Lower Bounds for Non-interactive Zero-Knowledge

Hoeteck Wee*

Computer Science Division
University of California, Berkeley
hoeteck@cs.berkeley.edu

Abstract. We establish new lower bounds and impossibility results for non-interactive zero-knowledge proofs and arguments with set-up assumptions.

- For the common random string model, we exhibit a lower bound for the trade-off between hardness assumptions and the length of the random string for non-interactive zero-knowledge *proofs*. This generalizes a previous result ruling out non-interactive zero-knowledge proofs for non-trivial languages with a random string of length $O(\log n)$.
- In the registered public key model, we show that there does not exist a non-interactive zero-knowledge *proof* for a non-trivial language.
- In the bare public key model with fully nonuniform simulation wherein the size of the simulator is also allowed to depend on the size of the distinguisher and the distinguishing gap, there does not exist a non-interactive zero-knowledge *proof* for an NP-complete language, unless the polynomial hierarchy collapses. On the other hand, there is a non-interactive zero-knowledge *argument* for all of NP with a fully nonuniform simulator.

Our negative results complement upper bounds and feasibility results from previous work.

Keywords: Non-interactive zero-knowledge, set-up assumptions, lower bounds.

1 Introduction

The seminal notion of *zero-knowledge proofs*, namely proofs that yield no knowledge beyond the validity of the assertion proved, was introduced by Goldwasser, Micali and Rackoff [GMR89]. Formally, an interactive protocol is zero-knowledge if there exists a simulator that can simulate the behavior of every, possibly malicious, verifier without access to the prover, such that its output is indistinguishable from the output of the verifier after having interacted with the honest prover.

Minimizing the number of rounds is an important goal in design of zero-knowledge proof systems. A lower bound was established by Goldreich and Oren [GO94], who showed that at least three rounds of interaction are necessary to achieve auxiliary-input zero-knowledge. To understand and overcome this limitation, recent work has focused on both impossibility and feasibility results for weaker notions of zero-knowledge. The relaxations include limiting the power of malicious verifiers [BLV06], limiting prover resources [DS02], quasipolynomial-time simulation [P03, BP04], and witness indistinguishability [DN00, BOV03].

* Part of this work was completed while visiting IBM T.J. Watson Research Center and IPAM, Los Angeles.

1.1 Non-interactive Zero-Knowledge with Set-Up

A different way to bypass the lower bound on interaction is to introduce set-up assumptions. This approach was initiated by Blum, Feldman and Micali [BFM88], who showed how to realize a non-interactive zero-knowledge protocol for NP, comprising a single message from the prover to the verifier. In this work, we will focus on set-up assumptions with a “public key” flavor, presented in decreasing order in the amount of the trust the prover and verifier needs to put in the set-up:

- *Common random string (CRS) model*. This is the original model proposed by Blum et al., wherein both the prover and the verifier receive a truly random string from a trusted party. A slight relaxation of this model is the *common reference string model*, wherein both parties receive a random string chosen accordingly to some polynomial-time samplable distribution.
- *Registered public key model*. Barak et al. [BCNP04] introduced the registered public key model as a relaxation of the CRS model under which general multi-party computations can still be securely realized within the UC framework. In addition, they showed how to realize non-interactive zero-knowledge in this model.¹ We will restrict ourselves to the special case wherein only the verifier registers a “public key” with a “registration authority”. An honest verifier upon registration receives a randomly generated public key, and does not need to keep the secret data used for key generation, whereas a cheating verifier may register any public key of its choice, but must provide the secret data associated with the key to the registration authority. Prior to participating in the protocol, the prover obtains the verifier’s key from the registration authority.
- *Bare public key model*. In the bare public key model introduced by Canetti et al. [CGGM00], the verifier again has a public key that has been registered prior to interacting with the prover. Here, there is no trusted “registration authority” that verifies (and enforces) properties of the registered key. In particular, an honest verifier registers a randomly generated key, whereas a cheating verifier may register any arbitrary key, possibly even a malformed one.

We stress that in each of these models, the proofs are publicly verifiable - verification does not require verifier’s secret key. Note that the bare public key model imposes the strongest requirements on the simulator (minimal trust requirements) whereas the common reference string model imposes the weakest requirement (maximal trust requirements: a non-interactive protocol that is zero-knowledge in the bare public key model is also zero-knowledge in the registered public key model, and a protocol satisfying the latter is zero-knowledge in the common reference string model. We refer the reader to [CGGM00, BCNP04] for further cryptographic motivations for these set-up assumptions.

