

On Concurrent Zero-Knowledge with Pre-Processing*

(Extended abstract)

GIOVANNI DI CRESCENZO¹

RAFAIL OSTROVSKY²

¹ University of California San Diego,
and Telcordia Technologies, Inc.

E-mail: giovanni@cs.ucsd.edu, giovanni@research.telcordia.com

² Telcordia Technologies, Inc.

445 South Street, Morristown, NJ, 07960-6438, USA.

E-mail: rafail@research.telcordia.com

*Copyright © 1999, Telcordia Technologies, Inc. All Rights Reserved.

Abstract. Concurrent Zero-Knowledge protocols remain zero-knowledge even when many sessions of them are executed together. These protocols have applications in a distributed setting, where many executions of the same protocol must take place at the same time by many parties, such as the Internet. In this paper, we are concerned with the number of rounds of interaction needed for such protocols and their efficiency. Here, we show an efficient constant-round concurrent zero-knowledge protocol with preprocessing for all languages in NP, where both the preprocessing phase and the proof phase each require 3 rounds of interaction. We make no timing assumptions or assumptions on the knowledge of the number of parties in the system. Moreover, we allow arbitrary interleavings in both the preprocessing and in the proof phase. Our techniques apply to both zero-knowledge proof systems and zero-knowledge arguments and we show how to extend our technique so that polynomial number of zero-knowledge proofs/arguments can be executed after the preprocessing phase is done.

1 Introduction

The notion of zero-knowledge proof systems was introduced in the seminal paper of Goldwasser, Micali and Rackoff [18]. Since their introduction, zero-knowledge proofs have proven to be very useful as a building block in the construction of cryptographic protocols, especially after Goldreich, Micali and Wigderson [17] have shown that all languages in NP admit zero-knowledge proofs. Due to their importance, the efficiency of zero-knowledge protocols has received considerable attention. One of the aspects of efficiency is the number of rounds of interaction necessary for zero-knowledge proofs, and it was shown that there exist computational zero-knowledge proofs with a constant number of rounds for all languages in NP by Goldreich and Kahan [16].

Previous work. Recently, a lot of attention has been paid to the setting where many concurrent executions of the same protocol take place (say, on the Internet). For example, in the context of identification schemes, this was discussed by Beth and Desmedt [3]. In the context of zero-knowledge, Di Crescenzo constructed proof systems that remain zero-knowledge in a synchronous setting without timing or complexity assumptions [8]. The general case of concurrent

zero-knowledge in an asynchronous setting was considered by Dwork, Naor and Sahai [11]. They showed that assuming that there are some bounds on the relative speed of the processors, one can construct four-round zero-knowledge argument for any language in NP (an argument differs from a proof system in the fact that the soundness property holds only with respect to polynomial-time provers [4]). Dwork and Sahai improved this result by reducing the use of timing assumptions to a preprocessing phase [12]. (Preprocessing was first used in the context of zero-knowledge proofs by De Santis, Micali and Persiano [7].) Recently, Richardson and Kilian [24] presented concurrent zero-knowledge protocols for all languages in NP that do not use timing assumption but require more than constant number of rounds of interaction. More specifically, given a security parameter k , Richardson and Kilian’s protocol requires $O(k)$ rounds and allows $\text{poly}(k)$ concurrent executions. A negative result was given by Kilian, Petrank and Rackoff [20], who showed that if there exists a 4-round protocol that is concurrent (black-box simulation) zero-knowledge for some language L then L is in BPP.

Our results. We consider a distributed model with a preprocessing phase and a proof phase. In this setting, we show how, based on complexity assumptions, and after a three-round pre-processing stage, three-round concurrent zero-knowledge proofs (and arguments) can be constructed for all languages in NP. Our protocol does not require any timing assumptions nor knowledge of the number of parties in the system nor knowledge of the total number of concurrent executions. In the case of proof systems, our protocol is based on the existence of perfectly secure commitment schemes, it is computational zero-knowledge and applies to all public-coin zero-knowledge proof systems. In the case of arguments, our protocol is based on the intractability of computing discrete logarithms modulo primes and it is perfect zero-knowledge. The requirement that we make (which is different than in the work of Dwork and Sahai [12]) is that all the concurrent executions of the pre-processing subprotocols finish before the concurrent executions of the proof protocols begin, although we allow arbitrary interleaving both in the pre-processing and in the proof phase. A different interpretation of our result is that we do not make any timing assumptions, but require a single “synchronization barrier” between the pre-processing stage and the main proof stage of the protocol, where all parties finish the pre-processing stage before the main stage begins. We believe that this setting may be of interest in several applications, since the pre-processing stage does not need the knowledge of the theorem(s) to be proved in the main phase of the protocol.

Tools used. An important tool used in the construction of our schemes is that of equivocable commitment schemes, which we consider in two variants: computationally and perfectly equivocable. The computational variant was first discussed by Beaver in [1], and a first construction was given by Di Crescenzo, Ishai and Ostrovsky in [9] in the common random string model (using a scheme by Naor [21]). Here we present a computationally equivocable and a perfectly equivocable commitment schemes in the interactive model, based on bit commitment schemes by Naor [21] and Pedersen [23], respectively. These schemes might be of independent interest and may have further applications, such as for identification schemes. We remark that a somewhat weaker version of our results could be alternatively derived by considering in our model appropriate modifications of techniques based on coin flipping and non-interactive zero-knowledge proofs, as those used in [24, 12, 13], resulting with even smaller round-complexity. However,

our techniques apply to perfect zero-knowledge arguments and offer additional efficiency properties since they do not use general NP reductions. Another tool we use is that of straight-line simulation, as formally defined by Dwork and Sahai [12] and also used by Feige and Shamir [14].

2 Notations and Definitions

In this section we give basic notations, we recall the notions of interactive proof systems, zero-knowledge proof systems in the two-party model and formally define concurrent zero-knowledge proof systems with preprocessing.

Probabilistic setting. The notation $x \leftarrow S$ denotes the *random process* of selecting element x uniformly from set S . Similarly, the notation $y \leftarrow A(x)$, where A is an algorithm, denotes the random process of obtaining y when running algorithm A on input x , where the probability space is given by the random coins (if any) of algorithm A . By $\{R_1; \dots; R_n : v\}$ we denote the set of values v that a random variable can assume, due to the distribution determined by the sequence of random processes R_1, \dots, R_n . By $\text{Prob}[R_1; \dots; R_n : E]$ we denote the probability of event E , after the ordered execution of random processes R_1, \dots, R_n .

