-HBG allows us to compile NIZKs in the hidden-bits model into reusable MDV-NIZKs. Therefore, we reduce the task of constructing reusable MDV-NIZKs to that of constructing MDV-HBG.

⁶ The fact that the prover can adaptively choose com *after* seeing s_1, \dots, s_k is handled by simply taking a union bound over all possible choices of com .

DV-HBG from CDH. We show how to instantiate a designated-verifier DV-HBG based on the computational Diffie-Hellman (CDH) assumption to get our reusable DV-NIZK from CDH. Our construction relies on the ideas underlying the Cramer-Shoup (1-universal) hash-proof system [CS98,CS02] which can be thought of as an unconditionally secure reusable DV-NIZK for the “equality of two discrete logs” – i.e., given some public group elements g, h we define the language consisting of tuples (g', h') such that $\text{DLOG}_g(g') = \text{DLOG}_h(h')$. In particular, we think of the projection key of the hash-proof system as the CRS of the DV-NIZK, and the hashing key as the associated trapdoor. In the body of our paper, we give our full construction using the specific Cramer-Shoup instantiation, but for the introduction we will treat the Cramer-Shoup reusable DV-NIZK proof system as a black-box.

Our DV-HBG construction works as follows. Let \mathbb{G} be some cyclic group of order p and let g be a generator.

- The **Setup** algorithm chooses random group elements h_1, \dots, h_k . It also instantiates k copies of the Cramer-Shoup DV-NIZK with respect to the public group elements (g, h_i) respectively. The **crs** consists of g, h_1, \dots, h_k together with the k values $\{\text{crs}_i\}_{i \in [k]}$ of the Cramer-Shoup DV-NIZK. The trapdoor $\text{td} = \{\text{td}_i\}_{i \in [k]}$ consists of the k trapdoors for the Cramer-Shoup DV-NIZK.
- The **GenBits**(**crs**) algorithm chooses $y \leftarrow \mathbb{Z}_q$ and sets $\text{com} = g^y$. For $i = 1 \dots, k$, it sets $t_i = h_i^y$, $r_i = \text{hc}(t_i)$, where **hc** is a hardcore predicate (e.g., Goldreich-Levin [GL89]). Finally it sets $\pi_i = (t_i, \pi_i^{CS})$ where π_i^{CS} is a Cramer-Shoup proof that $\text{DLOG}_g(\text{com}) = \text{DLOG}_{h_i}(t_i)$.
- The **Verify** algorithm gets r_i and $\pi_i = (t_i, \pi_i^{CS})$ and checks that $r_i = \text{hc}(t_i)$ and that π_i^{CS} is a valid Cramer-Shoup proof using the corresponding trapdoor td_i .

For the statistically binding property we note that given **crs**, **com** the values $t_i = h_i^y$ and therefore also the hidden bits $r_i = \text{hc}(t_i)$ are completely determined. The prover cannot lie about t_i and therefore also about r_i by the unconditional reusable security of the Cramer-Shoup proof, and this holds even given oracle access to the Cramer-Shoup verifier. For the computational hiding property we rely on the fact that, given g, h_i, g^y , the CDH assumption ensures that h_i^y is computationally unpredictable and therefore $\text{hc}(h_i^y)$ is indistinguishable from uniform. This holds even given h_j, h_j^y for various random h_j since the distinguisher can sample such values himself by sampling $h_j = g^{x_j}$ and computing $h_j^y = (g^y)^{x_j}$.

MDV-HBG from One-More CDH. Finally, we show how to instantiate our malicious-designated-verifier MDV-HBG based on the one-more CDH assumption to get our reusable MDV-NIZK from one-more CDH. The construction and the security intuition are somewhat involved and so we present them in several stages.

Initial Attempt. As a first attempt, we can try to use the previous construction directly as an MDV-HBG. In particular, we can set the **crs** to only consist of the uniformly random values $\text{crs} = (h_1, \dots, h_k)$. The Cramer-Shoup DV-NIZKs then

naturally define \mathbf{pk}, \mathbf{sk} . Here it helps to be concrete about how the Cramer-Shoup DV-NIZK works. For each i , the Cramer-Shoup proof system defines $\mathbf{pk}_i = h_i^{a_i} g^{b_i}$ and the corresponding $\mathbf{sk}_i = (a_i, b_i)$. The MDV-HBG public keys and secret keys consist of these values $\mathbf{pk} = \{\mathbf{pk}_i\}, \mathbf{sk} = \{\mathbf{sk}_i\}$. Given a commitment $\mathbf{com} = g^y$, recall that the i 'th hidden bit is defined by taking a hardcore predicate $r_i = \mathbf{hc}(t_i)$ where $t_i = h_i^y$. The opening to the i 'th hidden bit consists of $t_i = h_i^y$ and the Cramer-Shoup proof $\pi_i^{CS} = \mathbf{pk}_i^y$.

Attack On Initial Attempt. Unfortunately, it's clear that the above is not secure as MDV-HBG. For example, if the malicious verifier chooses $\mathbf{pk}_i = h_j$ for $j \neq i$ then, by opening the i 'th hidden bit and giving a proof $\pi_i^{CS} = \mathbf{pk}_i^y = h_j^y$, the prover inadvertently also reveals the j 'th hidden bit! While the above is easily detectable, the malicious verifier can alternately set $\mathbf{pk}_i = h_j^x$ for a random x and still perform the same attack without being detectable. At the very least, we need to modify our solution to overcome this particular attack.

The Fix. We start with the above “base scheme”, which is not secure in the MDV setting, and show how to immunize against the above attack. To do so, we use the “base scheme” to generate ℓ “base hidden values” for some $\ell \gg k$ and then combine them carefully to create the k “actual hidden bits”. Recall that the base scheme defines a commitment g^y and the ℓ base hidden values are $t_j = h_j^y$. We can open any base value by giving the opening $\pi_j^{CS} = \mathbf{pk}_j^y$.

Instead of using the base values directly, we define each of the k “actual hidden bits” by combining together a small group of base values and applying a (Goldreich-Levin) hard-core predicate \mathbf{hc} . The groups are chosen via a pseudo-random mapping φ which maps each $i \in [k]$ to a small group $\varphi(i) \subseteq [\ell]$. In other words, the i 'th actual hidden bit is defined as $r_i = \mathbf{hc}(\{t_j : j \in \varphi(i)\})$. The mapping φ is chosen by the prover and is a part of \mathbf{com} . To open any actual hidden bit $i \in [k]$ the prover opens all of the base hidden values t_j^y and also provides the corresponding Cramer-Shoup proofs \mathbf{pk}_j^y for $j \in \varphi(i)$. Note that, since φ is a part of \mathbf{com} and we require \mathbf{com} to be short, it is important that φ has a short description size and therefore it must be a pseudo-random rather than truly random mapping. For concreteness, we set the number of based hidden values to $\ell = 3k\lambda$ and the group size to $|\varphi(i)| = \lambda$, where λ is the security parameter.

Intuition for the Fix. Intuitively, this prevents the above attacks for the following reason. Assume that the verifier can choose \mathbf{pk} maliciously so that the opening of any base value j can inadvertently also reveal some other base value $j' = \psi(j)$, where ψ is some mapping defined implicitly by the choice of \mathbf{pk} . Nevertheless, it is likely that each hidden bit i depends on some hidden value $j \in \varphi(i)$ that is not revealed even if we open all the other hidden bits $i' \neq i$. In particular, opening the bits $i' \neq i$ corresponds to giving out the base hidden value j' as well as the inadvertently opened values $\psi(j')$ for each $j' \in \varphi(i')$. But the entire set of revealed values $R = \{j', \psi(j') : j' \in \varphi(i'), i' \neq i\}$ is of size $|R| \leq 2k\lambda$ and $\varphi(i) \subset [\ell = 3k\lambda]$ appears to be a random and independent subset of size $|\varphi(i)| = \lambda$. Hence it is likely that $\varphi(i)$ contains some value $j \notin R$ which was not revealed. Here we crucially rely on the fact that φ is chosen (pseudo-

)randomly by the prover after the verifier chooses \mathbf{pk} which defines the mapping ψ .

The One-More CDH Assumption. While the above idea seems to immunize against the particular class of attacks we previously discussed, proving security against general attacks is more challenging. Nevertheless, we manage to do so under the “one-more CDH” assumption. The one-more CDH assumption considers an adversary who is given g, g^y, h_1, \dots, h_k along with an oracle $O_y(\cdot)$ which takes as input an arbitrary group element f and returns $O_y(f) = f^y$. It says that even if the adversary makes m arbitrary calls to the oracle O_y he cannot predict more than m of the values $\{h_j^y\}$.