1.2 Weak Nonuniform Zero Knowledge

A non-interactive zero-knowledge protocol in the bare public key model is essentially a 2-round zero-knowledge protocol without set-up assumptions, except the verifier’s first

¹ Specifically, Barak et al. demonstrated a non-interactive protocol that realizes the UC zero-knowledge functionality, which implies zero-knowledge.

message must be independent of the instance. This means that the result of Goldreich and Oren [GO94] also rules out non-interactive zero-knowledge argument systems in the bare public key model for languages outside BPP. Therefore, we will relax the zero-knowledge requirement for the bare public key model in the following ways (as has previously been done for general zero-knowledge in [DNRS03]):

- We allow the simulator to depend nonuniformly on the cheating verifier (namely that for every nonuniform probabilistic polynomial-time cheating verifier, there is a nonuniform probabilistic polynomial-time simulator), with the additional guarantee of a polynomial relation between the size of the verifier and that of the simulator.² The main difference from auxiliary-input zero knowledge is the latter guarantees a single (uniform) polynomial-time algorithm that on input a description of the verifier, outputs a description of the simulator.
- Next, we allow the size of the simulator to depend on that of the distinguisher and the distinguishing gap.³ In particular, this guarantee (by itself, without nonuniformity) implies quasipolynomial-time simulation [P03] (and therefore the security guarantee is stronger than that offered by quasipolynomial-time simulation).

We refer to this relaxation as *weak nonuniform zero knowledge*, and it is meaningful also in the standard model without set-up assumptions. Informally, weak nonuniform zero knowledge guarantees that an efficient verifier can approximate whatever he learns from interacting with the prover within any inverse polynomial factor at a price of a polynomial blow-up in the running time and a polynomial amount of help (corresponding to the nonuniformity). One can regard the simulator's nonuniformity as a measure of the *knowledge* leaked by an interaction with the honest prover.

As with the standard notion of zero knowledge, weak nonuniform zero knowledge implies witness indistinguishability (in the standard sense with a negligible distinguishing gap). More generally, weak nonuniform zero knowledge is sufficient in applications wherein the zero knowledge property is only used to construct an intermediate hybrid distribution in order to establish computational indistinguishability. This occurs for instance in the construction of non-malleable cryptographic schemes by Dolev, Dwork and Naor [DDN00] (as pointed out in [DNRS03]) and specifically, in their non-malleable bit commitment scheme and Sahai's CCA2-secure encryption scheme [S99]. We stress that nonuniform zero knowledge yields nonuniform reductions in the proof of security for these schemes.

An analogous relaxation for circuit obfuscation as formalized by Barak et al. [BG⁺01] was previously presented in [W05].⁴ Indeed, the only known positive results for obfuscation in the standard model [C97, W05] merely achieve this weaker

² Refer to [G01, Sec 4.3.3] for a discussion of this relaxation.

³ Dwork et al. [DNRS03] had earlier considered $S(V, T, D)$ zero-knowledge, the simulator is allowed to depend on both the verifier V and the distinguisher T (whereas we allow a dependency on the size of T but not T itself). A dependency on the distinguishing gap was introduced by Dwork, Naor and Sahai in their work on concurrent zero-knowledge [DNS04].

⁴ In the formalization proposed in [W05], the simulator is allowed to depend on the distinguisher, although it is easy to verify that the virtual black-box simulators for the constructions in [C97, W05] only need to depend on the size of the distinguisher and the distinguishing gap.

requirement. One of the motivations for this work is to understand whether this relaxation may also be meaningfully exploited in the context of zero-knowledge.

1.3 Our Contributions

We present lower bounds and impossibility results for non-interactive zero-knowledge proofs and arguments with set-up assumptions, along with matching upper bounds and feasibility results. Our main contributions are in the lower bounds; the protocols, apart from the one for the bare public key model, follow readily from previous work [KP98, P03, BCNP04]. We stress that understanding fundamental limitations do facilitate protocol design in narrowing down possible approaches. As a whole, our results complement known protocols to provide a clearer picture of the qualitative differences in the various set-up assumptions, as well as better insight into what can and cannot be realized in each of these models, and why.

Common random string model. We already know how to construct non-interactive zero-knowledge proof systems for NP in the common random string model [BFM88, FLS99]. We establish a lower bound on the trade-off between hardness assumptions and length of the common random string used in these constructions:

Informal Theorem. [Lower bound] In the common random string model, if there is a polynomial-time algorithm for CIRCUIT-SAT with $\ell(n)$ variables, then non-interactive zero-knowledge proof systems with a random string of length $\ell(n)$ only exist for languages in BPP.

Informal Theorem. [Upper bound] In the common random string model, under $\ell^{-1}(n)$ -hardness assumptions for enhanced trapdoor permutations, there is a non-interactive zero-knowledge proof system for all of NP with a CRS of length $\text{poly}(\ell(n))$.

The trade-off achieved in the upper bound is widely believed to be optimal (up to polynomial factors in the length of the CRS) but has not been formally stated; we provide and prove a formal statement to that effect. In the proof, we use probabilistic hashing techniques from [GS89] to address an issue related to randomness-efficient sampling [DI06]. We point out two special cases of our lower bound: to achieve a CRS of length $\text{poly}(\log n)$, sub-exponential hardness assumptions for CIRCUIT-SAT are indeed necessary. Also, if the CRS has length $O(\log n)$, then the language is in BPP (since there is a trivial CIRCUIT-SAT with an exponential dependency on the number of variables). This special case (which extends to readily to arguments, unlike the general case) is folklore and was stated without proof in [DDP97].

