Statistical ZAPs from Group-Based Assumptions

Conference Paper**Author(s):**

Couteau, Geoffroy; Katsumata, Shuichi; Sadeghi, Elahe; Ursu, Bogdan

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000521611>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

Lecture Notes in Computer Science 13042, https://doi.org/10.1007/978-3-030-90459-3_16

Funding acknowledgement:

724307 - Preparing Cryptography for Modern Applications (EC)

Statistical ZAPs from Group-Based Assumptions

Geoffroy Couteau¹, Shuichi Katsumata², Elahe Sadeghi³, and Bogdan Ursu⁴

¹ CNRS, IRIF, Université de Paris
geoffroy.couteau@ens.fr

² AIST, Japan
shuichi.katsumata@aist.go.jp

³ University of Virginia, United States
elaheasadeghi@virginia.edu

⁴ Department of Computer Science, ETH Zurich, Switzerland
bogdan.ursu@inf.ethz.ch

Abstract. We put forth a template for constructing statistical ZAPs for NP. Our template compiles NIZKs for NP in the hidden-bit model (which exist unconditionally) into statistical ZAPs using a new notion of *interactive hidden-bit generator* (IHBG), which adapts the notion of hidden-bit generator to the plain model by building upon the recent notion of statistically-hiding extractable commitments. We provide a construction of IHBG from the explicit hardness of the decision Diffie-Hellman assumption (where *explicit* refers to requiring an explicit upper bound on the advantage of any polynomial-time adversary against the assumption) and the existence of statistical ZAPs for a specific simple language, building upon the recent construction of dual-mode hidden-bit generator from (Libert et al., EUROCRYPT 2020). We provide two instantiations of the underlying simple ZAP:

- Using the recent statistical ZAP for the Diffie-Hellman language of (Couteau and Hartmann, CRYPTO 2020), we obtain statistical ZAPs for NP assuming (the explicit hardness of) DDH in $\mathbb{G}_1$ and kernel-DH in $\mathbb{G}_2$ (a *search* assumption which is weaker than DDH), where $(\mathbb{G}_1, \mathbb{G}_2)$ are groups equipped with an asymmetric pairing. This improves over the recent work of (Lombardi et al., EUROCRYPT 2020) which achieved a relaxed variant of statistical ZAP for NP, under a stronger assumption.
- Using the recent work of (Couteau et al., EUROCRYPT 2020), we obtain statistical ZAPs for NP assuming the explicit hardness of DDH, together with the assumption that no efficient adversary can break the key-dependent message one-wayness of ElGamal with respect to *efficient* functions over groups of size 2^λ with probability better than $\text{poly}(\lambda)/2^{(c+o(1))\cdot\lambda}$, denoted $2^{-c\lambda}$ -OW-KDM, for a constant $c = 1/2$, in pairing-free groups. Note that the latter is a search discrete-log-style falsifiable assumption, incomparable to DDH (in particular, it is not known to imply public-key encryption).

1 Introduction

Zero-knowledge proof systems, introduced in [GMR89], are a fundamental cryptographic primitive, allowing a prover to convince a verifier of the veracity of a

statement, while not divulging anything beyond whether the statement is true. Zero-knowledge proofs have countless applications. However, they suffer from strong lower bounds on the number of rounds of interactions required in their execution: they require at least three rounds of interactions [GO94]. Therefore, the dream result of proofs that consists of a single message from the prover to the verifier (NIZKs [BFM88]) can only be achieved when assuming a trusted setup. Due to the importance of round-efficient zero-knowledge proofs, a large effort has been devoted to the construction of such proofs; yet, this trusted setup is often undesirable.

Witness-indistinguishability (WI) [FS90] is a natural relaxation of zero-knowledge, and is one of the most widely used privacy notions in proof systems. It provides the following guarantee: if there exist two witnesses (w_0, w_1) for a statement $x \in \mathcal{L}$, the verifier should not be able to distinguish an honest prover using w_0 from an honest prover using w_1 . Witness-indistinguishable proofs can replace zero-knowledge proofs in many of their applications. At the same time, their round complexity is not subject to *any* known lower bounds.

ZAPs. The work of Dwork and Naor [DN00] introduced (and constructed) ZAPs, which are two-message public-coin WI proof systems. These proof systems have several advantages: being public-coin, they are publicly verifiable (the validity of the proof can be verified solely by looking at the transcript). Furthermore, the first flow, which is just a uniformly random string, is inherently reusable for an arbitrary (polynomial) number of proofs on possibly different statements. ZAPs have proven to be important cryptographic primitives. By now, we have constructions of ZAPs from many standard assumptions, including trapdoor permutations (which is implied by factoring) [DN00], the decision linear assumption (DLIN) in bilinear maps [GOS06a], the (quasi-polynomial hardness of the) learning with error assumption [LVW19, GJJM20, BFJ⁺20], and also from more complex notions, such as indistinguishability obfuscation [BP15].

Statistical ZAP arguments. ZAPs were initially defined to satisfy unbounded soundness, and computational WI [DN00]. Statistical ZAP arguments provide the converse properties: computational soundness, and witness-indistinguishability against unbounded attackers. Unlike their computational WI counterpart, statistical ZAP arguments enjoy a very appealing property, that of *everlasting security*. Namely, soundness is an *online* security notion: as long as the prover cannot break soundness *at the time where it produces the proof*, security is guaranteed, even if the assumption it is based upon is later broken. On the other hand, WI and zero-knowledge should hold not only during the proof generation, but must continuously keep on holding in the future: compromising the assumptions underlying the WI property of proofs generated in the past at any point in the future would be sufficient to break privacy. Hence, targeting statistical privacy avoids being forced to assume the nonexistence of unforeseen cryptanalytic advances in the future.

Intriguingly, statistical ZAPs have proven much harder to construct than their computationally WI counterparts. In fact, for almost two decades after their introduction and until very recently, no construction of statistical ZAP argument was known, under any assumption. The situation changed very recently, with the construction of statistical ZAP arguments under the quasi-polynomial hardness of LWE , in two concurrent and independent works [GJJM20, BFJ⁺20]. Still, these results leave open the question of whether statistical ZAPs can be based on any of the other cryptographic assumptions that computational ZAPs can be based on, such as factoring or pairing-based assumptions.

The very recent work of [LVW20] comes very close to improving this state of affairs: they construct, from the quasi-polynomial hardness of the decision linear assumption in bilinear groups, *ZAPs with private randomness*. This primitive is essentially as versatile as a standard ZAP: while the verifier uses private coins, the proof remains publicly verifiable, and the first flow remains reusable. Yet, it still falls short of constructing true statistical ZAPs from pairing-based assumptions.

1.1 Our Result

In this work, we develop a new approach for constructing statistical ZAPs. At a high-level, our approach works by bootstrapping statistical ZAPs for simple languages to statistical ZAPs for NP, using a new primitive called *interactive hidden-bits generator* (IHBG), a plain-model variant of hidden-bits generators, which have been recently introduced in [CH19, QRW19, KNY19, LPWW20] for constructing NIZKs for NP from different assumptions. We provide two instantiations of our framework (in groups with or without pairings in the publicly verifiable setting), and obtain:

- **Statistical ZAPs in pairing groups.** A statistical ZAP argument for NP, assuming the explicit hardness⁵ of the DDH assumption in $\mathbb{G}_1$ and of the kernel Diffie-Hellman assumption in $\mathbb{G}_2$, where $(\mathbb{G}_1, \mathbb{G}_2)$ are groups equipped with an asymmetric pairing. The kernel Diffie-Hellman assumption is a standard *search* assumption in bilinear groups [MRV15, KW15], which is implied by (and is qualitatively weaker than) the DDH assumption. This improves over [LVW20], both in terms of assumption (we rely on a qualitatively weaker assumption, since [LVW20] requires DDH both in $\mathbb{G}_1$ and $\mathbb{G}_2$) and of the primitive constructed (we achieve a true statistical ZAP argument, while [LVW20] achieves a relaxed variant).
- **Statistical ZAPs in pairing-free groups.** A statistical ZAP argument from NP, assuming explicit hardness of the DDH assumption in a *pairing-*

⁵ *Explicit hardness* in [BFJ⁺20] assumes that there exists an explicit bound μ on the advantage of any polynomial time adversary against the assumption. In particular, this is a weaker requirement than superpolynomial hardness, for any arbitrarily small superpolynomial function. We note that previous works on statistical ZAPs using quasi-polynomial hardness [LVW19, GJJM20, BFJ⁺20, LVW20] can instead use explicit hardness.

free group $\mathbb{G}$ with $\log |\mathbb{G}| \approx \lambda^{1/2}$, and the assumption that no polynomial-time adversary can break the OW-KDM security of ElGamal with respect to efficient functions with success probability significantly better than $2^{-\lambda/2}$, denoted as $2^{-\lambda/2}$ -OW-KDM security. Note that the best-known attack against such OW-KDM security of ElGamal succeeds with probability $\text{poly}(\lambda) \cdot 2^{-\lambda}$. While non-standard, this is a *falsifiable* search assumption, and there is an exponential gap between the required security margin and the best known attack. Under the same KDM assumption, but assuming only the standard polynomial hardness of DDH, we also obtain statistical NIZKs (NISZKs) for NP in the common reference string (CRS) model (settling for computational NIZKs, we can further relax DDH to computational Diffie-Hellman). This builds upon and improves over the recent work of [CKU20] which constructed *computational NIZK arguments* in the CRS model, under CDH and a stronger assumption: the $2^{-3\lambda/4}$ -OW-KDM-hardness of ElGamal.

In all the above, the (decisional or kernel) Diffie-Hellman assumption can be replaced by any of its standard generalizations, namely the decisional k -Lin [HK07] and kernel k -Lin assumptions, or even more generally any assumption from the family of the (decisional or kernel) matrix Diffie-Hellman assumptions [EHK⁺13, MRV15].

Relation to [JZ21]. In a breakthrough work (very recently accepted at Eurocrypt’21), Jain and Zhengzhong have solved the long-standing open problem of basing NIZKs on a well-studied assumption in pairing-free groups (the subexponential hardness of DDH). Furthermore, their work also achieves a statistical ZAP under the same assumption. We clarify the relation of our work to theirs.

The results presented in our work have been obtained concurrently and independently of those presented in [JZ21]. However, we were made aware of the existence and content of [JZ21] while it was submitted to Eurocrypt (through private communication), and before we had completed the write-up of our paper. The techniques developed in our work are unrelated to those in [JZ21], and our results are complementary:

- We show that explicit hardness of DDH (or superpolynomial hardness of DDH, for any arbitrarily small superpolynomial function) gives statistical ZAPs in the pairing setting, and two-round statistical WI arguments in the pairing-free setting. In contrast, [JZ21] relies on the subexponential hardness of DDH (but does not need pairings to achieve public verifiability).
- In the pairing-free setting, we also rely on an exponential search discrete-log-style hardness assumption, which is incomparable to subexponential DDH (albeit the latter is of course more standard). In particular, our assumption is falsifiable, holds in the generic group model, and is not known to imply public-key encryption.

Still, although our results have been achieved concurrently and independently of theirs, we cannot (and do not) claim to achieve the first construction of a

statistical ZAP from standard group-based assumptions, since their construction precedes ours.

1.2 Our Techniques

At the heart of our results is a construction of a new cryptographic primitive, which we call an *interactive hidden-bits generator* (IHBG). At a high level, an IHBG adapts the notion of hidden-bits generator (defined in the CRS model) recently introduced and studied in [CH19, QRW19, KNY19, LPWW20] to the plain model.

Dual-Mode Hidden-Bits Generators. More precisely, our starting point is the notion of a dual-mode hidden-bits generator (HBG) from [LPWW20]. In a dual-mode HBG, there are three algorithms: a CRS generation algorithm, a hidden-bits generator **GenBits**, and a verification algorithm **VerifyBit**. Given a CRS, the prover can, using **GenBits**, produce a *short commitment* c to a long, pseudorandom hidden-bit string ρ , as well as *openings* π_i to all the bits ρ_i of ρ . Then, **VerifyBit** takes as input the CRS, a short commitment, a position i , a value ρ_i , and an opening certificate π_i , and returns 0 or 1 depending on whether the opening is accepted. A dual-mode HBG must satisfy three properties:

- (Mode indistinguishability) the CRS can be generated in one of two modes, the *hiding* and the *binding* modes, which are computationally indistinguishable.
- (Hiding) when the CRS is in hiding mode, the value ρ_i at all non-opened positions i is *statistically hidden*, even given c and openings (ρ_j, π_j) at all other positions.
- (Extractable) when the CRS is in binding mode, there exists an efficient extractor which can extract from c a string ρ such that no efficient prover can produce accepting openings for $1 - \rho_i$, for any position i .

As shown in [LPWW20], and following related transformations in [CH19, QRW19, KNY19], a dual-mode HBG can be used to convert a NIZK for NP in the hidden-bits model (which exists unconditionally) into a *dual-mode* NIZK for NP in the CRS model (with statistical zero-knowledge when the HBG is used in hiding mode, and statistical soundness otherwise). These compilation techniques have their roots in the seminal works of Feige, Lapidot, and Shamir [FLS90] and of Dwork and Naor [DN00].

Interactive Hidden-Bits Generators. The statistical NIZKs by Libert et al. [LPWW20] crucially rely on the dual-mode feature of the HBG: the statistical binding property appears unavoidable to compile a NIZK in the hidden-bits model. Hence, obtaining statistical zero-knowledge is done by generating the CRS in hiding mode, but switching it to the binding mode when analyzing soundness. Of course, this standard technique is limited to the CRS model.

In an exciting recent work [KKS18], Kalai, Khurana, and Sahai, building upon previous results and ideas from [JKKR17, BGI⁺17, KS17], introduced an

elegant and clever approach to partially emulate this “dual-mode feature” of the CRS model, but in the plain model. At a high level, they rely on statistically-hiding commitment schemes, which have the property that with some (negligible but *not too small*) probability, they will become binding and extractable; furthermore, this event cannot be detected by the committer. This in turn allows to obtain statistical privacy (e.g. statistical witness indistinguishability), while allowing to use the extractability properties to show soundness, at the cost of having to rely on assumptions which rule out even inverse-superpolynomial distinguishing advantages. This approach proved fruitful and led to a successful line of work [LVW19, GJJM20, BFJ⁺20] on building statistical ZAPs in the plain model.

Intuitively, our notion of interactive hidden-bits generator simply adapts this technique to the notion of dual-mode hidden-bits generator. That is, an IHBG is a pair $(\text{GenBits}, \text{VerifyBit})$, similar to a dual-mode HBG, with the following core differences:

- GenBits takes as input a uniformly random string, which will correspond to the verifier message in the ZAP.
- The non-opened values remain statistically hidden with overwhelming probability over the coins of VerifyBit , for any (possibly malicious) choice of the random string.
- There exists a simulator which can produce simulated random coins (indistinguishable from true random coins) such that for any (possibly malicious) prover, with some not-too-small probability μ (e.g. inverse-superpolynomial) over the coins of the simulator, the hidden bit string ρ can be extracted from c .