Security Under One-More CDH. Our high level proof goes as follows. Assume that a malicious verifier gets to choose $\mathbf{pk} = \{\mathbf{pk}_j\}_{j \in [\ell]}$ maliciously after seeing $\text{crs} = \{h_j\}_{j \in [\ell]}$ and can break hiding. This means that for some $i \in [k]$, if the verifier gets a random com and openings to all the hidden bits *except* for the i 'th one, he can distinguish hidden bit i from uniform with non-negligible advantage. Since the i 'th hidden bit is defined by taking the Goldreich-Levin hardcore bit of the base hidden values $j \in \varphi(i)$, this means that the verifier can also predict all these values with non-negligible probability. So, if the verifier gets $\varphi, g, g^y, \{h_j^y, \mathbf{pk}_j^y : j \in \varphi(i'), i' \in [k] \setminus \{i\}\}$ then he can predict $\{h_j^y : j \in \varphi(i)\}$. Intuitively, we want to use such a verifier to break one-more CDH.

But in the above scenario, the verifier gets many more values raised to the y power than he is able to output. To get around this, we want to “rewind” the verifier run him on many different choices of φ to get more values $\{h_j^y : j \in \varphi(i)\}$ out of him. But each time we rewind we also need to provide him with the appropriate values $\{h_j^y, \mathbf{pk}_j^y : j \in \varphi(i'), i' \in [k] \setminus \{i\}\}$ so we are again getting fewer powers of y out than we need to put in, which appears to be self-defeating. If φ were truly random, we could get around this by freshly sample $\varphi(i)$ on each rewinding but keep $\varphi(i') : i' \in [k] \setminus i$ fixed – that way we would only need to give out some fixed $2k\lambda$ values $\{h_j^y, \mathbf{pk}_j^y : j \in \varphi(i'), i' \in [k] \setminus \{i\}\}$ but on each rewinding we get some additional fresh values $\{h_j^y : j \in \varphi(i)\}$ out of the verifier, and eventually we get more out than we put in which allows us to break one-more CDH.

Unfortunately φ needs to have a short description, and therefore can only be pseudorandom, in which case it's not clear how to freshly re-sample $\varphi(i)$ while keeping $\varphi(i') : i' \in [k] \setminus i$ fixed. We resolve this issue by using a special form of pseudorandom functions (PRFs) called “somewhere equivocal PRFs” [HJO+16] which essentially allow us to do exactly this while keeping the description of φ short. Furthermore, such somewhere equivocal PRFs were constructed from only one-way functions using the ideas of “distributed point functions” [GI14, BGI15] and therefore don't introduce any additional assumptions.

1.3 Concurrent works

Concurrently and independently of ours, the works of [CH19] and [KNYY19] present a similar construction of reusable DV-NIZKs from CDH, compiling

the hidden-bits NIZK of [FLS99] using the Cramer-Shoup hash-proof system [CS98,CS02,CKS08]. Additionally, they respectively obtain the following results:

- [CH19] gives a construction of NIZKs for all NP assuming LWE, along with a *non-interactive witness indistinguishable* (NIWI) proof for the Bounded Distance Decoding problem.
- [KNYY19] builds pre-processing NIZKs for all NP with *succinct* proofs, namely a pre-processing NIZK from DDH with proofs of size $|C| + \text{poly}(\lambda)$ (where C is a circuit checking the NP relation), and a designated-prover NIZK from (strong) assumptions over pairing-friendly groups, with proof size $|C| + \text{poly}(\lambda)$.

Meanwhile, our work introduces the notion of *malicious designated-verifier* NIZKs (MDV-NIZK), and presents a construction from the One-More CDH assumption.

Organization

Basic definitions and notations are given in Section 2. In Section 3 we introduce our new notion of Hidden Bits Generator (HBG). In Section 4 we show how to use an HBG to construct NIZKs. In Section 5 we construct a designated-verifier Hidden Bits Generator assuming CDH. A few extensions are mentioned in Section 7. In the full version of the paper, we additionally give a construction of a HBG from the CDH assumption over bilinear groups and we construct a HBG from (doubly-enhanced) trapdoor permutations.

2 Preliminaries

We will denote by λ the security parameter. The notation $\text{negl}(\lambda)$ denotes any function f such that $f(\lambda) = \lambda^{-\omega(1)}$, and $\text{poly}(\lambda)$ denotes any function f such that $f(\lambda) = \mathcal{O}(\lambda^c)$ for some $c > 0$.

We define the statistical distance between two random variables X and Y over some domain Ω as: $\mathbf{SD}(X, Y) = \frac{1}{2} \sum_{w \in \Omega} |X(w) - Y(w)|$. We say that two ensembles of random variables $X = \{X_\lambda\}$, $Y = \{Y_\lambda\}$ are *statistically indistinguishable*, denoted $X \stackrel{s}{\approx} Y$, if $\mathbf{SD}(X_\lambda, Y_\lambda) \leq \text{negl}(\lambda)$.

We say that two ensembles of random variables $X = \{X_\lambda\}$, and $Y = \{Y_\lambda\}$ are *computationally indistinguishable*, denoted $X \stackrel{c}{\approx} Y$, if, for all (non-uniform) PPT distinguishers Adv , we have $|\Pr[\text{Adv}(X_\lambda) = 1] - \Pr[\text{Adv}(Y_\lambda) = 1]| \leq \text{negl}(\lambda)$.

For a set X , integer k and sequence $x \in X^k$, we denote by x_i the i -th entry in the sequence, for any $i \in [k]$. For a subset $I \subset [k]$, we denote by $x_I = (x_i)_{i \in I}$ the subsequence of x in locations I .

For a probabilistic algorithm $\text{alg}(\cdot)$, we may explicit its internal randomness as follows: $\text{alg}(\cdot; \text{coins})$.

2.1 The Diffie-Hellman Assumption

A *group generator* $(\mathbb{G}, p, g) \leftarrow \text{GroupGen}(1^\lambda)$ is a PPT algorithm which, on input 1^λ , outputs the description of a cyclic group \mathbb{G} of order p , and a generator g of \mathbb{G} . We require that there are efficient algorithms running in time $\text{poly}(\lambda)$ to perform the group operation in \mathbb{G} and to test membership in \mathbb{G} . For notational simplicity, we will often shorten such an output (\mathbb{G}, p, g) to \mathbb{G} and assume that g, p are implicit. A *prime-order group generator* additionally ensures that p is prime.

Definition 2.1 (Computational Diffie-Hellman (CDH) assumption). *Let GroupGen be a group generator. We say that the Computational Diffie-Hellman (CDH) assumption holds relative to GroupGen if for all PPT algorithm \mathcal{A} , we have:*

$$\Pr \left[\mathcal{A}(\mathbb{G}, p, g, g^a, g^b) = g^{ab} : (\mathbb{G}, p, g) \leftarrow \text{GroupGen}(1^\lambda), (a, b) \xleftarrow{\$} \mathbb{Z}_p^2 \right] \leq \text{negl}(\lambda).$$

Given such a group generator satisfying the CDH assumption, we can consider an associated (randomized) *hard-core bit* $\text{hc} : \mathbb{G} \rightarrow \{0, 1\}$ such that for all PPT algorithm \mathcal{A} , we have:

$$\Pr \left[\mathcal{A}(\mathbb{G}, p, g, g^a, g^b, \tau) = \text{hc}(g^{ab}; \tau) : \begin{array}{l} \tau \xleftarrow{\$} \{0, 1\}^{L(\lambda)} \\ (\mathbb{G}, p, g) \leftarrow \text{GroupGen}(1^\lambda) \\ (a, b) \xleftarrow{\$} \mathbb{Z}_p^2 \end{array} \right] \leq 1/2 + \text{negl}(\lambda),$$

where the hard-core bit hc uses $L(\lambda)$ random coins.

Such a hard-core bit can be generically obtained, using the Goldreich-Levin construction [GL89].

2.2 Reusable Designated-Verifier NIZKs

In this section we define the notion of Reusable Designated-Verifier NIZKs (and obtain the standard notion of NIZK as a special case).