Registered public key model. For the registered public key model, we establish a separation between proof systems and argument systems:

Informal Theorem. [Impossibility] In the registered public key model, non-interactive zero-knowledge proof systems only exist for languages in BPP.

Informal Theorem. [Feasibility] In the registered public key model, under sub-exponential hardness assumptions for enhanced trapdoor permutations, there is a non-interactive zero-knowledge argument system for all of NP.

In the registered public key model, the only advantage the simulator has for generating accepting transcripts for YES instances is the secret key corresponding to the public key. However, a computationally unbounded adversary can easily sample a secret key corresponding to the public key, and then run the simulator. This will yield an accepting transcript for YES instances, but not for NO instances if the protocol is a proof system. As such, we can have non-trivial non-interactive argument systems but not proof systems in the registered public key model.

Bare public key model. For the weak non-uniform non-interactive zero-knowledge in the bare public key model, we also establish a separation between proof systems and argument systems. Our feasibility result shows that weak nonuniform simulation can indeed be meaningfully exploited in the context of zero-knowledge. Both of the following results refer to the weak non-uniform setting:

Informal Theorem. [Impossibility] In the bare public key model, non-interactive zero-knowledge proof systems only exist for languages in coNP/poly . In particular, there is no non-interactive zero-knowledge proof system for all of NP unless the polynomial hierarchy collapses [KL80].

Informal Theorem. [Feasibility] In the bare public key model, under sub-exponential hardness assumptions for enhanced trapdoor permutations, there is a non-interactive zero-knowledge argument system for all of NP.

We exploit derandomization via nonuniformity using the probabilistic method for both results.

We use the same protocol, namely a straight-forward adaptation of Pass's 2-round public-coin zero-knowledge argument [P03], for the feasibility results in the registered and bare public key models, with somewhat different simulators. One can therefore view the weak non-uniform zero-knowledge as a fall-back guarantee provided by Pass's protocol in the registered public key model: even if the assumption about the verifier's key being well-formed is not satisfied, the protocol still guarantees non-uniform zero-knowledge, which implies witness indistinguishability.

1.4 Additional Related Work

Lower bounds for non-black-box zero-knowledge. On the whole, our lower bounds and impossibility results make use of many of the insights and techniques from the work of Barak, Lindell and Vadhan [BLV06] on lower bounds for zero-knowledge, specifically, those for 2-round zero-knowledge protocols against uniform adversaries. The latter is an arguably less natural notion than non-interactive zero-knowledge protocols with set-up assumptions, while imposing the technical constraint of uniformity. Indeed, our results show that the ideas from [BLV06] are applicable to a more general setting.

Unconditional characterizations of NIZK. Another closely related work is that of Pass and Shelat [PS05] providing a systematic unconditional study of non-interactive proof systems in the common reference string model, and the secret parameter variant thereof,

wherein the verifier also has a secret key corresponding to the public key. A natural next project would be to extend their work to the registered public key and bare public key models and to argument systems (after all, our lower bounds indicate that the limitation to proof systems may be too restrictive), and to extend our work to secret parameter variants of the different models.

2 NIZK with Set-Up

2.1 Non-interactive Protocols with Set-Up

We consider a set-up phase, parameterized by a (deterministic) function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$, that represents a method for computing a public set-up key given a secret (and supposedly random) key. An protocol (P, V) is a *non-interactive proof system with set-up* for a language L if there is a relation R such that $L = L_R$, a set-up function f such that the following holds:

COMPLETENESS. If $(x, w) \in R$,

$$\Pr[\sigma \leftarrow f(U_k); \pi \leftarrow P(x, w, \sigma) : V(x, \sigma, \pi) = 1] \geq 2/3$$

SOUNDNESS. If $x \notin L$, then for every P^* ,

$$\Pr[\sigma \leftarrow f(U_k); \pi \leftarrow P^*(x, \sigma) : V(x, \sigma, \pi) = 1] \leq 1/3$$

We say that a protocol has *perfect completeness* if the expression $2/3$ is replaced by 1 , and *negligible soundness* if the expression $1/3$ is replaced by $\text{neg}(|x|)$. We say that (P, V) is a *non-interactive argument system with set-up* if the soundness condition is replaced by:

COMPUTATIONAL SOUNDNESS. If $x \notin L$, then for every nonuniform PPT P^* ,

$$\Pr[\sigma \leftarrow f(U_k); \pi \leftarrow P^*(x, \sigma) : V(x, \sigma, \pi) = 1] \leq 1/3$$

We note that our positive results satisfy the following stronger notion of soundness:

ADAPTIVE SOUNDNESS. For every P^* ,

$$\Pr[\sigma \leftarrow f(U_k); (x, \pi) \leftarrow P^*(\sigma) : x \notin L \text{ and } V(x, \sigma, \pi) = 1] \leq 1/3$$

We emphasize that in the formulations of completeness and soundness, both parties receive a randomly generated public key, and the verifier does not receive the secret randomness used to generate the key.