Defining IHBG and Statistical ZAPs for NP. The above is of course very informal. Formally defining an interactive hidden-bits generator requires some care. In particular, we observe that the definition of extractability for statistically hiding extractable commitments in [LVW19, GJJM20, BFJ⁺20] do not suffice in our setting. At a high level, this is because these definition roughly say the following: the event that the commitments become extractable happens with probability μ , and whenever this event happens, the extracted value are *guaranteed* to be correct.

However, this will not hold in our setting: given a tuple $(c, \{i, \pi_i\}_i)$ of a commitment and set of openings from a possibly malicious prover, the hidden-bit string ρ recovered by the extractor is correct if $\text{VerifyBit}(c, i, 1 - \rho_i, \pi_i) = \perp$ for all the opened positions i . Unfortunately, we can only guarantee that this will hold with overwhelming probability in our concrete construction, and not with probability 1. It turns out that, when building statistical ZAPs for NP, this is a crucial issue: in the soundness game of the ZAP construction from IHBG, the challenger will want extraction to succeed with probability μ *even when conditioning on other checks being successful*. A *guaranteed* correctness of extraction (conditioned on extraction succeeding) would ensure that this is the case, but an *overwhelming* probability of correctness does not, since conditioning on other events could arbitrarily change this probability.

To work around this issue, we adopt an approach closer in spirit to the definition of [LVW20]. We define μ -extractability as follows: an IHBG is μ -extractable if there exists an efficient simulator SimCoin and an efficient opener Open such that, for any PPT adversary $\mathcal{A}$ and any PPT distinguisher D , given simulated coins $(\tilde{r}, \tau) \leftarrow_r \text{SimCoin}$ (where τ is an associated trapdoor for the opener), and a tuple $(c, S, \rho_S^*, \{\pi_i\}_S, \text{st}) \leftarrow_r \mathcal{A}(\tilde{r})$ where c is a short commitment, S is a set of positions, ρ_S^* are the values which $\mathcal{A}$ opens the position to, the π_i are certificates of correct openings, and st is an arbitrary state, and letting $\rho \leftarrow \text{Open}(\tilde{r}, c, \tau)$, the probability p_1 that $\text{VerifyBit}(\tilde{r}, c, i, 1 - \rho_i, \pi_i)$ returns $\perp$ for all $i \in S$ and at the same time the distinguisher D , given st , outputs 1, should satisfy

$$p_1 \geq \mu(\lambda) \cdot (p_2 - \text{negl}(\lambda)),$$

where p_2 is the probability of the same event *without* the check that the procedure $\text{VerifyBit}(\tilde{r}, c, i, 1 - \rho_i, \pi_i)$ returns $\perp$ for all $i \in S$. That is, μ -extractability requires that *for any other efficient conditions that we were verifying*, the probability that these conditions are *still verified* and that simultaneously, extraction succeeded and produced a correct output, should not decrease by a factor more than μ compared to the initial probability. This strong security notion is the key to capture the intuition that the extraction should succeed with probability μ *essentially independently of everything else*.

Given this notion of μ -extractable IHBG, we provide a natural construction of statistical ZAP for NP, which follows the standard template of using the IHBG to compile an unconditional NIZK for NP in the hidden-bits model, and formally prove that the resulting construction is a ZAP.

Constructing IHBG. It remains to construct IHBG with a statistical hiding property, satisfying the strong μ -extractability notion defined above. The first natural idea is to rely on the construction of dual-mode HBG from [LPWW20], and to convert it into a plain model protocol by letting the verifier sample the CRS herself. However, this immediately runs into obstacles: nothing prevents the verifier from sampling the CRS in binding mode, breaking the statistical hiding property. To recover the statistical hiding property, we let the prover *tweak* the CRS sampled by the verifier in a way that simultaneously guarantee two things:

- With overwhelming probability over the coins of the prover, the tweaked CRS will be in hiding mode, yet
- The tweak comes from a superpolynomial-size set, and by successfully guessing the tweak in advance, a simulator can engineer the sampled CRS (in a way that is indistinguishable from sampling a CRS honestly) such that the tweaked CRS will be in *binding* mode.

To achieve these two features, we rely on an elegant linear-algebra trick. In order to explain the idea, we first recall the high-level template of the construction of dual-mode HBG described in [LPWW20]. Let m be the length of the hidden bit string. The LPWW construction works in a hard-discrete-log group $\mathbb{G}$ of order p with generator g . It has the following structure:

- The hiding CRS is $g^{\mathbf{A}}$, where $\mathbf{A}$ is a random *full-rank* matrix $\mathbf{A} \in \mathbb{Z}_p^{(m+1) \times (m+1)}$.
- The binding CRS is $g^{\mathbf{A}}$, where $\mathbf{A}$ is a random *rank-1* matrix in $\mathbb{Z}_p^{(m+1) \times (m+1)}$.

Under the DDH assumption, the two modes are indistinguishable. Let $\mathbf{a}_0, \dots, \mathbf{a}_m$ denote the columns of $\mathbf{A}$. To provide a short commitment to a pseudorandom length- m hidden bit string, the prover picks a random length- $(m+1)$ vector $\mathbf{y}$, and computes $c = g^{\mathbf{y}^\top \cdot \mathbf{a}_0}$. Then, the i -th hidden bit is defined to be $\rho_i = \text{HB}(g^{\mathbf{y}^\top \cdot \mathbf{a}_i})$, where $\text{HB}(\cdot)$ is a hardcore bit function (e.g. *a la* Goldreich-Levin). Eventually, to prove correct opening of ρ_i , given the commitment c and the CRS $g^{\mathbf{A}}$, the prover reveals $c_i = g^{\mathbf{y}^\top \cdot \mathbf{a}_i}$ and uses a NIZK to demonstrate the existence of a vector $\mathbf{y}$ such that $c = g^{\mathbf{y}^\top \cdot \mathbf{a}_0}$ and $c_i = g^{\mathbf{y}^\top \cdot \mathbf{a}_i}$ (from now on, we will call this language the LPWW language, $\mathcal{L}_{\text{LPWW}}$).

Observe that when the CRS is in binding mode, we have $\mathbf{a}_i = v_i \cdot \mathbf{a}_0$ for some value v_i (since $\mathbf{A}$ has rank 1), hence the above language becomes essentially a DDH language. Adapting existing statistical NIZKs for the DDH language suffices to guarantee extractability in binding mode. On the other hand, when the CRS is in hiding mode, where $\mathbf{A}$ has full rank, *any* number of openings (of which there is at most m) $g^{\mathbf{y}^\top \cdot \mathbf{a}_i}$ leak *statistically* no information about the unopened values (since $\mathbf{A}$ is of dimension $(m+1) \times (m+1)$). This is because for any possible choice of values for the unopened positions, there exists a unique vector $\mathbf{y}$ that coincides with all the opened and unopened values when $\mathbf{A}$ is full rank. Hence, this guarantees statistical hiding.

Now, the core idea to achieve statistical hiding and μ -extractability in our construction (where μ is some arbitrary fixed inverse-superpolynomial function) is to let the verifier sample and send $g^{\mathbf{A}}$ herself, but to let the prover *tweak* this sample as follows: let $\mathbf{I}_{m+1}$ denote the identity matrix in $\mathbb{Z}_p^{(m+1) \times (m+1)}$. The prover picks a small exponent α at random from a subset of $\mathbb{Z}_p$ of size $\approx 1/\mu$, e.g. by picking α as a random integer smaller than $[1/\mu]$, and using a natural encoding of integers in $\{0, \dots, p-1\}$ as elements of $\mathbb{Z}_p$. Then, the prover defines the tweaked CRS $g^{\mathbf{A}'}$ to be $g^{\mathbf{A} - \alpha \cdot \mathbf{I}_{m+1}}$, and uses this tweaked CRS in the dual-mode HBG construction of [LPWW20].⁶

To see why this tweak achieves exactly what we want, observe that the following holds:

- First, we show that with overwhelming probability $1 - (m+1)\mu$, the matrix $\mathbf{A}'$ has full rank. Indeed, if $\mathbf{A}'$ does *not* have full rank, it means that there is a nonzero vector $\mathbf{u}$ in the kernel of $\mathbf{A}'$. But then, $\mathbf{u} \cdot \mathbf{A}' = \mathbf{0}$ rewrites to $\mathbf{u} \cdot \mathbf{A} = \alpha \cdot \mathbf{u}$ – in equivalent terms, this means that α must be an *eigenvalue* of $\mathbf{A}$. But since $\mathbf{A}$ can have at most $m+1$ eigenvalues and α is randomly sampled from a set of size $1/\mu$, then this event can happen with probability at most $(m+1)\mu$.
- Second, we sketch why μ -extractability holds. First, the simulator will guess a value α' , and set $\mathbf{A} \leftarrow \mathbf{M} + \alpha' \cdot \mathbf{I}_{m+1}$, where $\mathbf{M}$ is a rank-1 matrix. Observe

⁶ There is an obvious additional necessary change: when proving correctness of an opening, the statistical NIZK for $\mathcal{L}_{\text{LPWW}}$ is replaced by a statistical ZAP for $\mathcal{L}_{\text{LPWW}}$.

that when the simulator guesses correctly, which happens with probability μ , it holds that $g^{\mathbf{A}'}$ is a binding CRS. Furthermore, under the assumption that no PPT adversary can distinguish DDH tuples from random tuples with probability better than $\mu \cdot \text{negl}(\lambda)$, the replacement of truly random coins by simulated coins will not be detected. Hence, when further assuming that the ZAP for $\mathcal{L}_{\text{LPWW}}$ guarantees a bound $\mu \cdot \text{negl}(\lambda)$ on the probability that a malicious PPT prover breaks soundness, we can extract with probability almost μ a correct hidden-bit string. In Section 3, we will formally prove that μ -extractability holds with respect to an arbitrary PPT distinguisher D .

Summing up, the above provides a construction of IHBG (which in turns implies statistical ZAPs for NP), assuming

- the hardness of DDH with distinguishing advantage $\mu \cdot \text{negl}(\lambda)$ for any PPT adversary and for any negligible functions μ and negl (an assumption in-between standard polynomial time hardness and superpolynomial time hardness, which is called *explicit hardness* in [BFJ⁺20]), and
- the existence of statistical ZAPs for $\mathcal{L}_{\text{LPWW}}$ with $\mu \cdot \text{negl}(\lambda)$ -soundness.

Instantiating the Statistical ZAPs for $\mathcal{L}_{\text{LPWW}}$. Looking ahead, the formal analysis of our construction actually requires a slightly exotic notion of soundness: $\mathcal{L}_{\text{LPWW}}$ is formally not a language, but a parametrized family of languages, and (adaptive) soundness must hold for parameters sampled uniformly at random from a specific subset of language parameters (which are those that correspond to $\mathbf{A}$ being of rank 1). We call a ZAP for the parameterized family of languages $\mathcal{L}_{\text{LPWW}}$ *IHBG-friendly* when it satisfies this notion of soundness. We provide two instantiations for the underlying IHBG-friendly statistical ZAP.

Using pairings. First, we observe that the recent work of Couteau and Hartmann [CH20] provides a statistical ZAP for the DDH language, which extends directly to an IHBG-friendly statistical ZAP for the $\mathcal{L}_{\text{LPWW}}$ language, under the standard kernel-DH assumption, in groups equipped with an asymmetric pairing. This leads to a statistical ZAP for NP under the explicit hardness of DDH in $\mathbb{G}_1$, and the explicit hardness of kernel-DH in $\mathbb{G}_2$, where $(\mathbb{G}_1, \mathbb{G}_2)$ are groups equipped with an asymmetric pairing.

Without pairings. Secondly, we revisit the recent construction of statistical NIZKs for the DDH language in pairing-free groups by Couteau, Katsumata, and Ursu [CKU20]. Their construction relies on the assumption that no PPT algorithm can break the one-wayness of ElGamal against key-dependent message (OW-KDM) attacks with respect to efficient functions (i.e., the assumption that no PPT adversary can recover m from an ElGamal encryption of m , even when m is some efficiently computable function of the ElGamal secret key) with probability better than $2^{-3\lambda/4+o(\lambda)}$ (note that the best known PPT attack against this assumption, in appropriate groups, succeeds with probability $2^{-\lambda+o(\lambda)}$; furthermore, the restriction of KDM hardness to *efficient* functions of the secret key

makes the assumption falsifiable “in spirit” – i.e., up to the negligible winning advantage). We denote this assumption the $2^{-3\lambda/4}$ -OW-KDM hardness of ElGamal. We adapt the CKU construction to the LPWW language. Along the way, we put forth a modification of their construction which significantly improves the underlying assumption: we only need to assume that no PPT adversary can break the OW-KDM hardness of ElGamal with probability better than $2^{-\lambda/2+o(\lambda)}$. This change directly improves the result of [CKU20]. With this instantiation, and observing that this statistical NIZK is also a statistical ZAP when the verifier can choose the CRS, we obtain a statistical ZAP for NP in pairing-free groups under the explicit hardness of DDH, and the $2^{-\lambda/2}$ -OW-KDM hardness of ElGamal (we note that the latter is incomparable to DDH: it is a search, discrete-logarithm-type assumption, which is not even known to imply public-key encryption).

1.3 A Direct Construction using Pairings

Eventually, we point out that if one is willing to rely on a stronger assumption, one of our two instantiations (the pairing-based instantiation) can be obtained from our techniques in a much more direct (and simple-in-hindsight) way, without going through the hidden-bit model. Specifically, the core idea for our IHBG construction is to modify the CRS of a dual-mode NIZK using a simple tweak, sampled from a small set by the prover, which guarantees that with overwhelming probability a maliciously sampled CRS will be in hiding mode (but it will be in binding mode in the case when the verifier guesses the tweak).

A similar tweak can be applied directly to the dual-mode NIZK of Groth, Ostrovsky, and Sahai [GOS06b] instantiated with Groth-Sahai commitments [GS08]. Briefly, a Groth-Sahai commitment is of the form $(1, g^m) \cdot \mathbf{u}^r \cdot \mathbf{v}^s$, where $\mathbf{u}, \mathbf{v}$ are two random vectors of length two, and $\cdot$ denotes the coordinate-wise product (we write $\mathbf{u}^r$ for (u_1^r, u_2^r) , where $\mathbf{u} = (u_1, u_2)$). When the vectors $(\mathbf{u}, \mathbf{v})$ are random, the commitments are perfectly hiding; when $\mathbf{v}$ is in the span of $\mathbf{u}$, they become perfectly binding. A GOS proof for circuit satisfiability, given a circuit C and a witness w such that $C(w) = 1$, works by committing to all bits of w , as well as to the bits on all wires during the evaluation of $C(w)$. Then, the proof proceeds by showing that all commitments commit to bits, that all gate relations are satisfied (which reduces to proving that a linear combination of the committed input and output bits – homomorphically computed from the commitments – is itself a bit), and that the output commitment contains 1. All these proofs can be reduced to pairing-product equations, hence can be proven with a Groth-Sahai NIZK [GS08].

