

# Composition of Zero-Knowledge Proofs with Efficient Provers<sup>\*</sup>

Eleanor Birrell<sup>1</sup> and Salil Vadhan<sup>2</sup>

<sup>1</sup> Department of Computer Science, Cornell University.  
eleanor@cs.cornell.edu

<sup>2</sup> School of Engineering and Applied Sciences and Center for Research on  
Computation and Society, Harvard University<sup>\*\*</sup>  
salil@seas.harvard.edu

**Abstract.** We revisit the composability of different forms of zero-knowledge proofs when the honest prover strategy is restricted to be polynomial time (given an appropriate auxiliary input). Our results are:

1. When restricted to efficient provers, the original Goldwasser–Micali–Rackoff (GMR) definition of zero knowledge (STOC ‘85), here called *plain zero knowledge*, is closed under a constant number of sequential compositions (on the same input). This contrasts with the case of unbounded provers, where Goldreich and Krawczyk (ICALP ‘90, SICOMP ‘96) exhibited a protocol that is zero knowledge under the GMR definition, but for which the sequential composition of 2 copies is not zero knowledge.
2. If we relax the GMR definition to only require that the simulation is indistinguishable from the verifier’s view by uniform polynomial-time distinguishers, with no auxiliary input beyond the statement being proven, then again zero knowledge is not closed under sequential composition of 2 copies.
3. We show that auxiliary-input zero knowledge with efficient provers is not closed under *parallel* composition of 2 copies under the assumption that there is a secure key agreement protocol (in which it is easy to recognize valid transcripts). Feige and Shamir (STOC ‘90) gave similar results under the seemingly incomparable assumptions that (a) the discrete logarithm problem is hard, or (b)  $UP \not\subseteq BPP$  and one-way functions exist.

## 1 Introduction

Composition has been one of the most active subjects of research on zero-knowledge proofs. The goal is to understand whether the zero-knowledge property is preserved when a zero-knowledge proof is repeated many times. The answers vary depending on the variant of zero knowledge in consideration and the

---

<sup>\*</sup> These results first appeared in the first author’s undergraduate thesis [5] and in the full version of the paper is available on the *Cryptology ePrint Archive* [6].

<sup>\*\*</sup> 33 Oxford Street, Cambridge, MA 02138. <http://seas.harvard.edu/~salil/>. Supported by NSF grant CNS-0831289.

form of composition (e.g. sequential, parallel, or concurrent). The study of composition was first aimed at reducing the soundness error of basic constructions of zero-knowledge proofs (via sequential or parallel composition), but was later also motivated by considering networked environments in which an adversary might be able to open several instances of a protocol (even concurrently).

Soon after Goldwasser, Micali, and Rackoff introduced the concept of zero-knowledge proofs [20], it was realized that composability is a subtle issue. In particular, this motivated a strengthening of the GMR definition, known as *auxiliary-input zero knowledge* [21, 19, 9], which was shown to be closed under sequential composition [19]. The need for this stronger definition was subsequently justified by a result of Goldreich and Krawczyk [16], who showed that the original GMR definition is not closed under sequential composition. Specifically, they exhibited a protocol that is *plain zero knowledge* when executed once, but fails to be zero knowledge when executed twice sequentially.

The starting point for our work is the realization that the Goldreich–Krawczyk protocol is not an entirely satisfactory counterexample, because the prover strategy is inefficient (i.e. super-polynomial time). Most cryptographic applications of zero-knowledge proofs require a prover strategy that can be implemented efficiently given an appropriate auxiliary input (e.g. NP witness). Prover efficiency can intuitively have an impact on the composability of zero-knowledge proofs, because an adversarial verifier may be able to use the extra computational power of one prover copy to “break” the zero-knowledge property of another copy. Indeed, known positive results on the parallel and concurrent composability of witness-indistinguishable proofs (a weaker variant of zero-knowledge proofs) rely on prover efficiency [9].

Thus, we revisit the sequential composability of plain zero knowledge, but restricted to efficient provers. Our first result is positive, and shows that such proofs *are* closed under any constant number of sequential compositions (in contrast to the Goldreich–Krawczyk result with unbounded provers). The case of a superconstant or polynomial number of compositions remains an interesting open question. This positive result refers to the standard formulation of plain zero knowledge, where the simulation and the verifier’s view are required to be indistinguishable by nonuniform polynomial-time distinguishers (or distinguishers that are given the prover’s auxiliary input in addition to the statement being proven).

We then consider the case where the distinguishers are uniform probabilistic polynomial-time algorithms, whose only additional input is the statement being proven. In this case, we obtain a negative result analogous to the one of Goldreich and Krawczyk, showing that zero knowledge is not closed under sequential composition of even 2 copies (assuming that  $\mathcal{NP} \not\subseteq \mathcal{BPP}$ ). Informally, these two results say that plain zero knowledge is closed under a constant number of sequential compositions if and only if the distinguishers are at least as powerful as the prover.

We also examine the *parallel* composability of *auxiliary-input* zero knowledge. Here, too, Goldreich and Krawczyk [16] gave a negative result that utilizes an

inefficient prover. Feige and Shamir [9], however, gave a negative result with an efficient prover, under the assumption that the discrete logarithm is hard, or more generally under the assumptions that  $\mathcal{UP} \not\subseteq \mathcal{BPP}$  and one-way functions exist. We are interested in whether the complexity assumption used by Feige and Shamir can be weakened. To this end, we provide a negative result under a seemingly incomparable assumption, namely that there exists a key agreement protocol (in which it is easy to recognize valid transcripts).