Interactive protocols. Following [18], an *interactive Turing machine* is a Turing machine with a public input tape, a public communication tape, a private random tape and a private work tape. An *interactive protocol* is a pair of interactive Turing machines sharing their public input tape and communication tape. The *transcript* of an execution of an interactive protocol (A, B) is a sequence containing the random tape of B and all messages appearing on the communication tape of A and B . The notation $(y_1, y_2) \leftarrow (A(x_1), B(x_2))(x)$ denotes the random process of running interactive protocol (A, B) , where A has private input x_1 , B has private input x_2 , x is A and B 's common input, y_1 is A 's output and y_2 is B 's output, where any of x_1, x_2, y_1, y_2, x can be an empty string; the notation $y \leftarrow (A(x_1), B(x_2))(x)$ is an abbreviation for $(y, y) \leftarrow (A(x_1), B(x_2))(x)$. We assume wlog that the output of both parties A and B at the end of an execution of protocol (A, B) contains a transcript of the communication exchanged between A and B during such execution. An *interactive protocol with preprocessing* is a pair of interactive protocols $((A_1, B_1), (A_2, B_2))$. The mechanics of an interactive protocol with preprocessing is divided in two phases, as follows. In a first phase, called the *preprocessing phase*, the first pair (A_1, B_1) is executed; at the end of this phase a string α is output by A_1 and given as private input to A_2 , and a string β is output by B_1 and given as private input to B_2 . Now, an input string x is given as common input to A_2 and B_2 , and the second pair (A_2, B_2) is executed. In this paper we will be concerned with two types of interactive protocols: *proof systems* and *arguments*, which we now describe.

Interactive proof systems and arguments. An interactive proof system for a language L is an interactive protocol in which, on input a string x , a computationally unbounded prover convinces a polynomial-time bounded verifier that x belongs to L . The requirements are two: completeness and soundness. Informally, completeness states that for any input $x \in L$, the prover convinces the verifier with very high probability. Soundness states that for any $x \notin L$ and any prover, the verifier is convinced with very small probability. A formal definition can be found in [18, 15]. An argument is defined similarly to a proof

system, the only difference being that the prover is assumed to be polynomially bounded (see [4]).

Zero-knowledge proof systems in the two-party model. A zero-knowledge proof system for a language L is an interactive proof system for L in which, for any $x \in L$, and any possibly malicious probabilistic polynomial-time verifier V' , no information is revealed to V' that he could not compute alone before running the protocol. This is formalized by requiring, for each V' , the existence of an efficient simulator $S_{V'}$ which outputs a transcript ‘indistinguishable’ from the view of V' in the protocol. There exist three notions of zero-knowledge, according to the level of indistinguishability: computational, statistical and perfect. We refer the reader to [18, 15] for the definitions of computational, statistical and perfect indistinguishability between distributions and for the definitions of computational, statistical and perfect zero-knowledge proof systems.

The concurrent model for zero-knowledge proof systems. Informally speaking, the concurrent model describes a distributed model in which several parties can run concurrent executions of some protocols. In real-life distributed system, the communication is not necessarily synchronized; more generally, the model makes the worst-case assumption that the communication in the system happens in an *asynchronous* manner; this means that there is no fixed bound on the amount of time that a message takes to arrive to its destination. This implies that, for instances, the order in which messages are received during the execution of many concurrent protocols can be arbitrarily different from the order in which they were sent. Such variation in the communication model poses a crucial complication into designing a zero-knowledge protocol, since the possibly arbitrary interleaving of messages can help some adversary to break the zero-knowledge property. In fact, as a worst case assumption, one may assume that the ordering in which messages are received can be decided by the adversary. Now we proceed more formally.

We consider a distributed model, with two distinguished sets of parties: a set $\mathcal{P} = \{P_1, \dots, P_q\}$ of provers and a set $\mathcal{V} = \{V_1, \dots, V_q\}$ of verifiers, where P_i is connected to V_i , for $i = 1, \dots, q$. Let (A, B) be a zero-knowledge proof system. At any time a verifier V_i may decide to run protocol (A, B) ; therefore, for any fixed time, there may be several pairs of prover and verifier running (A, B) . The adversary \mathcal{A} is allowed to corrupt all the verifiers. Then \mathcal{A} can be formally described as a probabilistic polynomial time algorithm that, given the history so far, returns an index i and a message m_i , with the meaning that the (corrupted) verifier V_i , for $i \in \{1, \dots, q\}$, is sending message m_i to prover P_i . We assume wlog that P_i is required to send his next message after receiving m_i and before receiving a new message from \mathcal{A} . We now define the view of the adversary \mathcal{A} . First we define a *q-concurrent execution of (A, B)* as the possibly concurrent execution of q instances of protocol (A, B) , where all verifiers are controlled by \mathcal{A} . Also, we define the *q-concurrent transcript* of a q -concurrent execution of (A, B) as a sequence containing the random tapes of verifiers V_1, \dots, V_q and all messages appearing on the communication tapes of P_1, \dots, P_q and V_1, \dots, V_q , where the ordering of such messages is determined by the adversary corrupting all the verifiers. The notation $(T; y_{11}, y_{12}, \dots, y_{1q}, y_{2q}) \leftarrow ((P_1(p_1), V_1(v_1))(x_1), \dots, (P_q(p_q), V_q(v_q))(x_q))$ denotes the random process of running a q -concurrent execution of interactive protocol (A, B) , where each P_i has private input x_i , each V_i has private input v_i , x_i is P_i and V_i 's common input, y_{1i} is P_i 's output, y_{2i} is V_i 's output, and

T is the q -concurrent transcript of such execution (we assume wlog that the output of both parties P_i and V_i at the end of an execution of the interactive protocol (A,B) contains a transcript of the communication exchanged between P_i and V_i during such execution). Then the view of \mathcal{A} , denoted as $View_{\mathcal{A}}(\mathbf{x})$, is the transcript T output by the random process $(T; y_{11}, y_{12}, \dots, y_{1q}, y_{2q}) \leftarrow ((P_1(p_1), \mathcal{A}(v_1))(x_1), \dots, (P_q(p_q), \mathcal{A}(v_q))(x_q))$, where $\mathbf{x} = (x_1, \dots, x_q)$. Finally, a proof system (A,B) for language L is *concurrent zero-knowledge* if for any probabilistic polynomial time algorithm \mathcal{A} , there exists an efficient algorithm $S_{\mathcal{A}}$ such that for any polynomial $q(\cdot)$, and any x_1, \dots, x_q , where $q = q(n)$, and $|x_1| = \dots = |x_q| = n$, the two distributions $S_{\mathcal{A}}(\mathbf{x})$ and $View_{\mathcal{A}}(\mathbf{x})$ are indistinguishable.