Now, letting the verifier choose the CRS $(\mathbf{u}, \mathbf{v})$ themselves, the prover can sample a small tweak $z \leftarrow_r [1/\mu]$, and set the CRS to be $(\mathbf{u}', \mathbf{v}') = (\mathbf{u} \cdot (1, g^z), \mathbf{v} \cdot (1, g^z))$. For any adversarial choice of $(\mathbf{u}, \mathbf{v})$, $(\mathbf{u}', \mathbf{v}')$ will not be colinear except with negligible probability; on the other hand, with probability μ , the verifier can guess the tweak z and cause $(\mathbf{u}', \mathbf{v}')$ to be in binding mode. To make the analysis work, we need to rely on the same notion of μ -extractability which we defined previously. This direct approach leads to a statistical ZAP for NP in groups $(\mathbb{G}_1, \mathbb{G}_2)$ equipped with an asymmetric pairing, assuming the explicit

hardness of DDH in both $\mathbb{G}_1$ and $\mathbb{G}_2$, a slightly stronger assumption compared to the one we obtain when going through the hidden-bit model. While simple in hindsight, this construction was apparently missed in previous works: the recent work of [LVW20] achieved, under the same assumption, a strictly weaker result (a ZAPR argument for NP), using a considerably more involved and highly non-trivial construction.

2 Preliminaries

Due to page limitations, we provide the definitions of standard notations and cryptographic tools used throughout the paper in the full version of the paper [CKSU21].

2.1 Hardness Assumptions

Let DHGen be a deterministic algorithm that on input 1^λ returns a description $\mathcal{G} = (\mathbb{G}, p)$ where $\mathbb{G}$ is a cyclic group of prime order p . Let PGen be a deterministic algorithm that on input 1^λ returns a description $\mathcal{PG} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, p)$ where $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ are cyclic groups of prime order p equipped with a bilinear pairing operation $\bullet : \mathbb{G}_1 \times \mathbb{G}_2 \mapsto \mathbb{G}_T$. Below, we recall the definition of the decision Diffie-Hellman assumption in a cyclic group, as well as the definition of the kernel Diffie-Hellman assumption in a pairing group. Following [BFJ⁺20], we also consider the *explicit* hardness of the assumptions, where we say that an assumption has explicit μ -hardness if μ is an explicit bound on the advantage of any polynomial time adversary. Note that this notion of explicit hardness is stronger than standard polynomial hardness, but weaker than superpolynomial hardness⁷ for any superpolynomial factor.

Definition 1 (DDH Assumption). *We say that the decisional Diffie-Hellman (DDH) assumption holds relative to DHGen if for all PPT adversaries $\mathcal{A}$, it holds that $\text{Adv}^{\text{DDH}}(\mathcal{A}) \leq \text{negl}(\lambda)$, where*

$$\text{Adv}^{\text{DDH}}(\mathcal{A}) = |\Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{G}, g, g^\alpha, g^\beta, g^\gamma)] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{G}, g, g^\alpha, g^\beta, g^{\alpha\beta})]|.$$

Here, note that $\mathcal{G} \leftarrow \text{DHGen}(1^\lambda)$ and DHGen outputs a fixed group $\mathbb{G}$ per security parameter, and $g \leftarrow_r \mathbb{G}$, $\alpha, \beta, \gamma \leftarrow_r \mathbb{Z}_p$ are chosen uniformly. Furthermore, let $\mu(\lambda)$ be an efficiently computable function. We say that the μ -explicit hardness of the DDH assumption holds relative to DHGen , if $\text{Adv}^{\text{DDH}}(\mathcal{A}) \leq \mu(\lambda)$ for all PPT adversaries $\mathcal{A}$.

We now recall the definition of the kernel Diffie-Hellman assumption in a pairing group. The kernel DH assumption is a standard search assumption in bilinear groups, introduced in [MRV15] and used in several papers, e.g. [KW15]. In particular, kernel Diffie-Hellman in a group $\mathbb{G}_2$ is implied by (and is qualitatively weaker than) the DDH assumption in the same group.

⁷ We consider adversaries that run in superpolynomial time in case of superpolynomial hardness.

Definition 2 (Kernel DH Assumption). *We say that the kernel Diffie-Hellman ($\ker DH$) assumption holds relative to $PGen$ if for all PPT adversaries $\mathcal{A}$, it holds that $\text{Adv}^{\ker DH}(\mathcal{A}) \leq \text{negl}(\lambda)$, where*

$$\text{Adv}^{\ker DH}(\mathcal{A}) = \Pr \left[\begin{array}{l} \mathcal{PG} \leftarrow PGen(1^\lambda), \\ (g_1, g_2) \leftarrow_r \mathbb{G}_1 \times \mathbb{G}_2, e \leftarrow_r \mathbb{Z}_p, : (u, v) \in \ker((1, e)^\top) \wedge v \neq 0 \\ (g_1^u, g_1^v) \leftarrow \mathcal{A}(1^\lambda, \mathcal{PG}, g_1, g_2, g_2^e) \end{array} \right].$$

Furthermore, let $\mu(\lambda)$ be an efficiently computable function. We say that the μ -explicit hardness of the kernel DH assumption holds relative to $PGen$, if $\text{Adv}^{\ker DH}(\mathcal{A}) \leq \mu(\lambda)$ for all PPT adversaries $\mathcal{A}$.

To see why the above is implied by DDH in $\mathbb{G}_2$, observe that on input $(g, g^\alpha, g^\beta, g^\gamma)$, an adversary against DDH can run the kernel DH adversary on input (g_1, g, g^α) , where $g_1 \leftarrow_r \mathbb{G}_1$ and e is implicitly set as α . It then gets a vector (g_1^u, g_1^v) in $\mathbb{G}_1^2$ from the kernel DH adversary such that (u, v) is in the kernel of $(1, \alpha)$. Now, if $(g, g^\alpha, g^\beta, g^\gamma)$ is a DDH tuple, then (u, v) is also in the kernel of $(g^\beta, g^\gamma) = (g, g^\alpha)^\beta$, and this can be checked efficiently given (g_1^u, g_1^v) with the help of the pairing operation.

Remark 3 (Extensions to Matrix Diffie-Hellman). For the sake of concreteness and simplicity, we state our results in this paper in terms of the DDH and kernel DH assumptions. However, all our results can be generalized to hold under the standard generalizations of the Diffie-Hellman assumption, namely the decisional k -Lin [HK07] and kernel k -Lin assumptions, or even more generally any assumption from the family of the (decisional or kernel) matrix Diffie-Hellman assumptions [EHK⁺13, MRV15].

One-Way KDM Security of ElGamal. The last hardness assumption we will use in this work states, in essence, that no PPT adversary can recover m given an ElGamal encryption of m , even when m might be an efficiently computable function of the ElGamal secret key, with probability significantly better than $2^{-c \cdot \lambda}$ for some constant $c < 1$ (where λ is the logarithm of the group size). Note that the best known attack against this falsifiable search assumption succeeds with probability $\text{poly}(\lambda)/2^\lambda$. To formally introduce the assumption, we introduce a natural secret-key variant of ElGamal (which suffices for our construction and leads to a more conservative assumption compared to the public-key variant).

Definition 4 (Secret-Key ElGamal). *Let $\tilde{\mathbb{G}} = \{\tilde{\mathbb{G}}_\lambda\}_\lambda$ be an ensemble of groups where each group $\tilde{\mathbb{G}}_\lambda$ is of order q such that $\lceil \log q \rceil \approx \lambda$. The natural (secret-key) variant of additive ElGamal with message space $\mathbb{Z}_q$ consists of the following three PPT algorithms.*

- $\text{Setup}(1^\lambda)$: The setup algorithm outputs a public-parameter $\tilde{G} \leftarrow_r \tilde{\mathbb{G}}_\lambda$ and a secret key $k \leftarrow_r \mathbb{Z}_q$.
- $\text{Enc}_{\tilde{G}}(k, m)$: The encryption algorithm samples $\tilde{R} \leftarrow_r \tilde{\mathbb{G}}$ and outputs a ciphertext $\tilde{C} = (\tilde{R}, \tilde{R}^k \cdot \tilde{G}^m)$.

- $\text{HalfDec}(k, \tilde{\mathbf{C}})$: The half decryption algorithm parses $\tilde{\mathbf{C}}$ as $(\tilde{C}_0, \tilde{C}_1)$ and outputs $\tilde{C}_1/\tilde{C}_0^k$.

Throughout the paper, we omit the subscript when the meaning is clear. Note that the scheme does not allow for full decryption, but only for decryption “up to discrete logarithm”: for every $(\tilde{G}, k, m)$, it holds that $\text{HalfDec}(k, \text{Enc}_{\tilde{G}}(k, m)) = \tilde{G}^m$. One important property of the scheme is that it enjoys the notion of *universality*. Informally, the notion claims that the ciphertexts are not associated with a specific key, but rather, could have been an output of *any* key.

Definition 5 (Universality). For all $\lambda \in \mathbb{N}$, $\tilde{G} \in \tilde{\mathbb{G}}_\lambda$, and $k^* \in \mathbb{Z}_q$, the ciphertexts of ElGamal satisfies

$$\{\tilde{\mathbf{C}} : (k, m) \leftarrow_r \mathbb{Z}_q^2, \tilde{\mathbf{C}} \leftarrow_r \text{Enc}_{\tilde{G}}(k, m)\} = \{\tilde{\mathbf{C}} : m \leftarrow_r \mathbb{Z}_q, \tilde{\mathbf{C}} \leftarrow_r \text{Enc}_{\tilde{G}}(k^*, m)\} = \mathcal{U}_{\tilde{\mathbb{G}}^2}.$$

Definition 6 (OW-KDM Security). Let $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of sets of functions where each $\mathcal{F}_\lambda = \{F_u\}_u$ is a family of (possibly randomized) efficiently-computable functions. We say that ElGamal satisfies (one-query) δ -hard OW-KDM security with respect to $\mathcal{F}$ if for every $F_u \in \mathcal{F}_\lambda$, superpolynomial function s , and every (non-uniform) PPT adversary $\mathcal{A}$, it holds that

$$\Pr_{\substack{(\tilde{G}, k) \leftarrow_r \tilde{\mathbb{G}}_\lambda \times \mathbb{Z}_q \\ m \leftarrow F_u(\tilde{G}, k) \\ \tilde{\mathbf{C}} \leftarrow_r \text{Enc}_{\tilde{G}}(k, m)}} [\mathcal{A}(\tilde{G}, \tilde{\mathbf{C}}) = m] \leq s(\lambda) \cdot \delta(\lambda).$$

When ElGamal satisfies δ -hard OW-KDM security for $\delta(\lambda) = 2^{-(c+o(1)) \cdot \lambda}$ for some constant $c \in (0, 1]$, we say it is $2^{-c\lambda}$ -OW-KDM secure or more simply, strong OW-KDM secure.

The strong OW-KDM security of ElGamal was introduced in [CCRR18]. However, this work considered an extreme variant of the notion with $c = 1$ (that is, $2^{-\lambda}$ -OW-KDM), and where security was required to hold with respect to *all functions* (even inefficient ones). The more conservative variant (with $c < 1$ and a restriction to efficiently computable functions) was introduced in [CKU20], which used it (with constant $c = 3/4$) to build correlation-intractable hash functions. In this work, we will rely on an even more conservative variant with $c = 1/2$.

2.2 ZAP

ZAP [DN00, DN07] is a public-coin two-move witness indistinguishable non-interactive argument. In this work, we focus on statistical ZAPs where witness indistinguishability holds unconditionally.

Definition 7 (ZAP). A ZAP system Π_{ZAP} for an NP language $\mathcal{L} = \{\mathcal{L}_\lambda\}_\lambda$ with corresponding relation $\mathcal{R} = \{\mathcal{R}_\lambda\}_\lambda$ with public-coin length $\ell(\lambda)$ is a tuple of PPT algorithms (Prove, Verify) defined as follows.

$\text{Prove}(r, x, w) \rightarrow \pi$: The proving algorithm is given the public-coin $r \in \{0, 1\}^\ell$, a statement x , and a witness w , and outputs a proof π .
 $\text{Verify}(r, x, \pi) \rightarrow \top$ or $\perp$: The verification algorithm is given the public-coin $r \in \{0, 1\}^\ell$, a statement x , and a proof π , and outputs $\top$ for acceptance or $\perp$ for rejection.

We additionally require the following properties to hold.

Correctness: For any $\lambda \in \mathbb{N}$, $r \in \{0, 1\}^\ell$ and $(x, w) \in \mathcal{R}_\lambda$, we have $\Pr[\text{Verify}(r, x, \text{Prove}(r, x, w)) = \top] = 1$.

(Non-Adaptive) Computational Soundness: For any $\lambda \in \mathbb{N}$, PPT adversary $\mathcal{A}$, and any statement $x \notin \mathcal{L}_\lambda$, we have

$$\Pr[r \leftarrow \{0, 1\}^\ell, \pi \leftarrow_r \mathcal{A}(r, x) : \text{Verify}(r, x, \pi) = \top] \leq \text{negl}(\lambda).$$

(Adaptive) Statistical Witness Indistinguishability: For any $\lambda \in \mathbb{N}$ and unbounded adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, we have

$$\left| \Pr \left[\begin{array}{l} (r, x, w_0, w_1, \text{st}) \leftarrow_r \mathcal{A}_0(1^\lambda) : \quad \mathcal{A}_1(\text{st}, \pi_0) = 1 \\ \pi_0 \leftarrow_r \text{Prove}(r, x, w_0) \quad \wedge \quad (x, w_0) \in \mathcal{R}_\lambda \\ \quad \wedge \quad (x, w_1) \in \mathcal{R}_\lambda \end{array} \right] - \Pr \left[\begin{array}{l} (r, x, w_0, w_1, \text{st}) \leftarrow_r \mathcal{A}_0(1^\lambda) : \quad \mathcal{A}_1(\text{st}, \pi_1) = 1 \\ \pi_1 \leftarrow_r \text{Prove}(r, x, w_1) \quad \wedge \quad (x, w_0) \in \mathcal{R}_\lambda \\ \quad \wedge \quad (x, w_1) \in \mathcal{R}_\lambda \end{array} \right] \right| \leq \text{negl}(\lambda).$$

Remark 8 (On Adaptive Soundness). In this work, we construct a ZAP that is *non-adaptive* computationally sound and adaptive *statistical* witness indistinguishable. This security property is in alignment with all the recent ZAPs (or ZAP with private randomness) [GJJM20, BFJ⁺20, LVW20]. Constructing ZAPs satisfying *adaptive* soundness and *statistical* witness indistinguishability seems to be difficult, where the former stipulates that the adversary can choose the statement $x \notin \mathcal{L}$ after it sees the public-coin r . Although we do not have any formal proofs of nonexistence of such ZAPs, we do have some evidence indicating the difficulty of obtaining them. In the context of NIZKs satisfying statistical zero-knowledge (NISZKs), Pass [Pas13] shows that there is no black-box reduction from the adaptive soundness of NISZK to a falsifiable assumption [Nao03, GW11].

2.3 NIZKs in the Hidden-Bits Model

We recall the notion of a NIZK in the hidden-bits model [FLS99].

Definition 9. A non-interactive proof system Π_{HBM} in the hidden-bits model for an NP language $\mathcal{L} = \{\mathcal{L}_\lambda\}_\lambda$ with corresponding relation $\mathcal{R} = \{\mathcal{R}_\lambda\}_\lambda$ with hidden-bits length $m(\lambda)$ is a pair of PPT algorithms $(\text{Prove}, \text{Verify})$ defined as follows.