## 2 Definitions and Preliminaries

### 2.1 Interactive Proofs

Given two interactive Turing machines – a prover  $P$  and a verifier  $V$  – we consider two types of interactive protocols: proofs of language membership (interactive proofs) and proofs of knowledge. In each case, both parties receive a common input  $x$ , and  $P$  is trying to convince  $V$  that  $x \in L$  for some language  $L$ . We will allow  $P$  to have an extra “auxiliary input” or “witness”  $y$ . We use the notation  $(P, V)$  to denote an interactive protocol and the notation  $\langle P(x, y), V(x) \rangle$  to denote the verifier  $V$ ’s view of that protocol with inputs  $(x, y)$  and  $x$  respectively. The choices for  $y$  will be given by a relation of the following kind:

**Definition 2.1 (Poly-balanced Relation).** *A binary relation  $R$  is poly-balanced if there exists a polynomial  $p$  such that for all  $(x, y) \in R$ ,  $|y| \leq p(|x|)$ . The language generated by such a relation is denoted  $L_R = \{x : (x, y) \in R\}$ .*

Observe that we don’t require  $R$  to be polynomial-time verifiable, so every language  $L$  is generated by such a relation, for example the relation  $R = \{(x, y) : |y| = |x| \text{ and } x \in L\}$ .

**Definition 2.2 (Interactive Proof).** *We say that an interactive protocol  $(P, V)$  is an interactive proof system for a language  $L$  if there exists a poly-balanced relation  $R$  such that  $L = L_R$  and the following properties hold:*

- (Verifier Efficiency): *The verifier  $V$  runs in time at most  $\text{poly}(|x|)$  on input  $x$ .*
- (Completeness): *If  $(x, y) \in R$  then the verifier  $V(x)$  accepts with probability 1 after interacting with the prover  $P(x, y)$  on common input  $x$  and prover auxiliary input  $y$ .*
- (Soundness): *There exists a function  $s(n) \leq 1 - 1/\text{poly}(n)$  (called the soundness error) for which it holds that for all  $x \notin L$  and for all prover strategies  $P^*$ , the verifier  $V(x)$  accepts with probability at most  $s(|x|)$  after interacting with  $P^*$  on common input  $x$  and prover auxiliary input  $y$ .*

**Definition 2.3 (Proof of Knowledge).** *Let  $R$  be a poly-balanced relation. Given an interactive protocol  $(P, V)$ , we let  $p(x, y, r)$  be the probability that  $V$  accepts on common input  $x$  when  $y$  is  $P$ ’s auxiliary input and  $r$  is the random input generated by  $P$ ’s random coin flips. Let  $P_{x,y,r}$  be the function such that*

$P_{x,y,r}(\bar{m})$  is the message sent by  $P$  after receiving messages  $\bar{m}$ . An interactive protocol  $(P(x,y), V(x))$  is an interactive proof of knowledge for the relation  $R$  if the following three properties hold:

- (Verifier Efficiency): The verifier  $V$  runs in time at most  $\text{poly}(|x|)$  on input  $x$ .
- (Completeness): If  $(x,y) \in R$ , then  $V$  accepts after interacting with  $P$  on common input  $x$ .
- (Extraction): There exists a function  $s(n) \leq 1 - 1/\text{poly}(n)$  (called the soundness error), a polynomial  $q$ , and a probabilistic oracle machine  $K$  such that for every  $x,y,r \in \{0,1\}^*$ ,  $K$  satisfies the following condition: if  $p(x,y,r) > s(|x|)$  then on input  $x$  and with access to oracle  $P_{x,y,r}$  machine  $K$  outputs  $w$  such that  $(x,w) \in R$  within an expected number of steps bounded by  $q(|x|)/(p(x,y,r) - s(|x|))$ .

Observe that extraction implies soundness, so a proof of knowledge for  $R$  is also an interactive proof for  $L_R$ .

Although the above definitions require a polynomial-time verifier, neither places any restriction on the computational power of the prover  $P$ . In keeping with the standard model of “realistic” computation, we sometimes prefer to limit the computational resources of both parties to polynomial time. Specifically, we add the additional requirement that there exists a polynomial  $p$  such that the prover  $P(x,y)$  runs in time  $p(|x|, |y|)$  where  $x$  is the common input and  $y$  is the prover’s auxiliary input. We refer to such protocols as *efficient* or *efficient-prover* proofs.

## 2.2 Zero Knowledge

In keeping with the literature, we define zero knowledge in terms of the indistinguishability of the output distributions.

**Definition 2.4 (Uniform/Nonuniform Indistinguishability).** *Two ensembles of probability distributions  $\{\Pi_1(x)\}_{x \in S}$  and  $\{\Pi_2(x)\}_{x \in S}$  are uniformly (resp. nonuniformly) indistinguishable if for every uniform (resp. nonuniform) probabilistic polynomial-time algorithm  $D$ , there exists a negligible function  $\mu$  such that for every  $x \in S$ ,*

$$\left| \Pr[D(1^{|x|}, \Pi_1(x)) = 1] - \Pr[D(1^{|x|}, \Pi_2(x)) = 1] \right| \leq \mu(|x|),$$

where the probability is taken over the samples of  $\Pi_1(x)$  and  $\Pi_2(x)$  and the coin tosses of  $D$ .

Often, definitions of computational indistinguishability give the distinguisher the index  $x$  (not just its length). This makes no difference for nonuniform distinguishers – since they can have  $x$  hardwired in – but it does matter for uniform distinguishers. Indeed, we will see that zero-knowledge proofs demonstrate different properties under composition depending on how much information the distinguisher is given about the inputs.

Also, uniform indistinguishability is usually not defined with a universal quantifier over  $x \in S$ , but instead with respect to all polynomial-time samplable distributions on  $x \in S$  (e.g. [2][12]). We use the above definition for simplicity, but our results also extend to the usual definition.

For the purposes of this paper, we consider two different definitions of zero knowledge. The first, which has primarily been of interest for historical reasons, is the one originally introduced by Goldwasser, Micali, and Rackoff [20]:

**Definition 2.5 (Plain Zero Knowledge).** *An interactive proof system  $(P, V)$  for a language  $L = L_R$  is plain zero knowledge (with respect to nonuniform distinguishers) if for all probabilistic polynomial-time machines  $V^*$ , there exists a probabilistic polynomial-time algorithm  $M_{V^*}$  that on input  $x$  produces an output probability distribution  $\{M_{V^*}(x)\}$  such that  $\{M_{V^*}(x)\}_{(x,y) \in R}$  and  $\{\langle P(x, y), V^*(x) \rangle\}_{(x,y) \in R}$  are nonuniformly indistinguishable.*

As is standard, the above definition refers to *nonuniform* distinguishers (which can have  $x, y$  and any additional information depending on  $x, y$  hard-wired in as nonuniform advice). However, it is also natural to consider *uniform* distinguishers. In this setting, it is important to differentiate between the case where the distinguisher is only given the single verifier input  $x$  and the case where the distinguisher is given both  $x$  and the prover's auxiliary input  $y$ .

