

Making Zero-Knowledge Provers Efficient

Mihir Bellare* Erez Petrank†

September 14, 1995

Abstract

We look at the question of how powerful a prover must be to give a zero-knowledge proof.

We present the first unconditional bounds on the complexity of a statistical ZK prover. The result is that if a language possesses a statistical zero-knowledge then it also possesses a statistical zero-knowledge proof in which the prover runs in probabilistic, polynomial time with an NP oracle. Previously this was only known given the existence of one-way permutations.

Extending these techniques to protocols of knowledge complexity $k(n) > 0$, we derive bounds on the time complexity of languages of “small” knowledge complexity.

Underlying these results is a technique for efficiently generating an “almost” random element of a set $S \in \mathcal{NP}$. Specifically, we construct a probabilistic machine with an NP oracle which, on input 1^n and $\delta > 0$ runs in time polynomial in n and $\lg \delta^{-1}$, and outputs a random string from a distribution within distance δ of the uniform distribution on $S \cap \{0, 1\}^n$.

* High Performance Computing and Communications, IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598, USA. e-mail: mihir@watson.ibm.com

† Department of Computer Science, Technion, Haifa, Israel. e-mail : erez@cs.technion.ac.il

1 Introduction

The investigation of the properties of interactive and zero-knowledge proofs [19] has already yielded many powerful results with diverse applications in complexity theory and cryptography. Here we continue this investigation by looking at the complexity of the prover in an interactive proof system, focusing in particular on the case where the proof system is statistical zero-knowledge. As we hope the following will indicate, bounding the complexity of zero-knowledge provers is not only interesting in its own right but also provides novel paradigms for some more classical complexity theoretic tasks.

Let us now proceed to describe our results.

1.1 Making ZK Provers Efficient without Assumptions

Suppose a language L has an interactive proof (P, V) . The prover P in such a proof is a (possibly probabilistic) function which given the common input and history of the interaction so far, returns the next message to send to the verifier V . The “complexity of the prover” refers to the computational complexity of this function.

The complexity of the prover has been addressed by Lund, Fortnow, Karloff and Nisan [23], Beigel, Bellare, Feigenbaum and Goldwasser [4] and (for multi-prover proofs [8]) by Babai, Fortnow and Lund [3]. It is also related to the notion of program checking of Blum and Kannan [9].

Here we are interested in the particular case where the interactive proof (P, V) must also be statistical zero-knowledge[†] (statistical means that the zero-knowledge is in a strong sense: see §2 for definitions). One of the things that makes the question of the complexity of a statistical ZK prover intriguing is that the additional burden the ZK constraint puts on the prover appears to be quite large. In fact, no obvious upper bound on the power of a statistical ZK prover seems apparent. This is in contrast to the case of (plain) interactive proofs where a polynomial space prover always suffices. We note that this is true even though we know that languages in statistical ZK are of (relatively) low complexity: results of Fortnow [12] and Aiello-Håstad [1] imply that $\mathcal{SZK} \subseteq \mathcal{AM} \cap \text{co-AM} \subseteq \Sigma_2^P \cap \Pi_2^P$ (where \mathcal{SZK} denotes the class of languages possessing statistical ZK interactive proofs of membership).

We do know that coins are necessary: Oren [24] and Goldreich and Oren [15] show that any ZK prover for a non-trivial language must be probabilistic.

Upper bounds on the complexity of a statistical ZK prover have been established by making use of (unproven) complexity assumptions. The first such result was that of Bellare, Micali and Ostrovsky [5] who showed that any language which possesses a statistical ZK proof also possesses a statistical ZK proof whose prover is a probabilistic, polynomial time (PPT) machine with a Σ_2^P oracle, provided the discrete log problem is hard.[‡] Since then, both the complexity and the assumption have been reduced: on the one hand Ostrovsky [25] showed that the Σ_2^P oracle may be replaced by an NP one, and on the other hand Ostrovsky, Venkatesan and Yung [26] show that any one-way permutation suffices.

[†] The analogous question for computational ZK is easily resolved as a corollary of the results in [7, 21].

[‡] [5] only claims a bound of probabilistic polynomial space, but the improvement to PPT with a Σ_2^P oracle is immediate on combining their construction with the results of Jerrum, Valiant and Vazirani [22] on uniform generation.

To summarize, the cumulative effect of all this work was to establish that any language which possesses a statistical ZK proof also possesses a statistical ZK proof whose prover is a PPT machine with an NP oracle, given that one-way permutations exist. The question that remained was whether it was possible to do away with the cryptographic assumption entirely.

Using different techniques, we answer this question in the affirmative. Namely, we establish the following unconditional result.

Theorem 1.1 *Let L be a language possessing a statistical zero-knowledge proof system. Then L has a statistical zero-knowledge proof system in which the prover is a probabilistic, polynomial time machine with access to an NP oracle.*

This theorem is proved by a transformation. Given a statistical ZK proof system (P, V) for L we show how to construct a PPT oracle machine P_{eff} such that $(P_{\text{eff}}^{\text{NP}}, V)$ is a statistical ZK proof system for L . Other than the unconditionality, our transformation offers some other improvements over previous ones. Some of these are

- We preserve the number of rounds ([5] blows them up by a constant factor and [26, 25] by a polynomial factor).
- We can preserve *perfect* ZK if we allow the prover to have a $\Sigma_2^{\mathcal{P}}$ oracle instead of an NP one (in previous solutions, even if the original proof system had been perfect ZK, the transformed one would be statistical).

As in some of the results in [15, 24], one of the ideas in our proof is to make (appropriate) use of the “auxiliary” inputs that the definition of ZK provides (cf. [11, 15, 19, 24, 29]). Previous solutions [5, 25, 26] did not exploit this feature of ZK.

We note that auxiliary inputs are important to the definition of ZK: Goldreich and Krawczyk [13] show that without auxiliary inputs in the definition, ZK would not even be closed under sequential composition.

Another idea of our proof is to build a “universal” verifier which, for our purpose, captures the behaviors of all possible cheating verifiers. We believe that it is an interesting open question to determine whether some such universal verifier can effectively eliminate the need to consider all cheating verifiers in other ZK applications, and in particular be used to remove assumptions in other applications (cf. §7).

While previous solutions used only the fact that the given proof system (P, V) was statistical ZK with respect to the honest verifier, our solution takes advantage of the full power of the ZK definition (that is, we use the fact that (P, V) is ZK with respect to all verifiers).

We remark that Theorem 1.1 implies that $\mathcal{SZK} \subseteq \mathcal{BPP}^{\text{NP}}$, providing an alternative way of showing that languages in SZK are of relatively low complexity (recall that the earlier and better results are due to Fortnow [12] and Aiello and Håstad [1]).

1.2 The Complexity of Non-Zero Knowledge Complexity

Knowledge complexity (KC) that is not zero was presented by Goldwasser, Micali, and Rackoff [18] and was formally formulated by Goldreich and Petrank [17]. The idea in the definition is that knowledge complexity measures the *computational* advantage gained in an interaction. One of the motivations to quantify KC originated from the hope that this new

approach may shed new light on interesting questions in complexity theory. In this paper, we make a step towards understanding how the time complexity of a language depends on its knowledge complexity.

Our bounds on the time complexity of the language are obtained by bounding the complexity of the prover (using techniques similar to those used to establish Theorem 1.1). Thus, bounding the power of the prover becomes a tool for bounding the complexity of the underlying language.

Our result involves a trade-off between the knowledge complexity and the number of rounds of the proof system. We are able to bound the time complexity of languages that have “short” interactive proofs of “small” knowledge complexity. More precisely, we show the following:

Theorem 1.2 *If L has an interactive proof of $g(n)$ rounds and of (statistical) knowledge complexity $k(n)$ satisfying $g(n)k(n) = O(\log n)$, then L is in $\mathcal{BPP}^{\text{NP}}$.*

Recently, the restriction on the number of rounds has been removed. Goldreich, Ostrovsky, and Petrank [16] used the same procedure as ours with a finer analysis to show that any language that has perfect knowledge complexity $O(\log n)$ is in $\mathcal{BPP}^{\text{NP}}$.

We should stress one important difference between ZK proofs and proofs of higher KC. The difference is that, for the latter, one expects composition to increase the knowledge complexity. One consequence of this is that reducing the error probability by standard amplification techniques (such as serial composition) will increase the knowledge complexity commensurately. So, in saying what is a proof of knowledge complexity $k(n)$, it is important to say what the error-probability is. We adopt the natural notion that the error-probability is a negligible (i.e. less than the reciprocal of any polynomial) function of n , and that is what the above theorem assumes.

Our result can be viewed as an extension to higher KC of the results of [12, 1] showing that languages of zero (statistical) knowledge complexity (i.e. statistical ZK languages) have “low” complexity (they show that $\mathcal{SZK} \subseteq \mathcal{AM} \cap \text{co-AM} \subseteq \mathcal{BPP}^{\text{NP}}$). Our techniques, however, are quite different. Generalizing the techniques of [12, 1] to the setting of higher KC would yield a weaker result than Theorem 1.2 because it seems one must assume that the error-probability of the original proof system was exponentially small.[†]

1.3 Efficient Almost Uniform Generation

A technique underlying these results is to generate almost random elements of a set in probabilistic, polynomial time with an NP oracle. As this might be of independent interest, let us describe the result in more detail.

Let $S \subseteq \{0,1\}^*$ be a set verifiable in polynomial time (i.e., $S \in \mathcal{NP}$), and let $S_n = S \cap \{0,1\}^n$. The uniform generation problem is to generate, on input 1^n , an element of S_n distributed uniformly at random. Jerrum, Valiant and Vazirani [22], using results of Stockmeyer [28] on approximate counting, showed that uniform generation can be done in probabilistic, polynomial time with a Σ_2^P oracle.

[†] Extending Fortnow’s techniques [12] one can show that if L has a $g(n)$ round, KC $k(n)$, proof system with error probability $2^{-O(n g(n))}$, and if $g(n)k(n) = O(\log n)$, then $L \in \text{co-AM}$.

For our applications we would like a lower complexity than PPT with a Σ_2^P oracle. On the other hand, we could tolerate a slight deviation from uniform of the output distribution. Accordingly, we consider the problem of *almost* uniform generation: on input 1^n and $\delta > 0$ generate a random element from a distribution within distance δ of the uniform distribution on S_n (the distance between distributions E_1 and E_2 is defined as $\frac{1}{2} \sum_x |\Pr_{E_1}[x] - \Pr_{E_2}[x]|$).

If $\delta = n^{-c}$ for some fixed constant c then techniques from Impagliazzo, Levin and Luby [20] can be used to do almost uniform generation in probabilistic, polynomial (in n) time with an NP oracle. However (for applications in this paper in particular) we would like to be able to achieve values of δ which are exponentially (in n) small, in the same complexity. We show that this can be done.