$\text{Prove}(\text{hb}, x, w) \rightarrow (I, \pi)$: The proving algorithm is given a random bit string $\text{hb} \in \{0, 1\}^m$ and a statement x , and a witness w as inputs, and outputs a subset $I \subseteq [m]$ together with a proof π .

$\text{Verify}(S, \text{hb}_S, x, \pi) \rightarrow \top$ **or** $\perp$: The verification algorithm is given a subset $S \subseteq [m]$, a string $\text{hb}_S \in \{0, 1\}^{|S|}$, a statement x and a proof π as inputs, and outputs $\top$ for acceptance or $\perp$ for rejection.

We additionally require the following properties to hold.

Correctness. For any $\lambda \in \mathbb{N}$, $(x, w) \in \mathcal{R}_\lambda$, any $\text{hb} \in \{0, 1\}^m$, and for $(I, \pi) \leftarrow_r \text{Prove}(\text{hb}, x, w)$, we have $\text{Verify}(x, \text{hb}_S, x, \pi) = \top$.

Statistical ε -Soundness. For any $\lambda \in \mathbb{N}$ and (possibly unbounded) adversary $\mathcal{A}$, we have

$$\Pr[\text{hb} \leftarrow_r \{0, 1\}^m, (x, S, \pi) \leftarrow_r \mathcal{A}(\text{hb}) : \text{Verify}(S, \text{hb}_S, x, \pi) = \top \wedge x \notin \mathcal{L}_\lambda] \leq \varepsilon.$$

Perfect Zero-Knowledge. For any $\lambda \in \mathbb{N}$ and any (possibly unbounded) stateful adversary $\mathcal{A}$, there exists a PPT⁸ zero-knowledge simulator Sim such that for every $(x, w) \in \mathcal{R}_\lambda$, the distributions $\{(S, \text{hb}_S, \pi) : \text{hb} \leftarrow_r \{0, 1\}^m, (S, \pi) \leftarrow_r \text{Prove}(\text{hb}, x, w)\}$ and $\{\text{Sim}_{\text{zk}}(x)\}$ are perfectly indistinguishable.

We use the following result regarding the existence of NIZKs in the hidden-bits model [FLS90].

Theorem 10 (NIZK for all of NP in the hidden-bits model). *Let $k = k(\lambda)$ be any positive integer-valued function. Then, unconditionally, there exists a non-interactive proof system Π_{HBM} for any NP language $\mathcal{L} = \{\mathcal{L}_\lambda\}_\lambda$ in the hidden-bits model that uses $\text{hb} = k \cdot \text{poly}(\lambda)$ hidden-bits with soundness error $\epsilon \leq 2^{-k \cdot \lambda}$, where poly is a polynomial function related to the NP language $\mathcal{L}$.*

2.4 Correlation-Intractable Hash Functions

Finally, we recall the definition of correlation-intractable hash functions (CIH). We also require a CIH to be *programmable*, which roughly means for any input-output pair (x, y) , we can efficiently find a key k such that $H(k, x) = y$. Due to page limitations, the formal definition is provided in the full version.

Definition 11 (Correlation Intractable Hash Function). *A collection $\mathcal{H} = \{H_\lambda : K_\lambda \times I_\lambda \mapsto O_\lambda\}_\lambda$ of (efficient) keyed hash functions is a $\mathcal{R}$ -correlation intractable hash (CIH) family, with respect to a parameterized relation ensemble $\mathcal{R} = \{\mathcal{R}_\lambda\}_\lambda = \{\{\mathcal{R}_{\lambda,t} \subseteq I_\lambda \times O_\lambda\}_{t \in T_\lambda}\}_\lambda$, if for every (non-uniform) PPT adversary $\mathcal{A}$ and $t \in T_\lambda$, it holds that*

$$\Pr_{\substack{k \leftarrow_r K_\lambda \\ x \leftarrow_r \mathcal{A}(k)}} [(x, H_\lambda(k, x)) \in \mathcal{R}_{\lambda,t}] \leq \text{negl}(\lambda).$$

Furthermore, let $\mu(\lambda)$ be an efficiently computable function. We say that the collection $\mathcal{H}$ satisfies $(\mu, \mathcal{R})$ -correlation intractability if the above probability is bounded by $\mu(\lambda)$ for all PPT adversaries $\mathcal{A}$.

⁸ Note that we can also relax the definition to allow for an *unbounded* zero-knowledge simulator.

3 Interactive Hidden-Bits Generating Protocol and ZAPs for NP

In this section, we formally define an *interactive hidden-bits generating* (IHBG) protocol. Our definition builds on the definition of a (dual-mode) hidden-bits generator from [QRW19, LPWW20] (and the similar notion of (designated-verifier) PRG [DN00, DN07, CH19]). The main difference is that we allow a two-round interaction between the hidden-bits generator and the verifier, while removing the common reference string. Below, we define a public-coin flavor of an IHBG protocol to allow for public verifiability and reusability of the message from the verifier.

3.1 Definition

We formalize the notion of an interactive hidden-bits generating (IHBG) protocol.

Definition 12 (Interactive Hidden-Bits Generating Protocol). *Let $s(\lambda)$ and $m(\lambda)$ be positive valued polynomials. An interactive hidden-bits generating (IHBG) protocol Π_{IHBG} with public-coin length $\ell(\lambda)$ is a tuple of efficient algorithms $(\text{GenBits}, \text{VerifyBit})$ defined as follows.*

$\text{GenBits}(1^\lambda, m, r) \rightarrow (\sigma, \rho, \{\pi_i\}_{i \in [m]})$: *The hidden-bits generator algorithm is given the security parameter 1^λ (in unary), a length m , a public-coin $r \in \{0, 1\}^\ell$ and outputs a commitment $\sigma \in \{0, 1\}^s$, a string $\rho \in \{0, 1\}^m$, and a set of proofs $\{\pi_i\}_{i \in [m]}$.*

$\text{VerifyBit}(r, \sigma, i, \rho_i, \pi_i) \rightarrow \top \text{ or } \perp$: *The verification algorithm is given a public-coin $r \in \{0, 1\}^\ell$, a commitment $\sigma \in \{0, 1\}^s$, a bit $\rho_i \in \{0, 1\}$, and a proof π_i , and outputs $\top$ for acceptance or $\perp$ for rejection.*

We additionally require the following properties to hold. Below, we assume that the security parameter is provided to all algorithms, and omit it for simplicity.

Correctness: For any $\lambda \in \mathbb{N}$, $j \in [m]$, and $r \in \{0, 1\}^\ell$, we have

$$\Pr[(\sigma, \rho, \{\pi_i\}_{i \in [m]}) \leftarrow_r \text{GenBits}(m, r) : \text{VerifyBit}(r, \sigma, j, \rho_j, \pi_j) = \top] = 1.$$

Succinctness: The commitment length s only depends on the security parameter, i.e., $s(\lambda) = \text{poly}(\lambda)$, and in particular, does not depend on the length m of the generated bits.

μ -Extractability: There exists a PPT public-coin simulator SimCoin and a deterministic polynomial-time open algorithm Open such that for all polynomial m , the following two conditions hold. For an intuitive explanation for μ -successful extraction, we refer the readers to the technical overview in Section 1.2.

- (Public-Coin Indistinguishability) for any PPT adversary $\mathcal{A}$, we have

$$\begin{aligned} & |\Pr[r \leftarrow_r \{0, 1\}^\ell : \mathcal{A}(m, r) = 1] \\ & - \Pr[(\tilde{r}, \tau) \leftarrow_r \text{SimCoin}(1^\lambda, m) : \mathcal{A}(m, \tilde{r}) = 1]| \leq \text{negl}(\lambda). \end{aligned}$$

- (μ -Successful Extraction) for any PPT adversary $\mathcal{A}$ and any PPT distinguisher D , we have

$$\begin{aligned} & \Pr \left[\begin{array}{l} (\tilde{r}, \tau) \leftarrow_r \text{SimCoin}(1^\lambda, m) \\ (\sigma, S, \rho_S^*, \{\pi_i\}_{i \in S}, \text{st}) \leftarrow_r \mathcal{A}(m, \tilde{r}) : \\ \rho \leftarrow \text{Open}(\tilde{r}, \sigma, \tau) \end{array} : D(\text{st}) = 1 \wedge \rho \in \{0, 1\}^m \wedge \forall i \in S, \text{VerifyBit}(\tilde{r}, \sigma, i, 1 - \rho_i, \pi_i) = \perp \right] \\ & \geq \mu(\lambda) \cdot \Pr \left[\begin{array}{l} (\tilde{r}, \tau) \leftarrow_r \text{SimCoin}(1^\lambda, m) \\ (\sigma, S, \rho_S^*, \{\pi_i\}_{i \in S}, \text{st}) \leftarrow_r \mathcal{A}(m, \tilde{r}) : D(\text{st}) = 1 \end{array} \right] - \mu(\lambda) \cdot \text{negl}(\lambda). \end{aligned}$$

Statistical Hiding: For all polynomial m , public-coin $r \in \{0, 1\}^\ell$, and all unbounded adversaries $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, there exists a (possibly unbounded) simulator Sim such that

$$\begin{aligned} & \left| \Pr \left[\begin{array}{l} (\sigma, \rho, \{\pi_i\}_{i \in [m]}) \leftarrow_r \text{GenBits}(m, r) : S \subseteq [m] \wedge \mathcal{A}_1(r, S, \sigma, \rho, \{\pi_i\}_{i \in S}) = 1 \end{array} \right] - \right. \\ & \left. \Pr \left[\begin{array}{l} \rho \leftarrow_r \{0, 1\}^m, S \leftarrow_r \mathcal{A}_0(\rho) \\ (\sigma, \{\pi_i\}_{i \in S}) \leftarrow_r \text{Sim}(m, r, S, \rho_S) : S \subseteq [m] \wedge \mathcal{A}_1(r, S, \sigma, \rho, \{\pi_i\}_{i \in S}) = 1 \end{array} \right] \right| \\ & \leq \text{negl}(\lambda). \end{aligned}$$

3.2 ZAPs for NP from Interactive Hidden-Bits Generating Protocols

Here, we construct a ZAP for NP based on an IHBG protocol and a NIZK in the hidden-bits model, where the latter exists unconditionally.

Building Block. Let $\mathcal{L}$ be an NP language and $\mathcal{R}$ be its corresponding relation.⁹ We construct a ZAP for $\mathcal{L}$ based on the following building blocks.

- $\Pi_{\text{IHBG}} = (\text{GenBits}, \text{VerifyBit})$ is an interactive hidden-bits generating protocol. We assume it has public-coin length $\ell(\lambda)$, commitment length $s(\lambda)$, and output length $m(\lambda)$ (i.e., $\rho \in \{0, 1\}^m$). We further assume it satisfies $\mu(\lambda)$ -extractability.
- $\Pi_{\text{HBM}} = (\text{HBM.Prove}, \text{HBM.Verify})$ is a NIZK in the hidden-bits model for $\mathcal{L}$. We assume the hidden-bits length is $m(\lambda)$ and it is statistically ε_{HBM} -sound, where $\varepsilon_{\text{HBM}} = 2^{-s(\lambda)} \cdot \mu(\lambda) \cdot \text{negl}(\lambda)$.¹⁰

Construction. The construction of a ZAP for $\mathcal{L}$ with public-coin length $\ell'(\lambda) = \ell(\lambda) + m(\lambda)$, denoted as Π_{ZAP} , is described as follows.

ZAP.Prove(r', x, w) : On input a public-coin $r' \in \{0, 1\}^{\ell'}$, a statement x and a witness w , parse it as $(r, \Delta) \leftarrow r'$ such that $r \in \{0, 1\}^\ell$ and $\Delta \in \{0, 1\}^m$. Then run $(\sigma, \rho, \{\pi_{\text{IHBG}, i}\}_{i \in [m]}) \leftarrow_r \text{GenBits}(1^\lambda, m, r)$ and compute an HBM proof $(S, \pi_{\text{HBM}}) \leftarrow_r \text{HBM.Prove}(\text{hb}, x, w)$, where $\text{hb} := \rho \oplus \Delta$. Finally, output $\pi_{\text{ZAP}} = (\sigma, S, \rho_S, \{\pi_{\text{IHBG}, i}\}_{i \in S}, \pi_{\text{HBM}})$.

⁹ Although $\mathcal{L}$ and $\mathcal{R}$ are parameterized by the security parameter λ , we omit them throughout the paper for better readability whenever the meaning is clear.

¹⁰ Here, m can be set sufficiently large for both Π_{IHBG} and Π_{HBM} so that the existence of Π_{HBM} is guaranteed unconditionally by Theorem 10.

$\text{ZAP.Verify}(r', x, \pi_{\text{ZAP}})$: On input a public-coin $r' \in \{0, 1\}^{\ell'}$, a statement x and a proof π_{ZAP} , parse it as $(r, \Delta) \leftarrow r'$ such that $r \in \{0, 1\}^{\ell}$ and $\Delta \in \{0, 1\}^m$, and $(\sigma, S, \rho_S, \{\pi_{\text{IHBG}, i}\}_{i \in S}, \pi_{\text{HBM}}) \leftarrow \pi_{\text{ZAP}}$. Then, output $\top$ if $\text{HBM.Verify}(S, \rho_S \oplus \Delta_S, x, \pi_{\text{HBM}}) = \top$ and $\text{VerifyBit}(r, \sigma, i, \rho_i, \pi_{\text{IHBG}, i}) = \top$ for all $i \in S$. Otherwise, output $\perp$.

3.3 Security

Correctness of our ZAP follows from a routine check. Below, we show our ZAP satisfies non-adaptive computational soundness and adaptive statistical witness indistinguishability in Theorems 13 and 14. Since the proof of witness indistinguishability is similar to those in [LPWW20], we provide the details in the full version of this paper [CKSU21].

Theorem 13 (Statistical Witness Indistinguishability). *If Π_{IHBG} is statistically hiding and Π_{HBM} has perfect zero-knowledge, then Π_{ZAP} is adaptive statistical witness indistinguishability.*

Theorem 14 (Soundness). *If Π_{IHBG} is μ -extractable and Π_{HBM} has statistical ε_{HBM} -soundness, where $\varepsilon_{\text{HBM}} = 2^{-s(\lambda)} \cdot \mu(\lambda) \cdot \text{negl}(\lambda)$, then Π_{ZAP} has non-adaptive computational soundness.*

Proof. Assume there exists a statement $x \notin \mathcal{L}$ and a PPT adversary $\mathcal{A}$ against the non-adaptive computational soundness of Π_{ZAP} with advantage ε . Below, we consider the following sequence of games between $\mathcal{A}$ and a challenger and denote E_i as the event that the challenger outputs 1.

Game₁: This is the real soundness game that proceeds as follows: The challenger first samples a public-coin $r' \leftarrow_r \{0, 1\}^{\ell'}$ and sends it to $\mathcal{A}$. $\mathcal{A}$ then outputs a proof π_{ZAP}^* and sends it to the challenger. The challenger outputs 1 if $\text{ZAP.Verify}(r', x, \pi_{\text{ZAP}}^*) = \top$, and outputs 0 otherwise. By definition $\Pr[E_1] = \varepsilon$.