**Definition 2.6.** *An interactive proof system  $(P, V)$  for a language  $L = L_R$  is plain zero knowledge with respect to  $V$ -uniform distinguishers if for all probabilistic polynomial-time machines  $V^*$ , there exists a probabilistic polynomial-time algorithm  $M_{V^*}$  that on input  $x$  produces an output probability distribution  $\{M_{V^*}(x)\}$  such that  $\{(x, M_{V^*}(x))\}_{(x,y) \in R}$  and  $\{(x, \langle P(x, y), V^*(x) \rangle)\}_{(x,y) \in R}$  are uniformly indistinguishable.*

**Definition 2.7.** *An interactive proof system  $(P, V)$  for a language  $L = L_R$  is plain zero knowledge with respect to  $P$ -uniform distinguishers if for all probabilistic polynomial-time machines  $V^*$ , there exists a probabilistic polynomial-time algorithm  $M_{V^*}$  that on input  $x$  produces an output probability distribution  $\{M_{V^*}(x)\}$  such that  $\{(x, y, M_{V^*}(x))\}_{(x,y) \in R}$  and  $\{(x, y, \langle P(x, y), V^*(x) \rangle)\}_{(x,y) \in R}$  are uniformly indistinguishable.*

The next definition of zero knowledge that we will consider is the more standard definition which incorporates an auxiliary input for the verifier.

**Definition 2.8 (Auxiliary-Input Zero Knowledge).** *An interactive proof system  $(P, V)$  for a language  $L$  is auxiliary-input zero knowledge if for every probabilistic polynomial-time machine  $V^*$  and every polynomial  $p$  there exists a probabilistic polynomial-time machine  $M_{V^*}$  such that the probability ensembles  $\{\langle P(x, y), V^*(x, z) \rangle\}_{(x,y) \in R, z \in \{0,1\}^{p(|x|)}}$  and  $\{M_{V^*}(x, z)\}_{(x,y) \in R, z \in \{0,1\}^{p(|x|)}}$  are nonuniformly indistinguishable.*

Observe that although this last definition is given only in terms of nonuniform indistinguishability, this is actually equivalent to requiring only uniform indistinguishability; any nonuniform advice used by the distinguisher can instead be incorporated into the verifier's auxiliary input  $z$ .

### 2.3 Composition

In this section, we explicitly state the definitions of sequential and parallel composition that will be used throughout this paper. These definitions can be applied to any of the definitions of zero knowledge given in the previous section.

**Definition 2.9.** *Given an interactive proof system  $(P, V)$  and a polynomial  $t(n)$ , we consider the  $t(n)$ -fold sequential composition of this system to be the interactive system consisting of  $t(n)$  copies of the proof executed in sequence. The  $i^{\text{th}}$  copy of the protocol is initialized after the  $(i - 1)^{\text{th}}$  copy has concluded. All copies of the protocol are initialized with the same inputs.*

We can extend our notion of zero knowledge to this setting in the natural way.

**Definition 2.10.** *An interactive proof  $(P, V)$  for the language  $L$  is sequential zero knowledge if for all polynomials  $t(n)$ , the  $t(n)$ -fold sequential composition of  $(P, V)$  is a zero knowledge proof for  $L$ .*

Note that although the verifiers in the different proof copies may be distinct entities and may in fact be honest, this definition implicitly assumes the worst case in which a single adversary controls all verifier copies. That is, it considers a sequential adversary (verifier) to be an interactive Turing machine  $V^*$  that is allowed to interact with  $t(n)$  independent copies of  $P$  (all on common input  $x$ ) in sequence.

Our definition of parallel composition is analogous to the above definition:

**Definition 2.11.** *Given an interactive proof system  $(P, V)$  and a polynomial  $t(n)$ , we consider the  $t(n)$ -fold parallel composition of this system to be the interactive system consisting of  $t(n)$  copies of the proof executed in parallel. Each message in the  $i^{\text{th}}$  round of a copy of the protocol must be sent before any message from the  $(i + 1)^{\text{th}}$  round. All copies of the protocol are initialized with the same inputs.*

We can again extend our notion of zero knowledge to this setting:

**Definition 2.12.** *An interactive proof  $(P, V)$  for the language  $L$  is parallel zero knowledge if for all polynomials  $t(n)$  the  $t(n)$ -fold parallel composition of  $(P, V)$  is a zero-knowledge proof for  $L$ .*

Thus a parallel adversary (verifier) is an interactive Turing machine  $V^*$  that is allowed to interact with  $t(n)$  independent copies of  $P$  (all on common input  $x$ ) in parallel. That is the  $i^{\text{th}}$  message in each copy is sent before the  $(i + 1)^{\text{th}}$  message of any copy of the protocol.

## 3 Sequential Zero Knowledge

### 3.1 Previous Results

In the area of sequential zero knowledge, there are two major results. The first is a negative result concerning the composition of plain zero-knowledge proofs.

**Theorem 3.1 (Goldreich and Krawczyk [16]).** *There exists a plain zero-knowledge proof (with respect to nonuniform distinguishers) whose 2-fold sequential composition is not plain zero-knowledge.*

The second significant result to emerge from the area concerns the composition of auxiliary-input zero-knowledge proofs. In this case it is possible to show that the zero-knowledge property is retained under sequential composition.

**Theorem 3.2 (Goldreich and Oren [19]).** *If  $(P, V)$  is auxiliary-input zero knowledge, then  $(P, V)$  is auxiliary-input sequential zero knowledge.*

These two results provide a context for our new results on sequential composition.

### 3.2 New Results

While Theorem 3.1 demonstrates that the original definition of zero knowledge is not closed under sequential composition, it relies on the fact that the prover can be computationally unbounded. In this section, we address the question: what happens when you compose *efficient-prover* plain zero-knowledge proofs? We obtain two results that partially characterize this behavior.