2.2 Non-interactive Zero-Knowledge

Since our main contributions are the negative results, we present the weakest possible notion of security (in particular, we consider non-adaptive zero-knowledge). Establishing lower bounds for weaker notions makes our results stronger. We only present definitions for zero-knowledge in the non-interactive setting.

ZERO-KNOWLEDGE IN COMMON REFERENCE STRING MODEL. There exists a PPT simulator S such that the following distributions are nonuniformly computationally indistinguishable:

$$\{\sigma \leftarrow f(U_k); \pi \leftarrow P(x, w, \sigma) : (\sigma, \pi)\}_{(x,w) \in R}$$

$$\text{and } \{(\sigma, \pi) \leftarrow S(x) : (\sigma, \pi)\}_{(x,w) \in R}$$

We refer to the special case where f is the identity as the Common Random String model.

ZERO-KNOWLEDGE IN REGISTERED PUBLIC KEY MODEL. There exists a PPT simulator S such that the following distributions are nonuniformly computationally indistinguishable:

$$\{P(x, w, f(r))\}_{(x,w) \in R, r \in \{0,1\}^k} \text{ and } \{S(x, r)\}_{(x,w) \in R, r \in \{0,1\}^k}$$

ZERO-KNOWLEDGE IN BARE PUBLIC KEY MODEL. There exists a PPT simulator S such that the following distributions are nonuniformly computationally indistinguishable:

$$\{P(x, w, \sigma)\}_{(x,w) \in R, \sigma \in \{0,1\}^{\text{poly}(k)}} \text{ and } \{S(x, \sigma)\}_{(x,w) \in R, \sigma \in \{0,1\}^{\text{poly}(k)}}$$

Note that zero-knowledge in the bare public key model implies zero-knowledge in the registered public key model, which in turn implies zero-knowledge in the common reference string model. Also, recall that in the definition of (auxiliary-input) zero-knowledge in the interactive setting, it suffices to consider deterministic cheating verifiers; for the same reason, once we have established the zero-knowledge property for a fixed public key in the registered public key and bare public key models, we derive the zero-knowledge property for any (adversarial) distribution over public keys.

3 Common Random String Model

3.1 Lower Bounds

Theorem 1. *If a language L has a non-interactive zero-knowledge proof system in the common string model with a CRS of length $\ell(n)$ (where n is the length of the instance) and there exists a probabilistic $\text{poly}(\ell^{-1}(\#variables), \text{circuitsize})$ algorithm for the CIRCUIT-SAT Problem, then $L \in \text{BPP}$.*

By zero-knowledge and soundness, the distribution of the simulated random strings is pseudorandom for YES instances and statistically far from uniform for NO instances. The idea is to use the CIRCUIT-SAT algorithm to design an efficient test that

- outputs 1 with high probability for samplable distributions over $\{0, 1\}^{\ell(n)}$ that are statistically far from uniform.
- outputs 1 with small probability on the uniform distribution over $\{0, 1\}^{\ell(n)}$.

The latter will correspond to YES instances and the former will correspond to NO instances. The difficulty lies in that the sampling algorithm may use $\text{poly}(n)$ bits of randomness, so we cannot directly test if the input lies in the support of the sampling distribution. To overcome this, we use pairwise independent sampling to reduce the randomness complexity of the sampling algorithm. This is inspired by the Goldwasser-Sipser protocol for proving lower bounds on set sizes [GS89]; the formal analysis is also very similar.

Proof (sketch). Suppose L has a non-interactive zero-knowledge proof system (P, V) in the CRS model with a CRS of length $\ell = \ell(n)$ and a simulator S . We modify the proof system to satisfy the following additional properties:

- The completeness and soundness errors are both at most $1/64$. This can be achieved using randomness-efficient error reduction while increasing the CRS by an additive $O(1)$ bits [DDP02], although naive parallel repetition with a $O(1)$ multiplicative increase is fine too.
- On every input x , the simulator S always outputs accepting transcripts, and the distinguishing error for YES instances is at most $1/32$.