Game₂: This game is identical to the previous game except that the public-coin $r' \in \{0, 1\}^{\ell'}$ is sampled differently. Let SimCoin be the PPT public-coin simulator of the IHBG protocol Π_{IHBG} . Then, in this game, the challenger first runs $(\tilde{r}, \tau) \leftarrow_r \text{SimCoin}(m)$ and samples $\Delta \leftarrow_r \{0, 1\}^m$, where $\tilde{r} \in \{0, 1\}^{\ell}$, and outputs the simulated public-coin $\tilde{r}' := (\tilde{r}, \Delta) \in \{0, 1\}^{\ell'}$. The rest is defined the same as in the previous game.

Game₃: This game is identical to the previous game except that the challenger checks an additional condition regarding π_{ZAP}^* output by $\mathcal{A}$. Let Open be the efficient deterministic open algorithm of the IHBG protocol Π_{IHBG} . Then, in this game, when $\mathcal{A}$ outputs π_{ZAP}^* , the challenger first parses

$$(\sigma^*, S^*, \rho_{S^*}^*, \{\pi_{\text{IHBG}, i}^*\}_{i \in S^*}, \pi_{\text{HBM}}^*) \leftarrow \pi_{\text{ZAP}}^*$$

and runs $\rho \leftarrow \text{Open}(\tilde{r}, \sigma^*, \tau)$. It then outputs 1 if $\text{ZAP.Verify}(r', x, \pi_{\text{ZAP}}^*) = \top$, $\rho \in \{0, 1\}^m$, and $\rho_{S^*}^* = \rho_{S^*}$, and 0 otherwise.

The following Lemmas 15 to 17 establish $\Pr[E_1] = \varepsilon \leq \text{negl}(\lambda)$, thus completing the proof.

Lemma 15. *If Π_{HBM} is μ -extractable for all PPT adversary, then we have $|\Pr[E_1] - \Pr[E_2]| \leq \text{negl}(\lambda)$, hence $\Pr[E_2] \geq \varepsilon - \text{negl}(\lambda)$.*

Proof. The only difference between the two games is how the public-coin is generated. Let us consider the following adversary $\mathcal{B}$ against the public-coin indistinguishability of Π_{HBM} : $\mathcal{B}$ receives $r \in \{0, 1\}^\ell$ from its challenger and samples $\Delta \leftarrow_r \{0, 1\}^m$. It then invokes $\mathcal{A}$ on input $r' = (r, \Delta)$, and outputs 1 if the proof π_{ZAP} output by $\mathcal{A}$ satisfies $\text{ZAP.Verify}(r', x, \pi_{\text{ZAP}}) = \top$, and 0 otherwise. Since $\mathcal{B}$ perfectly simulates **Game**₁ (resp. **Game**₂) when $r \leftarrow_r \{0, 1\}^\ell$ (resp. $(r, \tau) \leftarrow_r \text{SimCoin}(m)$), we have $|\Pr[E_1] - \Pr[E_2]| \leq \text{negl}(\lambda)$.

Lemma 16. *If Π_{HBM} is μ -extractable for all PPT adversary, then we have $\Pr[E_3] \geq \mu(\lambda) \cdot (\Pr[E_2] - \text{negl}(\lambda))$.*

Proof. This follows from the μ -successful extractability of Π_{HBM} . Let us consider the following adversary $\mathcal{B}$ and distinguisher $\mathcal{D}$ against the μ -successful extractability: $\mathcal{B}$ on input m and $\tilde{r}$ invokes $\mathcal{A}$ and simulates the challenger in **Game**₂. When $\mathcal{A}$ outputs a forgery $\pi_{\text{ZAP}}^* = (\sigma^*, S^*, \rho_{S^*}^*, \{\pi_{\text{HBM},i}^*\}_{i \in S^*}, \pi_{\text{HBM}}^*)$, $\mathcal{B}$ outputs $(\sigma^*, S^*, \rho_{S^*}^*, \{\pi_{\text{HBM},i}^*\}_{i \in S^*}, \text{st})$, where $\text{st} = (\tilde{r}, \pi_{\text{ZAP}}^*)$; $\mathcal{D}$ on input st , checks if $\text{ZAP.Verify}(\tilde{r}, x, \pi_{\text{ZAP}}^*) = \top$, and outputs 1 if so and outputs 0 otherwise. Observe that the probability $\mathcal{D}$ outputs 1 is the same as the probability that event E_2 occurs. Below, we relate the probability that event E_3 occurs with the left hand side equation of μ -successful extractability.

The only difference between **Game**₂ and **Game**₃ is the check that $\rho \in \{0, 1\}^m$ and $\rho_{S^*}^* = \rho_{S^*}$. Now, consider a variant **Game**₃' of **Game**₃ where, instead of checking $\rho_{S^*}^* = \rho_{S^*}$, the challenger checks that for all $i \in S^*$, it holds that

$$\text{VerifyBit}(\tilde{r}, \sigma^*, i, 1 - \rho_i, \pi_{\text{HBM},i}^*) = \perp.$$

Let E_3' be the event that the challenger outputs 1 in this variant. Observe that if event E_3' occurs then so does event E_3 . Indeed, whenever the challenger outputs 1 in E_3' , it holds in particular that

$$\begin{aligned} \forall i \in S^*, \text{VerifyBit}(\tilde{r}, \sigma^*, i, \rho_i^*, \pi_i^*) &= \top, \text{ and} \\ \forall i \in S^*, \text{VerifyBit}(\tilde{r}, \sigma^*, i, 1 - \rho_i, \pi_i^*) &= \perp. \end{aligned}$$

The latter implies that it can never hold, for any $i \in S^*$, that $\rho_i^* = 1 - \rho_i$; hence, since we check $\rho \in \{0, 1\}^m$ in both events, whenever E_3' happens, it further holds that $\rho_{S^*}^* = \rho_{S^*}$ and E_3 therefore holds as well. In other terms,

$$\Pr[E_3] \geq \Pr[E_3'].$$

Therefore, by applying the μ -successful extractability of Π_{HBM} with respect to $\mathcal{B}$ and $\mathcal{D}$, since the only difference between **Game**₂ and **Game**₃' is the check that $\rho \in \{0, 1\}^m$ and $\text{VerifyBit}(\tilde{r}, \sigma^*, i, 1 - \rho_i, \pi_{\text{HBM},i}^*) = \perp$, we get

$$\Pr[E_3'] \geq \mu(\lambda) \cdot (\Pr[E_2] - \text{negl}(\lambda)),$$

which concludes the proof of Lemma 16.

Lemma 17. *If Π_{HBM} is statistical ε_{HBM} -sound, then we have $\Pr[\text{E}_3] \leq \mu(\lambda) \cdot \text{negl}(\lambda)$.*

Proof. Let $(\sigma^*, S^*, \rho_{S^*}^*, \{\pi_{\text{HBM},i}^*\}_{i \in S^*}, \pi_{\text{HBM}}^*) \leftarrow \pi_{\text{ZAP}}^*$ be $\mathcal{A}$'s output. When the challenger outputs 1 (i.e., event E_3 occurs), we have $\rho_{S^*}^* = \rho_{S^*}$, where $\rho \leftarrow \text{Open}(\tilde{r}, \sigma^*, \tau)$, and $\text{HBM.Verify}(S^*, \rho_{S^*}^* \oplus \Delta_{S^*}, x, \pi_{\text{HBM}}^*) = \top$. For an any $S^* \subseteq [m]$ and ρ_{S^*} , if $\Delta \leftarrow_r \{0, 1\}^m$ is sampled uniformly at random, then $\rho_{S^*} \oplus \Delta_{S^*}$ is distributed uniformly random. Then, by soundness of Π_{HBM} , for a fixed ρ_{S^*} we have

$$\Pr[\text{HBM.Verify}(S^*, \rho_{S^*} \oplus \Delta_{S^*}, x, \pi_{\text{HBM}}^*) = \top] \leq \varepsilon_{\text{HBM}},$$

where the probability is taken over the randomness of Δ , $\mathcal{A}$, and the challenger, conditioned on $\mathcal{A}$ outputting $\rho_{S^*}^*$ that is consistent with ρ_{S^*} . Here, we do not include the condition $x \notin \mathcal{L}$ in the above equation since we consider non-adaptive soundness for Π_{ZAP} .

If we fix an arbitrary $(\tilde{r}, \tau)$, then for any commitment $\sigma \in \{0, 1\}^s$ the output of $\rho \leftarrow \text{Open}(\tilde{r}, \sigma, \tau)$ is uniquely defined since Open is deterministic. Let us denote the unique ρ as ρ^σ . Then, taking a union bound over all possible commitments $\sigma \in \{0, 1\}^s$, we have

$$\begin{aligned} \Pr[\exists \sigma \in \{0, 1\}^\ell \text{ s.t. } \text{HBM.Verify}(S^*, \rho_{S^*}^\sigma \oplus \Delta_{S^*}, x, \pi_{\text{HBM}}^*) = \top] &\leq 2^s \cdot \varepsilon_{\text{HBM}} \\ &= \mu(\lambda) \cdot \text{negl}(\lambda). \end{aligned}$$

Thus, we conclude $\Pr[\text{E}_3] \leq \mu(\lambda) \cdot \text{negl}(\lambda)$.

Putting everything together, this gives $\mu(\lambda) \cdot (\varepsilon - \text{negl}(\lambda)) \leq \mu(\lambda) \cdot \text{negl}(\lambda)$, which implies $\varepsilon \leq \text{negl}(\lambda)$. This concludes the proof.

4 The LPWW Language $\mathcal{L}_{\text{LPWW}}$

To instantiate the generic construction of statistical ZAP for NP given in Section 3, we will construct an IHBG which builds upon the dual-mode hidden-bit generator of Libert, Passelègue, Wee, and Wu [LPWW20]. In this section, we first recall the specific *parameterized* language considered by [LPWW20] (denoted as the LPWW language $\mathcal{L}_{\text{LPWW}}$). We then introduce some tools related to this parameterized language: a specific type of statistical ZAP for $\mathcal{L}_{\text{LPWW}}$, which we call *IHBG-friendly statistical ZAP*, and a Σ -protocol for $\mathcal{L}_{\text{LPWW}}$.

4.1 Definition

Formally, we denote by $\mathcal{L}_{\text{LPWW}} := \{\mathcal{L}_{\text{LPWW},\lambda}\}_\lambda$ the following family of parametrized languages: let $\mathbb{G}$ be a cyclic group of prime order p . We implicitly fix a vector length $d \in \mathbb{N}$ and a generator $g \in \mathbb{G}$ for each security parameter λ .¹¹ Let a set

¹¹ To be precise, $g \in \mathbb{G}$ will be sampled for each security parameter λ and the family of parameterized language $\mathcal{L}_{\text{LPWW},\lambda}$ is defined with respect to such generator g . For better readability, we may make the random sampling of g implicit when the context is clear.

of parameter space Λ_λ be $(\mathbb{G}^d \setminus \{\mathbf{1}\})^2$, where $\mathbf{1} := g^{\mathbf{0}}$ for $\mathbf{0} \in \mathbb{Z}_p^d$. Then, for any parameter $\text{par} = (g^{\mathbf{v}}, g^{\mathbf{w}}) \in \Lambda_\lambda$, we define $\mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}} = \{\mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}}\}_{\text{par} \in \Lambda_\lambda}$ such that $\mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}}$ is the following parametrized language:

$$\mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}} := \left\{ (g^s, g^u) \in \mathbb{G}^2 \mid \exists \mathbf{y} \in \mathbb{Z}_p^d \text{ s.t. } g^{\mathbf{y}^\top \mathbf{v}} = g^s \wedge g^{\mathbf{y}^\top \mathbf{w}} = g^u \right\}.$$

Let $\text{Col}(\mathbb{G}^d) \subset \Lambda_\lambda$ denote the set of elements of the form $(g^{\mathbf{v}}, g^{\alpha \mathbf{v}})$ for some $\mathbf{v} \neq \mathbf{0}$ and $\alpha \in \mathbb{Z}_p^*$, that is, the exponents form colinear vectors over $(\mathbb{Z}_p)^d$. Observe that for any $\text{par} \in \text{Col}(\mathbb{G}^d)$, $\mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}}$ is a non-trivial Diffie-Hellman-style language (hence, $\mathcal{L}_{\text{LPWW}}^{\text{par}}$ is a sparse subset of Λ_λ); however, for any $\text{par} \in \Lambda_\lambda \setminus \text{Col}(\mathbb{G}^d)$, $\mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}}$ is actually equal to $\mathbb{G}^2$ (hence, $\mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}}$ is a trivial language). Below, we may omit the security parameter and use the shorthand $\mathcal{L}_{\text{LPWW}} = \{\mathcal{L}_{\text{LPWW}}^{\text{par}}\}_{\text{par} \in \Lambda}$ when the meaning is clear.

4.2 IHBG-Friendly Statistical ZAPs for the LPWW Language $\mathcal{L}_{\text{LPWW}}$

Looking ahead, our construction of IHBG in Section 5 will rely at its core on an adaptively secure statistical ZAP for the family of parametrized languages $\mathcal{L}_{\text{LPWW}} = \{\mathcal{L}_{\text{LPWW}}^{\text{par}}\}_{\text{par} \in \Lambda}$. More precisely, the statistical ZAP which we will use in our construction satisfies a variant of the standard notion of adaptive computational soundness (which we defined for a single language in Section 2): we require adaptive computational soundness to hold with respect to parameters par *sampled uniformly from* $\text{Col}(\mathbb{G}^d) \subset \Lambda$ (recall that $\text{Col}(\mathbb{G}^d)$ is the subset of parameters such that $\mathcal{L}_{\text{LPWW}}^{\text{par}}$ is nontrivial). In contrast, adaptive statistical witness indistinguishability must hold *even for adversarially chosen parameters* $\text{par} \in \Lambda$ (hence, in a sense, WI is doubly-adaptive: with respect to the statement, and with respect to the language parameters). We call a statistical ZAP with these properties an *IHBG-friendly* statistical ZAP for $\mathcal{L}_{\text{LPWW}}$. We provide a formal definition below.

Definition. We formally introduce the notion of IHBG-friendly statistical ZAP for the family of parametrized languages $\mathcal{L}_{\text{LPWW}}$.

Definition 18 (IHBG-Friendly Statistical ZAP for $\mathcal{L}_{\text{LPWW}}$). Let $\Lambda_\lambda = (\mathbb{G}^d \setminus \{\mathbf{1}\})^2$ be the parameter space for any $\lambda \in \mathbb{N}$ and consider the family of parameterized NP languages $\mathcal{L}_{\text{LPWW}} = \{\mathcal{L}_{\text{LPWW}, \lambda}\}_\lambda = \{\{\mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}}\}_{\text{par} \in \Lambda_\lambda}\}_\lambda$, with associated witness relation $\mathcal{R}_{\text{LPWW}} = \{\mathcal{R}_{\text{LPWW}, \lambda}\}_\lambda = \{\{\mathcal{R}_{\text{LPWW}, \lambda}^{\text{par}}\}_{\text{par} \in \Lambda_\lambda}\}_\lambda$. Then, an IHBG-friendly ZAP system Π_{ZAP} for $\mathcal{L}_{\text{LPWW}}$ with with public-coin length $\ell(\lambda)$ is a tuple of PPT algorithms (Prove, Verify) defined as follows.