Theorem 1.3 *Let $S \in \mathcal{NP}$. Then there is a probabilistic oracle machine A which on input 1^n and $\delta > 0$ runs in time polynomial in n and $\lg \delta^{-1}$, and has the property that the distribution $A^{\text{NP}}(1^n, \delta)$ is within distance δ of the uniform distribution on S_n .*

Remark 1.4 *The special case in which S is decidable in polynomial time (i.e. $S \in \mathcal{P}$) is important on its own and specifically, throughout this paper we use Theorem 1.3 only for sets in \mathcal{P} .*

In Theorem 3.2 we actually prove something a little stronger: the almost uniform generation is “universal” (in the sense that A does not depend on S but rather gets a description of S as an input).

This result is established by combining techniques from Jerrum, Valiant and Vazirani [22] and Stockmeyer [28] with Carter-Wegman universal hash function [10] based techniques for estimating set sizes (Sipser [27]). The details are in §3.

2 Preliminaries

If E_1 and E_2 are probability spaces then the distance between them, denoted $d(E_1, E_2)$, is

$$\frac{1}{2} \sum_{x \in \{0,1\}^*} |\Pr_{E_1}[x] - \Pr_{E_2}[x]|.$$

We say that E_1 and E_2 are δ -close if the distance between them is $\leq \delta$.

An ensemble over $L \subseteq \{0,1\}^*$ is a collection $\{E(x, a)\}_{(x,a) \in L \times \{0,1\}^*}$ of probability spaces indexed by $L \times \{0,1\}^*$. We extend the definition of distance to ensembles, with the distance between ensembles $\mathcal{E}_1 = \{E_1(x, a)\}_{(x,a) \in L \times \{0,1\}^*}$ and $\mathcal{E}_2 = \{E_2(x, a)\}_{(x,a) \in L \times \{0,1\}^*}$ being the function $d^*: \mathbb{N} \rightarrow \mathbb{N}$ defined by

$$d^*(n) = \max_{x \in L, |x|=n} \sup_{a \in \{0,1\}^*} d(E_1(x, a), E_2(x, a)).$$

A function $\epsilon: \mathbb{N} \rightarrow \mathbb{N}$ is negligible if for any constant $c > 0$ there is an integer n_c such that $\epsilon(n) \leq n^{-c}$ for all $n \geq n_c$. We say that the ensembles \mathcal{E}_1 and \mathcal{E}_2 are *statistically indistinguishable* if the distance between them is negligible.

Next we define interactive proofs [19]. We begin by specifying the parties involved: the verifier and the prover. We have chosen to be very specific in our definitions, because this will simplify the exposition of our theorems and proofs.

Note that our verifiers take an auxiliary input (in addition to the common input and history of interaction) — intuitively, it captures a possible history of past interactions of the verifier.

In the following, $g, l, q: \mathbb{N} \rightarrow \mathbb{N}$ are polynomially bounded, polynomial time computable functions. We write g, l, q for $g(n), l(n), q(n)$ respectively whenever n is understood.

Definition 2.1 *A verifier is a function V with the following properties. Suppose $x \in \{0, 1\}^n$, $\alpha_1, \beta_1, \dots, \alpha_t, \beta_t$ are $l(n)$ bit strings, $a \in \{0, 1\}^*$ and $R \in \{0, 1\}^{q(n)}$. Then V on input x (the common input), a (auxiliary input), $\alpha_1\beta_1 \dots \alpha_t\beta_t$ (the conversation so far), and R (random tape of V) returns*

- *A $l(n)$ bit string α_{t+1} (next message to prover) if $t < g(n)$*
- *A bit (accept or reject) if $t = g(n)$.*

Moreover V runs in time polynomial in n . We refer to g, l, q as the number of rounds, message length, and random tape length of V , respectively.

Notice that by definition a verifier message must be a l bit string, and a verifier must output some message for each round (it cannot halt in the middle of the protocol). Such conditions may be assumed without loss of generality anyway, and it will simplify our exposition and proofs to instead build them into the definition.

Definition 2.2 *A prover is a (probabilistic) function P with the following property. Suppose $x \in \{0, 1\}^n$ and $\alpha_1, \beta_1, \dots, \alpha_{t-1}, \beta_{t-1}, \alpha_t$ are $l(n)$ bit strings, $t \leq g$. Then P on input x (the common input) and $\alpha_1\beta_1 \dots \alpha_{t-1}\beta_{t-1}\alpha_t$ (the conversation so far) returns a $l(n)$ bit string β_t (next message to verifier). We refer to g, l as the number of rounds and the message length of P , respectively.*

Again, note that a prover message in any round must be of length l . When we speak of a prover P interacting with a verifier V , or couple them as a pair (P, V) , it is to be understood that both parties have the same number of rounds and same message length.

The probability that a verifier V accepts (resp. rejects) at the end of an interaction with a prover P , on common input x and auxiliary input a for V , is the probability that $V(x, a, \alpha_1\beta_1 \dots \alpha_g\beta_g, R) = 1$ (resp. 0) in the experiment given by choosing $R \in \{0, 1\}^{q(n)}$ at random and setting $\alpha_t = V(x, a, \alpha_1\beta_1 \dots \alpha_{t-1}\beta_{t-1}, R)$, $\beta_t = P(x, \alpha_1\beta_1 \dots \alpha_t)$ for $t = 1, \dots, g$ (here g is the number of rounds, l the message length, and r the random tape length of V).

Definition 2.3 *We say that a prover/verifier pair (P, V) is an interactive proof for a language L if*

- (1) *The probability that V rejects at the end of the interaction with P on a common input $x \in L$ is negligible, and*
- (2) *For any prover \hat{P} and common input $x \notin L$ the probability that V accepts at the end of the interaction with \hat{P} is negligible*

(the auxiliary input of V is set to the empty string λ in both cases). The first condition is called the completeness condition and the second the soundness condition.

Next we define zero-knowledge (ZK) proofs [19]. We remark that the auxiliary inputs (cf. [11, 15, 19, 24, 29]) are crucial to make the definition meaningful. For one thing, without them, the composition of zero-knowledge protocols is not necessarily a zero-knowledge protocol as we would like it to be [13]. We recall also that any ZK proof which is black box simulation ZK (as all known ones are) is ZK in the auxiliary input model [15, 24].

Let us first define simulators. In the following let $p: \mathbb{N} \rightarrow \mathbb{N}$ be a polynomially bounded, polynomial time computable function.

Definition 2.4 *A simulator is an algorithm S which, on inputs $x \in \{0, 1\}^n$, $a \in \{0, 1\}^*$ and $r \in \{0, 1\}^{p(n)}$, outputs, in time polynomial in n , a pair of strings $(R, \alpha_1\beta_1 \dots \alpha_{g(n)}\beta_{g(n)})$, where $R \in \{0, 1\}^{q(n)}$ and $\alpha_1, \beta_1, \dots, \alpha_{g(n)}, \beta_{g(n)} \in \{0, 1\}^{l(n)}$. We refer to p as the number of coins the simulator tosses. We refer to g, l, q as the number of rounds, message length, and random tape length of S , respectively.*

Note that we have not yet said anything about a simulator's ability to produce the view of the verifier. We do this in the next definition. For us, a simulator is simply an algorithm whose output format conforms to certain parameters. The purpose of this is again to simplify later exposition by building into the definition those features that could be assumed without loss of generality anyway.

If V is a verifier and we talk of S being a simulator for V it is to be understood that these parties have the same number of rounds, message length, and random tape length. If S is a simulator and p the number of its coin tosses then we let $S(x, a)$ denote the probability space which to each string s assigns the probability that $S(x, a, r) = s$ when $r \in \{0, 1\}^{p(x)}$ is chosen at random.

Fix a (cheating) verifier \hat{V} , common input x , and auxiliary input a for \hat{V} . We let $View_{(P, \hat{V})}(x, a)$ denote the probability distribution on pairs $(R, \alpha_1\beta_1 \dots \alpha_g\beta_g)$ which is given by picking at random $R \in \{0, 1\}^{q(n)}$ and then setting $\alpha_t = V(x, a, \alpha_1\beta_1 \dots \alpha_{t-1}\beta_{t-1}, R)$ and $\beta_t = P(x, \alpha_1\beta_1 \dots \alpha_t)$ for $t = 1, \dots, g$. Intuitively, this captures \hat{V} 's view of the interaction: his own coin tosses and the transcript of the conversation with the prover.

Definition 2.5 *Let P be a prover, \hat{V} a verifier, and S a simulator for \hat{V} . We call S a statistical (resp. perfect) ZK P -simulator for \hat{V} if the ensembles $\{View_{(P, \hat{V})}(x, a)\}_{(x, a) \in L \times \{0, 1\}^*}$ and $\{S_{\hat{V}}(x, a)\}_{(x, a) \in L \times \{0, 1\}^*}$ are statistically indistinguishable (resp. equal).*

Definition 2.6 *We say that (P, V) is statistical (resp. perfect) zero-knowledge if for each verifier \hat{V} there is a statistical (resp. perfect) ZK P -simulator for \hat{V} .*

We call \hat{V} a cheating verifier.

Finally we also summarize the definition of the knowledge complexity measure (greater or equal to zero) which we use in the sequel. This definition is due to [17], and they call it the "fraction" version. For intuition and motivation we refer the reader to the original paper [17].

Let S be a simulator, p the number of its coin tosses, and $C \subseteq \{0, 1\}^{p(n)}$ a subset of its coin tosses, then we let $S(x, a, C)$ denote the probability space which to each string s assigns the probability that $S(x, a, r) = s$ when r is chosen at random from C . In the following let $k: \mathbb{N} \rightarrow \mathbb{N}$.

Definition 2.7 Let P be a prover, \hat{V} a verifier, and $S_{\hat{V}}$ a simulator for \hat{V} . Let p be the number of coin tosses of $S_{\hat{V}}$. We call $S_{\hat{V}}$ a *statistical (resp. perfect) KC $k(n)$ P -simulator* for \hat{V} if for each n and each $x \in L \cap \{0, 1\}^n$ there is a set $SUCC_{x,a} \subseteq \{0, 1\}^{p(n)}$ such that

- (1) $2^{-p(n)} |SUCC_{x,a}| \geq 2^{-k(n)}$ ($SUCC_{x,a}$ has density at least $2^{-k(n)}$)
- (2) The ensembles $\{\text{View}_{(P, \hat{V})}(x, a)\}_{(x,a) \in L \times \{0,1\}^*}$ and $\{S_{\hat{V}}(x, a, SUCC_{x,a})\}_{(x,a) \in L \times \{0,1\}^*}$ are statistically indistinguishable (resp. equal).

We call $SUCC_{x,a}$ the *success set* of S for x .

Definition 2.8 (Knowledge Complexity – Fraction Version) We say that (P, V) has *statistical (resp. perfect) knowledge complexity $k(n)$ for a language L* if for any verifier \hat{V} there exists a *statistical (resp. perfect) KC $k(n)$ P -simulator $S_{\hat{V}}$ for \hat{V} .*

PPT abbreviates “probabilistic, polynomial time.”