Let r denote the number of random bits used by S , and let S_1 be S with the output truncated to just the simulated CRS. Consider the following algorithm M for deciding L : on input x ,

1. Run $S(x)$ for n independent iterations to obtain transcripts (σ_i, π_i) , $i = 1, 2, \dots, n$. In addition, pick n independent pairwise-independent hash functions $h_i : \{0, 1\}^{\ell-2} \rightarrow \{0, 1\}^r$.
2. Reject if for the majority of $i = 1, 2, \dots, n$, we have $(\sigma_i, x, h_i) \in L_{\text{aux}}$, where

$$L_{\text{aux}} = \left\{ (\sigma, x, h) \mid \exists u \in \{0, 1\}^{\ell-2} \text{ s.t. } S_1(x; h(u)) = \sigma \right\} \in \text{BPP}$$

For each i , we show that $\Pr_{h_i, \sigma_i}[(\sigma_i, x, h_i) \in L_{\text{aux}}]$ is small for $x \in L$ and large for $x \notin L$:

- $x \in L$. By a union bound,

$$\Pr_{h_i, \sigma_i}[(\sigma_i, x, h_i) \in L_{\text{aux}}] \leq \Pr_{h_i}[(U_\ell, x, h_i) \in L_{\text{aux}}] + \frac{1}{32} \leq \frac{2^{\ell-2}}{32} + \frac{1}{32} < \frac{1}{3}$$

- $x \notin L$. By soundness, $|S_1(x; \{0, 1\}^r)| \leq \frac{1}{64} \cdot 2^\ell$. Let A be the set of “low probability” strings in $S_1(x; \{0, 1\}^r)$, that is,

$$A = \left\{ \sigma : \Pr[S_1(x; U_r) = \sigma] \leq \frac{1}{2^{\ell-4}} \right\}$$

A union bound yields

$$\Pr[\sigma_i \in A] \leq \frac{1}{64} \cdot 2^\ell \cdot \frac{1}{2^{\ell-4}} = \frac{1}{4}$$

On the other hand, for the “high probability” strings, a standard analysis via the Chebyshev inequality yields

$$\Pr_{h_i}[(\sigma_i, x, h_i) \notin L_{\text{aux}} \mid \sigma_i \notin \Lambda] \leq \frac{1}{4}$$

Hence,

$$\Pr_{h_i, \sigma_i}[(\sigma_i, x, h_i) \in L_{\text{aux}}] \geq \frac{1}{2}$$

Hence, M is a BPP algorithm for deciding L . □

3.2 Upper Bounds

The following result follows from a variant of the Kilian-Petrank non-interactive zero-knowledge proof system for NP in the CRS model [KP98] (alluded to in [GOS06]) wherein the length on the random string depends polynomially on the security parameter (and not the length of the instance). The idea is to rewrite the input as a conjunction of a polynomial number of constant-sized statements and prove each of these statements using the same CRS (as in [FLS99]).

Proposition 1 ([KP98, GOS06]). *Suppose there exist enhanced trapdoor permutations secure against $\ell^{-1}(n)^{\omega(1)}$ -size circuits. Then, there exists a non-interactive zero-knowledge proof system for NP in the common random string model wherein the CRS has length $O(\ell(n)^3)$ (where n is the length of the instance). In addition, the proof system has perfect completeness, negligible adaptive soundness error and an efficient prover.*

4 Registered Public Key Model

4.1 Impossibility Results

Theorem 2. *If a language L has a non-interactive zero-knowledge proof system in the registered public key model, then $L \in \text{BPP}$.*

Proof. Consider the following algorithm M for deciding L : on input x ,

1. pick $r \leftarrow U_k$.
2. accept iff $V(x, f(r), S(x, r))$ accepts.

Completeness and zero-knowledge guarantees that for all $x \in L$, M accepts with probability at least $2/3 - \text{neg}(|x|)$. Next, consider a (unbounded) cheating prover P^* that for all $x \notin L$ and all σ , outputs π such that $V(x, \sigma, \pi) = 1$ if such a π exists, and $\perp$ otherwise. Then, for all $x \notin L$ and $\sigma \in f(\{0, 1\}^k)$,

$$\Pr[V(x, \sigma, P^*(x, \sigma)) = 1] \geq \Pr_{r: f(r)=\sigma}[V(x, \sigma, S(x, r)) = 1]$$

Averaging over σ , we obtain, for all $x \notin L$,

$$\begin{aligned} \Pr[M(x) = 1] &= \Pr_r[V(x, f(r), S(x, r)) = 1] \\ &\leq \Pr_r[V(x, f(r), P^*(x, f(r))) = 1] \\ &\leq 1/3 \quad (\text{by soundness}) \end{aligned}$$

Hence, M is a BPP algorithm for deciding L . $\square$

4.2 Feasibility Results

Indeed, by relaxing the soundness requirement to computational soundness, Barak et al. constructed a non-interactive UC zero-knowledge protocol in the registered public key model [BCNP04]. The protocol requires that the prover also has a public key in order to achieve additional guarantees required by universal composability. We observe that it is not necessary for the prover to register a key if zero-knowledge is our only goal (but paying the price of subexponential hardness assumptions); in particular, we may use the variant of Pass's protocol [P03] shown in Fig 1.

Proposition 2. *Suppose there exist enhanced trapdoor permutations secure against 2^{n^δ} -size circuits for some constant $\delta > 0$. Then, there exists a non-interactive zero-knowledge argument system for NP in the registered public key model. In addition, the argument system has perfect completeness, negligible adaptive soundness error and an efficient prover.*

Note that our negative results do not extend to the secret parameter model. There, Pass, Shelat, Vaikuntanathan [PSV06] constructed a non-interactive zero-knowledge proof system for all of NP assuming the existence of standard trapdoor permutations (or any semantically secure encryption scheme).