Prove(par, r, x, w) $\rightarrow \pi$: The proving algorithm is given the parameters $\text{par} \in \Lambda_\lambda$, the public-coin $r \in \{0, 1\}^\ell$, a statement x , and a witness w , and outputs a proof π .

$\text{Verify}(\text{par}, r, x, \pi) \rightarrow \top \text{ or } \perp$: *The verification algorithm is given the parameters $\text{par} \in \Lambda_\lambda$, the public-coin $r \in \{0, 1\}^\ell$, a statement x , and a proof π , and outputs $\top$ for acceptance or $\perp$ for rejection.*

We additionally require the following properties to hold.

Correctness: For any $\lambda \in \mathbb{N}$, $r \in \{0, 1\}^\ell$, $\text{par} \in \Lambda_\lambda$, and $(x, w) \in \mathcal{R}_{\text{LPWW}, \lambda}^{\text{par}}$, we have

$$\Pr[\text{Verify}(\text{par}, r, x, \text{Prove}(\text{par}, r, x, w)) = \top] = 1.$$

(Adaptive) Computational ϵ_{sound} -Soundness w.r.t. Colinear Parameters: For any $\lambda \in \mathbb{N}$ and PPT adversary $\mathcal{A}$, we have

$$\Pr \left[\text{par} \leftarrow_r \text{Col}(\mathbb{G}^d), r \leftarrow \{0, 1\}^\ell, (x, \pi) \leftarrow_r \mathcal{A}(\text{par}, r) : \begin{array}{l} x \notin \mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}} \wedge \\ \text{Verify}(\text{par}, r, x, \pi) = \top \end{array} \right] \leq \epsilon_{\text{sound}}.$$

(Doubly-Adaptive) Statistical Witness Indistinguishability: For any $\lambda \in \mathbb{N}$ and unbounded adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, we have

$$\begin{aligned} & \left| \Pr \left[\begin{array}{l} (r, \text{par}, x, w_0, w_1, \text{st}) \leftarrow_r \mathcal{A}_0(1^\lambda) \\ \pi \leftarrow_r \text{Prove}(\text{par}, r, x, w_0) \end{array} : \begin{array}{l} \text{par} \in \Lambda_\lambda \wedge \mathcal{A}_1(\text{st}, \pi) = 1 \\ \wedge (x, w_0) \in \mathcal{R}_{\text{LPWW}, \lambda}^{\text{par}} \\ \wedge (x, w_1) \in \mathcal{R}_{\text{LPWW}, \lambda}^{\text{par}} \end{array} \right] \right| \\ & - \left| \Pr \left[\begin{array}{l} (r, \text{par}, x, w_0, w_1, \text{st}) \leftarrow_r \mathcal{A}_0(1^\lambda) \\ \pi \leftarrow_r \text{Prove}(\text{par}, r, x, w_1) \end{array} : \begin{array}{l} \text{par} \in \Lambda_\lambda \wedge \mathcal{A}_1(\text{st}, \pi) = 1 \\ \wedge (x, w_0) \in \mathcal{R}_{\text{LPWW}, \lambda}^{\text{par}} \\ \wedge (x, w_1) \in \mathcal{R}_{\text{LPWW}, \lambda}^{\text{par}} \end{array} \right] \right| \leq \text{negl}(\lambda). \end{aligned}$$

Building IHBG-Friendly Statistical ZAPs for $\mathcal{L}_{\text{LPWW}}$. In Section 6, we will provide two constructions of an IHBG-friendly statistical ZAPs for $\mathcal{L}_{\text{LPWW}}$, one in pairing groups (Theorem 25), and one in pairing-free groups (Theorem 28). Both constructions are obtained by compiling the Σ -protocol for $\mathcal{L}_{\text{LPWW}}$ described in Section 4.3 into an IHBG-friendly statistical ZAP for $\mathcal{L}_{\text{LPWW}}$. Below, we give an overview of the main lemmas regarding our two constructions whose proofs are provided in Section 6.

Pairing-Based Construction. The pairing-based construction builds upon the Couteau-Hartmann compiler from [CH20], which relies on the hardness of the kernel Diffie-Hellman assumption in a group $\mathbb{G}_2$ (more generally, it can be based on the kernel k -Lin assumption in $\mathbb{G}_2$ for any k), a standard search assumption (which is implied in particular by DDH in $\mathbb{G}_2$) introduced in [MRV15] and used in several works on pairing-based NIZKs, e.g. [KW15].

Lemma 19. *Let $(\mathbb{G}_1, \mathbb{G}_2)$ be bilinear-map groups equipped with an asymmetric pairing (implicitly parameterized by the security parameter λ). There exists an IHBG-friendly adaptive statistical ZAP for the family of parametrized languages $\mathcal{L}_{\text{LPWW}}$ over $\mathbb{G}_1$ which satisfies adaptive computational ϵ_{sound} -soundness w.r.t. colinear parameters, and doubly-adaptive statistical witness indistinguishability, assuming the explicit ϵ_{sound} -hardness of the kernel Diffie-Hellman assumption in $\mathbb{G}_2$.*

Pairing-Free Construction. The pairing-free construction builds upon the compiler of [CKU20]. The work of [CKU20] build a correlation intractable hash function under the $2^{-3\lambda/4}$ -OW-KDM security of ElGamal, which suffices to compile the above Σ -protocol into a statistical ZAP. We refine their approach and achieve a similar result under a weaker assumption, by managing to reduce the constant $3/4$ to $1/2$, that is, rely on the $2^{-\lambda/2}$ -OW-KDM security of ElGamal. We note that the best known attack against this falsifiable search assumption succeeds with probability $\text{poly}(\lambda)/2^\lambda$.

Lemma 20. *Let $\mathbb{G}$ be a group of order p such that $\lambda \approx 2\lceil \log p \rceil^2$. There exists an IHBG-friendly adaptive statistical ZAP for the family of parametrized languages $\mathcal{L}_{\text{LPWW}}$ over $\mathbb{G}$ which satisfies adaptive computational $\varepsilon_{\text{sound}}$ -soundness w.r.t. colinear parameters for any $\varepsilon_{\text{sound}} = 2^{-o(\lceil \log p \rceil^2)}$, and doubly-adaptive statistical witness indistinguishability, assuming the $2^{-\lambda/2}$ -OW-KDM hardness of ElGamal over another group $\tilde{\mathbb{G}}$ of size $|\tilde{\mathbb{G}}| \approx 2^\lambda$.*

4.3 Σ -protocols for the LPWW Language $\mathcal{L}_{\text{LPWW}}$

To construct our IHBG-friendly statistical ZAPs for $\mathcal{L}_{\text{LPWW}}$, we rely on a Σ -protocol for the family of parameterized language $\mathcal{L}_{\text{LPWW}}$. To this end, we need to first extend the standard definition of Σ -protocols for a single language to a family of parameterized languages. As the definition is a natural extension of the standard definition, we provide the detail in the full version of this paper [CKSU21].

We now provide a Σ -protocol for $\mathcal{L}_{\text{LPWW}}$ as follows. Fix some parameters $\text{par} = (g^{\mathbf{v}}, g^{\mathbf{w}}) \in \Lambda = (\mathbb{G}^d \setminus \{\mathbf{1}\})^2$ (implicitly parameterized by the security parameter λ). To match with the notations which we will use later when building an IHBG, we denote the dimension d in $\mathcal{L}_{\text{LPWW}}^{\text{par}}$ by $m+1$. We consider a statement $(\hat{X}, \hat{Y}) := (g^{\hat{x}}, g^{\hat{y}}) \in \mathcal{L}_{\text{LPWW}}^{\text{par}}$ and let $\mathbf{y} \in \mathbb{Z}_p^{m+1}$ be the prover witness (i.e., $\mathbf{y}$ is any vector over $\mathbb{Z}_p^{m+1}$ such that $\mathbf{y}^\top \mathbf{v} = \hat{x}$ and $\mathbf{y}^\top \mathbf{w} = \hat{y}$). Let $n \in \mathbb{N}$ be any positive integer. Then, a Σ -protocol for $\mathcal{L}_{\text{LPWW}} = \{\mathcal{L}_{\text{LPWW}}^{\text{par}}\}_{\text{par} \in \Lambda}$ is provided in Figure 1. Correctness can be checked by routine calculation. Below, we prove perfect witness indistinguishability and adaptive soundness. Due to page limitation, the proof is provided in the full version of this paper [CKSU21].

Fig. 1. Σ -protocol with statement $(\hat{X}, \hat{Y}) \in \mathcal{L}_{\text{LPWW}}^{\text{par}}$ where $\text{par} := (\mathbf{g}, \mathbf{h}) = (g^{\mathbf{v}}, g^{\mathbf{w}})$.

Lemma 21 (Perfect Witness Indistinguishability). *The IHBG-friendly Σ -protocol for the family of parametrized languages $\mathcal{L}_{\text{LPWW}} = \{\mathcal{L}_{\text{LPWW}}^{\text{par}}\}_{\text{par} \in \Lambda}$ in Figure 1 satisfies perfect witness indistinguishability.*

Lemma 22 (Adaptive Soundness). *The IHBG-friendly Σ -protocol for the family of parametrized languages $\mathcal{L}_{\text{LPWW}} = \{\mathcal{L}_{\text{LPWW}}^{\text{par}}\}_{\text{par} \in \Lambda}$ in Figure 1 satisfies adaptive $(\frac{1}{p-1})^{n-1}$ -soundness.*

5 Interactive Hidden-Bits Generating Protocols from the Explicit Hardness of DDH and an IHBG-Friendly Statistical ZAPs for $\mathcal{L}_{\text{LPWW}}$

In this section, we construct an IHBG protocol based on explicit μ -hardness of the DDH assumption (over a pairing-free group, for a negligible function μ arbitrarily close to an inverse polynomial function) and an IHBG-friendly statistical ZAP for the language $\mathcal{L}_{\text{LPWW}}$, defined in Section 4, which is naturally induced from the (non-interactive) hidden-bits generator of Libert et al. [LPWW20].

5.1 Constructing the IHBG Protocol

Building Block. Our construction is parametrized by λ and $\mu(\lambda)$, and relies on the following building blocks:

- $\mathcal{H} = \{\mathcal{H}_\lambda\}_\lambda = \{\{H : \mathbb{G} \mapsto \{0, 1\}\}_H\}_\lambda$ is a family of universal hash functions with description size of at most $O(\log_2 p)$ bits, where $\mathbb{G}$ and p are implicitly parameterized by the security parameter.
- $\Pi_{\text{ZAP}} = (\text{ZAP.Prove}, \text{ZAP.Verify})$ is an IHBG-friendly ZAP for the parametrized family of languages $\mathcal{L}_{\text{LPWW}} = \{\mathcal{L}_{\text{LPWW}, \lambda}\}_\lambda = \{\{\mathcal{L}_{\text{LPWW}, \lambda}^{\text{par}}\}_{\text{par} \in \Lambda_\lambda}\}_\lambda$ with public-coin length $\ell'(\lambda)$, satisfying adaptive computational ϵ_{sound} -soundness w.r.t. colinear parameters for $\epsilon_{\text{sound}} = \frac{\mu(\lambda)}{m(\lambda)}$ and doubly-adaptive statistical witness indistinguishability. Here, we set the vector length parameter $d(\lambda)$ in $\mathcal{L}_{\text{LPWW}, \lambda}$ to $m(\lambda) + 1$, where $m(\lambda)$ is the polynomial output bit length of the IHBG protocol defined below.

Construction. The construction of an IHBG protocol denoted as Π_{IHBG} is described as follows. The commitment length is at most $s(\lambda) = \lceil \log_2 p \rceil + O(\log_2 p)$ where $(\mathbb{G}, p) \leftarrow \text{DHGen}(1^\lambda)$ (note that DHGen guarantees in particular $p > \lambda^{\omega(1)}$, which is needed to use the uniformity property of H). The output bit length $m(\lambda)$ is an arbitrary large enough fixed polynomial $\text{poly}(\lambda)$, and the public-coin length $\ell(\lambda)$ is $m \cdot \ell' + (m+2) \cdot \lceil \log_2 p \rceil$. We rely on one more parameter $\nu(\lambda)$ and require the parameters to satisfy the following conditions:

- In order to prove statistical hiding, $m(\lambda) \cdot \mu(\lambda)$ must be negligible; this holds by setting $\mu(\lambda)$ to be a negligible function.
- For technical reasons in the hybrid games, we need a negligible gap between ν and μ ; that is, $\nu(\lambda)$ is a negligible function satisfying $\mu(\lambda) = \nu(\lambda) \cdot \text{negl}(\lambda)$.

- We also need $1/\mu(\lambda)$ (and hence $1/\nu(\lambda)$) to be small compared to p (otherwise, assuming explicit μ -hardness of DDH over $\mathbb{G}$ does not make sense: a polynomial time attack with $O(1/p)$ advantage against DDH trivially exists). In particular, $\mu(\lambda)$ can be set as an arbitrary close to an inverse polynomial, i.e., $\lambda^{-\omega(1)}$. Here, since $1/\nu(\lambda)$ is small compared to p , any element $z \in [1/\nu(\lambda)]$ can be seen as an element of $\mathbb{Z}_p$.

We proceed with the description of the scheme. In the following we may omit the dependency on λ for better readability when the context is clear.

GenBits($1^\lambda, m, r$) : On input the security parameter 1^λ , bit length m , and a public-coin $r \in \{0, 1\}^\ell$, parse $((r_{\text{ZAP},i})_{i \in [m]}, g, g^{\mathbf{M}}) \leftarrow r$, where $g \in \mathbb{G}$ and $\mathbf{M} := (\mathbf{v}|\mathbf{w}_1| \dots |\mathbf{w}_m) \in \mathbb{Z}_p^{(m+1) \times (m+1)}$.¹² Then sample $z \leftarrow_r [1/\nu]$, and compute $g^{\mathbf{M}-z \cdot \mathbf{I}_{m+1}}$, where we denote $\mathbf{M}' := \mathbf{M} - z \cdot \mathbf{I}_{m+1} = (\mathbf{v}'|\mathbf{w}'_1| \dots |\mathbf{w}'_m) \in \mathbb{Z}_p^{(m+1) \times (m+1)}$. Further sample a random hash function $H \leftarrow_r \mathcal{H}$ and a uniformly random seed $\mathbf{y} \leftarrow_r \mathbb{Z}_p^{m+1}$, and compute a commitment $g^s \leftarrow g^{\mathbf{y}^\top \mathbf{v}'}$, openings $g^{u_i} \leftarrow g^{\mathbf{y}^\top \mathbf{w}'_i}$, and the hidden bits $\rho_i \leftarrow H(g^{u_i})$ for all $i \in [m]$. For each $i \in [m]$, set the language parameter $\text{par}_i := (g^{\mathbf{v}'}, g^{\mathbf{w}'_i})$, statement $x_i := (g^s, g^{u_i})$, and witness $w := \mathbf{y}$ for membership to the parametrized language $\mathcal{L}_{\text{LPWW}}^{\text{par}_i}$, and compute $\pi_{\text{ZAP},i} \leftarrow_r \text{ZAP.Prove}(\text{par}_i, r_{\text{ZAP},i}, x_i, w)$ and set $\pi_i = (g^{u_i}, \pi_{\text{ZAP},i})$. Finally, output the commitment $\sigma := (H, g^s, z) \in \mathcal{H} \times \mathbb{G} \times [1/\nu]$, string $\rho := (\rho_i)_{i \in [m]} \in \{0, 1\}^m$ and the set of proofs $\{\pi_i\}_{i \in [m]}$.