3 Efficient Almost Uniform Generation

In this section, we build our main technical tool: the efficient almost uniform generator. We believe that this tool is important in its own sake and can be useful in other settings as well. The generator samples “almost uniformly” from a given set $S \in \mathcal{NP}$. The generator is universal in the sense that it is not built for a specific set S , but is given a description of S in the input.

The description of S is given to the sampler as an encoding $\langle M \rangle$ of a machine M that verifies S . Namely, M has two inputs (r, w) and the set S is equal to:

$$S = \{r : \exists w \text{ s.t. } M \text{ accepts the input } (r, w)\}.$$

For any set S in \mathcal{NP} there exists a machine M that runs in time polynomial in $|r|$ and corresponds to S in the above manner. We define $M_n \stackrel{\text{def}}{=} S \cap \{0, 1\}^n$ and we denote by $\mathcal{U}[B]$ the uniform distribution over a finite set B .

Definition 3.1 A *universal almost uniform generator* is a (probabilistic) machine A which has the property that $A(1^n, \delta, \langle M \rangle)$ and $\mathcal{U}[M_n]$ are δ -close, for all $n \in \mathbb{N}$, $\delta > 0$ and TMs M .

Our result is the following.

Theorem 3.2 *There is a probabilistic, oracle machine U such that U^{NP} is a universal almost uniform generator, and moreover, the running time of U^{NP} on inputs $1^n, \delta, \langle M \rangle$ is polynomial in $n, \lg \delta^{-1}$ and $T_M(n)$.*

Here $T_M(n)$ denotes the maximum, over all $r \in \{0, 1\}^n$, $w \in \{0, 1\}^*$, of the running time of M on input (r, w) .

A very important property of the procedure of Theorem 3.2 is that its running time is polynomial in $\lg \delta^{-1}$ rather than polynomial in δ^{-1} , so that in time polynomial in n we can achieve exponentially small (in n) error.

We remark that even the special case in which we sample a set $S \in \mathcal{P}$ is important, and in fact, we use Theorem 3.2 throughout this paper only for sets which are decidable in polynomial time.

The proof of Theorem 3.2 is derived by combining techniques from [22, 28, 27]. First, we build an approximate counter. The approximate counter of Stockmeyer [28] took polynomial time with a Σ_2^P oracle. We begin by showing that its “almost” version can be implemented in PPT with an NP oracle. Our construction uses hashing techniques similar to the ones used by Sipser [27]. Once we have this approximate counter, we use the PPT reduction of Jerrum, Valiant and Vazirani [22] from uniform generation to approximate counting. The key observation now is the strength of the [22] reduction: it is capable of sampling distributions within an exponentially small distance from the uniform distribution, when given a primitive for estimating set sizes (namely approximate counting) that is only accurate to within the reciprocal of a polynomial. The details follow.

We use universal hash functions [10]. Let $H_{k,b}$ denote the set of all affine transformations $h_{A,\alpha} : \{0, 1\}^k \rightarrow \{0, 1\}^b$ given by $h_{A,\alpha}(y) = Ay + \alpha$ where A is a b by k matrix over $\text{GF}[2]$ and α is a b -vector over $\text{GF}[2]$, and the arithmetic is over $\text{GF}[2]$. The following lemma follows from Lemma 4.1 in Aiello-Håstad [1], which in turn is based on ideas of Babai [2] and Sipser [27].

Lemma 3.3 *Let b, k be positive integers, $\delta > 0$ and $S \subseteq \{0, 1\}^k$. Let $l = 16 \lg(2\delta^{-1})$. Select random and independent hash functions $h_1, \dots, h_{2l} \in H_{k,b}$ and set $z \stackrel{\text{def}}{=} 0^b$. Let the random variable X equal the number of indices i for which $h_i^{-1}(z) \cap S \neq \emptyset$. Then the probability that $X > l$ is*

- (1) $\geq 1 - \delta$ if $|S| \geq 4 \cdot 2^b$
- (2) $\leq \delta$ if $|S| \leq 2^b/4$.

We remark that the choice $z = 0^b$ is arbitrary and selecting z to be any constant string in $\{0, 1\}^b$ will do as well. The notion of approximate counting to within a factor $(1 + \epsilon)$ with success probability at most $(1 - \delta)$ is as follows (cf. [28]).

Definition 3.4 *A universal approximate counter is a (probabilistic) machine C which, on inputs $1^n, \epsilon > 0, \delta > 0, \langle M \rangle$ outputs an estimate v which, with probability $\geq 1 - \delta$, satisfies $|M_n|/(1 + \epsilon) \leq v \leq |M_n|(1 + \epsilon)$.*

Let us now establish the result on approximate counting that we need.

Theorem 3.5 *There exists a probabilistic, oracle machine C such that C^{NP} is a universal approximate counter, and moreover, the running time of C^{NP} on inputs $1^n, \epsilon, \delta, \langle M \rangle$ is polynomial in $n, \epsilon^{-1}, \lg \delta^{-1}$ and $T_M(n)$.*

Proof: The algorithm C is given in Figure 1. The idea behind it is as follows.

We first look at Lemma 3.3 and derive an algorithm to approximate a set S (given as $\langle M \rangle$) to within a constant factor with success probability of $1 - \delta$. We note that the computational bottle neck in running this algorithm is having to decide, given z , h , and $\langle M \rangle$, whether $h_i^{-1}(z) \cap S \neq \emptyset$. We shall have to use the NP-oracle to make these decisions. All other steps in the algorithm can be performed in polynomial time.

$C^{\text{NP}}(1^n, \epsilon, \delta, \langle M \rangle)$
 Choose t such that $8^{1/t} \leq 1 + \epsilon$
 { **Comment** t will be polynomial in ϵ^{-1} }
 Let $k = nt$ and let $N = M^t$ be the machine defined by $N(y_1, \dots, y_t) = 1$ iff $M(y_1) = \dots = M(y_t) = 1$ and all the y_i have the same length.
 { **Comment** $N_k = M_n^t$ }
 $b \leftarrow 0$; $l \leftarrow 16 \lg(2k\delta^{-1})$
repeat
 $b \leftarrow b + 1$
 Pick at random $h_1, \dots, h_{2l} \in H_{k,b}$ and set $z = 0^b$. Use the NP oracle to determine, for each i , whether or not $h_i^{-1}(z) \cap N_k = \emptyset$, and let X be the number of indices i for which $h_i^{-1}(z) \cap N_k \neq \emptyset$.
until ($X \leq l$ or $b = k$)
 { **Comment** $\Pr[2^{b-1}/4 < |N_k| < 4 \cdot 2^b] \geq 1 - \delta$ }
 Let $\alpha = 2^b$.
 { **Comment** $\Pr[|N_k|/4 < \alpha < 8|N_k|] \geq 1 - \delta$ }
 Output $\alpha^{1/t}$
 { **Comment** $\Pr[|M_n|/(1 + \epsilon) < \alpha < |M_n|(1 + \epsilon)] \geq 1 - \delta$ }

Figure 1: Universal Approximate Counting in PPT with an NP Oracle

Using the setting of Lemma 3.3 the approximation algorithm is as follows. We enumerate over all possible $b = 1, 2, \dots, n$, and select the first b for which the experiment defined in Lemma 3.3 fails. Namely, running the experiment with $b-1$ yielded the result $X > l$ whereas running it with b yielded $X \leq l$. We conclude that with probability at least $1 - n \cdot \delta$ it holds that

$$\frac{2^{b-1}}{4} \leq |S| \leq 4 \cdot 2^b. \quad (1)$$

Namely, with probability at least $1 - n\delta$, in all the experiments we performed in the algorithm (we always perform at most n experiments), only events with probability greater than $1 - \delta$ (see Lemma 3.3) occurred.

Next, we note that a PPT machine with an NP oracle can decide whether $h_i^{-1}(z) \cap S \neq \emptyset$ since an NP oracle can determine whether there exists a couple (x, w) for which $h(r) = z$ and M accepts (r, w) .

Last, we have to improve the quality of the approximation. To this end, we define a machine N which runs M on t inputs and accepts iff M accepts all the t inputs. The set defined by the machine N is S^t which is also a set in \mathcal{NP} . To be done, we note that approximating the size of S^t to within a factor of 8 implies an approximation of the size of S to within $8^{1/t}$. Thus, we can get an approximation for the size of S to within a factor of $1 + \epsilon$ by setting t to be polynomial in ϵ and applying our constant-factor approximation on S^t (or on the machine N). The protocol is given in Figure 1. ■

Notice that the estimates from the above theorem are only accurate to within the reciprocal of a polynomial (that is, ϵ^{-1} must be bounded by a polynomial in n). We now use the reduction of [22] which is strong enough in this sense. Given an approximation to within a factor of $1 + \epsilon$ (for an ϵ which is a reciprocal of a polynomial) it samples a distribution which is very close to the uniform distribution. Specifically, they show the following:

Theorem 3.6 [22]: *If there exists a universal approximate counter to within a factor of $(1 + \epsilon)$ (that never fails) then there exists a universal (exact) uniform generator which, when given access to the approximate counter as an oracle, runs in polynomial time in n and $\frac{1}{\epsilon}$.*

For further reference, let $Q(n, \frac{1}{\epsilon})$ be a polynomial which bounds the number of oracle queries that the generator guaranteed in Theorem 3.6 makes to the approximate counter. We are going to run this generator with our universal approximate counter. The problem is that the generator expects an approximation procedure which is always correct, while our procedure fails to approximate the given set (to within $1 + \epsilon$) with positive probability (which is only bounded by δ).

We partition the analysis (of running the generator with our approximator) into two cases: In one case, our approximate counter succeeds in all $Q(n, \frac{1}{\epsilon})$ queries. This happens with probability at least $1 - Q(n, \frac{1}{\epsilon}) \cdot \delta$, and in this case, the generator outputs a uniform element in the set as required. The other possible case, is that the approximate counter makes an error in one of the queries. In this case, it is not possible to tell which distribution is output by the generator. However, since with probability at least $1 - Q(n, \frac{1}{\epsilon}) \cdot \delta$ the generator samples uniformly, we get that its output distribution is at distance at most $Q(n, \frac{1}{\epsilon}) \cdot \delta$ from the uniform distribution. Since the running time of the approximation counter is polynomial in the logarithm of $\frac{1}{\delta}$, then the polynomial multiplicative factor $Q(n)$, has no significant effect on the complexity of the procedure. Setting $\delta' \stackrel{\text{def}}{=} \frac{\delta}{Q(n)}$ we get:

Corollary 3.7 *If there exists a universal approximate counter to within a factor of $(1 + \epsilon)$, C^{NP} , which runs in time polynomial in $n, 1/\epsilon, \log \frac{1}{\delta'}$ and $T_M(n)$, and which fails with probability at most δ' , then there exists a universal almost uniform generator, U^{NP} , which runs in time polynomial in $n, T_M(n)$ and $\log \frac{1}{\delta'}$.*

Combinning Corollary 3.7 with Theorem 3.5, we get Theorem 3.2 and we are done.