Concurrent zero-knowledge proof systems with preprocessing. We now consider a variant of the above distributed model, in which we would like to consider concurrent zero-knowledge proof systems with preprocessing. In this variant a concurrent execution of an interactive protocol with preprocessing $((A1,B1),(A2,B2))$ is divided into two phases: a preprocessing phase and a proof phase. In the preprocessing phase there is a concurrent execution (in the sense defined before) of the preprocessing pair $(A1,B1)$ and in the proof phase there is a concurrent execution of the proof pair $(A2,B2)$. The requirements are, as before, completeness, soundness and concurrent zero-knowledge. Now we give a formal definition of concurrent zero-knowledge proof systems with preprocessing.

Definition 1. Let $(A,B) = ((A1,B1),(A2,B2))$ be an interactive protocol with preprocessing. We say that (A,B) is a concurrent (computational) (statistical) (perfect) zero-knowledge proof system with preprocessing for language L if the following holds:

1. **Completeness:** For any $x \in L$, it holds that $\text{Prob}[(\alpha, \beta) \leftarrow (A1, B1)(1^{|x|}); (t, (t, out)) \leftarrow (A2(\alpha), B2(\beta))(x) : out = \text{ACCEPT}] \geq 1 - 2^{-|x|}$.
2. **Soundness:** For any $x \notin L$, and any $(A1', A2')$, it holds that $\text{Prob}[(\alpha, \beta) \leftarrow (A1', B1)(1^{|x|}); (t, (t, out)) \leftarrow (A2'(\alpha), B2(\beta))(x) : out = \text{ACCEPT}] < 2^{-|x|}$.
3. **Concurrent Zero-Knowledge:** For each probabilistic polynomial time algorithm $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists an expected polynomial time simulator algorithm $S_{\mathcal{A}}$ such that for any polynomial $q(\cdot)$, for each $x_1, \dots, x_q \in L$, where $|x_1| = \dots = |x_q| = n$ and $q = q(n)$, the two distributions $S_{\mathcal{A}}(\mathbf{x})$ and $View_{\mathcal{A}}(\mathbf{x})$ are (computationally) (statistically) (perfectly) indistinguishable, where

$$View_{\mathcal{A}}(\mathbf{x}) = \{(T_1, \alpha_1, \beta_1, \dots, \alpha_q, \beta_q) \leftarrow ((P_1, \mathcal{A}_1), \dots, (P_q, \mathcal{A}_1))(1^n); \\ (T_2, \cdot, \cdot, \dots, \cdot, \cdot) \leftarrow ((P_1(\alpha_1), \mathcal{A}_2(\beta_1))(x_1), \dots, (P_q(\alpha_q), \mathcal{A}_2(\beta_q))(x_q)) : \\ (T_1, T_2, \beta_1, \dots, \beta_q)\}$$

denotes the view of \mathcal{A} on input $\mathbf{x} = (x_1, \dots, x_q)$, and where P_1, \dots, P_q run algorithms $A1$ in the preprocessing phase and $A2$ in the proof phase.

3 Cryptographic tools and a class of languages

We briefly review some cryptographic tools as commitment schemes and coin flipping protocols. Then we discuss a class of languages having a certain type of zero-knowledge protocols, which will be used in this paper.

Commitment schemes. Informally speaking, a *commitment scheme* (Alice, Bob) is a two-phase interactive protocol between two probabilistic polynomial time parties Alice and Bob, called the sender and the receiver, respectively, such that the following is true. Alice commits to his bit b in the first phase (called the commitment phase); in the second phase (called the decommitment phase) Alice convinces Bob of the value of the bit b Alice had committed to in the first phase (if Bob is not convinced, he outputs a special string \perp). A commitment scheme has three requirements. First, if Alice and Bob behave honestly, then at the end of the decommitment phase Bob is convinced that Alice had committed to bit b with high probability (this is the *correctness* requirement). Then, no matter which polynomial-time computable strategy Bob uses in the commitment phase, Bob is not able to guess such bit with probability significantly better than its a priori probability at the end of such phase (this is the *security* property). Finally, for any strategy played by Alice, the probability that he can later decommit both as 0 and as 1 is negligible (this is the *binding* property). There are two main variants of commitment schemes in the literature, with respect to the type of security guaranteed: *computationally secure* (i.e., the security property holds against a polynomial time bounded receiver, while the binding property holds against an unrestricted sender) and *perfectly secure* (i.e., security holds against an unrestricted receiver, while binding holds against polynomial time bounded sender). In this paper we will consider both variants.

Coin flipping protocols. A coin flipping protocol [5] is a protocol among two parties, Alice and Bob, who exchange messages. At the end of the protocol, both parties have a bit that is uniformly distributed, no matter how each of them tries to influence its distribution. Standard constructions for coin flipping protocols are obtained by using bit commitment schemes, as follows. Alice commits to a random bit a , Bob sends a random bit b to Alice, then Alice reveals her bit a ; the result of the coin flipping protocol is $c = a \oplus b$.

Three-round public-coin honest-verifier zero-knowledge protocols. We will consider a class of languages having a special type of zero-knowledge protocol; namely, a protocol having the following properties: 1) the protocol has three rounds; namely, the prover sends a message (called *first message*) to the verifier, the verifier sends a message (called *challenge*) to the prover, the prover sends a final message (called *answer*) to the verifier who decides whether to accept or not; 2) it is public-coin; namely, the challenge consists of some random bits sent by the verifier; 3) it is honest-verifier zero-knowledge; namely, the zero-knowledge requirement holds only with respect to a verifier which follows its program. In the literature there are several examples of protocols satisfying the above properties, and the class of languages having such protocols is quite large. For instance, all languages in NP have a computational zero-knowledge proof system (e.g., [17, 5]) with these properties, and all random self-reducible languages and formula compositions over them have perfect zero-knowledge proof systems with these properties ([18, 17, 25, 6]).