VerifyBit($r, \sigma, i, \rho_i, \pi_i$) : Parse $((r_{\text{ZAP},i})_{i \in [m]}, g, g^{\mathbf{M}}) \leftarrow r$, $(H, g^s, z) \leftarrow \sigma$, $(g^{u_i}, \pi_{\text{ZAP},i}) \leftarrow \pi_i$, and compute $g^{\mathbf{M}'} \leftarrow g^{\mathbf{M}-z \cdot \mathbf{I}_{m+1}}$. Then, set the language parameter as $\text{par}_i := (g^{\mathbf{v}'}, g^{\mathbf{w}'_i})$ and the statement as $x_i := (g^s, g^{u_i})$. Check $\rho_i = H(g^{u_i})$ and $\text{ZAP.Verify}(\text{par}_i, r_{\text{ZAP},i}, x_i, \pi_{\text{ZAP},i}) = \top$. Output $\top$ if both check passes and otherwise output $\perp$.

Succinctness The length of the commitment $\sigma = (H, g^s, z)$ only depends on the security parameter, and in particular, independent of m . This is because g^s requires $\lceil \log_2 p \rceil$ bits, z requires $\lceil \log_2(1/\nu(\lambda)) \rceil \leq \lceil \log_2 p \rceil$ and the description of the universal hash function H requires at most $O(\log_2 p)$ bits.

5.2 Security

Correctness of our IHBG protocol can be verified by a routine check. Below, we show our IHBG protocol satisfies extractability and statistical hiding in the following Theorems 23 and 24. Due to page limitations, the proof is provided in the full version of this paper [CKSU21], and we only give a proof sketch in the main body.

Theorem 23 (Extractability). *Consider $\mu(\lambda)$ an efficiently computable function, $\varepsilon_{\text{sound}} = \frac{\nu(\lambda)}{m(\lambda)}$, and a negligible function $\nu(\lambda)$ such that $\mu(\lambda) = \nu(\lambda) \cdot \text{negl}(\lambda)$.*

¹² Note the algorithm only has knowledge of the encodings $g^{\mathbf{M}}$, and does not know the discrete logarithms $\mathbf{M}$ themselves.

If the IHBG-friendly ZAP for $\mathcal{L}_{\text{LPWW}}$ is adaptively computational $\varepsilon_{\text{sound}}$ -sound w.r.t. colinear parameters and the DDH assumption is μ -explicitly hard, then IHBG satisfies ν -extractability.

Proof Sketch Recall that the hidden-bits generator receives $g^{\mathbf{M}}$ and m first flows $(r_{\text{ZAP},i})_{i \in [m]}$ of the underlying ZAP. Verification is performed with respect to a matrix $\mathbf{M}'$, which the verifier computes as $g^{\mathbf{M}'} \leftarrow g^{\mathbf{M} - z\mathbf{I}_{m+1}}$ (where the value z is part of the commitment σ outputted by the hidden-bits generator). Intuitively, what this means is that a malicious hidden-bits generator can only influence $g^{\mathbf{M}'}$ with its choice for z . In our proof, the SimCoin simulator will randomly pick $\tilde{z}$ (its guess for z), generate a matrix $\mathbf{M}''$ of rank 1 and compute $g^{\mathbf{M}} := g^{\mathbf{M}'' + \tilde{z}\mathbf{I}_{m+1}}$. Public-coin indistinguishability follows from polynomial DDH; when encoded in the exponent, rank 1 matrices are indistinguishable from full-rank ones.

When guessing z is successful (which happens with probability ν), the matrix $\mathbf{M}'$ will be equal to $\mathbf{M}''$ (of rank 1), and the commitment uniquely determines the hidden-bits string ρ (moreover, we show that there exists an efficient, deterministic algorithm **Open** which extracts ρ).

Extractability requires more work, because we need to remove the problematic extra checks that **Open** succeeds and that $\text{VerifyBit}(\tilde{r}, \sigma, i, 1 - \rho_i, \pi_i) = \perp$ for all $i \in S$. The probability we end up with should not be too far from our starting point (see extractability in Definition 12). In our hybrids, we first switch the real coins to simulated ones. Then, for simulated coins we know that the underlying ZAP satisfies soundness - so we can remove the checks $\text{VerifyBit}(\tilde{r}, \sigma, i, 1 - \rho_i, \pi_i) = \perp$ by relying on the soundness of the underlying ZAP. The extra check that opening works is removed by a statistical argument. Finally, we revert back from simulated coins to real ones. Since all these steps are conditioned on our initial guess of z being correct, we require the explicit hardness of DDH and polynomial DDH does not suffice. For a formal proof, please see the full version of this paper [CKSU21].

Theorem 24 (Statistical Hiding). *If the IHBG-friendly ZAP for $\mathcal{L}_{\text{LPWW}}$ is doubly-adaptive statistically witness indistinguishable, the hash function family $\mathcal{H}$ is universal, and $\nu(\lambda)$ is negligible, then Π_{IHBG} is statistically hiding.*

6 IHBG-Friendly Statistical ZAPs for $\mathcal{L}_{\text{LPWW}}$

In this section, we provide two instantiations for the IHBG-friendly statistical ZAP used in the construction of IHBG from the previous section, one in pairing groups, and one in pairing-free groups. These constructions and their analysis constitute the proofs of Lemma 19 and Lemma 20.

6.1 First Construction: a Statistical ZAP for $\mathcal{L}_{\text{LPWW}}$ in Pairing Groups

For this construction, we employ the Couteau-Hartmann compiler from [CH20]. The high-level idea of the compiler is very simple: assume that the family of

parametrized languages $\mathcal{L}_{\text{LPWW}} = \{\mathcal{L}_{\text{LPWW}}^{\text{par}}\}_{\text{par} \in \Lambda}$ is defined over a group $\mathbb{G}_1$, such that there exists another group $\mathbb{G}_2$ and an asymmetric pairing from $\mathbb{G}_1 \times \mathbb{G}_2$ to a target group $\mathbb{G}_T$. Let $g_2 \in \mathbb{G}_2$ be a generator of $\mathbb{G}_2$. Then, the Couteau-Hartmann compiler converts a Σ -protocol with linear answer for the target language into a statistical ZAP by parsing the random message of the verifier as a pair (g_2, g_2^e) , where e is seen as some random verifier challenge for the Σ -protocol. The compiled ZAP is constructed by computing the first flow of the Σ -protocol normally, and the last flow (which is a linear function of the challenge e with coefficients known to the prover) “in the exponent of g_2 ” using (g_2, g_2^e) . The verification step is carried out using a pairing. Below, we adapt this compiler to the family of parameterized languages $\mathcal{L}_{\text{LPWW}}$ and prove its security.

Construction. Let $(\mathbb{G}_1, \mathbb{G}_2)$ be elliptic curves equipped with an asymmetric pairing $\bullet : \mathbb{G}_1 \times \mathbb{G}_2 \mapsto \mathbb{G}_T$, where $\mathbb{G}_1$ and $\mathbb{G}_2$ both have prime order p . We extend the definition of $\bullet$ to vectors in the conventional manner. Let g_1 be a generator of $\mathbb{G}_1$ and d be a vector length parameter. Let $\text{par} = (\mathbf{g}, \mathbf{h}) \in \Lambda = (\mathbb{G}_1^d \setminus \{\mathbf{1}\})^2$ be the language parameters. We will rely on the Σ -protocol from Section 4.3 with repetition parameter $n = 1$. In particular, we do not require to rely on the adaptive soundness of the Σ -protocol (i.e., Lemma 22) to achieve adaptive soundness (looking ahead, higher value of n (i.e., adaptive soundness of the Σ -protocol) will only be useful in our pairing-free instantiation). The construction of a ZAP for $\mathcal{L}_{\text{LPWW}}$ over $\mathbb{G}_1$ with public coin length $\ell = 2\lceil \log |\mathbb{G}_2| \rceil$, denoted as Π_{ZAP} , is described as follows.

- **ZAP.Prove**(par, r', x, w) : On input parameters $\text{par} = (\mathbf{g}, \mathbf{h}) \in \Lambda$, a public coin $r \in \{0, 1\}^\ell$, a statement $x := (X, Y) \in \mathcal{L}_{\text{LPWW}}^{\text{par}}$ and a witness $w := \mathbf{y} \in \mathbb{Z}_p^d$ such that $(X, Y) = ((\mathbf{g}^\top)^\mathbf{y}, (\mathbf{h}^\top)^\mathbf{y})$, parse r as $(g_2, g_2^e) \in \mathbb{G}_2^2$ and proceed as follows:
 - Pick $\mathbf{z} \leftarrow_r \mathbb{Z}_p^d$ and set $(R, S) \leftarrow ((\mathbf{g}^\top)^\mathbf{z}, (\mathbf{h}^\top)^\mathbf{z})$. Note that this corresponds to computing the first flow of the prover in the Σ -protocol from Section 4.3, with $n = 1$.
 - Set $g_2^{\mathbf{d}} \leftarrow (g_2^e)^\mathbf{y} \cdot g_2^\mathbf{z}$. Note that this corresponds to computing the last flow of the prover in the Σ -protocol from Section 4.3, *in the exponent domain* of $\mathbb{G}_2$.
 - Output $\pi_{\text{ZAP}} = (R, S, g_2^{\mathbf{d}})$.
- **ZAP.Verify**($\text{par}, r, x, \pi_{\text{ZAP}}$) : On input parameters $\text{par} = (\mathbf{g}, \mathbf{h}) \in \Lambda$, a public coin $r \in \{0, 1\}^\ell$, a statement $x = (X, Y)$, and a proof π_{ZAP} , parse π_{ZAP} as $(R, S, g_2^{\mathbf{d}})$, and parse r as $(g_2, g_2^e) \in \mathbb{G}_2^2$. Check that $\mathbf{g}^\top \bullet g_2^{\mathbf{d}} = (X \bullet g_2^e) \cdot (R \bullet g_2)$ and $\mathbf{h}^\top \bullet g_2^{\mathbf{d}} = (Y \bullet g_2^e) \cdot (S \bullet g_2)$. Note that this corresponds to executing the verification procedure of the Σ -protocol from Section 4.3 (with $n = 1$), but using the pairings to emulate the exponentiations of $(\mathbf{g}^\top, \mathbf{h}^\top)$ and (X, Y) (which are all over $\mathbb{G}_1$) by $\mathbf{d}$ and e respectively, since the latter are now only known in the exponent of g_2 .

We prove Lemma 19. Namely, we show our IHBG-friendly ZAP for $\mathcal{L}_{\text{LPWW}}$ satisfies doubly-adaptive perfect witness indistinguishability and adaptive computational $\varepsilon_{\text{sound}}$ -soundness w.r.t. colinear parameters. Due to page limitations,

we provide them in the full version of this paper [CKSU21]. Plugging this IHBG-friendly adaptive statistical ZAP for $\mathcal{L}_{\text{LPWW}}$ into the construction of IHBG of Section 5 and combining it with the construction of statistical ZAP for NP from any IHBG from Section 3, we get our first main theorem:

Theorem 25 (Statistical ZAPs in Pairing Groups). *Assume that the explicit μ -hardness of the DDH assumption holds in a group $\mathbb{G}_1$, and the explicit (μ/m) -hardness of the kernel Diffie-Hellman assumption holds in a group $\mathbb{G}_2$, where $(\mathbb{G}_1, \mathbb{G}_2)$ are groups equipped with a bilinear pairing, m is the output length of the IHBG protocol, and for any negligible function μ (which can be arbitrarily close to an inverse polynomial function). Then there exists an adaptive statistically witness indistinguishable ZAP for NP with non-adaptive computational soundness.*

6.2 Second Construction: a Statistical ZAP for $\mathcal{L}_{\text{LPWW}}$ in Pairing-Free Groups

A Correlation-Intractable Hash Function for $\mathcal{R}_{\text{LPWW}}$. Let λ be the security parameter. We consider a group $\tilde{\mathbb{G}}$ of order $q(\lambda)$ with $\lceil \log q \rceil \approx \lambda$. Let $\text{Trunc} : \tilde{\mathbb{G}} \mapsto \{0, 1\}^{\lambda/2}$ be the function which, on input a group element $\tilde{G} \in \tilde{\mathbb{G}}$, parses it as a $\lceil \log q \rceil$ -bit string and returns the first $\lambda/2$ bits of its input. We consider the following hash function $\mathbf{H} : \tilde{\mathbb{G}}^2 \times \mathbb{Z}_q \mapsto \{0, 1\}^{\lambda/2}$ based on secret key ElGamal:

- Sampling the key: sample $(\tilde{G}, k, m) \leftarrow_r \tilde{\mathbb{G}} \times \mathbb{Z}_q^2$ and set the hash key as $\tilde{\mathbf{C}} \leftarrow_r \text{Enc}_{\tilde{G}}(k, m)$. Note that the key distribution is exactly the uniform distribution over $\tilde{\mathbb{G}}^2$ due to universality (see Definition 5).
- Evaluating $\mathbf{H}(\tilde{\mathbf{C}}, \cdot) : \mathbf{H}(\tilde{\mathbf{C}}, x) = \text{Trunc}(\text{HalfDec}(x, \tilde{\mathbf{C}}))$.