Definition 2.2 (Reusable DV-NIZKs). *Let L be an NP language with witness relation R_L . A Reusable Designated-Verifier Non-Interactive Zero-Knowledge (DV-NIZK) Proof for L is a tuple of PPT algorithms $(\text{Setup}, \mathcal{P}, \mathcal{V})$ where:*

- $\text{Setup}(1^\lambda, 1^n)$: On input the security parameter λ and statement length n , outputs a common reference string crs and a trapdoor td ;
- $\mathcal{P}(\text{crs}, x, w)$: On input a common reference string crs , a statement x of length n and a witness w , outputs a proof π ;
- $\mathcal{V}(\text{crs}, \text{td}, x, \pi)$: On input a common reference string crs , a trapdoor td , a statement x and a proof π , outputs **accept** or **reject**,

such that they satisfy the following properties:

- **Completeness:** We require that for all $(x, w) \in R_L$, we have:

$$\Pr \left[\mathcal{V}(\text{crs}, \text{td}, x, \pi) = \text{accept} : \begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^{|x|}) \\ \pi \leftarrow \mathcal{P}(\text{crs}, x, w) \end{array} \right] = 1;$$

- **Statistical Soundness:** Let n and Q be any polynomials, and let $\tilde{\mathcal{P}}$ be any (computationally unbounded) cheating prover that makes at most $Q(\lambda)$ queries to an oracle $\mathcal{V}(\text{crs}, \text{td}, \cdot, \cdot)$ which takes as input (x, π) , and outputs $\mathcal{V}(\text{crs}, \text{td}, x, \pi)$. We require that:

$$\Pr \left[\mathcal{V}(\text{crs}, \text{td}, x, \pi) = \text{accept} \wedge x \notin L : \begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^{n(\lambda)}) \\ (x, \pi) \leftarrow \tilde{\mathcal{P}}^{\mathcal{V}(\text{crs}, \text{td}, \cdot, \cdot)}(\text{crs}) \end{array} \right] \leq \text{negl}(\lambda);$$

- **Zero-Knowledge (Selective):** We require that there exists a PPT simulator Sim such that for any PPT stateful⁷ adversary \mathcal{A} , the two following distributions are computationally indistinguishable:

$\text{EXP}^{\text{Real}}(1^\lambda) :$ $(x, w) \leftarrow \mathcal{A}(1^\lambda)$ <i>where</i> $(x, w) \in R_L$ $(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^{ x })$, $\pi \leftarrow \mathcal{P}(\text{crs}, x, w)$ Output $\mathcal{A}(\text{crs}, \text{td}, \pi)$	$\text{EXP}^{\text{Ideal}}(1^\lambda) :$ $(x, w) \leftarrow \mathcal{A}(1^\lambda)$ <i>where</i> $(x, w) \in R_L$ $(\text{crs}, \text{td}, \pi) \leftarrow \text{Sim}(1^\lambda, x)$ Output $\mathcal{A}(\text{crs}, \text{td}, \pi)$
--	---

Our basic definition only considers selective ZK where the statement being proven is chosen ahead of time, prior to seeing the CRS. In Section 4.1 we also consider a stronger notion of adaptive ZK.

Our definition of designated-verifier NIZK coincides with that of standard (publicly verifiable) NIZK if the trapdoor td is empty.

Definition 2.3. A publicly-verifiable NIZK is a reusable designated-verifier NIZK where the trapdoor td output by Setup is an empty string.

Remark 2.4 (Bounding the number of queries to the Verify oracle). Notice that for soundness we only allow the unbounded cheating prover to make a *polynomial* number of queries to $\mathcal{V}(\text{crs}, \text{td}, \cdot, \cdot)$. One would ideally allow the unbounded cheating prover to make *arbitrarily* many queries to \mathcal{V} (matching more closely the publicly-verifiable setting, where a cheating prover can indeed query the verification algorithm on arbitrarily many inputs). It turns out that any DV-NIZK satisfying this stronger notion can be generically turned into a publicly-verifiable one. This is because the cheating prover can query all possible proofs to \mathcal{V} for any $x \notin L$; and therefore soundness can only hold if there are no valid proof of any false statement (with overwhelming probability over the choice of crs), in which case soundness also holds when the prover is given the trapdoor. Therefore this is essentially the best requirement one can hope for as a meaningful notion of reusable DV-NIZKs which is weaker than publicly-verifiable ones.

⁷ Throughout this paper we follow the convention that whenever a stateful adversary \mathcal{A} is invoked with some inputs it also produces some state which it gets as input on the next invocation.

Remark 2.5 (Single-Theorem vs. Multi-Theorem Zero-Knowledge). The definition of ZK above is often referred to as “single-theorem ZK” since it only requires zero-knowledge to hold for a *single* statement. However, there is a generic compiler from single-theorem ZK to multi-theorem ZK where zero-knowledge holds polynomially many statements via the “OR trick” [FLS99]. We note that the very same transformation directly applies to both the selective and adaptive ZK setting and also both the publicly-verifiable and the designated-verifier setting.

2.3 NIZKs in the Hidden-Bits Model

We now recall the definition of a NIZK in the hidden-bits model:

Definition 2.6 (NIZK in the Hidden-Bits Model). *Let L be an NP language and n be an integer. A Non-Interactive Zero-Knowledge Proof in the Hidden-Bits Model for L is given by a pair of PPT algorithms $(\mathcal{P}, \mathcal{V})$, and a polynomial $k(\lambda, n)$, where:*

- $\mathcal{P}(1^\lambda, r, x, w)$: On input string $r \in \{0, 1\}^{k(\lambda, n)}$, a statement x of size $|x| = n$ and a witness w , output a set of indices $I \subseteq [k]$ and proof π .
- $\mathcal{V}(1^\lambda, I, r_I, x, \pi)$: On input a subset $I \subseteq [k]$, a string r_I , a statement x and a proof π , outputs **accept** or **reject**,

such that they satisfy the following properties:

- **Completeness:** We require that for all $x \in L$ of size $|x| = n$ with witness w we have:

$$\Pr \left[\mathcal{V}(1^\lambda, I, r_I, x, \pi) = \text{accept} : \begin{array}{l} r \xleftarrow{\$} \{0, 1\}^{k(\lambda, n)} \\ (I, \pi) \leftarrow \mathcal{P}(1^\lambda, r, x, w) \end{array} \right] = 1;$$

- **Soundness:** We require that for all polynomial $n = n(\lambda)$, and all unbounded cheating prover $\tilde{\mathcal{P}}$, we have:

$$\Pr \left[\begin{array}{l} \mathcal{V}(1^\lambda, I, r_I, x, \pi) = \text{accept} \\ \wedge x \notin L \\ \wedge |x| = n \end{array} : \begin{array}{l} r \xleftarrow{\$} \{0, 1\}^{k(\lambda, n)} \\ (x, \pi, I) \leftarrow \tilde{\mathcal{P}}(1^\lambda, r) \end{array} \right] \leq \text{negl}(\lambda);$$

- **Zero-Knowledge:** We require that there exists an efficient simulator Sim such that for any adversary \mathcal{A} the two following distributions are statistically indistinguishable:

$$(I, r_I, \pi) \overset{\$}{\approx} (I', r'_I, \pi')$$

where $(x, w) \leftarrow \mathcal{A}(1^\lambda)$, $r \leftarrow \{0, 1\}^{k(\lambda, |x|)}$, $(I, \pi) \leftarrow \mathcal{P}(1^\lambda, r, x, w)$, $(I', r'_I, \pi') \leftarrow \text{Sim}(1^\lambda, x)$.

When clear from context, we will omit 1^λ as an argument to the algorithms defined above.

Remark 2.7 (Amplifying soundness). Let $\ell(\lambda, n)$ be a polynomial. Then, given any NIZK in the hidden-bits model, we can build one with soundness $2^{-\ell(\lambda, n)} \cdot \text{negl}(\lambda)$. This is simply done by running $\ell(\lambda, n)$ copies of the NIZK in parallel, and where the new verification algorithm accepts a proof if and only if all of the executions accept. Note that doing so requires to use $k \cdot \ell(\lambda, n)$ hidden bits instead of k initially.

Theorem 2.8 ([FLS99], see also [Gol01, Section 4.10.2]). *Every $L \in \text{NP}$ has a NIZK in the Hidden-Bits Model.*

3 Hidden-Bits Generator

In this section, we define our new notion of Hidden-Bits Generator (HBG). For simplicity, we first define a publicly verifiable version of HBG and then extend the definition to a designated-verifier version (DV-HBG).

Definition 3.1 (Hidden-Bits Generator). *A Hidden-Bits Generator (HBG) is given by a set of PPT algorithms (Setup, GenBits, Verify):*

- **Setup**($1^\lambda, 1^k$): *Outputs a common reference string crs.*
- **GenBits**(crs): *Outputs a triple $(\text{com}, r, \{\pi_i\}_{i \in [k]})$, where $r \in \{0, 1\}^k$.*
- **Verify**(crs, com, i, r_i, π_i): *Outputs accept or reject, where $i \in [k]$.*

We require any Hidden-Bits Generator to satisfy the following properties:

Correctness: *We require that for every polynomial $k = k(\lambda)$ and for all $i \in [k]$, we have:*

$$\Pr \left[\text{Verify}(\text{crs}, \text{com}, i, r_i, \pi_i) = \text{accept} : \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^k) \\ (\text{com}, r, \pi_{[k]}) \leftarrow \text{GenBits}(\text{crs}) \end{array} \right] = 1.$$

Succinct Commitment: *We require that there exists some set $\mathcal{COM}(\lambda)$ and some constant $\delta < 1$ such that $|\mathcal{COM}(\lambda)| \leq 2^{k^\delta \text{poly}(\lambda)}$, and such that for all crs output by **Setup**($1^\lambda, 1^k$) and all com output by **GenBits**(crs) we have $\text{com} \in \mathcal{COM}(\lambda)$. Furthermore, we require that for all $\text{com} \notin \mathcal{COM}(\lambda)$, **Verify**(crs, com, \cdot, \cdot) always outputs **reject**.⁸*

⁸ The set of commitments \mathcal{COM} should not be thought of as the set of all valid commitments (and indeed it may contain commitments not in the support of **GenBits**). In particular, the simplest way to satisfy this property is to bound the bit-length of **com** and have the verifier reject commitments that are too large. Note that additional structural properties about **com** can be checked by the **Verify** algorithm.