First we show that the Goldreich and Krawczyk result (Theorem 3.1) cannot be extended to efficient-prover plain zero-knowledge proofs. Indeed, we show that such proofs are closed under a *constant* number of compositions.

**Theorem 3.3.** *If  $(P, V)$  is an efficient-prover plain zero-knowledge proof system with respect to nonuniform (resp.,  $P$ -uniform) distinguishers then for every constant  $k$ , the  $k$ -fold sequential composition of  $(P, V)$  is also plain zero knowledge w.r.t. nonuniform (resp.,  $P$ -uniform) distinguishers.*

We leave the case of a super-constant number of compositions as an intriguing open problem.

Next we consider the case of  $V$ -uniform distinguishers, and we show that such protocols are *not* closed under 2-fold sequential composition with efficient provers.

**Theorem 3.4.** *If  $\mathcal{NP} \not\subseteq \mathcal{BPP}$  then there exists an efficient-prover plain zero-knowledge proof with respect to  $V$ -uniform distinguishers whose 2-fold composition is not plain zero knowledge with respect to  $V$ -uniform distinguishers.*

Informally, Theorems 3.3 and 3.4 say that plain zero knowledge is closed under a constant number of sequential compositions if and only if the distinguishers are at least as powerful as  $P$ .

**Proof of Theorem 3.3.** We now prove that efficient-prover plain zero-knowledge is closed under  $O(1)$ -fold sequential composition.

*Proof.* Let  $(P_k, V_k)$  denote the sequential composition of  $k$  copies of  $(P, V)$ . We prove by induction on  $k$  that  $(P_k, V_k)$  is plain zero knowledge with respect to nonuniform (resp.,  $P$ -uniform) distinguishers.

$(P_1, V_1)$  is zero knowledge by assumption.

Assume for induction that  $(P_{k-1}, V_{k-1})$  is zero knowledge, and consider the interactive protocol  $(P_k, V_k)$ . Let  $V_k^*$  be some sequential verifier strategy for interacting with  $P_k$ , and let  $V_{k-1}^*$  denote the sequential verifier that emulates  $V_k^*$ 's interactions with the first  $k-1$  copies of the the proof system  $(P, V)$  and then halts. Since  $(P_{k-1}, V_{k-1})$  is zero knowledge, there exists a simulator  $M_{k-1}$  that successfully simulates  $V_{k-1}^*$ .

Define  $H_k^*$  to be the “hybrid” verifier strategy (for interaction with  $P$ ) that consists of running the simulator  $M_{k-1}$  to obtain a simulated view  $v$  of the first  $k-1$  interactions, and then emulates  $V_k^*$  (starting from the simulated view  $v$ ) in the  $k$ th interaction. Since  $(P, V)$  is plain zero knowledge, there exists a polynomial-time simulator  $M_k$  for this verifier strategy.

We now show that  $M_k$  is also a valid simulator for  $(P_k, V_k^*)$ . Since by induction  $(P_{k-1}, V_{k-1})$  is plain zero knowledge versus nonuniform (resp.,  $P$ -uniform) distinguishers, the ensembles  $\Pi_1(x, y) = (x, y, \langle P_{k-1}(x, y), V_{k-1}^*(x) \rangle)$  and  $\Pi_2(x, y) = (x, y, M_{k-1}(x))$  are nonuniformly (resp., uniformly) indistinguishable when  $(x, y) \in R$ . Consider the function  $f(x, y, v) = (x, y, v')$  that emulates  $V_k^*$  starting from view  $v$  in one more interaction with  $P(y)$  to obtain view  $v'$ . Since  $f$  is polynomial-time computable, we have that  $f(\Pi_1(x, y))$  and  $f(\Pi_2(x, y))$  are also nonuniformly (resp., uniformly) indistinguishable. Observe that  $f(\Pi_1(x, y)) = (x, y, \langle P_k(x, y), V_k^*(x) \rangle)$  and  $f(\Pi_2(x, y)) = (x, y, M_k(x))$  therefore  $M_k$  is a valid simulator for  $(P_k, V_k^*)$  and hence  $(P_k, V_k)$  is plain zero knowledge with respect to nonuniform (resp.,  $P$ -uniform) distinguishers.  $\square$

In this proof, we implicitly rely on the fact that the number of copies  $k$  is a constant. It is possible that the running time of the simulation is  $\Theta(n^{g(k)})$  for some growing function  $g$ , and hence super-polynomial for nonconstant  $k$ .

Note that this result doesn't conflict with either Theorem 3.1 (in which the prover was allowed to use exponential time and was therefore able to distinguish between a simulated interaction and a real interaction) or Theorem 3.4 (in which the prover is polynomial time but the distributions are only indistinguishable to a  $V$ -uniform distinguisher, so the prover was still able to distinguish between a simulated interaction and a real interaction). Instead, it demonstrates that when neither party has more computational resources than the distinguisher, it is possible to prove a sequential closure result for plain zero knowledge, albeit restricted to a constant number of compositions.

**Proof of Theorem 3.4.** We now prove Theorem 3.4, showing that plain zero knowledge with respect to  $V$ -uniform distinguishers is *not* closed under

sequential composition. Our proof of Theorem 3.4 is a variant of the Goldreich-Krawczyk [16] proof of Theorem 3.1, so we begin by reviewing their construction.

*Overview of the Goldreich-Krawczyk Construction [16].* In the proof of Theorem 3.1, the key to constructing a zero-knowledge protocol that breaks under sequential composition lies in taking advantage of the difference in computational power between the unbounded prover and the polynomial-time verifier. The proof requires the notion of an *evasive pseudorandom ensemble*. This is simply a collection of sets  $S_i \subseteq \{0,1\}^{p(i)}$  such that each set is pseudorandom and no polynomial-time algorithm can generate an element of  $S_i$  with non-negligible probability. The existence of such ensembles was proven by Goldreich and Krawczyk in [17]. Using this, Goldreich and Krawczyk [16] construct a protocol such that in the first sequential copy, the verifier learns some element  $s \in S_{|x|}$ . In the second iteration, the verifier uses this  $s$  (whose membership in  $S_{|x|}$  can be confirmed by the prover) to extract information from  $P$ . A polynomial-time prover would be unable to generate or verify  $s \in S_{|x|}$ , therefore the result inherently relies on the super-polynomial time allotted to the prover.