4 The Simulation Based Prover

A basic tool in our results is the construction, from the simulator for a protocol, of a new prover based on this simulator. Here we describe a key component of this construction which we call a simulation based prover, and develop techniques to implement it efficiently. Similar ideas were used first in [5] and then later in [25]. These papers however did not incorporate auxiliary inputs into the notion, a difference that is important to some of our results.

Definition 4.1 *If S is a simulator of g rounds and message length l (cf. Definition 2.4) and $S(x, a, r) = (R, \alpha_1\beta_1 \dots \alpha_g\beta_g)$ then for each $t = 1, \dots, g$ we let $S_t(x, a, r) = \alpha_1\beta_1 \dots \alpha_{t-1}\beta_{t-1}\alpha_t$.*

Definition 4.2 *Let S be a simulator (of g rounds and message length l), and let p be the number of its coin tosses. We denote by P_S the probabilistic function whose output (next message to verifier), on input x (common input), a (auxiliary input) and $\alpha_1\beta_1 \dots \alpha_t$ (conversation so far), is given by the following experiment:*

- (1) *Pick at random a string r from the set $\{r \in \{0,1\}^{p(n)} : S_t(x,a,r) = \alpha_1\beta_1 \dots \alpha_t\}$*
- (2) *Then compute $S(x,a,r)$ — this has the form*

$$(R, \alpha_1\beta_1 \dots \alpha_t\beta'_t \dots \alpha'_g\beta'_g)$$

for some $\beta'_t, \dots, \alpha'_g, \beta'_g \in \{0,1\}^l$

- (3) *Finally, output the string β'_t .*

We call P_S the simulation based prover for S .

Note that P_S is not quite a “prover” in the sense of Definition 2.2 because it has too many arguments: a prover does not take an auxiliary input. We choose to name it as we have because we will, later, use it as a basis to construct “true” provers.

Intuitively, if S is a (statistical) ZK P -simulator for V , then P_S is attempting to find P 's replies in the interaction of P with V . He does this by using S to appropriately sample the view of V in its interaction with P . The role played by the auxiliary input is more subtle, and we will discuss it later. It should be noted, though, that in interacting with P the verifier V may have some auxiliary input of its own, and this is not necessarily the same as the auxiliary input to P_S .

We will denote the output of P_S on input x, a and $\alpha_1\beta_1 \dots \alpha_t$ by $P_S(x, a, \alpha_1\beta_1 \dots \alpha_t)$, but keep in mind the function is probabilistic so that $P_S(x, a, \alpha_1\beta_1 \dots \alpha_t)$ is actually a distribution on l bit strings.

The complexity of computing P_S is the question we address next.

Exact computation of P_S is clearly possible in probabilistic polynomial space, and this is easily improved to PPT with a Σ_2^P oracle by using the results of [22] on uniform generation.

Let us now turn to approximations: we will be interested in computing a distribution which is δ -close to P_S . If $\delta = n^{-c}$ for some constant c then, following [25], this can be done in PPT with an NP oracle by using techniques from [20]. That result will not, however, suffice for the applications in this paper: we need to be able to achieve values of δ which are exponentially small (in n). The following Theorem says we can do this, still in PPT with an NP oracle.

Theorem 4.3 *Let S be a simulator (of g rounds and message length l). Then there exists a probabilistic oracle machine T with the following property. Suppose $x \in \{0,1\}^n, a \in \{0,1\}^*$, $\alpha_1, \beta_1, \dots, \alpha_t$ are $l(n)$ bit strings ($t \leq g$), and $\delta > 0$. Then the probability spaces $T^{\text{NP}}(x, a, \alpha_1\beta_1 \dots \alpha_t, \delta)$ and $P_S(x, a, \alpha_1\beta_1 \dots \alpha_t)$ are δ -close. Moreover, T runs in time polynomial in n and $\lg \delta^{-1}$.*

Note the running time of T is polynomial in $\lg \delta^{-1}$ (rather than δ^{-1}), which is why we can achieve exponentially small error.

The proof of this theorem is straightforward (given the implementation of *universal almost uniform generation* in PPT with an NP oracle, i.e., Theorem 3.2): follow the algorithm of Definition 4.2, using Theorem 3.2 to implement the first step.

5 Making ZK Provers Efficient Without Assumptions

Let us begin with some background.

5.1 The Problem

We are given a statistical ZK interactive proof system (P, V) for L . We wish to construct a new statistical ZK interactive proof system $(\overline{P}, \overline{V})$ for L whose prover is efficient. Let us begin by reviewing previous solutions to the problem and see whence arise the complexity assumptions that we wish to eliminate.

The solution of [5] begins with the observation that it is easy to construct an efficient prover P_h such that the system (P_h, V) is statistical ZK with respect to the honest verifier V (that is, there exists a P_h -simulator for V). In our present terminology, this prover P_h is simply the simulation based prover P_{S_V} (more precisely, $P_h(x, \alpha_1 \beta_1 \dots \alpha_t) = P_{S_V}(x, \lambda, \alpha_1 \beta_1 \dots \alpha_t)$). The problem is that (P_h, V) is not statistical ZK with respect to all verifiers as the definition of statistical ZK requires (that is, there does not necessarily exist a P_h -simulator for $\hat{V} \neq V$). Indeed, sampling the output of S_V in order to reply to a verifier \hat{V} does not yield a statistical ZK proof with respect to \hat{V} .

What if, when talking to \hat{V} , the prover should sample from the outputs of the simulator $S_{\hat{V}}$ for this particular cheating verifier \hat{V} ? That is, when talking to \hat{V} , the prover would be running $P_{S_{\hat{V}}}$. This would yield the right distribution, but cannot be (efficiently) achieved. For we want to construct a single efficient prover. How could this prover know which (of the infinitely many) provers $P_{S_{\hat{V}}}$ to run at any point in time? He cannot tell who he is talking to, and, anyway, cannot keep an infinite number of pieces of code handy.

In [5] the problem of the cheating verifier is surmounted by constructing a “compiler” which, given a proof system for L which is statistical ZK with respect to the honest verifier, transforms it into a new proof system which is statistical ZK with respect to *any* (cheating) verifier, while preserving the power of the prover. They then apply this compiler to transform (P_h, V) . It is the compiler that uses a complexity assumption.

To summarize, their solution is in two steps: given a statistical ZK proof system (P, V) for L ,

Step 1: Construct an efficient prover P_h such that (P_h, V) is a proof system for L which is statistical ZK with respect to the honest verifier V

Step 2: Using a complexity assumption, “compile” the (P_h, V) proof into a proof $(\overline{P}, \overline{V})$ in which \overline{P} is of the same complexity as P_h but now $(\overline{P}, \overline{V})$ is statistical ZK with respect to *all* verifiers (as required by the definition of ZK).

All subsequent work on the question (until now) has followed the same two step paradigm, and (because of Step 2) has thereby used complexity assumptions. Below we will see that the problem of the cheating verifier can be viewed differently in this context, and the two step paradigm can be avoided. The complexity assumption is then unnecessary.

5.2 Our Solution

Our goal is to construct a PPT oracle machine P_{eff} such that $(P_{\text{eff}}^{\text{NP}}, V)$ is a statistical zero-knowledge interactive proof system for L .

The first step is to define a single *universal* verifier who, in essence, can “simulate” the behavior of all verifiers as far as the prover is concerned. The key in the definition of the universal verifier V^* (as in several proofs in [15, 24]) is to make appropriate use of the auxiliary inputs.

Definition 5.1 *The universal verifier V^* (of g rounds and message length l) is defined as follows. Let $x \in \{0, 1\}^n$ and let $\alpha_1, \beta_1, \dots, \alpha_{t-1}, \beta_{t-1}, \alpha_1^*, \dots, \alpha_i^*$ be $l(n)$ bit strings, $t, i \leq g$. On input x (common input), $\alpha_1^* \dots \alpha_i^*$ (auxiliary input), $\alpha_1 \beta_1 \dots \alpha_{t-1} \beta_{t-1}$ (conversation so far), and R (random tape) the message $V^*(x, \alpha_1^* \dots \alpha_i^*, \alpha_1 \beta_1 \dots \alpha_{t-1} \beta_{t-1}, R)$ that V^* computes (and sends to the prover) is*

$$\begin{cases} \alpha_t^* & \text{if } 1 \leq t \leq i \\ 0^l & \text{otherwise.} \end{cases}$$

At the end of the interaction, V^* accepts.

In other words, the universal verifier’s response in the t -th round is the t -th string written on his auxiliary input as long as there are $\geq t$ strings on his auxiliary input, and 0^l otherwise. Note that this response of V^* is independent of the conversation $\alpha_1 \beta_1 \dots \alpha_{t-1} \beta_{t-1}$ so far, and that V^* is deterministic (random tape length 0). Also note that at the end of the protocol V^* blindly accepts.

Since (P, V) is ZK we know that there exists a ZK P -simulator S^* for V^* . Note that outputs of S^* have the form $(\lambda, \alpha_1 \beta_1 \dots \alpha_g \beta_g)$ where λ is the empty string.

Definition 5.2 *Let V^* be the universal verifier and S^* a (statistical) ZK P -simulator for V^* . We define P^* by*

$$P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = P_{S^*}(x, \alpha_1 \dots \alpha_t, \alpha_1 \beta_1 \dots \alpha_t),$$

where P_{S^*} is the simulation based prover for S^* . We call P^* a **universal prover** for (P, V) .

It is to be understood that if (P, V) is perfect ZK then S^* is a perfect ZK P -simulator. Let us now try to give some intuition for these definitions and their use.

Fix a common input x , and suppose we want to compute $P(x, \alpha_1 \beta_1 \dots \alpha_t)$, the prover’s response when the history is $\alpha_1 \beta_1 \dots \alpha_t$. The idea is that since this response $P(x, \alpha_1 \beta_1 \dots \alpha_t)$ depends *only* on the history $\alpha_1 \beta_1 \dots \alpha_t$, it does not matter how the verifier is computing his messages as long as the history is indeed $\alpha_1 \beta_1 \dots \alpha_t$. In particular, we may as well assume P is talking to the universal verifier with the latter’s auxiliary input set to $\alpha_1 \alpha_2 \dots \alpha_t$. But the result of P talking to V^* (on that auxiliary input) is mimicked by $S^*(x, \alpha_1 \alpha_2 \dots \alpha_t)$, and P_{S^*} just samples this last distribution to appropriately extract the prover’s next message. So we can compute P using P^* . In fact, if (P, V) were perfect (rather than statistical) ZK then P^* would equal P (recall that both are probabilistic functions, so that this equality means that on any input the distribution given by both functions is identical). That is, the new simulation based prover is just another way of looking at the original prover P . When (P, V) is statistical (rather than perfect) ZK, P^* and P are sufficiently close. The following Lemma states this more precisely.