Statistical Binding: *There exists an (inefficient) deterministic algorithm $\text{Open}(1^k, \text{crs}, \text{com})$ such that for every polynomial $k = k(\lambda)$, on input 1^k , crs and com , the algorithm outputs r such that for every (potentially unbounded) cheating prover $\tilde{\mathcal{P}}$:*

$$\Pr \left[\begin{array}{l} r_i^* \neq r_i \\ \wedge \text{Verify}(\text{crs}, \text{com}, i, r_i^*, \pi_i) = \text{accept} \end{array} : \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^k) \\ (\text{com}, i, r_i^*, \pi_i) \leftarrow \tilde{\mathcal{P}}(\text{crs}) \\ r \leftarrow \text{Open}(1^k, \text{crs}, \text{com}) \end{array} \right] \leq \text{negl}(\lambda).$$

Computationally Hiding: *We require that for all polynomial $k = k(\lambda)$ and $I \subseteq [k]$, the two following distributions are computationally indistinguishable:*

$$\begin{array}{c} (\text{crs}, \text{com}, I, r_I, \pi_I, r_I) \\ \overset{c}{\approx} \\ (\text{crs}, \text{com}, I, r_I, \pi_I, r'_I), \end{array}$$

where $\text{crs} \leftarrow \text{Setup}(1^\lambda, 1^k)$, $(\text{com}, r, \pi_{[k]}) \leftarrow \text{GenBits}(\text{crs})$ and $r' \overset{\$}{\leftarrow} \{0, 1\}^k$.

Designated-Verifier Hidden-Bits Generator. We define the Designated-Verifier version of a Hidden-Bits Generator (DV-HBG) similarly, but with the following differences:

- $\text{Setup}(1^\lambda, 1^k)$: Now outputs (crs, td) , where td is a trapdoor associated to the crs ;
- $\text{Verify}(\text{crs}, \text{td}, \text{com}, i, r_i, \pi_i)$ takes the trapdoor td as an additional input, and outputs accept or reject as before;
- For *Statistical Binding*, the cheating prover $\tilde{\mathcal{P}}$ can now make a polynomial number of oracle queries to $\text{Verify}(\text{crs}, \text{td}, \dots)$. We require that for any such $\tilde{\mathcal{P}}$:

$$\Pr \left[\begin{array}{l} r_i^* \neq r_i \\ \wedge \text{Verify}(\text{crs}, \text{td}, \text{com}, i, r_i^*, \pi_i) = \text{accept} \end{array} : \begin{array}{l} (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^k) \\ (\text{com}, i, r_i^*, \pi_i) \leftarrow \tilde{\mathcal{P}}^{\text{Verify}(\text{crs}, \text{td}, \dots)}(\text{crs}) \\ r \leftarrow \text{Open}(1^k, \text{crs}, \text{com}) \end{array} \right] \leq \text{negl}(\lambda).$$

- For *Computational Hiding*, we require that the distributions are indistinguishable given the associated trapdoor td :

$$(\text{crs}, \text{td}, \text{com}, I, r_I, \pi_I, r_I) \overset{c}{\approx} (\text{crs}, \text{td}, \text{com}, I, r_I, \pi_I, r'_I),$$

where $(\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^k)$, $(\text{com}, r, \pi_{[k]}) \leftarrow \text{GenBits}(\text{crs})$ and $r' \overset{\$}{\leftarrow} \{0, 1\}^k$.

4 From Hidden-Bits Generator to NIZKs

We now prove that we can combine any (DV-)HBG with a NIZK in the Hidden-Bits model to get a (Reusable DV-)NIZK in the CRS model. Recall that our

basic notion of NIZKs considered selective version of ZK where the statement to be proven is chosen prior to seeing the CRS. In Section 4.1 we will then extend our compiler to the adaptive ZK setting.

Theorem 4.1. *Suppose there exists a Hidden-Bits Generator, then there exists a publicly verifiable NIZK. Suppose there exists a designated-verifier Hidden-Bits Generator (DV-HBG), then there exists a reusable designated-verifier NIZK (reusable DV-NIZK).*

For simplicity, we first consider the publicly verifiable version of Theorem 4.1, the minor differences that are needed to extend it to the designated-verifier setting are discussed in the full version of the paper.

Construction. Let L be an NP language and n be an integer. Let $(\text{Setup}^{\text{BG}}, \text{GenBits}, \text{Verify})$ be a hidden-bits generator (Definition 3.1), where $|\mathcal{COM}| = |\mathcal{COM}(\lambda)| \leq 2^{k^\delta p(\lambda)}$ for some polynomial p and constant $\delta < 1$. (where k is the number of hidden bits generated).

Given a NIZK in the hidden-bits model for L using $k' = k'(\lambda, n)$ hidden bits (which exists unconditionally by Theorem 2.8), by Remark 2.7, there exists, for all polynomial $q(\lambda, n)$ (which we will set later), a NIZK in the hidden-bits model $(\mathcal{P}^{\text{HB}}, \mathcal{V}^{\text{HB}})$ using $k = k' \cdot q(\lambda, n)$ hidden bits with soundness-error $2^{-q(\lambda, n)} \cdot \text{negl}(\lambda)$.

Consider the following candidate NIZK $(\text{Setup}^{\text{ZK}}, \mathcal{P}, \mathcal{V})$ in the CRS model:

- $\text{Setup}^{\text{ZK}}(1^\lambda, 1^n)$: Compute $\text{crs}^{\text{BG}} \leftarrow \text{Setup}^{\text{BG}}(1^\lambda, 1^k)$, sample $s \xleftarrow{\$} \{0, 1\}^k$ and output:

$$\text{crs} = (\text{crs}^{\text{BG}}, s);$$

- $\mathcal{P}(\text{crs}, x, w)$: Compute $(\text{com}, r^{\text{BG}}, \pi_{[k]}) \leftarrow \text{GenBits}(\text{crs}^{\text{BG}})$. Set $r_i = r_i^{\text{BG}} \oplus s_i$ for all $i \in [k]$, and run the hidden-bits prover to get $(I \subseteq [k], \pi^{\text{HB}}) \leftarrow \mathcal{P}^{\text{HB}}(r, x, w)$. Output:

$$\Pi = (I, \pi^{\text{HB}}, \text{com}, r_I, \pi_I).$$

- $\mathcal{V}(\text{crs}, x, \Pi = ((I, \pi^{\text{HB}}, \text{com}, r_I, \pi_I)))$: Compute $r_i^{\text{BG}} = r_i \oplus s_i$ for all $i \in [k]$. Accept if for all $i \in I$, $\text{Verify}(\text{crs}^{\text{BG}}, \text{com}, i, r_i^{\text{BG}}, \pi_i)$ accepts, and if $\mathcal{V}^{\text{HB}}(I, r_I, x, \pi^{\text{HB}})$ also accepts.

We refer the reader to the full version of the paper for a proof that $(\text{Setup}^{\text{ZK}}, \mathcal{P}, \mathcal{V})$ is a NIZK, and how to extend this construction to the designated-verifier setting.

4.1 Adaptive ZK

Our default definition of (reusable designated-verifier) NIZKs considers a *selective* version of the zero-knowledge property, where the statement x is chosen before the CRS. We also consider a stronger *adaptive* zero-knowledge property, where the statement x can depend adaptively on the CRS. Let us begin by defining adaptive ZK.

Definition 4.2 (Adaptive ZK). A (reusable designated-verifier) NIZK satisfies adaptive Zero-Knowledge (adaptive ZK) if the following holds. We require that there exists a stateful PPT simulator Sim such that for any stateful PPT adversary \mathcal{A} the two following distributions are computationally indistinguishable:

$$\begin{array}{ll}
 \text{EXP}^{\text{Real}}(1^\lambda) : & \text{EXP}^{\text{Ideal}}(1^\lambda) : \\
 1^n \leftarrow \mathcal{A}(1^\lambda) & 1^n \leftarrow \mathcal{A}(1^\lambda) \\
 (\text{crs}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n) & (\text{crs}, \text{td}) \leftarrow \text{Sim}(1^\lambda, 1^n) \\
 (x, w) \leftarrow \mathcal{A}(\text{crs}, \text{td}) & (x, w) \leftarrow \mathcal{A}(\text{crs}, \text{td}) \\
 \text{where } (x, w) \in R_L, |x| = n & \text{where } (x, w) \in R_L, |x| = n \\
 \pi \leftarrow \mathcal{P}(\text{crs}, x, w) & \pi \leftarrow \text{Sim}(x) \\
 \text{Output } \mathcal{A}(\pi) & \text{Output } \mathcal{A}(\pi)
 \end{array}$$

The compiler of Theorem 4.1 can be extended to the adaptive setting:

Theorem 4.3. Suppose there exists a Hidden-Bits Generator, then there exists a publicly verifiable NIZK with adaptive ZK security. Suppose there exists a designated-verifier Hidden-Bits Generator (DV-HBG), then there exists a reusable designated-verifier NIZK (DV-NIZK) with adaptive ZK security.