Correlation-Intractability of $\mathbf{H}$. Fix a parameter $n \in \mathbb{N}$. Consider a group $\mathbb{G}$ of order $p(\lambda)$ with $\lceil \log p \rceil \approx \lambda/2n$. Fix a parameter $t \in \mathbb{Z}_p^*$ and define the set of parameters $\Lambda^t := \{(g^{\mathbf{v}}, g^{t \cdot \mathbf{v}})\}_{\mathbf{v} \in \mathbb{Z}_p^{m+1} \setminus \{\mathbf{0}\}} \subset \Lambda = (\mathbb{G}^d \setminus \{\mathbf{1}\})^2$ implicitly parameterized by the security parameter λ . Define $\mathcal{R}_{\text{LPWW}}^{\text{sparse}} = \{\mathcal{R}_{\text{LPWW}, t}^{\text{sparse}}\}_{t \in \mathbb{Z}_p^*}$ to be the natural sparse relation associated to the Σ -protocol of Section 4.3 for the parametrized family of languages $\mathcal{L}_{\text{LPWW}}$, with repetition parameter n . That is,

$$\mathcal{R}_{\text{LPWW}, t}^{\text{sparse}} := \{(\alpha, \beta) \in \mathbb{G}^{2n} \times (\mathbb{Z}_p^*)^n : \exists x, \gamma, \text{par} \in \Lambda_t \text{ s.t. } x \notin \mathcal{L}_{\text{LPWW}}^{\text{par}} \wedge V(x, \alpha, \beta, \gamma) = \top\},$$

where $\alpha := \{(R_i, S_i)\}_{i \in [n]}$, $\beta := \{e_i\}_{i \in [n]}$, and $\gamma := \{\mathbf{d}_i\}_{i \in [n]}$ in Figure 1. Here, the above relation can also be described alternatively using the following (inefficient) randomized function:

$$f_t(\alpha; z) : \begin{cases} \mathbb{G}^{2n} \times \mathbb{Z}_p^* \mapsto (\mathbb{Z}_p^*)^n \\ ((R_i, S_i)_{i \in [n]}, z) \rightarrow (z, ((\log_{(R_1^t/S_1)}(R_i^t/S_i)) \cdot z)_{i \in [2, n]}) \end{cases}.$$

Given this function, it is straightforward (albeit tedious) to check that the relation rewrites to

$$\mathcal{R}_{\text{LPWW},t}^{\text{sparse}} = \{(\alpha, \beta) \in \mathbb{G}^{2n} \times (\mathbb{Z}_p^*)^n : \exists z \in \mathbb{Z}_p^*, f_t(\alpha; z) = \beta\}.$$

The following is the main contribution of this section. Due to page limitation, we provide the proof of the following theorem in the full version of this paper [CKSU21].

Theorem 26. *Assume that ElGamal satisfies $2^{-\lambda/2}$ -OW-KDM security with respect to efficient functions. Let $\mathcal{R}_{\text{LPWW}}^{\text{sparse}} = \{\mathcal{R}_{\text{LPWW},\lambda}^{\text{sparse}}\}_\lambda = \{\{\mathcal{R}_{\text{LPWW},\lambda,t}^{\text{sparse}}\}_{t \in \mathbb{Z}_p^*}\}_\lambda$ be the family of parameterized sparse relation induced by $\mathcal{L}_{\text{LPWW}}$. Then the hash family $\mathcal{H} = \{H : \tilde{\mathbb{G}}^2 \times \mathbb{Z}_q \mapsto \{0,1\}^{\lambda/2}\}_\lambda$ satisfies $(\varepsilon, \mathcal{R}_{\text{LPWW}}^{\text{sparse}})$ -correlation intractability for every negligible function ε satisfying $\varepsilon(\lambda) = 2^{-o(\lambda)}$.*

Remark 27. Theorem 26 should be compared to Theorem 24 from [CKU20]: in [CKU20], the authors restricted their attention to a Σ -protocol with only two parallel repetitions (the language we consider is also different, but this does not matter for the conclusion – both the DDH language from [CKU20] and the LPWW language could be used in their construction). As a consequence, they could only build a correlation-intractable hash function for their relation from the $2^{-3\lambda/4}$ -OW-KDM hardness of ElGamal. By considering the general case of n parallel repetitions, and adjusting n appropriately, we significantly strengthen their conclusion and manage to rely on the $2^{-\lambda/2}$ -OW-KDM hardness of ElGamal. By Definition 6, this means that no PPT adversary has significantly better advantage than $2^{-(1/2+o(1))\cdot\lambda}$, where the $o(1)$ in the exponent can be made smaller than $1/\lambda^\varepsilon$ for any constant $\varepsilon < 1$. Beyond this simple generalization, our analysis is essentially identical to that of [CKU20]; we provide it below for the sake of completeness.

IHBG-Friendly Statistical ZAP for $\mathcal{L}_{\text{LPWW}}$ in Pairing-Free Groups. Equipped with the above correlation-intractable hash function, we are now ready to give our construction of our IHBG-friendly statistical ZAP. We note that this construction will actually satisfy a stronger soundness notion than required for an IHBG-friendly ZAP: adaptive computational soundness will hold for *any* parameters (an not just only for parameters sampled uniformly from $\text{Col}(\mathbb{G}^d)$).¹³ Let $\mathbb{G}$ be a group of order p , and let $\tilde{\mathbb{G}}$ be a group of order q such that $\lceil \log q \rceil \approx \lambda \approx 2\lceil \log p \rceil^2$. Let Π_Σ be the Σ -protocol for $\mathcal{L}_{\text{LPWW}}$, with repetition parameter $n = \lceil \log p \rceil$. Let P_1, P_2 and V be the corresponding algorithms for the first and second move of the prover and the verifier, respectively. let $H : \tilde{\mathbb{G}}^2 \times \mathbb{Z}_q \mapsto \{0,1\}^{\lambda/2}$ be the correlation intractable hash function constructed above.

Construction. The construction of an IHBG-friendly statistical ZAP for $\mathcal{L}_{\text{LPWW}}$ with public coin length $\ell = 2\lceil \log q \rceil$, denoted as Π_{ZAP} , is described as follows.

¹³ We defined the weaker soundness notion since that was all we required to construct the IHBG protocol, and moreover, it was what we could construct from the kernel DH assumption.

- **ZAP.Prove**(par, r, x, w) : On input parameters $\text{par} = (\mathbf{g}, \mathbf{h})$, a public coin $r \in \{0, 1\}^\ell$, a statement $x := (X, Y) \in \mathcal{L}_{\text{PWW}}$ and a witness $w := \mathbf{y}$ such that $(X, Y) = ((\mathbf{g}^\top)^\mathbf{y}, (\mathbf{h}^\top)^\mathbf{y})$, run $\alpha \leftarrow_r P_1(\text{par}, x, w)$ and compute $\beta = H(r, \alpha)$, where r provides the description of the CIH hash H . Parse β as an element of $(\mathbb{Z}_p^*)^n$, and further run $\gamma \leftarrow_r P_2(\text{par}, x, w, \alpha, \beta)$. Finally, output $\pi_{\text{ZAP}} = (\alpha, \gamma)$.
- **ZAP.Verify**($\text{par}, r, x, \pi_{\text{ZAP}}$) : On input parameters $\text{par} = (\mathbf{g}, \mathbf{h})$, a public coin $r \in \{0, 1\}^\ell$, a statement x , and a proof π_{ZAP} , parse π_{ZAP} as $(\alpha, \gamma) \leftarrow \pi_{\text{ZAP}}$. Then, compute $\beta = H(r, \alpha)$ and output $\top$ if $V(\text{par}, x, \alpha, \beta, \gamma) = \top$. Otherwise, output $\perp$.

We prove Lemma 20. Namely, we show our IHBG-friendly ZAP for $\mathcal{L}_{\text{PWW}}$ satisfies doubly-adaptive perfect witness indistinguishability and adaptive computational $\varepsilon_{\text{sound}}$ -soundness w.r.t. colinear parameters. Due to page limitations, we provide them in the full version of this paper [CKSU21]. Plugging this IHBG-friendly adaptive statistical ZAP for $\mathcal{L}_{\text{PWW}}$ into the construction of IHBG of Section 5 and combining it with the construction of statistical ZAP for NP from any IHBG from Section 3, we get our second main theorem:

Theorem 28 (Statistical ZAPs in Pairing-Free Groups). *Assume that the explicit μ -hardness of the DDH assumption holds in a group $\mathbb{G}$ of order p for any negligible function μ (which can be arbitrarily close to an inverse polynomial function), and that the $2^{-\lambda/2}$ -OW-KDM security of ElGamal holds over a group $\mathbb{G}$ of order q such that $\lceil \log q \rceil \approx \lambda \approx 2\lceil \log p \rceil^2$. Then there exists an adaptive statistically witness indistinguishable ZAP for NP with non-adaptive computational soundness.*

7 Acknowledgements

Geoffroy Couteau was supported by ANR SCENE. Shuichi Katsumata was supported by JST CREST Grant Number JPMJCR19F6. Bogdan Ursu was partly supported by ERC PREP-CRYPTO Grant 724307.

References

- BFJ⁺20. S. Badrinarayanan, R. Fernando, A. Jain, D. Khurana, and A. Sahai. Statistical ZAP arguments. In *EUROCRYPT 2020, Part III, LNCS* 12107, pages 642–667. Springer, Heidelberg, May 2020.
- BFM88. M. Blum, P. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th ACM STOC*, pages 103–112. ACM Press, May 1988.
- BGI⁺17. S. Badrinarayanan, S. Garg, Y. Ishai, A. Sahai, and A. Wadia. Two-message witness indistinguishability and secure computation in the plain model from new assumptions. In *ASIACRYPT 2017, Part III, LNCS* 10626, pages 275–303. Springer, Heidelberg, December 2017.

- BP15. N. Bitansky and O. Paneth. ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation. In *TCC 2015, Part II, LNCS* 9015, pages 401–427. Springer, Heidelberg, March 2015.
- CCRR18. R. Canetti, Y. Chen, L. Reyzin, and R. D. Rothblum. Fiat-Shamir and correlation intractability from strong KDM-secure encryption. In *EUROCRYPT 2018, Part I, LNCS* 10820, pages 91–122. Springer, Heidelberg, April / May 2018.
- CH19. G. Couteau and D. Hofheinz. Designated-verifier pseudorandom generators, and their applications. In *EUROCRYPT 2019, Part II, LNCS* 11477, pages 562–592. Springer, Heidelberg, May 2019.
- CH20. G. Couteau and D. Hartmann. Shorter non-interactive zero-knowledge arguments and ZAPs for algebraic languages. In *CRYPTO 2020, Part III, LNCS* 12172, pages 768–798. Springer, Heidelberg, August 2020.
- CKSU21. G. Couteau, S. Katsumata, E. Sadeghi, and B. Ursu. Statistical ZAPs from group-based assumptions. Cryptology ePrint Archive, Report 2021/688, 2021. <https://ia.cr/2021/688>.
- CKU20. G. Couteau, S. Katsumata, and B. Ursu. Non-interactive zero-knowledge in pairing-free groups from weaker assumptions. In *EUROCRYPT 2020, Part III, LNCS* 12107, pages 442–471. Springer, Heidelberg, May 2020.
- DN00. C. Dwork and M. Naor. Zaps and their applications. In *41st FOCS*, pages 283–293. IEEE Computer Society Press, November 2000.
- DN07. C. Dwork and M. Naor. Zaps and their applications. *SIAM J. Comput.*, 36(6):1513–1543, 2007.
- EHK⁺13. A. Escala, G. Herold, E. Kiltz, C. Ràfols, and J. Villar. An algebraic framework for Diffie-Hellman assumptions. In *CRYPTO 2013, Part II, LNCS* 8043, pages 129–147. Springer, Heidelberg, August 2013.
- FLS90. U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st FOCS*, pages 308–317. IEEE Computer Society Press, October 1990.
- FLS99. U. Feige, D. Lapidot, and A. Shamir. Multiple noninteractive zero knowledge proofs under general assumptions. *SIAM J. Comput.*, 29(1):1–28, 1999.
- FS90. U. Feige and A. Shamir. Witness indistinguishable and witness hiding protocols. In *22nd ACM STOC*, pages 416–426. ACM Press, May 1990.
- GJJM20. V. Goyal, A. Jain, Z. Jin, and G. Malavolta. Statistical zaps and new oblivious transfer protocols. In *EUROCRYPT 2020, Part III, LNCS* 12107, pages 668–699. Springer, Heidelberg, May 2020.
- GMR89. S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.
- GO94. O. Goldreich and Y. Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, December 1994.
- GOS06a. J. Groth, R. Ostrovsky, and A. Sahai. Non-interactive zaps and new techniques for NIZK. In *CRYPTO 2006, LNCS* 4117, pages 97–111. Springer, Heidelberg, August 2006.
- GOS06b. J. Groth, R. Ostrovsky, and A. Sahai. Perfect non-interactive zero knowledge for NP. In *EUROCRYPT 2006, LNCS* 4004, pages 339–358. Springer, Heidelberg, May / June 2006.
- GS08. J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *EUROCRYPT 2008, LNCS* 4965, pages 415–432. Springer, Heidelberg, April 2008.

- GW11. C. Gentry and D. Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *43rd ACM STOC*, pages 99–108. ACM Press, June 2011.
- HK07. D. Hofheinz and E. Kiltz. Secure hybrid encryption from weakened key encapsulation. Cryptology ePrint Archive, Report 2007/288, 2007. <https://eprint.iacr.org/2007/288>.
- JKKR17. A. Jain, Y. T. Kalai, D. Khurana, and R. Rothblum. Distinguisher-dependent simulation in two rounds and its applications. In *CRYPTO 2017, Part II, LNCS 10402*, pages 158–189. Springer, Heidelberg, August 2017.
- JZ21. A. Jain and J. Zhengzhong. Non-interactive zero knowledge from sub-exponential ddh. In *EUROCRYPT 2021*, 2021.
- KKS18. Y. T. Kalai, D. Khurana, and A. Sahai. Statistical witness indistinguishability (and more) in two messages. In *EUROCRYPT 2018, Part III, LNCS 10822*, pages 34–65. Springer, Heidelberg, April / May 2018.
- KNYY19. S. Katsumata, R. Nishimaki, S. Yamada, and T. Yamakawa. Designated verifier/prover and preprocessing NIZKs from Diffie-Hellman assumptions. In *EUROCRYPT 2019, Part II, LNCS 11477*, pages 622–651. Springer, Heidelberg, May 2019.
- KS17. D. Khurana and A. Sahai. How to achieve non-malleability in one or two rounds. In *58th FOCS*, pages 564–575. IEEE Computer Society Press, October 2017.
- KW15. E. Kiltz and H. Wee. Quasi-adaptive NIZK for linear subspaces revisited. In *EUROCRYPT 2015, Part II, LNCS 9057*, pages 101–128. Springer, Heidelberg, April 2015.
- LPWW20. B. Libert, A. Passelègue, H. Wee, and D. J. Wu. New constructions of statistical NIZKs: Dual-mode DV-NIZKs and more. In *EUROCRYPT 2020, Part III, LNCS 12107*, pages 410–441. Springer, Heidelberg, May 2020.
- LVW19. A. Lombardi, V. Vaikuntanathan, and D. Wichs. 2-message publicly verifiable WI from (subexponential) LWE. Cryptology ePrint Archive, Report 2019/808, 2019. <https://eprint.iacr.org/2019/808>.
- LVW20. A. Lombardi, V. Vaikuntanathan, and D. Wichs. Statistical ZAPR arguments from bilinear maps. In *EUROCRYPT 2020, Part III, LNCS 12107*, pages 620–641. Springer, Heidelberg, May 2020.
- MRV15. P. Morillo, C. Ràfols, and J. L. Villar. Matrix computational assumptions in multilinear groups. Cryptology ePrint Archive, Report 2015/353, 2015. <https://eprint.iacr.org/2015/353>.
- Nao03. M. Naor. On cryptographic assumptions and challenges (invited talk). In *CRYPTO 2003, LNCS 2729*, pages 96–109. Springer, Heidelberg, August 2003.
- Pas13. R. Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC 2013, LNCS 7785*, pages 334–354. Springer, Heidelberg, March 2013.
- QRW19. W. Quach, R. D. Rothblum, and D. Wichs. Reusable designated-verifier NIZKs for all NP from CDH. In *EUROCRYPT 2019, Part II, LNCS 11477*, pages 593–621. Springer, Heidelberg, May 2019.