Lemma 5.3 *Let $x \in L \cap \{0,1\}^n$, let $\alpha_1, \alpha_2, \dots, \alpha_t$ be any $l(n)$ bit strings (the verifier messages in the conversation so far). Then*

(1) *If (P, V) is statistical ZK then:*

$$\sum_{\beta_1, \dots, \beta_{t-1}} \Pr[\text{View}_{(P, V^*)}(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t].$$

$$\sum_{\beta_t \in \{0,1\}^{l(n)}} |\Pr[P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t] - \Pr[P(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t]| \leq 2\delta$$

(2) *If (P, V) is perfect ZK then for any $\beta_1, \beta_2, \dots, \beta_{t-1}$ it holds that $P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = P(x, \alpha_1 \beta_1 \dots \alpha_t)$.*

Note that the equality guaranteed in the second part of the lemma is an equality of two distribution spaces on strings in $\{0,1\}^{l(n)}$.

Proof: We prove Part (1). The proof of Part (2) is a simplification of this proof for the more restricted case of perfect ZK. By hypothesis, the distance between $\text{View}_{(P, V^*)}(x, \alpha_1 \alpha_2 \dots \alpha_t)$ and $S(x, \alpha_1 \alpha_2 \dots \alpha_t)$ is at most δ and thus δ is also a bound on the distance between t -round prefixes of these two distributions. We denote by $S(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t \beta_t$ the event in which S , on input $x, \alpha_1 \dots \alpha_t$, outputs a conversation which has $\alpha_1 \beta_1 \dots \alpha_t \beta_t$ as a prefix.

Let us bound the distance mentioned in the lemma. All the following summations on the strings β_i , $1 \leq i \leq t$, are summations over $\beta_i \in \{0,1\}^{l(n)}$.

$$\sum_{\beta_1, \dots, \beta_{t-1}} \Pr[\text{View}_{(P, V^*)}(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t].$$

$$\sum_{\beta_t} |\Pr[P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t] - \Pr[P(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t]|$$

$$\leq \sum_{\beta_1, \dots, \beta_t} \left| \Pr[\text{View}_{(P, V^*)}(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t] \cdot \Pr[P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t] \right.$$

$$\left. - \Pr[\text{View}_{(P, V^*)}(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t] \cdot \Pr[P(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t] \right|$$

$$\leq \sum_{\beta_1, \dots, \beta_t} \left| \Pr[S(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t] \cdot \Pr[P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t] \right.$$

$$\left. - \Pr[S(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t] \cdot \Pr[P(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t] \right.$$

$$\left. + \Pr[\text{View}_{(P, V^*)}(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t] \cdot \Pr[P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t] \right.$$

$$\left. - \Pr[\text{View}_{(P, V^*)}(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t] \cdot \Pr[P(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t] \right|$$

Since $P^*(x, \alpha_1 \beta_1 \dots \alpha_t)$ is defined to behave exactly like the simulator on the prover step at round t given the history so far and the verifier messages $\alpha_1 \dots \alpha_t$ as auxiliary input, we may rewrite the first term as $\Pr[S(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t \beta_t]$. Also, since P is the prover in the interaction (P, V^*) , we may rewrite the last term as: $-\Pr[\text{View}_{(P, V^*)}(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t \beta_t]$. Thus, we get:

$$\leq \sum_{\beta_1, \dots, \beta_t} \left| \Pr[S(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t \beta_t] - \Pr[\text{View}_{(P, V^*)}(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t \beta_t] \right|$$

$$+ \left| \Pr[S(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t] - \Pr[\text{View}_{(P, V^*)}(x, \alpha_1 \dots \alpha_t) = \alpha_1 \beta_1 \dots \alpha_t] \right|$$

$$\cdot \Pr[P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t]$$

$$\leq 2\delta$$

and we are done. ■

Corollary 5.4 *Let (P, V) be statistical ZK, let $x \in L \cap \{0, 1\}^n$, and let \hat{V} be any (possibly cheating) verifier. Then*

$$\sum_{\alpha_1 \beta_1, \dots, \alpha_t} \Pr[\text{View}_{(P, \hat{V})}(x) = \alpha_1 \beta_1 \dots \alpha_t] \cdot \sum_{\beta_t \in \{0, 1\}^{l(n)}} |\Pr[P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t] - \Pr[P(x, \alpha_1 \beta_1 \dots \alpha_t) = \beta_t]| < 2\delta$$

Proof: The second part of Lemma 5.3 guarantees the above inequality for all possible strings $\alpha_1 \alpha_2 \dots \alpha_t$. Namely, the view of (P, V^*) with auxiliary input $(\alpha_1 \dots \alpha_t)$ contains only verifier messages which are the predetermined $\alpha_1 \dots \alpha_t$. Thus, the inequality also holds when the strings $\alpha_1 \alpha_2 \dots \alpha_t$ are produced by any distribution space generated by any verifier \hat{V} . Note that the distribution on $\beta_1 \dots \beta_{t-1}$ does not change since we use the same prover P as in the lemma. ■

Note that δ is negligible. Combining this Corollary with Theorem 4.3 yields the desired conclusion.

Theorem 5.5 *Suppose L has a statistical ZK interactive proof system (P, V) . Then there is PPT oracle machine P_{eff} such that $(P_{\text{eff}}^{\text{NP}}, V)$ is a statistical ZK interactive proof system for L .*

Proof: Let $d^*(n)$ equal

$$\max_{x \in L, |x|=n} \sup_{a \in \{0, 1\}^*} d(S^*(x, a), \text{View}_{(P, V^*)}(x, a)).$$

We know that this function is negligible.

Let P^* be a universal prover for (P, V) . Let \hat{V} be (any) verifier. Let $x \in L \cap \{0, 1\}^n$ and $a \in \{0, 1\}^*$. We begin by showing that Corollary 5.4 implies that the distance between $\text{View}_{(P, \hat{V})}(x, a)$ and $\text{View}_{(P^*, \hat{V})}(x, a)$ is at most $2g(n)d^*(n)$. Note that since d^* is negligible, so is $2gd^*$.

We bound the difference between these two distributions by an induction on the round number. Namely, the induction claim is that in round t , $1 \leq t \leq g(n)$ it holds that

$$\sum_{\alpha_1 \beta_1 \dots \alpha_t \beta_t} \left| \Pr[\text{View}_{(P, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \alpha_t \beta_t] - \Pr[\text{View}_{(P^*, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \alpha_t \beta_t] \right| \leq t \cdot 2d^*(n).$$

The verifier behaves in exactly the same manner in both interactions and thus for any prefix α_1 both distributions are equal. Combining this with Corollary 5.4 we get that the claim holds for $t = 1$.

Let us show the induction step. The induction hypothesis implies that

$$\sum_{\alpha_1 \beta_1 \dots \alpha_{t-1} \beta_{t-1}} \left| \Pr[\text{View}_{(P, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \alpha_{t-1} \beta_{t-1}] - \Pr[\text{View}_{(P^*, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \alpha_{t-1} \beta_{t-1}] \right| \leq (t-1) \cdot 2d^*(n).$$

Since the verifier steps behave the same in both distributions we also have

$$\sum_{\alpha_1 \beta_1 \dots \alpha_t} \left| \Pr[View_{(P, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \beta_{t-1} \alpha_t] - \Pr[View_{(P^*, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \beta_{t-1} \alpha_t] \right| \leq (t-1) \cdot 2d^*(n). \quad (2)$$

Let us now show that the induction claims hold for round t .

$$\begin{aligned} & \sum_{\alpha_1 \beta_1 \dots \alpha_t \beta_t} \left| \Pr[View_{(P, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \alpha_t \beta_t] - \Pr[View_{(P^*, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \alpha_t \beta_t] \right| \\ &= \sum_{\alpha_1 \beta_1 \dots \alpha_t \beta_t} \left| \Pr[View_{(P, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \beta_{t-1} \alpha_t] \cdot \Pr[P(\alpha_1 \beta_1 \dots \alpha_t) = \beta_t] \right. \\ & \quad \left. - \Pr[View_{(P^*, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \beta_{t-1} \alpha_t] \cdot \Pr[P^*(\alpha_1 \beta_1 \dots \alpha_t) = \beta_t] \right| \end{aligned}$$

Similarly to the proof of Corollary 5.4 this can be bounded by:

$$\begin{aligned} & \leq \sum_{\alpha_1 \beta_1 \dots \alpha_t \beta_t} \Pr[View_{(P, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \beta_{t-1} \alpha_t] \cdot \\ & \quad \left| \Pr[P(\alpha_1 \beta_1 \dots \alpha_t) = \beta_t] - \Pr[P^*(\alpha_1 \beta_1 \dots \alpha_t) = \beta_t] \right| \\ & \quad + \left| \Pr[View_{(P, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \beta_{t-1} \alpha_t] - \Pr[View_{(P^*, \hat{V})}(x, a) = \alpha_1 \beta_1 \dots \beta_{t-1} \alpha_t] \right| \end{aligned}$$

Using Equation (2) and Corollary 5.4 we are done.

The negligible distance of $View_{(P, \hat{V})}(x, a)$ and $View_{(P^*, \hat{V})}(x, a)$ implies that (P^*, V) is a statistical ZK interactive proof system for L . To see this first note that setting \hat{V} in the above statement to be the honest verifier V implies that the completeness condition holds for (P^*, V) . Since V is unchanged the soundness of course still holds. And if $S_{\hat{V}}$ is a statistical ZK P -simulator for \hat{V} then it is also a statistical ZK P^* -simulator for \hat{V} .

Finally we note that by Theorem 4.3 a distribution within negligible distance of $P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = P_{S^*}(x, \alpha_1 \dots \alpha_t, \alpha_1 \beta_1 \dots \alpha_t)$ is computable in PPT with an NP oracle (denote this machine by $P_{\text{eff}}^{\text{NP}}$). Setting $P_{\text{eff}}^{\text{NP}}$ to be the prover instead of P^* does not foil the statistical ZK property of the interaction, since for any \hat{V} and any auxiliary input $a \in \{0, 1\}^*$ it holds that $View_{(P^*, \hat{V})}(x, a)$ and $View_{(P_{\text{eff}}^{\text{NP}}, \hat{V})}(x, a)$ have negligible distance. ■

We remark that perfect ZK can be preserved at the computational price of allowing P_{eff} to have a Σ_2^P oracle instead of an NP one:

Theorem 5.6 *Suppose L has a perfect ZK interactive proof system (P, V) . Then there is PPT oracle machine P_{eff} such that $(P_{\text{eff}}^{\Sigma_2^P}, V)$ is a perfect ZK interactive proof system for L .*

Proof: Use the same argument as in the proof of the previous theorem (except of course use Lemma 5.3 (2) rather than Corollary 5.4), but implement P_{S^*} using the uniform generation algorithm of [22] (rather than the almost uniform generation of Theorem 3.2). ■

It is possible to extend these techniques to derive bounds on the complexity of the prover in a proof of knowledge complexity $k(n) > 0$ (cf. Definition 2.8) while preserving

the KC. However, the bounds we get are not nearly as good: a straightforward application of our techniques shows only that if L has an interactive proof of knowledge complexity $k(n) \leq n^{O(1)}$ then L has an interactive proof (P, V) of knowledge complexity $k(n)$ in which P is a probabilistic, exponential time machine with a NEXP oracle. Not only is the complexity much greater than for ZK, but we don't know how to do better even if we assume $k(n)$ is just 1 (rather than zero). Deriving better bounds on the complexity of non-zero KC provers while preserving the KC is another open question.