5 Bare Public Key Model

5.1 Weak Nonuniform Zero Knowledge

As noted in the introduction, the lower bound of Goldreich and Oren [GO94] also extends to the bare public key model:

Theorem 3 (implicit in [GO94]). *If a language L has a non-interactive zero-knowledge argument system in the bare public key model, then $L \in \text{BPP}$.*

As such, we will focus on weak nonuniform zero-knowledge in the bare public model. We say that a nonuniform PPT A has size s if the running time and the length of the nonuniform advice for A is bounded by s . Two distributions A, B are (s, ϵ) -indistinguishable if for every nonuniform PPT D of size s , $|\Pr[D(A) = 1] - \Pr[D(B) = 1]| < \epsilon$. Unlike the uniform setting, we need to define zero-knowledge for distributions over public keys chosen by an adversarial verifier V^* .

WEAK NON-UNIFORM ZERO-KNOWLEDGE IN BARE PUBLIC KEY MODEL. There exists a polynomial p such that for every function $s(n) = n^{O(1)}$ and $\epsilon(n) = 1/n^{O(1)}$, and for every nonuniform PPT V^* of size s , there exists a nonuniform PPT S of size $p(n, s, 1/\epsilon)$ such that for all sufficiently large n and for all $(x, w) \in R$ with $|x| = n$, the following distributions

$$\{(\tau, \sigma) \leftarrow V^*(1^n); (\tau, \sigma, P(x, w, \sigma))\} \text{ and } \{(\tau, \sigma, \pi) \leftarrow S(x); (\tau, \sigma, \pi)\}$$

are $(s(n), \epsilon(n))$ -indistinguishable.

We stress once again that the definition allows for the size of the simulator to depend on s , an upper bound on the sizes of the malicious verifier and the distinguisher, and on ϵ , the distinguishing gap, although the dependency is determined by a *fixed* polynomial p .

5.2 Impossibility Results

Theorem 4. *If a language L has a weak nonuniform non-interactive zero-knowledge proof system in the bare public model, then $L \in \text{P/poly}$.*

Proof. The idea behind the proof is to use the probabilistic method to derandomize the verifier in the NIZK proof system and obtain a polynomial number of deterministic nonuniform verifiers with some randomness hardwired into it. We then use the nonuniform simulators for these verifiers to decide the language.

Fix an input length n , and by the probabilistic method, there exists a set $A \subseteq \{0, 1\}^{\text{poly}(n)}$ of polynomial size satisfying the following properties:

- for all $x \in L \cap \{0, 1\}^n$ and a fixed witness w for each x ,

$$\left| \Pr_{\alpha \in \{0, 1\}^{\text{poly}(n)}} [V(x, f(\alpha), P(x, w, f(\alpha))) = 1] - \Pr_{\alpha \in A} [V(x, f(\alpha), P(x, w, f(\alpha))) = 1] \right| < \frac{1}{12}$$

where the probabilities are also taken over the coin tosses of the prover.

- for all $x \in \{0, 1\}^n \setminus L$,

$$\left| \Pr_{\alpha \in \{0, 1\}^{\text{poly}(n)}} [\exists \pi : V(x, f(\alpha), \pi) = 1] - \Pr_{\alpha \in A} [\exists \pi : V(x, f(\alpha), \pi) = 1] \right| < \frac{1}{12}$$

Now, for each $r \in A$, consider the malicious verifier V_r^* with r hardwired into it and sends $f(r)$ as its public key, and the class of distinguishers $\{D_{x,r} \mid x \in \{0, 1\}^n\}$ that on input a transcript $(r', \sigma'\pi)$ accepts iff $r' = r$ and $V(x, \pi) = 1$. Let S_r denote the nonuniform PPT simulator for V_r^* with distinguishing probability $\frac{1}{12}$ and which fools $\{D_{x,r} \mid x \in L \cap \{0, 1\}^n\}$. Hence, for all $x \in \{0, 1\}^n$:

$$\begin{aligned} \text{YES instance:} \quad & \Pr_{r \in A} [D_{x,r}(S_r(x)) = 1] > \frac{2}{3} - \frac{1}{12} - \frac{1}{12} = \frac{1}{2} \\ \text{NO instance:} \quad & \Pr_{r \in A} [D_{x,r}(S_r(x)) = 1] < \frac{1}{3} + \frac{1}{12} = \frac{5}{12} \end{aligned}$$

where the probabilities are also taken over the coin tosses of S_r . By hardwiring A and $\{S_r \mid r \in A\}$ as nonuniform advice, we obtain $L \in \text{BPP/poly} = \text{P/poly}$. $\square$

Remark 1. The analogous result in [BLV06] requires that the proof system has either perfect completeness or an efficient prover.

5.3 Feasibility Results

The idea is to derandomize the adversary and the distinguisher and hardwire the trapdoor information about the public key into the simulator.