We refer to the full version of the paper for the proof of Theorem 4.3.

5 Designated-Verifier Hidden-Bits Generator from CDH

Let $(\mathbb{G}, p, g) \leftarrow \text{GroupGen}(1^\lambda)$ be a prime-order group generator so that \mathbb{G} is a group of prime order p , with a generator g . Let hc be the corresponding Goldreich-Levin [GL89] hard-core bit. Let us define the following hidden-bits generator:

- $\text{Setup}(1^\lambda, 1^k)$: Let $(\mathbb{G}, p, g) \leftarrow \text{GroupGen}(1^\lambda)$. For all $i \in [k]$, pick random $a_i, b_i \xleftarrow{\$} \mathbb{Z}_p$ and $h_i \xleftarrow{\$} \mathbb{G}$ and compute:

$$f_i = h_i^{a_i} \cdot g^{b_i}.$$

Sample some random coins γ matching the randomness used by $\text{hc}(\cdot)$. Output:

$$\left(\text{crs} = (\mathbb{G}, \{(h_i, f_i)\}_{i \in [k]}, \gamma), \text{td} = \{(a_i, b_i)\}_{i \in [k]} \right).$$

- $\text{GenBits}(\text{crs})$: Pick a random $y \leftarrow \mathbb{Z}_p$, and compute for all $i \in [k]$: $t_i = h_i^y$ and $u_i = f_i^y$. Output:

$$\begin{aligned}
 \text{com} &= s = g^y, \\
 \{r_i &= \text{hc}(t_i; \gamma)\}_{i \in [k]}, \\
 \{\pi_i &= (u_i, t_i)\}_{i \in [k]}.
 \end{aligned}$$

- Verify($\text{crs}, \text{td} = \{(a_i, b_i)\}, \text{com} = s, i, r_i, \pi_i = (u_i, t_i)$) : Compute:

$$\rho_i = t_i^{a_i} \cdot s^{b_i},$$

and accept if and only if $\rho_i = u_i$, and $r_i = \text{hc}(t_i; \gamma)$.

Theorem 5.1. *The triple (Setup, GenBits, Verify) is a Designated-Verifier Hidden-Bits Generator under CDH.*

We refer to the full version of the paper for a proof of Theorem 5.1.

Combining Theorems 4.3 and 5.1, and Remark 2.5, we obtain the following:

Theorem 5.2 (Reusable DV-NIZK from CDH). *Under the CDH assumption, there exists a reusable DV-NIZK for all NP with statistical soundness, and adaptive, multi-theorem zero-knowledge (Definitions 2.2, 4.2).*

6 Malicious-Designated-Verifier NIZKs

In this section we consider a strengthening of designated-verifier NIZKs to the malicious-designated-verifier setting (MDV-NIZK). In this setting, the trusted setup consists solely of a common random string (CRS). Given the CRS, the (potentially malicious) verifier generates a public key pk along with a secret key sk . The rest of the protocol is otherwise similar to the previous setting: any prover can use the CRS along with the newly generated public key to build non-interactive proofs of (many) NP statements, which can be verified using the corresponding secret key. The main difference is that we require zero-knowledge to hold against malicious verifiers, who can generate arbitrarily malformed public keys pk .

6.1 More Preliminaries

Reusable Malicious-Designated-Verifier NIZK

Definition 6.1 (Reusable Malicious-Designated-Verifier NIZK (MDV-NIZK)). *Let L be an NP language with witness relation R_L . A Reusable Malicious-Designated-Verifier NIZK (MDV-NIZK) for L is a tuple of PPT algorithms (Setup, KeyGen, \mathcal{P} , \mathcal{V}) where:*

- Setup($1^\lambda, 1^n$): outputs a common random string crs ;
- KeyGen(crs): outputs a public key pk along with an associated secret key sk ;
- $\mathcal{P}(\text{crs}, \text{pk}, x, w)$: outputs a proof π ;
- $\mathcal{V}(\text{crs}, \text{sk}, \text{pk}, x, \pi)$: Outputs **accept** or **reject**.

We require those algorithms to satisfy the same completeness and statistical soundness properties as Reusable DV-NIZKs (see Definition 2.2) with direct modifications to match the new syntax above, where now (crs, pk) together act in place of what was previously just the crs . The requirement for zero-knowledge is strengthened to the following:

Malicious Zero-Knowledge (Adaptive): We require that there exists a PPT simulator Sim such that for any PPT stateful adversary \mathcal{A} , the two following distributions are computationally indistinguishable:

$$\begin{array}{ll}
 \text{EXP}^{\text{Real}}(1^\lambda) : & \text{EXP}^{\text{Ideal}}(1^\lambda) : \\
 1^n \leftarrow \mathcal{A}(1^\lambda) & 1^n \leftarrow \mathcal{A}(1^\lambda) \\
 \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^n) & \text{crs} \leftarrow \text{Sim}(1^\lambda, 1^n) \\
 (x, w, \text{pk}) \leftarrow \mathcal{A}(\text{crs}) & (x, w, \text{pk}) \leftarrow \mathcal{A}(\text{crs}) \\
 \text{where } (x, w) \in R_L, |x| = n & \text{where } (x, w) \in R_L, |x| = n \\
 \pi \leftarrow \mathcal{P}(\text{crs}, \text{pk}, x, w) & \pi \leftarrow \text{Sim}(\text{pk}, x) \\
 \text{Output } \mathcal{A}(\pi) & \text{Output } \mathcal{A}(\pi)
 \end{array}$$

Remark 6.2 (Single-Theorem vs. Multi-Theorem Zero-Knowledge). As in Definition 2.2, the definition above only captures *single-theorem* zero-knowledge. However the same ‘‘Or trick’’ of [FLS99] as in Remark 2.5 allows to generically compile any MDV-NIZK with single-theorem, adaptive (resp. selective) ZK into one satisfying *multi-theorem*, adaptive (resp. selective) ZK.

One-More CDH

We will use in this section a strengthening of the CDH assumption called *One-More CDH*. Intuitively, it states that given a set of challenge elements $\{h_j = g^{b_j}\}$ and the ability to make m queries to an oracle that raises arbitrary elements to some hidden exponent $a \in \mathbb{Z}_p$, it is hard to guess *more* than m of the values $h_j^{a_j} = g^{ab_j}$.

Definition 6.3 (One-More Computational Diffie-Hellman assumption (One-More CDH)). Let GroupGen be a group generator. Let $\ell = \ell(\lambda)$ and $m = m(\lambda)$ be polynomials. Consider, for any PPT \mathcal{A} , the following experiment:

- $$\begin{array}{l}
 \text{Exp}^{\text{One-More CDH}}(1^\lambda) \\
 1. (\mathbb{G}, p, g) \leftarrow \text{GroupGen}(1^\lambda) \\
 2. (g^a, \{g^{b_i}\}_{i \leq \ell}) \xleftarrow{\$} \mathbb{G}^{1+\ell} \\
 3. L \leftarrow \mathcal{A}^{\mathcal{O}_a(\cdot)}(\mathbb{G}, p, g, g^a, \{g^{b_i}\}_{i \leq \ell}) \\
 4. \text{Output } 1 \text{ if } \exists i_1 < \dots < i_{m+1} \in [\ell] \text{ such that } \forall j \leq m+1, g^{a \cdot b_{i_j}} \in L; \\
 \quad \text{Otherwise output } 0,
 \end{array}$$

where the oracle \mathcal{O}_a takes as input a group element $h \in \mathbb{G}$ and outputs h^a .

We say that the One-More CDH assumption holds relative to GroupGen^9 if for all PPT algorithm \mathcal{A} making at most m queries to \mathcal{O}_a , we have:

$$\Pr[\text{Exp}^{\text{One-More CDH}}(1^\lambda) = 1] \leq \text{negl}(\lambda).$$

Remark 6.4 (One-More CDH in Prior Works). A variety of previous works defined assumptions similar to the one above. To our knowledge, the first of this kind was introduced in the context of blind signatures in [Bol03], following the steps of [BNPS03] who first introduced One-More variants of the RSA and Discrete Log assumptions. More recently, another variant was used in the context of Oblivious PRFs (e.g. [JKK14]). The variant of [Bol03] requires the adversary to output one *single* guess for each target index $j \in J$, as opposed to a list of candidates L . As the adversary a-priori cannot test himself whether an element is correct, this makes it more difficult for the adversary to win the game and therefore the assumption of [Bol03] is *weaker* than our version in Definition 6.3. In [JKK14], on the other hand, the adversary is also given oracle access to a procedure that tests whether an element is a correct CDH output associated to some target index, but still has to output a single element for each target index. A direct reduction shows that this assumption is *at least as strong* as our variant: an adversary in the latter can call the oracle of [JKK14] on the whole list L to recover the matching indices.