*Overview of our Construction.* As in the Goldreich-Krawczyk construction, we take advantage of the difference in computational power between the two parties. However, since both are required to be polynomial-time machines, the only advantage that the prover has over the verifier is in the amount of nonuniform input each machine receives. The prover is allowed  $\text{poly}(|x|)$  bits of auxiliary input  $y$  whereas the verifier receives only the  $|x|$  bits from the common input  $x$ . In order to take advantage of this difference, we define *efficient bounded-nonuniform evasive pseudorandom ensembles*. Using the newly defined ensembles, we construct an analogous protocol; in the first iteration, the verifier learns some element of an efficient bounded-nonuniform evasive pseudorandom ensemble, and in the second it uses this information to extract otherwise unobtainable information from  $P$ .

**Definition 3.5.** *Let  $q$  be a polynomial and let  $S = \{S_1, S_2, \dots\}$  be a sequence of (non-empty) sets such that each  $S_n \subseteq \{0,1\}^n$ . We say that  $S$  is a efficient  $q(n)$ -nonuniform evasive pseudorandom ensemble if the following three properties hold:*

- (1) *For all probabilistic polynomial-time machines  $A$  with at most  $q(n)$  bits of nonuniformity,  $S_n$  is indistinguishable from the uniform distribution on strings of length  $n$ . That is, there exists a negligible function  $\epsilon$  such that for all sufficiently large  $n$ ,*

$$\left| \Pr_{x \in S_n} [A(x) = 1] - \Pr_{x \in U_n} [A(x) = 1] \right| \leq \epsilon(n).$$

- (2) *For all probabilistic polynomial-time machines  $B$  with at most  $q(n)$  bits of nonuniformity, it is infeasible for  $B$  to generate any element of  $S_n$  except*

with negligible probability. That is, there exists a negligible function  $\epsilon$  such that for all sufficiently large  $n$ ,

$$\Pr_{r \in \{0,1\}^{q(n)}} [B(x, r) \in S_n] \leq \epsilon(n).$$

- (3) There exists a polynomial  $p(n)$  and a sequence of strings  $\{\pi_n\}_{n \in \mathbb{N}}$  of length  $|\pi_n| = p(n)$  such that:
- (a) There exists a probabilistic polynomial-time machine  $D$  such that for all  $n \in \mathbb{N}$  and  $x \in \{0, 1\}^n$ ,  $D(\pi_n, x) = 1$  if  $x \in S_n$  and  $D(\pi_n, x) = 0$  else.
  - (b) There exists an expected probabilistic polynomial-time machine  $E$  such that for all  $n$   $E(\pi_n)$  is a uniformly random element of  $S_n$ .  
That is there exist efficient algorithms with polynomial-length advice for checking membership in the ensemble and for choosing an element uniformly at random.

This definition is similar in spirit to the notion of an evasive pseudorandom ensemble used by Goldreich and Krawczyk in the proof of Theorem 3.1. However, we add the additional requirement that a polynomial-time machine with an appropriate advice string  $\pi_n$  can identify and generate elements of the ensemble. In order for this to be possible, we relax the pseudorandomness and evasiveness requirements to only hold with respect to distinguishers with bounded nonuniformity rather than with respect to nonuniform distinguishers.

The introduction of this definition begs the question of whether or not such ensembles exist. Fortunately it turns out that they do.

**Theorem 3.6.** *There exists an efficient  $n/4$ -nonuniform evasive pseudorandom ensemble.*

The proof of this theorem appears in the full version [6]. It shows that if we select a hash function  $h_n : \{0, 1\}^n \rightarrow \{0, 1\}^{5n/16}$  from an appropriate pairwise independent family then with high probability  $S_n = h_n^{-1}(0^{5n/16})$  is an  $n/4$ -nonuniform evasive pseudorandom set. The pseudorandomness and evasiveness conditions (items (1) and (2)) are obtained by using pairwise independence and taking a union bound over all algorithms with  $n/4$  bits of nonuniformity. The efficiency condition (item (3)) is obtained by taking  $h_n$  to be from a standard family (e.g.,  $h_n(x) =$  the first  $5n/16$  bits of  $a \cdot x + b$ ) and taking  $\pi_n$  to be the descriptor of  $h_n$  (e.g.,  $(a, b)$ ).

We use this result to demonstrate that efficient-prover plain zero-knowledge proofs with respect to  $V$ -uniform distinguishers are not closed under sequential composition. The construction is analogous to the one by Goldreich and Krawczyk, and can be found in the full version of the paper [6].

## 4 Parallel Zero Knowledge

### 4.1 Previous Results

There are two classic results that provide context for our new result concerning the parallel composition of efficient-prover zero-knowledge proof systems. In both

cases, the result applies to auxiliary-input (as well as plain) zero knowledge, and both results are negative.

The first result establishes the existence of non-parallelizable zero-knowledge proofs independent of any complexity assumptions.

**Theorem 4.1 (Goldreich and Krawczyk [16]).** *There exists an auxiliary-input zero knowledge proof whose 2-fold parallel composition is not auxiliary-input zero knowledge (or even plain zero knowledge with respect to nonuniform distinguishers).*

While this result demonstrates that zero knowledge is not closed under parallel composition in general, the proof (like that of Theorem 3.1) inherently relies on the unbounded computational power of the provers. Without the additional computational resources necessary to generate a string and test membership in an evasive pseudorandom ensemble, the prover would be unable to execute the defined protocol.

The second such result constructs an *efficient-prover* non-parallelizable zero-knowledge proof based on a zero-knowledge proof of knowledge of the discrete-logarithm relation.