4 Equivocable commitment schemes

We recall the notion of equivocable commitment schemes, by presenting two variants of them: computationally and perfectly equivocable commitment schemes. Then we present an example of a computationally equivocable commitment

scheme under the assumption of the existence of any one-way permutation and an example of a perfectly equivocal commitment scheme under the assumption of intractability of computing discrete logarithms.

Definition of equivocal commitment schemes. Informally speaking, a bit commitment scheme is equivocal if it satisfies the following additional requirement. There exists an efficient algorithm, called the simulator, which outputs a transcript leading to a ‘fake’ commitment such that: (a) the ‘fake’ commitment can be decommitted both as 0 and as 1, and (b) the simulated transcript is indistinguishable from a real execution of the protocol. The extension to equivocal string commitment schemes is straightforward. The two types of equivocability, computational and perfect, differ according to the type of indistinguishability in (b). A formal definition follows.

Definition 2. Let (A, B) be a commitment scheme. We say that (A, B) is computationally equivocal (resp., perfectly equivocal) if for any probabilistic polynomial time algorithm B' , there exists a probabilistic polynomial time algorithm M such that:

1. *Equivocability.* For all constants c , all sufficiently large n , any string $s \in \{0, 1\}^k$, where $k = k(n)$ for some polynomial $k(\cdot)$, it holds that $|p_0 - p_1| \leq n^{-c}$, where p_0, p_1 are, respectively,

$$\text{Prob}[(\alpha, \beta) \leftarrow M_{B'}(1^n, 1^k); (t, (t, v)) \leftarrow M_{B'(\beta)}(\alpha, s, 1^n) : v = s],$$

$$\text{Prob}[(\alpha, \beta) \leftarrow (A(s), B')(1^n, 1^k); (t, (t, v)) \leftarrow (A(\alpha, s), B'(\beta))(1^n, 1^k) : v = s].$$

(resp., it holds that $p_0 = p_1$.)

2. *Indistinguishability.* For any string $s \in \{0, 1\}^k$, where $k = k(n)$ for some polynomial $k(\cdot)$, the distributions $T(M)$ and $T(A)$ are computationally (resp., perfectly) indistinguishable, where

$$T(M) = \{(\alpha, \beta) \leftarrow M_{B'}(1^n, 1^k); (t, (t, s)) \leftarrow M_{B'(\beta)}(\alpha, s, 1^n) : (\beta, (t, s))\},$$

$$T(A) = \{(\alpha, \beta) \leftarrow (A(s), B')(1^n, 1^k); (t, (t, s)) \leftarrow (A(\alpha, s), B'(\beta))(1^n, 1^k) : (\beta, (t, s))\}.$$

Equivocal bit-commitment schemes have properties similar to chameleon or trapdoor commitment schemes (see, e.g., [4]); a main difference is that in equivocal commitment schemes one among the two requirements of binding and security holds with respect to a computationally unlimited adversary (as opposed to only polynomial-time bounded). Computationally equivocal bit-commitment schemes have been first discussed in [1], who observed the somewhat paradoxical requirement that such schemes need to satisfy. In [9] it was shown that the implementation in the common random string model of the bit commitment scheme from [21] is computationally equivocal. Here we show in the interactive model that a variation of the scheme in [21] is a computationally equivocal commitment scheme and a variation of the scheme in [23] is a perfectly equivocal commitment scheme.

A computationally equivocal commitment scheme. We will present a scheme (A, B) by combining the scheme in [21] with a coin flipping protocol; we describe the scheme for a single bit, since the extension to a many-bit string is straightforward. The bit commitment scheme in [21] consists in a 2-round

commitment phase and a 1-round decommitment phase. In the commitment phase, the receiver sends a $3n$ -bit random string R to the committer and the committer replies with some pseudo-random message com , where n is a security parameter. In the decommitment phase, the committer sends a decommitment message dec which allows, together with R and com , the receiver to be convinced that the committed bit was b . The variation we consider here consists in the following. First of all committer and receiver run a coin flipping protocol, whose output we denote by r . Then they continue the protocol as in [21], but using the output r of the coin flipping protocol as string R is used in the original protocol. A more formal description follows. We will denote by n a security parameter, and assume that the parties share a perfectly secure commitment scheme (C,D) and a pseudo-random generator G .

The Commitment Protocol (A,B)

Input to A: a bit b .

Commitment Phase:

1. B uniformly chooses string $u \in \{0,1\}^{3n}$;
B commits to it using scheme (C,D) (possibly interacting with A);
2. A uniformly chooses string $v \in \{0,1\}^{3n}$ and sends it to B;
3. B decommits u to A using scheme (C,D) (possibly interacting with A);
4. If u is not properly decommitted, A halts.
A uniformly chooses $s \in \{0,1\}^n$ and computes $z = G(s)$;
A sets $com = z$ if $b = 0$ or $com = z \oplus u \oplus v$ if $b = 1$, and sends com to B.

Decommitment Phase:

1. A sends s to B;
2. B outputs: 0 if $G(s) = com$, 1 if $G(s) = com \oplus u \oplus v$, \perp otherwise.

We obtain the following

Theorem 3. *Assuming the existence of pseudo-random generators and perfectly-secure commitment schemes, the protocol (A,B) is a computationally equivocal and computationally secure commitment scheme.*

Proof. The correctness and the security property of (A,B) directly follow from those of the scheme in [21]. Now we consider the binding property. Note that since B commits to u using a perfectly-secure commitment scheme, no infinitely powerful A' can guess u better than guessing at random. Therefore, for any A' , the distribution of $u \oplus v$ is uniform over $\{0,1\}^{3n}$, and we can directly apply the analysis of [21] to conclude that (A,B) satisfies the binding property. Now we show that (A,B) is computationally equivocal. Informally, we present an efficient simulator M , which can run the commitment phase in such a way that it can later decommit both as a 0 and as a 1.