Somewhere-Equivocable PRFs (SEPRFs)

We recall here the concept of Somewhere-Equivocable pseudorandom function (SEPRF)s, introduced in [HJO⁺16]. This is a function $\text{PRF}(K, \cdot)$ with two modes of generating a key. There is the standard key generation algorithm which generates a key K honestly. In addition, there is a way to generate a key K' that leaves a “hole” at some particular point x^* but defines the PRF output at all other points; later one can “plug the hole” to any value r by creating a key K^* which agrees with K' on all values other than x^* but on x^* it outputs r . For any x^* and a random r one cannot distinguish between an honestly generated key K and the key K^* created as above. Intuitively, the second mode of key generation ensures that the function $\text{PRF}(K^*, \cdot)$ outputs a truly random and independent value on some specific point x^* .

Definition 6.5 (1-Somewhere-Equivocable PRFs (1-SEPRFs) [HJO⁺16]).

A 1-Somewhere-Equivocable PRF (*1-SEPRF*) with input size s and output size d is a tuple of PPT algorithms $(\text{ObvGen}, \text{PRF}, \text{Sim}_1, \text{Sim}_2)$:

⁹ Later, we will also use the (mild) additional property that one can *obviously* sample uniform group elements in \mathbb{G} , so that the One-More CDH assumption holds even given the random coins used to sample the group elements in Step 2. (and in particular a and the b_i 's should be computationally hidden). Note that most standard groups (such as \mathbb{Z}_p^* or elliptic curves) allow to do so. Looking ahead, if such a property does not hold, the resulting MDV-NIZK (Theorem 6.9) will use a common *reference* string instead.

- $\text{ObvGen}(1^\lambda)$: outputs a key K such that $\text{PRF}(K, \cdot)$ maps $\{0, 1\}^s$ to $\{0, 1\}^d$;
- $\text{Sim}_1(x^*)$: on input $x^* \in \{0, 1\}^s$, outputs a key K and a state state ;
- $\text{Sim}_2(\text{state}, r)$: on input $r \in \{0, 1\}^d$, outputs a key K' .

such that the following properties hold:

Correctness: We have that for all $x^* \in \{0, 1\}^s$ and $r \in \{0, 1\}^d$, if $(K, \text{state}) \xleftarrow{\$} \text{Sim}_1(x^*)$ and $K' \xleftarrow{\$} \text{Sim}_2(\text{state}, r)$, then:

$$\begin{aligned} \text{PRF}(K, x) &= \text{PRF}(K', x) & \text{if } x \neq x^* \\ \text{PRF}(K', x^*) &= r. \end{aligned}$$

Equivocation security: For all PPT adversary \mathcal{A} we have:

$$\left| \Pr \left[\begin{array}{c} x^* \xleftarrow{\$} \mathcal{A}(1^\lambda) \\ K \xleftarrow{\$} \text{ObvGen}(1^\lambda) \\ \mathcal{A}(K) = 1 \end{array} \right] - \Pr \left[\begin{array}{c} x^* \xleftarrow{\$} \mathcal{A}(1^\lambda), r^* \xleftarrow{\$} \{0, 1\}^d \\ (K, \text{state}) \leftarrow \text{Sim}_1(x^*) \\ K' \xleftarrow{\$} \text{Sim}_2(\text{state}, r^*) \\ \mathcal{A}(K') = 1 \end{array} \right] \right| \leq \text{negl}(\lambda).$$

Claim ([HJO⁺16]). Assuming one-way functions exist, there exist 1-SEPRFs, with key size $\mathcal{O}(s \cdot d \cdot \lambda)$.

6.2 Reusable Malicious-Designated-Verifier HBG (MDV-HBG)

To define a reusable Malicious-Designated-Verifier Hidden-Bits Generator (MDV-HBG), we extend the definition of a DV-HBG in a manner analogous to the difference between DV-NIZKs and MDV-NIZKs. Namely, instead of having a trusted setup that generates a public crs along with a secret key sk for the verifier, we now only have the setup algorithm generate the crs and allow the (potentially malicious) verifier to generate pk, sk on his own via a new KeyGen algorithm. Furthermore, we want to ensure that the generated hidden bits only depend on crs but not on pk ; only the openings of the hidden bits can depend on pk .

Definition 6.6 (Reusable Malicious-Designated-Verifier HBG (MDV-HBG)). A Reusable Malicious-Designated-Verifier HBG is a tuple of PPT algorithms $(\text{Setup}, \text{KeyGen}, (\text{GenBits.Commit}, \text{GenBits.Prove}), \text{Verify})$:

- $\text{Setup}(1^\lambda, 1^k)$: outputs a common random string crs.
- $\text{KeyGen}(\text{crs})$: outputs a public key pk with an associated secret key sk .
- $\text{GenBits}(\text{crs}, \text{pk})$ is now split into two sub-procedures:
 - $\text{GenBits.Commit}(\text{crs})$: on input a crs, outputs a commitment com , some bits $r \in \{0, 1\}^k$ and a state state .
 - $\text{GenBits.Prove}(\text{crs}, \text{pk}, \text{state})$: on input a public key pk , a crs and a state state , produces proofs $\{\pi_i\}_{i \in k}$.

- It outputs $(\text{com}, r, \{\pi_i\}_{i \in [k]})$.
- $\text{Verify}(\text{crs}, \text{sk}, \text{com}, i, r_i, \pi_i)$: Outputs **accept** or **reject**.

We require an MDV-HBG to satisfy the following properties. The first three (correctness, succinctness of the commitments and statistical binding), are direct adaptations of Definition 3.1 to the new syntax:

Correctness: We require that for every polynomial $k = k(\lambda)$ and for all $i \in [k]$, we have:

$$\Pr \left[\text{Verify}(\text{crs}, \text{sk}, \text{com}, i, r_i, \pi_i) = \text{accept} : \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^k) \\ (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{crs}) \\ (\text{com}, r, \pi_{[k]}) \leftarrow \text{GenBits}(\text{crs}, \text{pk}) \end{array} \right] = 1.$$

Succinct Commitment: We require that there exists some set $\mathcal{COM}(\lambda)$ and some constant $\delta < 1$ such that $|\mathcal{COM}(\lambda)| \leq 2^{k^\delta \text{poly}(\lambda)}$, and such that for all crs output by $\text{Setup}(1^\lambda, 1^k)$ and all com output by $\text{GenBits}(\text{crs})$ we have $\text{com} \in \mathcal{COM}(\lambda)$. Furthermore, we require that for all $\text{com} \notin \mathcal{COM}(\lambda)$, $\text{Verify}(\text{crs}, \text{com}, \cdot, \cdot)$ always outputs **reject**.

Statistical Binding: There exists an (inefficient) deterministic algorithm $\text{Open}(1^k, \text{crs}, \text{com})$ such that for every polynomial $k = k(\lambda)$, on input 1^k , crs and com , the algorithm outputs r such that for every (potentially unbounded) cheating prover $\tilde{\mathcal{P}}$:

$$\Pr \left[\begin{array}{l} r_i^* \neq r_i \\ \text{Verify}(\text{crs}, \text{sk}, \text{com}, i, r_i^*, \pi_i) = \text{accept} \end{array} : \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^k) \\ (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{crs}) \\ (\text{com}, i, r_i^*, \pi_i) \leftarrow \tilde{\mathcal{P}}(\text{crs}, \text{pk}) \\ r \leftarrow \text{Open}(1^k, \text{crs}, \text{com}) \end{array} \right] \leq \text{negl}(\lambda).$$

The main conceptual difference with Definition 3.1 comes from the computational hiding property, which now captures security against malicious verifiers:

Computationally Hiding against Malicious Verifiers: Consider, for an integer k , a bit b , and a stateful PPT adversary \mathcal{A} , the following experiment:

- $\text{Exp}^{\text{Hiding}, b}(1^\lambda, 1^k)$
0. $I \subseteq [k] \leftarrow \mathcal{A}(1^k)$
 1. $\text{crs} \leftarrow \text{Setup}(1^\lambda, 1^k)$
 2. $\text{pk} \leftarrow \mathcal{A}(\text{crs})$
 3. Compute $(\text{com}, r, \{\pi_i\}_{i \in [k]}) \leftarrow \text{GenBits}(\text{crs}, \text{pk})$.
- Set for all $i \notin I$: $\begin{cases} \rho_i = r_i & \text{if } b = 0; \\ \rho_i \xleftarrow{\$} \{0, 1\} & \text{otherwise.} \end{cases}$
4. Output : $\beta \leftarrow \mathcal{A}(\text{crs}, \text{com}, I, r_I, \pi_I, \{\rho_i\}_{i \notin I})$

We require that for all polynomial $k = k(\lambda)$ and stateful PPT adversary \mathcal{A} :

$$\left| \Pr \left[\text{Exp}^{\text{Hiding},0}(1^\lambda) = 1 \right] - \Pr \left[\text{Exp}^{\text{Hiding},1}(1^\lambda) = 1 \right] \right| \leq \text{negl}(\lambda).$$

6.3 Reusable MDV-NIZK from MDV-HBG

We present here an analogue to Theorem 4.1 in the malicious-verifier setting.