**Theorem 4.2 (Feige and Shamir [9]).** *If the discrete logarithm assumption holds then there exists an efficient-prover auxiliary-input zero-knowledge proof whose 2-fold parallel composition is not auxiliary-input zero knowledge (or even plain zero knowledge with respect to  $V$ -uniform distinguishers).*

This proof relies on the very specific assumption that the discrete logarithm problem is intractable. However as Feige and Shamir observed [9], the only properties of this problem which are actually necessary are the fact that discrete logarithms are unique and that they have a zero-knowledge proof of knowledge. It is therefore natural to consider generalizing the result to proofs of language membership for any language  $L \in \mathcal{NP}$  with exactly one witness for each element  $x \in L$ . The class of such languages is known as  $\mathcal{UP}$ . Moreover, if one-way functions exist, then every problem in  $\mathcal{NP}$  (and hence in  $\mathcal{UP}$ ) has a zero-knowledge proof of knowledge [18]. Thus:

**Theorem 4.3 (Feige and Shamir [9]).** *If  $\mathcal{UP} \not\subseteq \mathcal{BPP}$  and one-way functions exist then there exists an efficient-prover auxiliary-input zero-knowledge proof whose 2-fold parallel composition is not auxiliary-input zero knowledge (or even plain zero knowledge with respect to  $V$ -uniform distinguishers).*

## 4.2 New Results

In this work, we broaden the complexity assumptions under which we have *efficient-prover* non-parallelizable zero-knowledge proofs under more general complexity assumptions. Specifically, we show that such protocols can be constructed from any key agreement protocol (satisfying an additional technical condition). Following the standard notion of key agreement, we introduce the following definition.

**Definition 4.4.** A key agreement protocol is an efficient protocol between two parties  $P_1, P_2$  with the following four properties:

- *Input:* Both parties have common input  $1^\ell$  which is a security parameter written in unary.
- *Output:* The outputs of both parties are  $k$ -bit strings (for some  $k = \text{poly}(\ell)$ ).
- *Correctness:* The parties have the same output with probability 1 (when they follow the protocol). This common output is called the key.
- *Secrecy:* No probabilistic polynomial time Turing machine  $E$  given  $1^\ell$  and the transcript of the protocol (messages between  $P_1, P_2$ ) can distinguish with non-negligible advantage the key from a uniformly distributed  $k$ -bit string. That is,  $\{(1^\ell, \text{transcript}(P_1, P_2), \text{output}(P_1, P_2))\}_{1^\ell: \ell \in \mathbb{N}}$  is nonuniformly indistinguishable from  $\{(1^\ell, \text{transcript}(P_1, P_2), U_k)\}_{1^\ell: \ell \in \mathbb{N}}$ .

For technical reasons, we impose an additional technical condition.

**Definition 4.5.** Let  $(P_1, P_2)$  be a key agreement protocol. We say that a pair  $(i, r) \in \{1, 2\} \times \{0, 1\}^*$  is consistent with a transcript  $t$  of messages if the messages from  $P_i$  in  $t$  are what  $P_i$  would have sent had its coin tosses been  $r$  and had it received the prior messages specified by  $t$ . We say that  $t$  is valid if there exist  $r_1, r_2$  such that  $t$  is consistent with both  $(1, r_1)$  and  $(2, r_2)$ ; that is,  $t$  occurs with nonzero probability when the honest parties  $P_1$  and  $P_2$  interact. We say that  $(P_1, P_2)$  has verifiable transcripts if there is a polynomial-time algorithm that can decide whether a transcript  $t$  is valid given  $t$  and any pair  $(i, r)$  consistent with  $t$ .

We note that many existing key agreement protocols have verifiable transcripts, including the Diffie-Hellman key exchange and the protocols constructed from any public-key encryption scheme with verifiable public keys.

Our main result on non-parallelizable zero knowledge proofs follows:

**Theorem 4.6.** *If key agreement protocols with verifiable transcripts exist then there exists an efficient-prover auxiliary-input zero-knowledge proof whose 2-fold parallel composition is not auxiliary-input zero knowledge (or even plain zero knowledge with respect to  $V$ -uniform distinguishers).*

The existence of secure key agreement protocols with verifiable transcripts seems incomparable to the assumption that  $\mathcal{UP} \not\subseteq \mathcal{BPP}$  which was used in Theorem 4.3.

#### **Proof of Theorem 4.6.**

*Proof.* By assumption, key agreement protocols with verifiable transcripts exist. We consider an occurrence of a key agreement protocol to consist of the coin tosses of the two parties ( $r_1, r_2$  respectively) together with the transcript  $t$  of messages exchanged between the parties during the protocol.

Define a language  $L = \{t : \exists(i, r_i) \text{ consistent with } t\}$ .  $L = L_R$  for the relation  $R = \{(t, (i, r_i)) : (i, r_i) \text{ is consistent with } t\}$ ; we do not claim or require that  $L \notin$

$\mathcal{BPP}$ . Observe that  $L \in \mathcal{NP}$ , so there exists an efficient-prover zero-knowledge proof of knowledge (ZKPOK) of a pair  $(i, r_i)$  that is consistent with  $t$  with error  $s(n) \leq 2^{-m}$  where  $m$  is the maximum length of a witness  $(i, r_i)$ [18]. If necessary, the required error can be achieved by sequential composition of any initial ZKPOK.

We can use this proof as a subprotocol for constructing the following interactive proof for the language  $L$ .  $V$  begins by sending the message  $c = 0$  to  $P$ . If  $c = 0$ , then  $P$  uses the ZKPOK to demonstrate that he knows  $(i, r_i)$  consistent with the transcript  $t$ . If  $c \neq 0$ ,  $V$  demonstrates knowledge of  $(j, r_j)$  using the same ZKPOK. If the proof is successful and the transcript is valid (which can be checked by  $P$  by our assumption of verifiable transcripts), then  $P$  shows in zero knowledge that he too knows a witness  $(i, r_i)$  and then sends the common key  $k$  to  $V$ .

The protocol is summarized below.

Step	$P(t, (i, r_i))$	$V(t)$
1		$c = 0$
		$\leftarrow$ $c$
2	if $c = 0$ : ZKPOK of $(i, r_i)$ consistent with $t$ $\rightarrow$	
		$\leftarrow$ if $c \neq 0$ : ZKPOK of $(j, r_j)$ consistent with $t$
3	if $c \neq 0$ : ZKPOK of $(i, r_i)$ consistent with $t$ $\rightarrow$	
4	if $c \neq 0$ , $V$ 's ZKPOK is successful, and $t$ is valid: send $k$ $\rightarrow$	

**Fig. 1.** A efficient-prover non-parallelizable zero-knowledge proof for  $L$ .

The described protocol is a zero-knowledge proof for the language  $L$ .