Theorem 5. *Suppose there exist enhanced trapdoor permutations secure against 2^{n^δ} -size circuits for some constant $\delta > 0$. Then, there exists a weak nonuniform non-interactive zero-knowledge argument system for NP in the bare public key model. Furthermore, the argument system has perfect completeness, negligible soundness error and an efficient prover.*

Proof. Let L be an NP-complete language for some relation R . Under the assumed trapdoor permutation family, we can construct the following primitives:

- a one-way permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ secure against 2^{n^δ} -sized circuits;
- a non-interactive (perfectly binding, computationally hiding) commitment scheme Com that can be broken (that is, recover the plaintext from the commitment) in time $2^{n^{\delta/2}}$; and
- a zap system [DN00], namely a 2-round public-coin witness-indistinguishable proof system for NP. For simplicity and ease of presentation, we present the protocol and analysis assuming the existence of a 1-round zap (e.g. [BOV03]); for a 2-round zap, we include the first round message as part of the public key.

The argument system for L is shown in Fig 1. The completeness property of this protocol follows from that of the zap system. To prove computational soundness, consider a nonuniform PPT cheating prover P^* that convinces the honest verifier to accept some $x \notin L$ with non-negligible probability. By adaptive soundness of the zap system, the commitment sent by P^* must contain the value $\pi^{-1}(\sigma)$, which can be extracted in time $2^{O(n^{\delta/2})}$. Hence, we derive from P^* a nonuniform algorithm running in time $2^{O(n^{\delta/2})}$ and inverts π with non-negligible probability, a contradiction.

To prove weak nonuniform zero-knowledge, fix s, ϵ , a nonuniform PPT V^* and an input length n . Consider the following distributions for each $(x, w) \in R$ with $|x| = n$:

- Hybrid H_1 . This is the distinguisher's view in an interaction with the honest prover.

$$\{(\tau, \sigma), \text{Com}(0^n), P_{\text{zap}}((x, \text{Com}(0^n)), (w, \perp)); (\tau, \sigma) \xleftarrow{R} V^*(U_s)\}$$

- Hybrid H_2 . This is the distinguisher's view when the prover commits to $\pi^{-1}(\sigma)$ instead of 0^n .

$$\{(\tau, \sigma), \text{Com}(\pi^{-1}(\sigma)), P_{\text{zap}}((x, \text{Com}(\pi^{-1}(\sigma))), (w, \perp)); (\tau, \sigma) \xleftarrow{R} V^*(U_s)\}$$

- Hybrid H_3 . We modify H_2 so that the prover uses $\pi^{-1}(\sigma)$ (and the private randomness used for the commitment) instead of w as the witness in the zap system.

$$\{(\tau, \sigma), \text{Com}(\pi^{-1}(\sigma)), P_{\text{zap}}((x, \text{Com}(\pi^{-1}(\sigma))), (\perp, \pi^{-1}(\sigma))); (\tau, \sigma) \xleftarrow{R} V^*(U_s)\}$$

Set-up function: $f(r) = \pi(r)$, where π is a permutation.

Common input: An instance $x \in \{0, 1\}^n$, public key σ .

Prover's private input: A witness $w \in \{0, 1\}^{\text{poly}(n)}$.

$P \rightarrow V$: Send $z = \text{Com}(0^n)$ and a zap proving the statement “ $x \in L$ OR z is a commitment to $\pi^{-1}(\sigma)$ ” using witness w .

Fig. 1. Variant of Pass's protocol [P03] for an NP-complete language L

Note that H_1 and H_2 are $(s, \epsilon/4)$ -indistinguishable by the hiding property of Com , and that H_2 and H_3 are $(s, \epsilon/4)$ -indistinguishable by witness indistinguishability of the zap system. Observe that for a fixed choice of x and coin tosses for Com and P_{zap} , a sample from the distribution H_3 may be computed as a deterministic function of the choice of random coin tosses for V^* . Hence, by the probabilistic method, there exists a set $\Lambda \subseteq \{0, 1\}^s$ of size $\Theta((s \log s + p'(n))/\epsilon^2)$ where p' is a fixed polynomial equal to $|x|$ plus the total randomness used by Com and P_{zap} , such that the following distribution H_4 is $(s, \epsilon/4)$ -indistinguishable from H_3 .