In the next section we will see however that if we don't want to preserve the KC then better bounds *are* achievable, and this has applications to bounding the complexity of low KC languages.

6 Bounds on the Complexity of Low KC languages

In this section we present a result which ties the knowledge complexity of an interactive proof for a language L to the time complexity of L . This result relies on an extension to positive KC of our techniques from the previous sections.

Throughout this paper, we use the definition of knowledge complexity in the *fraction* sense [17]. Loosely speaking, we say that a protocol (P, V) has knowledge complexity $k(n)$ if for any verifier \hat{V} there exists a simulator $S_{\hat{V}}$ with the following “good subspace” property. For any $x \in L$ and auxiliary input $a \in \{0, 1\}^*$ there exists a subspace (denoted $SUCC_{x,a}$) of the set $\{0, 1\}^{p(n)}$ of the simulator's possible random tapes which has density $\geq 2^{-k(n)}$ (i.e. if we choose a tape at random from $\{0, 1\}^{p(n)}$ then this tape is in $SUCC_{x,a}$ with probability $\geq 2^{-k(n)}$), and the output of the simulator on input x, a and a uniformly chosen random tape in $SUCC_{x,a}$ is statistically close to a random element of $View_{(P, \hat{V})}(x, a)$. The formal definition appears in §2.

There are two other definitions of non-zero KC: the oracle definition and the hint definition [17]. It is shown in [17] that the oracle definition is equal up to an additive constant to the fraction definition. The proof of the equivalence result there does not change the interactive proof, and specifically, the equivalence holds also when a restriction is imposed on the number of rounds in the interaction. Therefore, our results apply to the oracle definition as well.

Our results apply also to the hint definition. However, with this definition, it is easy to prove even stronger results, and such stronger results already appear in [17]. Nevertheless, several reasons are presented in [17] why not to regard the hint definition as an appropriate measure.

Recall (Definition 2.3) that we have adopted the convention that an interactive proof has a negligible error probability. Let us now proceed to our result.

6.1 Overview

We are given a $g(n)$ round interactive proof system (P, V) for L with statistical KC $k(n)$, and we suppose $g(n)k(n) = O(\log n)$. Our goal is to show that $L \in \mathcal{BPP}^{\text{NP}}$. We begin with an overview of the proof.

We prove more than required. We show that under the given conditions, there exists an interactive proof for L in which the prover is a PPT machine with an access to an NP oracle. As the prover has the power to run the verifier as well, this implies that $L \in \text{BPP}^{\text{NP}}$.

Let S be a statistical KC $k(n)$ P -simulator for the honest verifier V (cf. Definition 2.7), and P_S the simulation based prover for S (cf. Definition 4.2). We let P^* be the prover derived from P_S by setting the auxiliary input to the empty string: $P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = P_S(x, \lambda, \alpha_1 \beta_1 \dots \alpha_t)$ (intuitively, think of P^* as being P_S ; the difference is only a technicality). We show that (P^*, V) has the following “separability” property: (1) if $x \in L$ then V , interacting with P^* , accepts with a probability that is greater than $n^{-O(1)}$, whereas, (2) if $x \notin L$ then V , interacting with any prover, accepts with a negligible probability. This is Lemma 6.3 (below).

Next, we use Theorem 4.3 to show that there is PPT machine with an NP oracle that can compute a distribution close enough to P^* such that the separability property is maintained. This gives us a PPT with NP oracle prover \overline{P} such that (\overline{P}, V) has the separability property. Finally, we observe that standard amplification techniques can be used to reduce the error probability and transform (\overline{P}, V) into an interactive proof system (cf. Definition 2.3) for L without increasing the power of the prover.

Remark 6.1 We note that although the definition of knowledge complexity (Definition 2.8) guarantees a KC $k(n)$ P -simulator for *each* possible verifier, we use only the KC $k(n)$ P -simulator for the honest verifier V . Therefore the validity of the theorem (Theorem 6.4) requires only that the interactive proof has knowledge complexity $k(n)$ with respect to the honest verifier.

Remark 6.2 For proving the result of this section we do not need to use the universal verifier nor to exploit the auxiliary string of the verifier and simulator. Therefore, throughout this section we shall use only the honest verifier with an auxiliary string set to λ .

6.2 Our Results

The main lemma is the following.

Lemma 6.3 *Suppose (P, V) is a $g(n)$ round interactive proof for L , and S is a statistical KC $k(n)$ P -simulator for V . Suppose also that $g(n)k(n) = O(\log n)$. Define the prover P^* by*

$$P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = P_S(x, \lambda, \alpha_1 \beta_1 \dots \alpha_t)$$

where P_S is the simulation based prover for S . Then (P^*, V) satisfies the following “separability” property:

- (1) *There is a constant $c \geq 0$ such that for any $x \in L$ the probability that V accepts at the end of the interaction with P^* on common input x is $\geq |x|^{-c}$.*
- (2) *For any (cheating) prover \hat{P} and common input $x \notin L$ the probability that V accepts at the end of the interaction with \hat{P} is negligible.*

Note that we are making no claims about the knowledge complexity of the system (P^*, V) constructed in the above lemma. This is in contrast to our results in §5 where we were

trying to preserve the knowledge complexity (in the particular case where this knowledge complexity was zero).

The (technical) proof of Lemma 6.3 appears in §6.3. Let us now state and prove our theorem using Lemma 6.3.

Theorem 6.4 *Suppose L has a $g(n)$ round interactive proof with statistical knowledge complexity $k(n)$, and suppose also that $g(n)k(n) = O(\log n)$. Then L is in BPP^{NP} .*

Proof: By assumption there is a $g(n)$ round interactive proof (P, V) with statistical knowledge complexity $k(n)$. It suffices to show that there is a PPT machine P_{eff} such that $(P_{\text{eff}}^{\text{NP}}, V)$ is an interactive proof system for L .

Let S be a statistical KC $k(n)$ P -simulator for V . Let P_S be the corresponding simulation based prover, and let P^* be as in Lemma 6.3. Then (P^*, V) has the separability property. We let \overline{P} denote the PPT with NP oracle prover which is given by using T rather than P_S in the definition of P^* , where T is the machine that is guaranteed in Theorem 4.3, setting $\delta = 2^{-n}$ in that theorem so that the output of T is 2^{-n} -close to the output of P_S .

Let $x \in L$. Observe that $\text{View}_{(P^*, V)}(x, \lambda)$ and $\text{View}_{(\overline{P}, V)}(x, \lambda)$ are $g2^{-n}$ -close, and thus, on common input x , the difference between the probability that P^* convinces V to accept and the probability that \overline{P} convinces V to accept is at most $g2^{-n}$. Since $g2^{-n}$ is negligible, this means that the separability property still holds:

- (1) There is a constant $c \geq 0$ such that for any $x \in L$ the probability that V accepts at the end of the interaction with \overline{P} on common input x is $\geq |x|^{-c}$
- (2) For any (cheating) prover \hat{P} and common input $x \notin L$ the probability that V accepts at the end of the interaction with \hat{P} is negligible.

Next, standard amplification techniques can be used to turn (\overline{P}, V) into an interactive proof for L without increasing the complexity of the prover. Finally, we note that a PPT machine with an NP oracle can play the role of both parties \overline{P} and V and decide whether $x \in L$. ■

6.3 Proof of Lemma 6.3

Let us restate Lemma 6.3 and prove it.

Lemma 6.3 *Suppose (P, V) is a $g(n)$ round interactive proof for L , and S is a statistical KC $k(n)$ P -simulator for V . Suppose also that $g(n)k(n) = O(\log n)$. Define the prover P^* by*

$$P^*(x, \alpha_1 \beta_1 \dots \alpha_t) = P_S(x, \lambda, \alpha_1 \beta_1 \dots \alpha_t)$$

where P_S is the simulation based prover for S . Then (P^*, V) satisfies the following “separability” property:

- (1) There is a constant $c \geq 0$ such that for any $x \in L$ the probability that V accepts at the end of the interaction with P^* on common input x is $\geq |x|^{-c}$
- (2) For any (cheating) prover \hat{P} and common input $x \notin L$ the probability that V accepts at the end of the interaction with \hat{P} is negligible

Proof: We begin with some intuition.

Condition (2) follows from the soundness condition of the interactive proof (P, V) ; the concern is condition (1). Fix an $x \in L$. Recall that when the random tape of the simulator S is uniformly picked in $SUCC_{x,\lambda}$, the output distribution of S is statistically close to the distribution of the conversations between P and V . We show that in each round of the interaction our new prover P^* (using P_S) picks its answer using a random tape in $SUCC_{x,\lambda}$ with probability $\geq 2^{-k}$. Therefore with probability at least $2^{-kg} = n^{-O(1)}$, P_S picks all its answers using random tapes from $SUCC_{x,\lambda}$. Intuitively, he thus gains an advantage which is $2^{-kg} = n^{-O(1)}$ times the advantage of the original prover P . However this is a simplification, because we must deal with the fact that in all but the first round, P_S is not sampling uniformly from the space of all coin tosses. Rather, he is sampling conditional to the prefix of the conversation so far being some particular value and in this conditional space it is not always the case that $SUCC_x$ has density at least 2^{-k} . The formal argument follows.

For simplicity, we first present the proof under the assumption that S is a perfect (rather than just statistical) KC $k(n)$ P -simulator for V . The extension to the more general case follows thereafter.

Let p be the number of coin tosses of S . We call a string c a t -th round prover prefix if $c = \alpha_1\beta_1 \dots \alpha_t\beta_t$ for some l bit strings $\alpha_1, \beta_1, \dots, \alpha_t, \beta_t$. Throughout this proof, we do not make a real use of the auxiliary input and it is always set to the empty string. In the sequel, we omit the auxiliary from the notations and it should be clear that the empty string is always used. We define R_c to be the set of random tapes of the simulator that are consistent with a t -th round prover prefix c . Namely,

$$R_c \stackrel{\text{def}}{=} \{ r \in \{0, 1\}^{p(n)} : S_t(x, r) = c \}$$

(Recall that $S_t(x, r)$ denotes the t -th round prefix of the conversation $S(x, \lambda, r)$, cf. Definition 4.1.) So, for example: $R_\lambda = \{0, 1\}^{p(n)}$. Similarly, we define G_c to be the subset of R_c that contains only random strings from the good subspace $SUCC_x$:

$$G_c \stackrel{\text{def}}{=} R_c \cap SUCC_x$$

Thus, $G_\lambda = SUCC_x$ and if $SUCC_x$ has density $2^{-k(n)}$ then $G_\lambda/R_\lambda = 2^{-k(n)}$. Also, the summation over all possible t -prefixes yields $\sum_c G_c = G_\lambda$ and $\sum_c R_c = R_\lambda$.