The algorithm M . On input 1^n , M uniformly chooses $u \in \{0,1\}^{3n}$ and two seeds $s_0, s_1 \in \{0,1\}^n$, and computes $z_0 = G(s_0)$ and $z_1 = G(s_1)$. Then it runs the commitment phase as follows: it receives a commitment to v from B' , it sends u to B' and it receives the decommitment of v from B' ; at this point, M either halts (if decommitments were not appropriate) or rewinds B' to the state right after B' committed to v , sets $u' = z_0 \oplus z_1 \oplus v$ and sends u' to B' ; now M again receives the decommitment of v from B' ; at this point, M either halts (if decommitments were not appropriate) or continues the simulation, and the result of the coin

flipping protocol will be $u' \oplus v$; finally, in order to commit, M sends $com = z_0$ to B; in order to decommit as b , M sends s_b to B, for $b = 0, 1$.

M's output is computationally indistinguishable from a real execution of (A,B). Note that in M's output the result of the coin flipping protocol is equal to $G(s_0) \oplus G(s_1)$ while in an execution of (A,B) such output is random. In particular, the same holds for the string sent to the receiver during the coin flipping protocol. This can be used to show, exactly as already done in [9], that if there exists an efficient algorithm distinguishing M's output from a real execution of (A,B), then there exists an efficient algorithm that distinguishes uniformly distributed strings from outputs of G . This contradicts the assumption that G is a pseudorandom generator.

M can decommit both as 0 and as 1. During the execution of M there are two executions of the coin flipping protocol, one before and one after the rewinding of B'; we will call them the first and the second execution of the coin flipping protocol, respectively; we note that M can halt both during the first and during the second execution of the coin flipping protocol because of an inappropriate decommitment by B'. Instead, if the result of the second execution of the coin flipping protocol between M and B' is $u' \oplus v = z_0 \oplus z_1$, then s_0 and s_1 are valid decommitment keys as 0 and 1, respectively, for the commitment $com = z_0$. Therefore, the probability p_0 , as defined in item 1 of Definition 2 is equal to the probability that the output of the second execution of the coin flipping protocol is equal to $z_0 \oplus z_1$. Let us define probability q_1 as the probability that M does not halt in the first execution of the coin flipping protocol, and probability q_0 as the probability that M halts neither in the first nor in the second execution of the coin flipping protocol because of some inappropriate decommitment of B'. We now observe three facts.

The first fact is that $p_1 = q_1$. This follows by observing the probabilistic experiments in the definition of p_1 and q_1 are the same; specifically, notice that the string u is uniformly chosen and B' decommits properly in both experiments.

The second fact we observe is that $p_0 = q_0$. This follows by observing the probabilistic experiments in the definition of p_0 and q_0 are the same; specifically, observe that M is successful in creating a fake commitment if and only if M does not halt neither in the first execution nor in the second execution of the coin flipping protocol.

The third fact is that $q_1 - q_0$ is negligible. Assume not. Then B' could be used to efficiently distinguish between the random string u sent by M during the first execution of the coin flipping protocol and the string u' sent by M during the second execution (since B' correctly decommits in the first case and incorrectly in the second). This implies that B' can be, in turn, used to contradict the fact that G is a pseudo-random generator (exactly as in the proof that M's output is computationally indistinguishable from a real execution of (A,B)), thus giving a contradiction.

By combining the above three facts, we have that $|p_1 - p_0| \leq |q_1 - q_0|$, and is therefore negligible. \square

We remark that the perfectly-secure commitment scheme (C,D) used in the above construction can be implemented assuming the existence of any one-way permutation using the scheme in [22]. Therefore, the weakest assumption under which the described scheme (A,B) is computationally equivocal is the existence of one-way permutations.

A perfectly equivocal commitment scheme. Our perfectly equivocal

scheme is based on a perfectly-secure commitment scheme in [23]; this scheme uses some elementary number theoretic definition concerning discrete logarithms, which we briefly review.

Discrete logarithms. Let p, q be primes such that $p = 2q + 1$ and let G_q denote the only subgroup of Z_p^* of order q . We note that it can be efficiently decided whether an integer a is in Z_q , by checking that $a^q \equiv 1 \pmod{p}$. Moreover, any element of G_q different from 1 generates such subgroup. For any $a, b \in G_q$, if $b \neq 1$ the *discrete logarithm of a in base b* is the integer x such that $b^x = a \pmod{p}$.

The commitment scheme in [23]. Given primes p, q such that $p = 2q + 1$ and $g, h \in G_q$, the commitment scheme goes as follows. In order to commit to an integer $s \in Z_q$, the committer uniformly chooses $r \in Z_q$, computes $com = g^s h^r \pmod{p}$, and sends com to the receiver. In order to decommit com as m , the committer sends s, r to the receiver, who checks that $com = g^s h^r \pmod{p}$. The perfect security property of this scheme follows from the fact that com is uniformly distributed in G_q ; the computational binding property follows from the fact that if a committer is able to successfully decommit a string com both as s and s' , then he can compute the discrete logarithm of h in base g .

Our variation. We consider a protocol (A,B) that is a variation of the above scheme, consisting in running first a coin flipping protocol to choose $g, h \in G_q$. A more formal description follows. We will assume that the parties share a computationally secure commitment schemes (C,D).

The Commitment Protocol (A,B)

Input to A: an integer $s \in Z_q$.

Commitment Phase:

1. A uniformly chooses k -bit primes p, q such that $p = 2q + 1$ and sends them to B;
2. B checks that $p = 2q + 1$ and that p, q are primes and uniformly chooses $u_1, u_2 \in G_q$; B commits to u_1, u_2 using scheme (C,D) (possibly interacting with A);
3. A uniformly chooses $v_1, v_2 \in G_q$ and sends them to B;
4. B decommits u_1, u_2 to A using scheme (C,D) (possibly interacting with A);
5. If u_1, u_2 are not properly decommitted, A halts;
6. A and B set $g = u_1 v_1 \pmod{p}$ and $h = u_2 v_2 \pmod{p}$;
7. A uniformly chooses $r \in Z_q$, computes $com = g^s h^r \pmod{p}$ and sends com to B.