- Hybrid H_4 . We modify H_3 so that we replace V^* 's coin tosses with a random sample from Λ , where $\Lambda \subseteq \{0, 1\}^s$ is to be determined. We stress that Λ only depends on $|x|$, s , ϵ and not on x itself.

$$\{(\tau, \sigma), \text{Com}(\pi^{-1}(\sigma)), P_{\text{zap}}((x, \text{Com}(\pi^{-1}(\sigma))), (\perp, \pi^{-1}(\sigma)), \sigma);$$

$$(\tau, \sigma) = V^*(r'), r' \xleftarrow{R} \Lambda\}$$

By hardwiring Λ and $\{\pi^{-1}(\sigma) \mid r' \in \Lambda, (\tau, \sigma) = V^*(r')\}$ as nonuniform advice, we have a nonuniform PPT S of size $O(|\Lambda|n + s) = \text{poly}(n, s, 1/\epsilon)$ that on input x , computes the distribution H_4 , which is (s, ϵ) -indistinguishable from H_1 . $\square$

Acknowledgements

I am very grateful towards Ran Canetti and Vinod Vaikuntanathan for many interesting discussions on non-interactive zero-knowledge, and in particular, for raising the question of NIZKs with short CRS. I would also like to thank Moni Naor, Alon Rosen, Luca Trevisan and Salil Vadhan for helpful discussions on an earlier version of this work, Joe Kilian for clarifications regarding [KP98], and the anonymous referees for various feedback and pointers.

References

- [BCNP04] B. Barak, R. Canetti, J. B. Nielsen, and R. Pass. Universally composable protocols with relaxed set-up assumptions. In *Proc. 45th FOCS*, 2004.
- [BFM88] M. Blum, P. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications. In *Proc. 20th STOC*, 1988.

- [BG⁺01] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In *Proc. Crypto '01*, 2001.
- [BLV06] B. Barak, Y. Lindell, and S. Vadhan. Lower bounds for non-black-box zero knowledge. *JCSS*, 72(2):321–391, 2006.
- [BOV03] B. Barak, S. J. Ong, and S. Vadhan. Derandomization in cryptography. In *Proc. Crypto '03*, 2003.
- [BP04] B. Barak and R. Pass. On the possibility of one-message weak zero-knowledge. In *Proc. 1st TCC*, 2004.
- [C97] R. Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In *Proc. Crypto '97*, 1997.
- [CGGM00] R. Canetti, O. Goldreich, S. Goldwasser, and S. Micali. Resettable zero-knowledge. In *Proc. 32nd STOC*, 2000.
- [DI06] B. Dubrov and Y. Ishai. On the randomness complexity of efficient sampling. In *Proc. 38th STOC*, 2006.
- [DDN00] D. Dolev, C. Dwork, and M. Naor. Non-malleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000.
- [DDP97] A. De Santis, G. Di Crescenzo, and P. Persiano. Randomness-efficient non-interactive zero knowledge. In *Proc. 24th ICALP*, 1997.
- [DDP02] A. De Santis, G. Di Crescenzo, and G. Persiano. Randomness-optimal characterization of two NP proof systems. In *Proc. Random '02*, 2002.
- [DN00] C. Dwork and M. Naor. Zaps and their applications. In *Proc. 41st FOCS*, 2000.
- [DNRS03] C. Dwork, M. Naor, O. Reingold, and L. Stockmeyer. Magic functions. *JACM*, 50(6):852–921, 2003.
- [DNS04] C. Dwork, M. Naor, and A. Sahai. Concurrent zero-knowledge. *JACM*, 51(6):851–898, 2004.
- [DS02] C. Dwork and L. Stockmeyer. 2-round zero knowledge and proof auditors. In *Proc. 34th STOC*, 2002.
- [FLS99] U. Feige, D. Lapidot, and A. Shamir. Multiple noninteractive zero knowledge proofs under general assumptions. *SICOMP*, 29(1):1–28, 1999.
- [G01] O. Goldreich. *Foundations of Cryptography: Basic Tools*. Cambridge University Press, 2001.
- [GMR89] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [GO94] O. Goldreich and Y. Oren. Definitions and properties of zero-knowledge proof systems. *J. Cryptology*, 7(1):1–32, 1994.
- [GOS06] J. Groth, R. Ostrovsky, and A. Sahai. Perfect non-interactive zero knowledge for NP. In *Proc. Eurocrypt '06*, 2006.
- [GS89] S. Goldwasser and M. Sipser. Private coins versus public coins in interactive proof systems. *Advances in Computing Research*, 5:73–90, 1989.
- [KL80] R. M. Karp and R. J. Lipton. Some connections between nonuniform and uniform complexity classes. In *Proc. 12th STOC*, 1980.
- [KP98] J. Kilian and E. Petrank. An efficient noninteractive zero-knowledge proof system for NP with general assumptions. *J. Cryptology*, 11(1):1–27, 1998.
- [P03] R. Pass. Simulation in quasi-polynomial time and its application to protocol composition. In *Proc. Eurocrypt '03*, 2003.
- [PS05] R. Pass and A. Shelat. Unconditional characterizations of non-interactive zero-knowledge. In *Proc. Crypto '05*, 2005.

- [PSV06] R. Pass, A. Shelat, and V. Vaikuntanathan. Construction of a non-malleable encryption scheme from any semantically secure one. In *Proc. Crypto '06*, 2006.
- [s99] A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *Proc. 40th FOCS*, 1999.
- [w05] H. Wee. On obfuscating point functions. In *Proc. 37th STOC*, 2005.