We now analyze the interaction between P^* and V on input $x \in L$. Let c be the t -th round prover prefix of this interaction. Recall that in the $t+1$ -st round $P_S(x, c)$ picks r at random from R_c , computes $S(x, r)$, and outputs β'_{t+1} , where $S(x, r) = (R, \alpha_1\beta_1 \dots \alpha_{t+1}\beta'_{t+1} \dots \alpha'_g\beta'_g)$ (cf. Definition 4.2). Let us denote by r_t the string r that was chosen by P_S in the t -th round. For $t = 1, \dots, g$ let us denote by GOOD_t the event that until round t only random tapes from $SUCC_x$ were used. Namely, $r_j \in SUCC_x$ for all $j = 1, \dots, t$. Fix GOOD_0 to be some event with probability 1.

We will show that for each round $t = 0, \dots, g-1$ it is the case that

$$\Pr[\text{GOOD}_{t+1} | \text{GOOD}_t] \geq 2^{-k} . \tag{3}$$

It follows that $\Pr[\text{GOOD}_g] \geq 2^{-kg} \geq n^{-d}$, where d is a constant such that $g(n)k(n) \leq d \lg n$. We still have to argue that given GOOD_g the interaction (P^*, V) ends with V accepting with high probability. Since S is, by assumption, a perfect KC $k(n)$ P -simulator for the

interaction (P, V) we know that it outputs the distribution $View_{(P, V)}(x)$ when its random tape r is chosen uniformly from $SUCC_x$ (cf. Definition 2.7). On the other hand, we know that P_S chooses each (possible) random tape with equal probability, the event GOOD_g being true implies that r_t is uniformly distributed in $G_c = SUCC_x \cap R_c$ for all rounds t . Therefore, the subspace of $View_{(P^*, V)}(x)$ obtained by conditioning on the event GOOD_g equals $View_{(P, V)}(x)$. So the probability that V accepts in the interaction with P^* is at least $\Pr[\text{GOOD}_g]$ times the probability that V accepts in the interaction with P . Thus, the completeness of (P, V) implies that the probability that V accepts in the interaction with P^* is $\geq n^{-O(1)}$, as desired. It remains to justify Equation (3). By definition, if we sample the space of random tapes of the simulator once, we hit $SUCC_x$ with probability at least 2^{-k} . However, when we are sampling (uniformly) a subset of the possible random tapes of the simulator (the subset that is consistent with the conversation so far), it is not necessarily true that we hit $SUCC_x$ with probability $\geq 2^{-k}$. Still, Equation (3) can be derived by analyzing the prefix of the conversation (and thus the subset of the simulators possible coin tosses) as a stochastic variable.

Fix a round t and assume GOOD_t . We will now give a lower bound on the probability of GOOD_{t+1} . All probabilities in this proof are taken over the probability space defined by the interaction between the original verifier V with the simulation based prover P^* . The summation in the following equation is over all possible t -th round verifier prefixes, $c \in \{0, 1\}^{t \cdot 2^l}$, and over all possible verifier message in round $t + 1$, $\alpha \in \{0, 1\}^l$.

$$\begin{aligned}
q &\stackrel{\text{def}}{=} \Pr[\text{GOOD}_{t+1} | \text{GOOD}_t] \\
&= \sum_{c \circ \alpha} \Pr[c \circ \alpha \wedge \text{GOOD}_{t+1} | \text{GOOD}_t] \\
&= \sum_{c \circ \alpha} \Pr[\text{GOOD}_{t+1} | c \circ \alpha \wedge \text{GOOD}_t] \cdot \Pr[c \circ \alpha | \text{GOOD}_t] \\
&= \sum_{c \circ \alpha} \Pr[\text{GOOD}_{t+1} | c \circ \alpha] \cdot \Pr[c \circ \alpha | \text{GOOD}_t] \tag{4}
\end{aligned}$$

The last equality is true since conditioned on the history $c \circ \alpha$, the event GOOD_{t+1} does not depend on GOOD_t . Next, we note that the event that V chooses his next message to be α given the history c does not depend on the event GOOD_t . Therefore, partitioning the above summation into a summation on the history so far c and a summation on the next verifier message α , we get:

$$q = \sum_c \Pr[c | \text{GOOD}_t] \cdot \sum_\alpha \Pr[\alpha | c] \cdot \Pr[\text{GOOD}_{t+1} | c \circ \alpha]$$

Given that the simulator's coins are picked uniformly in $SUCC_x$, we know that the distribution output by the simulator is exactly equal to the distribution of the interaction (P, V) .[†] Thus, given GOOD_t , we know that the t -prefixes of the interaction (P^*, V) are distributed the same as the t -prefixes of the interaction (P, V) and these are distributed equally to the t -prefixes output by the simulation given that we pick the simulator coins uniformly in $SUCC_x$. Therefore, we may write the probability that the prefix c appears in the interaction (P^*, V) in this case as $|G_c|/|G_\lambda|$. Namely, the probability of the prefix c given GOOD_t

[†]Recall that we are now dealing with the more restricted case of *perfect* knowledge complexity.

is equal to the number of simulator random tapes in $SUCC_x$ that agree with this prefix, divided by the number of random tapes in $SUCC_x$. Also, we may replace the probability that V responds with the string α to the history c by $|G_{c\circ\alpha}|/|G_c|$. Last, we note that the probability of GOOD_{t+1} given the history $c \circ \alpha$ is exactly $|G_{c\circ\alpha}|/|R_{c\circ\alpha}|$ since P^* chooses r_{t+1} randomly in $R_{c\circ\alpha}$. To summarize, we may write:

$$\begin{aligned} q &= \sum_c \frac{|G_c|}{|G_\lambda|} \cdot \sum_\alpha \frac{|G_{c\circ\alpha}|}{|G_c|} \cdot \frac{|G_{c\circ\alpha}|}{|R_{c\circ\alpha}|} \\ &= \frac{1}{|G_\lambda|} \cdot \sum_{c\circ\alpha} \frac{|G_{c\circ\alpha}|^2}{|R_{h\circ\alpha}|} \end{aligned} \quad (5)$$

By the Cauchi-Schwartz inequality we know that for positive $x_i, y_i, 1 \leq i \leq n$, it holds that

$$\left(\sum_{i=1}^n \frac{x_i^2}{y_i} \right) \cdot \left(\sum_{i=1}^n y_i \right) \geq \left(\sum_{i=1}^n x_i \right)^2$$

Using this with $x_i = |G_{c\circ\alpha}|$ and $y_i = |R_{h\circ\alpha}|$, in Equation (5), we get:

$$\begin{aligned} q &\geq \frac{1}{|G_\lambda|} \cdot \frac{(\sum_{c\circ\alpha} |G_{c\circ\alpha}|)^2}{\sum_{c\circ\alpha} |R_{h\circ\alpha}|} \\ &= \frac{1}{|G_\lambda|} \cdot \frac{|G_\lambda|^2}{|R_\lambda|} \\ &= \frac{|G_\lambda|}{|R_\lambda|} \geq 2^{-k} \end{aligned}$$

Thus, we get that $\Pr[\text{GOOD}_{t+1} | \text{GOOD}_t] \geq 2^{-k}$ as claimed and we are done with the proof of Lemma 6.3.

6.3.1 The case of statistical knowledge complexity

The case of statistical knowledge complexity is a little more involved. Again we are going to show two things. First, that GOOD_g has a large enough probability (Claim 6.6 bellow), and second that given GOOD_g , V accepts with high probability (Claim 6.5 bellow). Our proof extends the proof for the case of perfect knowledge complexity and in fact, we only show that the difference between the perfect case and the statistical case is small enough.

Recall that in the statistical case it is only guaranteed that the view of V in the interaction (P, V) is statistically close to the output distribution of the simulator when its random tape is picked uniformly in the good subspace $SUCC_x$.

We are going to define a new prover P_G^* and a new verifier V_G^* who will help us analyse the protocol (P^*, V) which interests us. The prover P_G^* and the verifier V_G^* are fictitious entities which do not really appear in any of the interactions. Also, the verifier V_G^* will not necessarily be implementable as a probabilistic polynomial time machine. We define P_G^* to be the simulation based prover that uses only random tapes from $SUCC_x$. Recall that the simulation based prover on a given history so far c , picks uniformly a random tape $r \in R_c$ and outputs the next message of the prover in the intraction $S(x, \lambda, r)$. The prover P_G^* will

pick uniformly $r \in G_c$ instead. Thus, if we were considering perfect simulation, we would get that the prover P_G^* acts exactly like the original prover P . However, since the simulation is not perfect, we only get some closeness between P_G^* and P which we have to analyse. The verifier V_G^* is defined analogously. Namely, in response to a history so far c , it picks uniformly in G_c and answers according to the next verifier step in the conversation $S(x, \lambda, r)$. Note again that V_G^* does not act exactly like V only because the simulation is not perfect.

Consider now the interaction between (P_G^*, V_G^*) . This is exactly equal to the output distribution of the simulator given $SUCC_x$. Two facts follow from this observation. First, we get that the simulator S is a *perfect* KC $k(n)$ P-simulator for the interaction (P_G^*, V_G^*) . Second, by the property of the good subspace $SUCC_x$ we get that, the interaction (P_G^*, V_G^*) is statistically close to the original interaction (P, V) . It can be proven by induction that the interaction between P_G^* to V is statistically close both to the interaction (P, V) and to the interaction (P_G^*, V_G^*) . In general, if an interaction (A, B) that has g rounds is δ -close to an interaction (A', B') then it also hold that the interaction (A', B) is $g \cdot \delta$ -close to both. (This can be formally proven by a simple induction).

From the fact that (P_G^*, V_G^*) is statistically close to (P, V) we get that (P_G^*, V_G^*) accepts x with probability almost 1. Since the simulator S is a perfect simulator for the interaction (P_G^*, V_G^*) , we can use our proof for the perfect case and get that P^* convinces V_G^* with probability at least n^{-c} for some constant $c > 0$. However, we are interested in the probability that P^* convinces V on x .

Let us first show that given GOOD_g , V accepts x with high probability.

Claim 6.5 *Given GOOD_g , the probability that V accepts on input $x \in L$ while talking to P^* , is $\geq 1 - \epsilon$ for some negligible fraction ϵ .*

Proof: Conditioned on the event GOOD_g , we get that P^* is behaving exactly like P_G^* . Now, since (P_G^*, V) is statistically close to (P, V) and our claim follows. ■

Next, we show that the probability of the event GOOD_g in the interaction (P^*, V) is bigger than some polynomial fraction. Here, we are going to show that the difference between the perfect case and the statistical case is small in this respect.