Decommitment Phase:

1. A sends s, r to B;
2. B outputs: s if $com = g^s h^r \pmod{p}$, \perp otherwise.

We obtain the following

Theorem 4. *Assuming the existence of computationally-secure commitment schemes and the intractability of computing discrete logarithms modulo integers of the form $p = 2q + 1$, for p, q primes, (A,B) is a perfectly equivocal and perfectly secure commitment scheme.*

Proof. The correctness and perfect security properties directly follow from the analogue properties of the scheme in [23]. Now we consider the binding property. Note that since B commits to u_1, u_2 using a computationally secure commitment scheme, for any polynomial time A' , the distribution of $u_1 v_1 \pmod{p}$ and $u_2 v_2 \pmod{p}$ has only negligible distance from the uniform distribution. Therefore we can apply the analysis of [23] to conclude that (A,B) is computationally binding assuming that computing discrete logarithms modulo integers of the form

$p = 2q + 1$, for q prime, is intractable. The proof that (A,B) is perfectly equivocal can be obtained similarly as for the previous scheme. Similarly, we can construct a simulator M who will run the protocol as A does, learn the string u committed by B , rewind B and force the output of the coin flipping protocol to be a predetermined string z . Such string will be predetermined by M as the string generating $g, h \in G_q$ such that $g^x = h \pmod p$, for some $x \in G_q$ chosen by M . Later, a commitment com computed as $com = g^m h^r \pmod p$ by M , can be decommitted as $s \in G_q$ in the following way: M randomly chooses $r' \in G_q$ and sets $s = m + x(r - r') \pmod p$. It can be seen that (s, r') is a valid decommitment of com as s . The proof continues along the same lines as the proof of Theorem 3. We can define the same probabilities q_0, q_1 and similarly show that $p_0 = q_0, q_1 = q_2$ and $p_1 = q_1$, from which we obtain that $p_1 = p_0$. \square

We remark that the assumption of intractability of computing discrete logarithms implies the existence of computationally secure commitment schemes, therefore Theorem 4 holds under the only assumption of the intractability of computing discrete logarithms.

5 Concurrent zero-knowledge proofs with preprocessing

In this section we consider the problem of constructing concurrent zero-knowledge proofs with preprocessing. We present a transformation which applies to all languages having a public-coin honest-verifier zero-knowledge proof systems, including, in particular, all languages in NP. The transformation returns a concurrent zero-knowledge proof system with preprocessing, requiring therefore no timing assumptions nor requiring the parties to know the number of users in the system. Important properties of the resulting proof system include the fact that it has round, communication and computation complexity comparable to those of the starting proof system; specifically, our transformation does not require general NP reductions, and requires a 3-round preprocessing phase and a proof phase with a number of rounds equal to the number of rounds of the original proof system. For simplicity of description, we now present our result for all languages having a 3-round public-coin honest-verifier zero-knowledge proof system. Formally, we achieve the following

Theorem 5. *Let L be a language having a 3-round public-coin honest-verifier zero-knowledge proof system. Assuming the existence of pseudo-random generators and perfectly secure commitment schemes, there exists (constructively) a concurrent computational zero-knowledge proof system with preprocessing for L where both preprocessing and proof phases require 3 rounds of communication.*

The result underlying our transformation is actually stronger than what stated in the above theorem, providing, for example, additional efficiency properties; various remarks and extensions are discussed in Section 5.2. In the following, we start with an informal description of our transformation, then present a formal description of our concurrent zero-knowledge proof system and then prove its properties.

An informal description. We start by informally describing the main ideas behind our technique. First of all, we should observe that proving a protocol to be concurrent zero-knowledge becomes a problem in the presence of potentially bad interleavings among the polynomially many concurrent executions. In

particular, such interleavings may ask the simulator for too many (eventually, more than polynomial) rewinds in order to be able to succeed in simulating the protocol. Therefore, we use a technique that allows to simulate part of our protocol without performing any rewinding. Part of our protocol uses a special zero-knowledge property, similar to a technique used in [12] in a trusted-center setting and in [14] in the case of arguments. This type of zero-knowledge property, formally defined as *straight-line zero-knowledge*, does not require the simulator to rewind the adversary. Straight line zero-knowledge arguments have been implemented in [12], using either a trusted center or a preprocessing phase with timing assumptions. In this paper we separate the preprocessing phases of all protocols from the proof phases of all protocols and obtain, without using timing assumptions or trusted centers, that the proof phase of our protocol is straight-line zero-knowledge. This is achieved by using the tool of equivocable commitment. Namely, the prover uses a computationally equivocable commitment scheme to commit to some random string d , receives the challenge c from the verifier, decommits his string d and uses the string $c \oplus d$ as a challenge according to the original protocol (A,B). In this way, during the simulation, he can compute such commitment keys in a way such that he can later open them both as a 0 and as a 1. In particular, he will be able to set string d *after* he has seen the challenge c from the adversary. Namely, he will be able to set d in such a way that the string $d \oplus c$ matches a challenge consistent with the first message he has already sent and with the output of an execution of the simulator S that he had previously computed. The scheme will use the computationally equivocable commitment presented in Section 4; in particular, the coin flipping subprotocol used in that scheme will be run in the preprocessing phase of this protocol. Notice that the coin flipping protocol is not straight line simulatable, but its concurrent composition can be simulated in time at most quadratic of the number of executions, no matter which interleavings among them are chosen by the adversary.

Formal description. Let L be a language having a 3-round public-coin honest-verifier zero-knowledge proof system (A,B) and let x be the common input, where $|x| = n$. We denote by S the honest-verifier simulator associated to (A,B), and by (mes, c, ans) a transcript of an execution of (A,B), where mes is the first message, c is the challenge, ans is the answer, and $|c| = m$. Also, let (C,D) be a perfectly secure commitment scheme and let G be a pseudo-random generator. Now we give a formal description of the preprocessing phase subprotocol (P1,V1) and the proof phase subprotocol (P2,V2) of our concurrent zero-knowledge proof system with preprocessing (P,V).

The Proof System ((P1,V1),(P2,V2))

Input to P1 and V1: 1^n , where n is a positive integer.