Theorem 6.7. *Suppose there exists a MDV-HBG. Then there exists a reusable MDV-NIZK with adaptive ZK security.*

The proof of Theorem 6.7 is a simple adaptation of the one of Theorem 4.1. We refer the reader to the full version of the paper for more details.

6.4 MDV-HBG from One-More CDH

Notation. Let d, k and ℓ be integers, where ℓ is a power-of-two. Given a function $\varphi : [k] \rightarrow [\ell]^d$ and some index $i \in [k]$, we define, for some vector \mathbf{u} of dimension ℓ , the vector:

$$\mathbf{u}_{\varphi(i)} := (u_{\varphi(i)_1}, \dots, u_{\varphi(i)_d}).$$

In other words, we can think of $\varphi(i)$ as a set of *neighbors* of vertex $i \in [k]$ in the bipartite (multi-)graph $([k], [\ell])$. Furthermore if the vertices $j \in [\ell]$ are labelled with some element u_j , then $u_{\varphi(i)}$ denotes the *list* of labels associated to neighbors of i . Note that vertices in $[k]$ have d neighbors in $[\ell]$ (where there can be multiple occurrences of the same edge). We naturally extend this definition for *sets* of indices: for $I \subseteq [k]$, we define

$$\mathbf{u}_{\varphi(I)} := (u_{\varphi(i)_1}, \dots, u_{\varphi(i)_d})_{i \in I}.$$

Let hc be the Goldreich-Levin [GL89] hard-core bit (which, on input a bit-string $x \in \{0, 1\}^L$, uses randomness $r \xleftarrow{\$} \{0, 1\}^L$ and outputs $\text{hc}(x; r) := (\langle x, r \rangle, r)$).

Construction. Let $(\mathbb{G}, p, g) \leftarrow \text{GroupGen}(1^\lambda)$ be a prime-order group generator so that \mathbb{G} is a group of prime order p , with a generator g . For $\lambda, k \in \mathbb{N}$, let $\ell = \ell(\lambda, k)$ be the least power-of-two greater than $3k\lambda$ (i.e. $\ell = 2^{\lceil \log(3k\lambda) \rceil}$), and let $d = \lambda$. Let $(\text{ObvGen}, \text{PRF}, \text{Sim}_1, \text{Sim}_2)$ be a 1-SEPRF (as defined in Section 6.1) where $\text{ObvGen}(1^\lambda)$ outputs keys K such that $\text{PRF}(K, \cdot)$ maps $\{0, 1\}^{\lceil \log k \rceil}$ to $\{0, 1\}^{d \cdot \log \ell}$ (and in particular maps $[k]$ to $[\ell]$).

Let us define the following hidden-bits generator:

- **Setup** $(1^\lambda, 1^k)$: Let $(\mathbb{G}, p, g) \leftarrow \text{GroupGen}(1^\lambda)$. For all $j \in [\ell]$, pick $h_j \xleftarrow{\$} \mathbb{G}$.
Output:

$$\text{crs} = (\mathbb{G}, \{h_j\}_{j \in [\ell]}).$$

- KeyGen(crs): For all $j \in [\ell]$, pick random $a_j, b_j \xleftarrow{\$} \mathbb{Z}_p$, compute:

$$f_j = h_j^{a_j} \cdot g^{b_j},$$

and output:

$$\begin{aligned} \text{pk} &= \{f_j\}_{j \in [\ell]}, \\ \text{sk} &= \{(a_j, b_j)\}_{j \in [\ell]}. \end{aligned}$$

- GenBits(crs, pk):

- GenBits.Commit(crs): Pick a random $y \leftarrow \mathbb{Z}_p$ and set $s = g^y$. Compute for all $j \in [\ell]$: $t_j = h_j^y$. Sample some random coins γ matching the randomness used by $\text{hc}(\cdot)$ taking as input (the bit-representation of) elements in \mathbb{G}^d . Sample $K \leftarrow \text{ObvGen}(1^\lambda)$. Parsing the output of $\text{PRF}(K, \cdot)$ as d blocks of $\log \ell$ bits, this defines for all $i \in [k]$:

$$\varphi(i) := (\text{PRF}(K, i)_1, \dots, \text{PRF}(K, i)_d) \in [\ell]^d. \quad (1)$$

Compute for all $i \in [k]$: $r_i = \text{hc}((\mathbf{h}^y)_{\varphi(i)}; \gamma)$, where we recall that by definition $(\mathbf{h}^y)_{\varphi(i)} = (h_{\text{PRF}(K, i)_1}^y, \dots, h_{\text{PRF}(K, i)_d}^y)$. Output:

$$\begin{aligned} \text{com} &= (s, \gamma, K), \\ \{r_i\}_{i \in [k]}, \\ \text{state} &= (y, K). \end{aligned}$$

- GenBits.Prove(crs, pk, state): Parse pk as $\{f_j\}_{j \in [\ell]}$. The key K in state defines a function φ as per Equation 1. Compute for all $j \in [\ell]$: $t_j = h_j^y$ and $u_j = f_j^y$. Compute for all $i \in [k]$:

$$\pi_i = \{(t_j, u_j)\}_{j \in \varphi(i)}.$$

Output:

$$(\text{com}, r, \{\pi_i\}_{i \in [k]}).$$

- Verify(crs, sk, com, i, r_i, π_i): Parse $\text{sk} = \{(a_j, b_j)\}_{j \in [\ell]}$, $\text{com} = (s, \gamma, K)$, $\pi_i = \{(t_j, u_j)\}_{j \in \varphi(i)}$. Compute for $j \in \varphi(i)$ (where $\varphi(i)$ is defined as per Equation 1):

$$\rho_j = t_j^{a_j} \cdot s^{b_j},$$

and accept if and only if $\rho_j = u_j$ for all $j \in \varphi(i)$, and $r_i = \text{hc}(\{\mathbf{t}\}_{\varphi(i)}; \gamma)$.

Theorem 6.8. *Suppose that $(\text{ObvGen}, \text{PRF}, \text{Sim}_1, \text{Sim}_2)$ is a 1-SEPRF (Definition 6.5). Then, assuming the One-More CDH assumption holds (Definition 6.3), $(\text{Setup}, \text{GenBits}, \text{Verify})$ is a reusable Malicious-Designated-Verifier Hidden-Bits Generator (Definition 6.6).*

We refer the reader to the full version of this paper for a proof of Theorem 6.8.

Combining Claim 6.1, Theorems 6.7 and 6.8, and Remark 6.2, we obtain the following:

Theorem 6.9 (MDV-NIZK from One-More CDH). *Under the One-More CDH assumption (Definition 6.3), there exists a MDV-NIZK for all NP (Definition 6.1) with statistical soundness, and adaptive, multi-theorem zero-knowledge.*

7 Extensions

We informally describe two simple extensions of our construction.

Unbounded Statement Size. In our construction of (reusable DV-)NIZKs, we need to have a bound n on the size of the statements that can be proved and the size of the CRS depends on n . Ideally, we would have a fixed-size CRS which allows us to prove statements of arbitrary size. Indeed, we can achieve this using non-interactive statistically-binding commitments in the CRS model, which exist assuming OWFs [Nao90,Nao91]. Let us fix 3SAT as the NP-complete language. To prove that some 3CNF is satisfiable the prover commits to the satisfying assignments one variable at a time. Then he uses a (reusable DV-)NIZK scheme for each clause separately to show that the 3 relevant committed values satisfy the clause. Note that the size of the statements being proved by the underlying (reusable DV-)NIZK is independent of the size of the actual 3CNF formula. Therefore the above technique bootstraps a (reusable DV-)NIZK for statements of some fixed size which depends only on the security parameter to construct a (reusable DV-)NIZK for statements of arbitrary size.

Proof of Knowledge. While our basic construction is not a proof-of-knowledge it is easy to generically add this property assuming the existence of public-key encryption (PKE). We can add a public-key com of a PKE scheme to the CRS and have the prover encrypt the witness under com and then use the (reusable DV-)NIZK to prove that the ciphertext is an encryption of a valid witness for the statement. The extractor would choose com along with a corresponding decryption key sk and use it to extract the witness.