Claim 6.6 *Let $x \in L$ and consider the probability space of the interaction (P^*, V) on the input x . Then,*

$$\Pr[\text{GOOD}_g] \geq n^{-O(1)}$$

Proof: We prove that for every t , $0 \leq t \leq g - 1$,

$$\Pr[\text{GOOD}_{t+1} | \text{GOOD}_t] \geq 2^{-k} - \epsilon_1 \tag{6}$$

for some negligible fraction ϵ_1 . Thus we get that

$$\Pr[\text{GOOD}_g] = \prod_{t=0}^{g-1} \Pr[\text{GOOD}_{t+1} | \text{GOOD}_t] \geq 2^{-kg} - \epsilon$$

for some negligible fraction ϵ , and we are done.

So let us show that Equation (6) holds. From now and on we are going to consider now two probability spaces: the probability space induced by the interaction (P^*, V) and the

probability space induced by the interaction (P^*, V_G^*) . Therefore, we shall explicitly write for each probability which space is considered.

Since S is a perfect simulator for the interaction (P^*, V_G^*) , our proof for the perfect case implies

$$\Pr_{(P^*, V_G^*)} [\text{GOOD}_{t+1} | \text{GOOD}_t] \geq 2^{-k}. \quad (7)$$

We are going to show that when V participates in the interaction instead of V_G^* , this property does not change significantly. Specifically, we will show that

$$\epsilon_1 \stackrel{\text{def}}{=} \Pr_{(P^*, V_G^*)} [\text{GOOD}_{t+1} | \text{GOOD}_t] - \Pr_{(P^*, V)} [\text{GOOD}_{t+1} | \text{GOOD}_t] \quad (8)$$

is a negligible fraction. This is enough since Equation (6) follows from Equation (7) and Equation (8).

Again, we sum over all t -round prefixes $c \in \{0, 1\}^{t \cdot 2l}$ and over all possible verifier responses $\alpha \in \{0, 1\}^l$:

$$\begin{aligned} \epsilon_1 &= \sum_{c \circ \alpha} \Pr_{(P^*, V_G^*)} [\text{GOOD}_{t+1} | c \circ \alpha] \cdot \Pr_{(P^*, V_G^*)} [c \circ \alpha | \text{GOOD}_t] \\ &\quad - \Pr_{(P^*, V)} [\text{GOOD}_{t+1} | c \circ \alpha] \Pr_{(P^*, V)} [c \circ \alpha | \text{GOOD}_t] \end{aligned}$$

Since the probability of the event GOOD_{t+1} conditioned on the prefix $c \circ \alpha$ depends only on the behavior of the prover P^* which is identical in both distributions, we may write:

$$\begin{aligned} \epsilon_1 &= \sum_{c \circ \alpha} \Pr_{(P^*, V_G^*)} [\text{GOOD}_{t+1} | c \circ \alpha] \cdot \left(\Pr_{(P^*, V_G^*)} [c \circ \alpha | \text{GOOD}_t] - \Pr_{(P^*, V)} [c \circ \alpha | \text{GOOD}_t] \right) \\ &\leq \sum_{c \circ \alpha} \Pr_{(P^*, V_G^*)} [c \circ \alpha | \text{GOOD}_t] - \Pr_{(P^*, V)} [c \circ \alpha | \text{GOOD}_t] \\ &= \sum_{c \circ \alpha} \Pr_{(P_G^*, V_G^*)} [c \circ \alpha] - \Pr_{(P_G^*, V)} [c \circ \alpha] \quad (9) \end{aligned}$$

The last equality holds since conditioning on GOOD_t in all relevant prover steps is equivalent to the prover being P_G^* . Finally, since the interaction (P_G^*, V) is statistically close to (P_G^*, V_G^*) , we get that the expression in Equation (9) is negligible and we are done. ■

7 Conclusions and Open Problems

We showed that if L has a statistical ZK proof then it has a statistical ZK proof with a prover who runs in PPT with an NP oracle. This was previously only known given complexity assumptions. Our first question is whether one can remove assumptions from other similar problems. In particular, can one unconditionally establish any of the following?

- (1) If L has a statistical ZK proof then it has a statistical ZK proof with perfect completeness (i.e. the verifier accepts with probability 1 when $x \in L$; cf. [14])
- (2) If L has a statistical ZK proof then it has a statistical ZK proof with a black-box simulator.

- (3) If L has an interactive proof which is statistical ZK with respect to the honest verifier, then it has a statistical ZK interactive proof.

We recall that these results are known given complexity assumptions [5].

Second, on the subject of the power of the prover. Our bound of PPT with NP oracle does not depend on the complexity of the language. Can one find “tighter” relationships between the complexity of L and the complexity of a statistical ZK prover for L ? For example, what can one say about statistical ZK in the model of [4] where the prover is PPT with an oracle for L ?

Acknowledgments

We thank Oded Goldreich and Avi Wigderson for helpful remarks and proof simplifications.

References

- [1] W. AIELLO AND J. HÅSTAD. Perfect Zero-Knowledge can be Recognized in Two Rounds. *Proceedings of the 28th Annual IEEE Symposium on the Foundations of Computer Science*, IEEE (1987).
- [2] L. BABAI. Trading Group Theory for Randomness. *Proceedings of the 17th Annual ACM Symposium on the Theory of Computing*, ACM (1985).
- [3] L. BABAI, L. FORTNOW AND C. LUND. Non-Deterministic Exponential Time has Two-Prover Interactive Protocols. *Proceedings of the 31st Annual IEEE Symposium on the Foundations of Computer Science*, IEEE (1990).
- [4] R. BEIGEL, M. BELLARE, J. FEIGENBAUM AND S. GOLDWASSER. Languages that are Easier than their Proofs. *Proceedings of the 32nd Annual IEEE Symposium on the Foundations of Computer Science*, IEEE (1991).
- [5] M. BELLARE, S. MICALI AND R. OSTROVSKY. The (True) Complexity of Statistical Zero-Knowledge. *Proceedings of the 22nd Annual ACM Symposium on the Theory of Computing*, ACM (1990).
- [6] M. BELLARE AND E. PETRANK. Making Zero-Knowledge Provers Efficient. Technical Report, Computer Science Department, Technion, Haifa, Israel.
- [7] M. BEN-OR, S. GOLDWASSER, O. GOLDREICH, J. HÅSTAD, J. KILIAN, S. MICALI AND P. ROGAWAY. Everything Provable is Provable in Zero-Knowledge. *Advances in Cryptology — Proceedings of CRYPTO 88*, Lecture Notes in Computer Science 403, Springer-Verlag (1989). S. Goldwasser, ed.
- [8] M. BEN-OR, S. GOLDWASSER, J. KILIAN AND A. WIGDERSON. Multi-Prover Interactive Proofs: How to Remove Intractability Assumptions. *Proceedings of the 20th Annual ACM Symposium on the Theory of Computing*, ACM (1988).
- [9] M. BLUM AND S. KANNAN. Designing Programs that Check their Work. *Proceedings of the 21st Annual ACM Symposium on the Theory of Computing*, ACM (1989).
- [10] L. CARTER AND M. WEGMAN. Universal Classes of Hash Functions. *J. Computer and System Sciences* **18**, 143–154 (1979).

- [11] U. FEIGE. Interactive Proofs. M.Sc Thesis, Weizmann Institute of Science. August 1987.
- [12] L. FORTNOW. The Complexity of Perfect Zero-Knowledge. *Advances in Computing Research* (ed. S. Micali) Vol. 18 (1989).
- [13] O. GOLDREICH AND H. KRAWCZYK. On the Composition of Zero-Knowledge Proof Systems. *Proceedings of ICALP 90*.
- [14] O. GOLDREICH, Y. MANSOUR AND M. SIPSER. Interactive Proof Systems: Provers that never Fail and Random Selection. *Proceedings of the 28th Annual IEEE Symposium on the Foundations of Computer Science*, IEEE (1987).
- [15] O. GOLDREICH AND Y. OREN. Definitions and Properties of Zero-Knowledge Proof Systems. Technical Report #570, Technion (1989).
- [16] O. GOLDREICH AND E. PETRANK. Quantifying Knowledge Complexity. *Proceedings of the 32nd Annual IEEE Symposium on the Foundations of Computer Science*, IEEE (1991). O. GOLDREICH, R. OSTROVSKY, AND E. PETRANK. Computational Complexity and Knowledge Complexity. *26th ACM Symp. on Theory of Computation*, May 1994. pp. 534-543.
- [17] O. GOLDREICH AND E. PETRANK. Quantifying Knowledge Complexity. *Proceedings of the 32nd Annual IEEE Symposium on the Foundations of Computer Science*, IEEE (1991).
- [18] S. GOLDWASSER, S. MICALI, AND C. RACKOFF. The Knowledge Complexity of Interactive Proofs. *Proceedings of the 17th Annual ACM Symposium on the Theory of Computing*, ACM (1985).
- [19] S. GOLDWASSER, S. MICALI, AND C. RACKOFF. The Knowledge Complexity of Interactive Proofs. *SIAM J. Comput.* **18** (1), 186-208 (February 1989).
- [20] R. IMPAGLIAZZO, L. LEVIN AND M. LUBY. Pseudo-Random Generation from One-Way Functions. *Proceedings of the 21st Annual ACM Symposium on the Theory of Computing*, ACM (1989).
- [21] R. IMPAGLIAZZO AND M. YUNG. Direct Minimum-Knowledge computations. *Advances in Cryptology — Proceedings of CRYPTO 87*, Lecture Notes in Computer Science 293, Springer-Verlag (1987).
- [22] M. JERRUM, L. VALIANT AND V. VAZIRANI. Random Generation of Combinatorial Structures from a Uniform Distribution. *Theoretical Computer Science* **43**, 169-188 (1986).
- [23] C. LUND, L. FORTNOW, H. KARLOFF AND N. NISAN. Algebraic Methods for Interactive Proof Systems. *Proceedings of the 31st Annual IEEE Symposium on the Foundations of Computer Science*, IEEE (1990).
- [24] Y. OREN. On The Cunning Power of Cheating Verifiers: Some Observations About Zero Knowledge Proofs. *Proceedings of the 28th Annual IEEE Symposium on the Foundations of Computer Science*, IEEE (1987).
- [25] R. OSTROVSKY. One-Way Functions, Hard on Average Problems, and Statistical Zero-Knowledge Proofs. Structures 1991.
- [26] R. OSTROVSKY, R. VENKATESAN AND M. YUNG. On the Complexity of Asymmetric Games. Manuscript (1990).
- [27] M. SIPSER. A Complexity Theoretic Approach to Randomness. *Proceedings of the 15th Annual ACM Symposium on the Theory of Computing*, ACM (1983).

- [28] L. STOCKMEYER. The Complexity of Approximate Counting. *Proceedings of the 15th Annual ACM Symposium on the Theory of Computing*, ACM (1983).
- [29] M. TOMPA AND H. WOLL. Random Self-Reducibility and Zero-Knowledge Proofs of Possession of Information. *Proceedings of the 28th Annual IEEE Symposium on the Foundations of Computer Science*, IEEE (1987).