Instructions for P1 and V1 (preprocessing phase):

1. V1 uniformly chooses $u_1, \dots, u_m \in \{0,1\}^{3n}$ and commits to them using scheme (C,D) (possibly interacting with P1);
2. P1 uniformly chooses $v_1, \dots, v_m \in \{0,1\}^{3n}$ and sends them to V1;
3. V1 decommits u_1, \dots, u_m using scheme (C,D) (possibly interacting with P1);
4. If u_1, \dots, u_m are not properly decommitted, P1 halts;
5. P1 sets α equal to the transcript so far and outputs: α ;

6. V1 sets β equal to the transcript so far and outputs: β .

Input to P2: x, α , where $|x| = n$. **Input to V2:** x, β , where $|x| = n$.

Instructions for P2 and V2 (proof phase):

1. P2 sets $mes = A(x)$ and uniformly chooses $d_1, \dots, d_m \in \{0, 1\}$;
2. For $i = 1, \dots, m$,
 - P2 uniformly chooses $s_i \in \{0, 1\}^n$ and computes $z_i = G(s_i)$;
 - P2 sets $com_i = z_i$ if $d_i = 0$ or $com_i = z_i \oplus u_i \oplus v_i$ if $d_i = 1$;
 - P2 sends mes, com_1, \dots, com_m to B.
3. V2 uniformly chooses $b \in \{0, 1\}^m$ and sends b to P2;
4. P2 decommits d_1, \dots, d_m by sending s_1, \dots, s_m to V2;
5. P2 sets $d = d_1 \circ \dots \circ d_m$, $c = a \oplus d$, $ans = A(x, mes, c)$ and sends ans to V2;
6. V2 checks that $com_i = G(s_i)$ if $d_i = 0$ and $com_i = G(s_i) \oplus u_i \oplus v_i$ if $d_i = 1$;
if any of these checks fails then V2 halts;
7. If $B(x, mes, c, ans) = \text{ACCEPT}$ then V2 outputs: ACCEPT
else V2 outputs: REJECT.

We remark that the protocol does not require any information about the input to be known at preprocessing phase, other than its length.

5.1 Properties of our protocol

We now prove the properties of the described protocol. Clearly the verifier's algorithms V1 and V2 can be performed in probabilistic polynomial time. The completeness requirement directly follows from the correctness of the equivocal commitment scheme and the completeness of the protocol (A,B). The soundness requirement directly follows from the perfect security of the commitment scheme (C,D) and the analysis in [21]. Now we see that the requirement of concurrent zero-knowledge is satisfied.

Concurrent Zero-Knowledge. In the following we describe an efficient simulator Sim which interacts with a probabilistic polynomial-time adversary \mathcal{A} who corrupts all verifiers V_1, \dots, V_q . We will show that for all $x_1, \dots, x_q \in L$, the distributions $\text{Sim}_{\mathcal{A}}(\mathbf{x})$ and $\text{View}_{\mathcal{A}}(\mathbf{x})$ are computationally indistinguishable, where $\mathbf{x} = (x_1, \dots, x_q)$.

The simulator Sim. We start with an informal description of Sim. The first step of the simulator Sim is to run the preprocessing phase of protocol (P,V), where Sim runs algorithm P1 and \mathcal{A} runs V1. This step is done to check whether \mathcal{A} decommits all the commitments in a proper way. If this is the case, then Sim continues the simulation; otherwise, namely, if there exists at least one commitment that is not open properly from \mathcal{A} , Sim outputs the transcript seen until then and halts. In this step Sim also keeps track of which verifier sent its commitment message first, call it V_1 , and which were its decommitments.

The second step of the simulator Sim consists in performing the simulation of the concurrent executions in the preprocessing phase. The algorithm Sim repeats for each verifier V_i , $i \in \{1, \dots, q\}$, the following four substeps, given a 'current' verifier V_i , and the value of its decommitted strings. First, it rewinds the adversary \mathcal{A} until right after V_i had sent its commitments. Then, using its

knowledge of the decommitted strings from V_i , he sends some pseudo-random strings which will allow later to send commitments that can be later opened both as 0 and as 1. Now, Sim will continue the simulation of the preprocessing phase by running the (polynomial-time) algorithm P1 and waiting for the next verifier sending its commitment message. Once, he has found such verifier, he will continue the simulation of the preprocessing phase by running P1 and waiting for the decommitments by such verifier, who becomes then the current verifier in the next execution of these four substeps. If at any time, a verifier sends an inappropriate decommitment, Sim just halts.

The third step of the simulator Sim consists in performing the simulation of the concurrent executions in the proof phase. Note that after the second step is over, if Sim has not halted, in all proofs Sim can compute commitments that he can open both as 0 and as 1. Therefore, the proof can now be simulated without rewindings of any verifier, since, after seeing the message from the verifier, Sim can set the challenge any value it likes by just properly decommitting the equivocable commitments.

A formal description is in Figure 5.1. We have the following

Lemma 6. *For all probabilistic polynomial time algorithms \mathcal{A} , the algorithm $\text{Sim}_{\mathcal{A}}$ runs in expected polynomial time.*

Proof. Let us first consider the simulation of the preprocessing phase. From its formal description, we first observe that Sim runs once the preprocessing phase by running algorithm P1; this step can be run in polynomial time since algorithm P1 also runs in polynomial time. Then we observe that Sim rewinds the algorithm \mathcal{A} at most $q(n)$ times, that is, a number of times polynomial in n , since q is a polynomial. During each rewinding Sim clearly runs at most a polynomial number of steps, since G is computable in polynomial time. Finally, it is easy to see that the simulation of the proof phase also takes at most a polynomial number of steps. \square

Lemma 7. *For all probabilistic polynomial time algorithms \mathcal{A} , all polynomials q , and all $x_1, \dots, x_q \in L$, the distribution $\text{Sim}_{\mathcal{A}}(\mathbf{x})$ is computationally indistinguishable from $\text{View}_{\mathcal{A}}(\mathbf{x})$, where $\mathbf{x} = (x_1, \dots, x_q)$.*

Sketch of Proof. We consider three cases according to whether the adversary decommits its commitments in a proper way in various steps of algorithm Sim. The first case we consider is the one in which the adversary \mathcal{A} always decommits its commitments in a proper way. All messages sent by \mathcal{A} are clearly equally distributed in the simulation and in the proof since they are computed in the same way. Now, let us consider the messages from the provers. We observe that the triple $(\text{mes}_i, c_i, \text{ans}_i)$ is output by the simulator S for (A, B) and therefore its distribution is computationally indistinguishable from the distribution of the same triple in the proof. The only remaining messages to consider are the strings $v_{i,j}$ in the preprocessing phase, the commitments $z_{i,j,\cdot}$ and the associated decommitments $s_{i,j,\cdot}$ in the proof phase. From the description of Sim and P1, we can see that the distributions of such strings are clearly different; for instance, the strings $v_{i,j}$ are random in a real proof and pseudo-random in the simulated execution. However, notice that each triple $(v_{i,j}, z_{i,j,\cdot}, s_{i,j,\cdot})$ is part of the transcript of an execution of the computationally equivocable commitment scheme

The Algorithm Sim

Input to Sim: x_1, \dots, x_q , where $|x_i| = n$.