Acknowledgments

Research supported by NSF grants CNS-1314722, CNS-1413964, CNS-1750795 and the Alfred P. Sloan Research Fellowship. The second author was supported in part by the Israeli Science Foundation (Grant No. 1262/18). We thank Geoffroy Couteau, Dennis Hofheinz, Shuichi Katsumata, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa for sharing their manuscripts [CH19,KNYY19] and for helpful discussions.

References

- BFM88. Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th Annual ACM Symposium on Theory of Computing*, pages 103–112, Chicago, IL, USA, May 2–4, 1988. ACM Press.
- BGI15. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EURO-CRYPT 2015, Part II*, volume 9057 of *Lecture Notes in Computer Science*,

- pages 337–367, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- BNPS03. Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Se-manko. The one-more-RSA-inversion problems and the security of Chaum’s blind signature scheme. *Journal of Cryptology*, 16(3):185–215, June 2003.
- Bol03. Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In Yvo Desmedt, editor, *PKC 2003: 6th International Workshop on Theory and Practice in Public Key Cryptography*, volume 2567 of *Lecture Notes in Computer Science*, pages 31–46, Miami, FL, USA, January 6–8, 2003. Springer, Heidelberg, Germany.
- BY93. Mihir Bellare and Moti Yung. Certifying cryptographic tools: The case of trapdoor permutations. In Ernest F. Brickell, editor, *Advances in Cryptology – CRYPTO’92*, volume 740 of *Lecture Notes in Computer Science*, pages 442–460, Santa Barbara, CA, USA, August 16–20, 1993. Springer, Heidelberg, Germany.
- CC18. Pyrros Chaidos and Geoffroy Couteau. Efficient designated-verifier non-interactive zero-knowledge proofs of knowledge. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 193–221, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.
- CCRR18. Ran Canetti, Yilei Chen, Leonid Reyzin, and Ron D. Rothblum. Fiat-Shamir and correlation intractability from strong KDM-secure encryption. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, pages 91–122, 2018.
- CH19. Geoffroy Couteau and Dennis Hofheinz. Towards non-interactive zero-knowledge proofs from CDH and LWE. In *EUROCRYPT*, 2019.
- CHK03. Ran Canetti, Shai Halevi, and Jonathan Katz. A forward-secure public-key encryption scheme. In Eli Biham, editor, *Advances in Cryptology – EURO-CRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 255–271, Warsaw, Poland, May 4–8, 2003. Springer, Heidelberg, Germany.
- CKS08. David Cash, Eike Kiltz, and Victor Shoup. The twin Diffie-Hellman problem and applications. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 127–145, Istanbul, Turkey, April 13–17, 2008. Springer, Heidelberg, Germany.
- CL17. Ran Canetti and Amit Lichtenberg. Certifying trapdoor permutations, re-visited. *IACR Cryptology ePrint Archive*, 2017:631, 2017.
- CS98. Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In Hugo Krawczyk, editor, *Advances in Cryptology – CRYPTO’98*, volume 1462 of *Lecture Notes in Computer Science*, pages 13–25, Santa Barbara, CA, USA, August 23–27, 1998. Springer, Heidelberg, Germany.
- CS02. Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 45–64, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer, Heidelberg, Germany.

- Dam93. Ivan Damgård. Non-interactive circuit based proofs and non-interactive perfect zero-knowledge with preprocessing. In Rainer A. Rueppel, editor, *Advances in Cryptology – EUROCRYPT’92*, volume 658 of *Lecture Notes in Computer Science*, pages 341–355, Balatonfüred, Hungary, May 24–28, 1993. Springer, Heidelberg, Germany.
- DDN91. Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *23rd Annual ACM Symposium on Theory of Computing*, pages 542–552, New Orleans, LA, USA, May 6–8, 1991. ACM Press.
- DFN06. Ivan Damgård, Nelly Fazio, and Antonio Nicolosi. Non-interactive zero-knowledge from homomorphic encryption. In Shai Halevi and Tal Rabin, editors, *TCC 2006: 3rd Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pages 41–59, New York, NY, USA, March 4–7, 2006. Springer, Heidelberg, Germany.
- DMP88. Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Non-interactive zero-knowledge proof systems. In Carl Pomerance, editor, *Advances in Cryptology – CRYPTO’87*, volume 293 of *Lecture Notes in Computer Science*, pages 52–72, Santa Barbara, CA, USA, August 16–20, 1988. Springer, Heidelberg, Germany.
- DMP90. Alfredo De Santis, Silvio Micali, and Giuseppe Persiano. Non-interactive zero-knowledge with preprocessing. In Shafi Goldwasser, editor, *Advances in Cryptology – CRYPTO’88*, volume 403 of *Lecture Notes in Computer Science*, pages 269–282, Santa Barbara, CA, USA, August 21–25, 1990. Springer, Heidelberg, Germany.
- DN00. Cynthia Dwork and Moni Naor. Zaps and their applications. In *41st Annual Symposium on Foundations of Computer Science*, pages 283–293, Redondo Beach, CA, USA, November 12–14, 2000. IEEE Computer Society Press.
- FLS99. Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM J. Comput.*, 29(1):1–28, September 1999.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology – CRYPTO’86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194, Santa Barbara, CA, USA, August 1987. Springer, Heidelberg, Germany.
- GI14. Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 640–658, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
- GL89. Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *21st Annual ACM Symposium on Theory of Computing*, pages 25–32, Seattle, WA, USA, May 15–17, 1989. ACM Press.
- GMR85. Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th Annual ACM Symposium on Theory of Computing*, pages 291–304, Providence, RI, USA, May 6–8, 1985. ACM Press.
- GMR89. Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.

- Gol01. Oded Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, Cambridge, UK, 2001.
- Gol04. Oded Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004.
- Gol11. Oded Goldreich. *Basing Non-Interactive Zero-Knowledge on (Enhanced) Trapdoor Permutations: The State of the Art*, pages 406–421. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- GOS06. Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for NP. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 339–358, St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Heidelberg, Germany.
- GR13. Oded Goldreich and Ron D. Rothblum. Enhancements of trapdoor permutations. *J. Cryptology*, 26(3):484–512, 2013.
- HJO⁺16. Brett Hemenway, Zahra Jafargholi, Rafail Ostrovsky, Alessandra Scafuro, and Daniel Wichs. Adaptively secure garbled circuits from one-way functions. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part III*, volume 9816 of *Lecture Notes in Computer Science*, pages 149–178, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.
- JKK14. Stanislaw Jarecki, Aggelos Kiayias, and Hugo Krawczyk. Round-optimal password-protected secret sharing and T-PAKE in the password-only model. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014, Part II*, volume 8874 of *Lecture Notes in Computer Science*, pages 233–253, Kaoshiung, Taiwan, R.O.C., December 7–11, 2014. Springer, Heidelberg, Germany.
- KMO90. Joe Kilian, Silvio Micali, and Rafail Ostrovsky. Minimum resource zero-knowledge proofs (extended abstract). In Gilles Brassard, editor, *Advances in Cryptology – CRYPTO’89*, volume 435 of *Lecture Notes in Computer Science*, pages 545–546, Santa Barbara, CA, USA, August 20–24, 1990. Springer, Heidelberg, Germany.
- KNYY19. Shuichi Katsumata, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa. Designated verifier/prover and preprocessing NIZKs from Diffie-Hellman assumptions. In *EUROCRYPT, 2019*.
- KW18. Sam Kim and David J. Wu. Multi-theorem preprocessing NIZKs from lattices. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 733–765, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.
- LS91. Dror Lapidot and Adi Shamir. Publicly verifiable non-interactive zero-knowledge proofs. In Alfred J. Menezes and Scott A. Vanstone, editors, *Advances in Cryptology – CRYPTO’90*, volume 537 of *Lecture Notes in Computer Science*, pages 353–365, Santa Barbara, CA, USA, August 11–15, 1991. Springer, Heidelberg, Germany.
- Nao90. Moni Naor. Bit commitment using pseudo-randomness. In Gilles Brassard, editor, *Advances in Cryptology – CRYPTO’89*, volume 435 of *Lecture Notes in Computer Science*, pages 128–136, Santa Barbara, CA, USA, August 20–24, 1990. Springer, Heidelberg, Germany.
- Nao91. Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, 1991.

- NY90. Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd Annual ACM Symposium on Theory of Computing*, pages 427–437, Baltimore, MD, USA, May 14–16, 1990. ACM Press.
- PsV06. Rafael Pass, abhi shelat, and Vinod Vaikuntanathan. Construction of a non-malleable encryption scheme from any semantically secure one. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 271–289, Santa Barbara, CA, USA, August 20–24, 2006. Springer, Heidelberg, Germany.
- SW14. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 475–484, New York, NY, USA, May 31 – June 3, 2014. ACM Press.