**Efficient-Prover Interactive Proof:** The fact that this protocol is an interactive proof follows directly from the fact that the subprotocol is (by assumption) a proof of knowledge. Completeness and soundness follow from completeness and extraction properties of the ZKPOK that  $P$  conducts in Step 2 or Step 3 respectively. Prover and verifier efficiency likewise follow from the respective properties of the ZKPOK subprotocol.

**Zero Knowledge:** Given any verifier strategy  $V^*$  we can construct a simulator  $M_{V^*}$ .  $M_{V^*}$  begins by randomly choosing and fixing the coin tosses of the verifier  $V^*$ , and then runs the verifier  $V^*$  in order to obtain its first message  $c$ . If  $c = 0$ ,  $M_{V^*}$  then emulates the simulator for the ZKPOK to simulate Step 2. It then does nothing for Step 3. If  $c \neq 0$ , then  $M_{V^*}$  simulates the ZKPOK in Step 2 by following the correct “verifier” protocol and running  $V^*$  in order to simulate the “prover” half of the protocol.  $M_{V^*}$  then simulates Step 3 using the simulator

for the subprotocol. The expected time of all of these steps is polynomial; this follows directly from the running time of the simulators provided by the various subprotocols.

Finally, the simulator proceeds to Step 4. If  $c = 0$  then there is no message sent in Step 4. If  $c \neq 0$  and the ZKPOK in Step 2 was unsuccessful, then there is again no message sent in Step 4. If  $c \neq 0$  and the proof in Step 2 was successful, then  $M_{V^*}$  runs the following two extraction techniques in parallel, halting when one succeeds: First, it attempts to extract some  $(j, r_j)$  consistent with  $t$  by employing the extractor  $K$  using  $V^*$ 's strategy from Step 2 as an "oracle." Second it attempts to learn some witness  $(j, r_j)$  by trying each of the  $2^m$  possible witnesses in sequence. If  $M_{V^*}$  has successfully found a witness, it uses  $(j, r_j)$  together with the transcript  $t$  to determine whether  $t$  is valid and then to determine the common key  $k$  by emulating the actions of one party and responding to the "messages" from the other party as described in the transcript  $t$ . This key  $k$  is then used to simulate Step 4.

The indistinguishability and expected polynomial running time of the simulation follow from those of the ZKPOK simulator, except for the simulation of Step 4 in the case  $c \neq 0$ . To analyze this, let  $p$  be the probability that  $V^*$  succeeds in the ZKPOK in Step 2. If  $p > 2 \cdot 2^{-m}$ , then there exists such an extractor  $K$  that extracts a witness  $(j, r_j)$  in expected time  $q(|x|)/(p - s(|x|))$ . Since this occurs with probability  $p$ , the expected time for this case is bounded by  $(p \cdot q(|x|))/(p - s(|x|)) \leq (p \cdot q(|x|))/(p - 2^{-m}) \leq (p \cdot q(|x|))/(p/2) \leq 2q(|x|) = \text{poly}(|x|)$ . If  $p \leq 2 \cdot 2^{-m}$  then the brute force technique will find a witness in expected time  $p \cdot 2^m \leq 2 = \text{poly}(|x|)$ . Checking  $t$ 's validity takes polynomial time by assumption, and determining  $k$  takes time  $\Theta(|x|)$ , therefore the entire simulation runs in expected polynomial time.

The indistinguishability of the final step of this simulation relies on the fact that the transcript  $t$  is valid. Therefore, by the correctness of the key agreement protocol, the same key will be computed using the extracted witness  $(j, r_j)$  as with the prover's witness  $(i, r_i)$  even if they are not the same, so the simulation is polynomially indistinguishable from  $V^*$ 's view of the interactive protocol.

**Parallel Execution:** Consider now two executions,  $(\tilde{P}_1, \tilde{V})$  and  $(\tilde{P}_2, \tilde{V})$  in parallel. A cheating verifier  $V^*$  can always extract some witness  $w \in \{(1, r_1), (2, r_2)\}$  from  $\tilde{P}_1$  and  $\tilde{P}_2$  using the following strategy: in Step 1,  $V^*$  sends  $c = 0$  to  $\tilde{P}_1$  and  $c = 1$  to  $\tilde{P}_2$ . Now  $V^*$  has to execute the protocol  $(P, V)$  twice: once as a verifier talking to the prover  $\tilde{P}_1$ , and once as a prover talking to the verifier  $\tilde{P}_2$ . This he does by serving as an intermediary between  $\tilde{P}_1$  and  $\tilde{P}_2$ , sending  $\tilde{P}_1$ 's messages to  $\tilde{P}_2$ , and  $\tilde{P}_2$ 's messages to  $\tilde{P}_1$ . Now  $\tilde{P}_2$  willfully sends  $k$  to  $\tilde{V}$  (which, by the secrecy property of the key agreement protocol,  $\tilde{V}$  is incapable of computing on his own).  $\square$

## 5 Conclusions and Open Problems

We view our results as pointing out the significance of prover efficiency, as well as the power of the distinguishers, in the composability of zero-knowledge proofs.

Indeed, we have shown that with prover efficiency, the original GMR definition enjoys a greater level of composability than without. Nevertheless, the now-standard notion of auxiliary input zero knowledge still seems to be the appropriate one for most purposes. In particular, we still do not know whether plain zero knowledge is closed under a super-constant number of compositions. We also have not considered the case that different statements are being proven in each of the copies, much less (sequential) composition with arbitrary protocols. For these, it seems likely that auxiliary input zero knowledge, or something similar, is necessary.

One way in which our negative result on sequential composition (of plain zero knowledge with respect to  $V$ -uniform distinguishers, Theorem 3.4) can be improved is to provide an example where the prover’s auxiliary inputs are defined by a relation that can be decided in polynomial time (in contrast to our construction, where the prover’s auxiliary input contains the advice string  $\pi_{4n}$ , which may be hard to recognize).

For the parallel composition of auxiliary-input zero knowledge with efficient provers, it remains open to determine whether a negative result can be proven under a more general assumption such as the existence of one-way functions. The methods of Feige and Shamir [9] (Theorem 4.3) can be generalized to replace the assumption  $\mathcal{UP} \not\subseteq \mathcal{BPP}$  with the assumption that there is a problem in  $\mathcal{NP}$  for which the witnesses have a “uniquely determined feature” [22] that is hard to compute. That is, there is a poly-balanced, poly-time relation  $R$ , an efficiently computable  $f$ , and a function  $g$  such that (a) if  $(x, w) \in R$ , then  $f(x, w) = g(x)$ , and (b) there is no probabilistic polynomial-time algorithm that computes  $g(x)$  correctly for all  $x \in L_R$ . (The assumption that  $\mathcal{UP} \not\subseteq \mathcal{BPP}$  corresponds to the case that  $f(x, w) = w$ . In general, we allow the witnesses for  $x$  to have a “unique part,” namely  $g(x)$ , which is still hard to compute.) Our result (Theorem 4.6) can be viewed as constructing such an  $R$ ,  $f$ , and  $g$  from a key agreement protocol.

Our construction complements that of Haitner, Rosen, and Shaltiel [22] — they consider the parallel repetition of natural zero-knowledge proofs (such as 3-Coloring [18] or Hamiltonicity [7]), and argue that “certain black-box techniques” cannot prove that a feature  $g(x)$  will remain hard to compute by the verifier (on average). In contrast, we consider the parallel repetition of a contrived zero-knowledge proof and show that a cheating verifier can always learn a certain hard-to-compute feature  $g(x)$ .

## Acknowledgments

We thank the TCC 2010 reviewers for helpful comments.

## References

1. B. Barak. How to go beyond the Black-Box Simulation Barrier. In *42nd IEEE Symposium on Foundations of Computer Science*, pages 106-115, 2001.

2. B. Barak, Y. Lindell, S. Vadhan. Lower Bounds for Non-Black-Box Zero Knowledge. *Proc. of the 44th IEEE Symposium on the Foundation of Computer Science*, pages 384-393, 2003.
3. M. Bellare and O. Goldreich. On defining proofs of knowledge. In *Advances in Cryptology – CRYPTO '92*, volume 740 of *Lecture Notes in Computer Science*, pages 390-420. Springer-Verlag, 1993, 16-20 August 1992.
4. M. Ben-Or, O. Goldreich, S. Goldwasser, J. Hastad, J. Kilian, S. Micali, and P. Rogaway. Everything provable is provable in zero-knowledge. In *Proceedings of Advances in Cryptology Crypto88. Lecture Notes in Computer Science, Vol. 403*. Springer-Verlag, New York, 1990, pp. 37-56.
5. E. Birrell. Composition of Zero-Knowledge Proofs. Undergraduate Thesis. Harvard University, 2009.
6. E. Birrell and S. Vadhan. Composition of Zero Knowledge Proofs with Efficient Provers. Cryptology eprint archive, 2009.
7. M. Blum. How to prove a theorem so no one else can claim it. In *Proceedings of the International Congress of Mathematicians*, pages 1444-1451, 1987.
8. W. Diffie and M. Hellman. New Directions in Cryptography. *IEEE Trans. on Info. Theory*, IT-22 Nov. 1976, pages 644-654.
9. U. Feige and A. Shamir. Witness Indistinguishability and Witness Hiding Protocols. In *22nd ACM Symposium on the Theory of Computing*, pages 416-426, 1990.
10. U. Feige and A. Shamir. Zero-Knowledge Proofs of Knowledge in Two Rounds. In *Crypto '89*, Springer-Verlag LNCS Vol. 435, pages 526-544, 1990.
11. O. Goldreich. *Foundations of Cryptography - Basic Tools*. Cambridge University Press, 2001.
12. O. Goldreich. A Uniform Complexity Treatment of Encryption and Zero Knowledge. *Journal of Cryptology*, Vol. 6, No. 1, pages 21-53, 1993.
13. O. Goldreich. Zero-Knowledge twenty years after its invention. Cryptology ePrint Archive, Report 2002/186, 2002. <http://eprint.iacr.org/>
14. O. Goldreich, S. Goldwasser, and S. Micali. How to Construct Random Functions. *Journal of the Association for Computing Machinery*, Vol. 33, No. 4, pages 792-807, 1986.
15. O. Goldreich and A. Kahan. How to Construct Constant-Round Zero-Knowledge Proof Systems for NP. *Journal of Cryptology*. Vol. 9, No. 2, pages 167-189, 1996
16. O. Goldreich and H. Krawczyk. On the Composition of Zero-Knowledge Proof Systems. *SIAM Journal on Computing*, Vol. 25, No. 1, February 1996, pages 169-192. Preliminary version in *ICALP '90*.
17. O. Goldreich and H. Krawczyk. Sparse Pseudorandom Distributions. *Random Structures & Algorithms*, Vol. 3, No. 2, pages 163-174, 1992.
18. O. Goldreich, S. Micali, and A. Wigderson. Proofs that Yield Nothing but their Validity or All Languages in NP have Zero-Knowledge Proof Systems. *Journal of the ACM*, Vol. 38, No. 1, pages 691-729, 1991.
19. O. Goldreich and Y. Oren. Definitions and Properties of Zero-Knowledge Proof Systems. *Journal of Cryptology*, Vol. 7. No. 1, pages 1-32, 1994.
20. S. Goldwasser, S. Micali, and C. Rackoff. Knowledge Complexity of Interactive Proofs. *Proc. 17th STOC*, 1985. pages 291-304.
21. S. Goldwasser, S. Micali, and C. Rackoff. The Knowledge Complexity of Interactive Proof Systems. *SIAM Journal on Computing*, Vol. 18, pages 186-208, 1989.
22. I. Haitner, A. Rosen, and R. Shaltiel. On the (Im)possibility of Arthur-Merlin Witness Hiding Protocols. In *Theory of Cryptography Conference*, 2009.
23. S. Vadhan. *Pseudorandomness*. To appear in *Foundations and Trends in Theoretical Computer Science*, 2010.