Instructions for Sim (preprocessing phase):

1. Run the preprocessing subprotocol (P1,V1), playing as P1 and interacting with \mathcal{A} who plays as V1; (if \mathcal{A} ever decommits inappropriately, output the transcript obtained until then and halt;)
 call V_1 the first verifier sending its commitments (according to scheme (C,D));
 set $i = 1$ and let $u_{1,1}, \dots, u_{1,m} \in \{0, 1\}^{3n}$ be the strings decommitted from V_1 ;
2. repeat while $i \leq q$:
 rewind \mathcal{A} until right after V_i has sent its commitments;
 uniformly choose $s_{i,j,0}, s_{i,j,1} \in \{0, 1\}^n$, set $z_{i,j,0} = G(s_{i,j,0})$, $z_{i,j,1} = G(s_{i,j,1})$,
 and $pr_{i,j} = z_{i,j,0} \oplus z_{i,j,1}$, for $j = 1, \dots, m$;
 set $v_{i,j} = pr_{i,j} \oplus u_{i,j}$ and send $v_{i,j}$ to V_i , for $j = 1, \dots, m$;
 get a message *advmes* from \mathcal{A} ;
 if *advmes* is a decommitment message from V_i then
 if the decommitment is not appropriate then halt;
 if $i = q$ then exit from the repeat loop else get a message *advmes* from \mathcal{A} ;
 if *advmes* is a commitment message from a verifier $\neq V_i$ then
 call such verifier V_{i+1} and set $i = i + 1$;
 repeat
 get a message *advmes* from \mathcal{A} ;
 if *advmes* is a commitment message from a verifier $\neq V_i$ then
 uniformly choose $v_{i,1}, \dots, v_{i,m} \in \{0, 1\}^{3n}$ and send them to \mathcal{A} ;
 if *advmes* is a decommitment message from a verifier $\neq V_i$ then
 if the decommitment is not appropriate then halt;
 until *advmes* is the decommitment message from V_i ;
 if the decommitment by V_i is not appropriate then halt;
 let $u_{i,1}, \dots, u_{i,m} \in \{0, 1\}^{3n}$ be the strings decommitted from V_i ;
3. for $i = 1, \dots, m$, let α_i be a transcript of the communication between Sim and V_i ,
 let $\beta_i = (r_i \circ \alpha_i)$, where r_i is the random tape used by V_i and let T_1 be the entire transcript between all P_i 's and all V_i 's so far.

Instructions for Sim (proof phase):

1. Repeat
 get a message *advmes* from \mathcal{A} ;
 if *advmes* = 'start i -th proof' then
 uniformly choose $rand_i$ and let $(mes_i, c_i, ans_i) = S(rand_i, x_i)$;
 send $z_{i,1,0}, \dots, z_{i,m,0}$ to \mathcal{A} ;
 if *advmes* = $b_i \in \{0, 1\}^m$ then
 send *ans_i* and let d_{ij} be the j -th bit of $c_i \oplus b_i$;
 decommit each $z_{i,j,0}$ as d_{ij} by sending to \mathcal{A} $s_{i,j,0}$ if d_{ij} is 0 or $s_{i,j,1}$ if d_{ij} is 1;
2. let T_2 be the entire transcript between all P_i 's and all V_i 's in this phase;
3. output: $(T_1, T_2, \beta_1, \dots, \beta_m)$.

Figure 1: The simulator algorithm Sim

presented in Section 4. Therefore, we can directly apply Theorem 3 to conclude that the distribution of such triple in the real proof and the distribution of such triple in the simulation are computationally indistinguishable.

The second case we consider is the one in which the adversary \mathcal{A} decommits its commitments in step 1 in a proper way but decommits in a non proper way in one of the remaining steps of Sim. Notice that in this case Sim halts without output and therefore the two distributions differ. However we now see that this happens with negligible probability. Assume not; then \mathcal{A} is able to distinguish between the first simulated execution (in which all the $u_{i,j}$'s are random) and the second simulated execution (in which the $u_{i,j}$'s are pseudo-random). In the proof of Theorem 3 we showed that in the case of a single execution of a single bit protocol, this fact implies the existence of an efficient distinguisher between a random string and the output of G , thus contradicting the assumption that G is a pseudo-random generator. A modification of that proof can be used here too, in the concurrent setting, where the modification consists in using a hybrid argument to take care of the fact that there are several strings $u_{i,j}$ that are random in one execution and pseudo-random in the second execution.

The last case we need to consider is the one in which the adversary \mathcal{A} does not decommit in a proper way its commitments in step 1 of algorithm Sim. Note that if this happens the simulator Sim halts by outputting the transcript so far; this is the same view as that of \mathcal{A} during the real proof (since the prover also halts in this case); therefore, the simulation is perfect in this case. \square

5.2 Remarks and Extensions

We present some remarks and extensions of our main result, as given in Theorem 5, concerning minimal complexity assumptions, efficiency, allowing any polynomial number of proofs, and extending to any public-coin protocol.

Efficiency. In terms of communication and computation complexity, when applied to several 3-round public-coin honest-verifier zero-knowledge protocols in the literature (e.g., [17, 5, 18, 25, 6]) our proof system has efficiency comparable to that of the original protocol. In particular, contrarily to previously given concurrent zero-knowledge proof system, the construction of our protocol does not require any NP reduction, which could potentially blow up the parameters. Considering that our protocol is also a proof of knowledge, this may have applications to identification schemes. Moreover, in many examples for which the original computational zero-knowledge proof system already requires a commitment scheme (e.g., [17, 5]), an optimization can be performed: such protocols can be made concurrent zero-knowledge in our model by only implementing the commitment scheme using our equivocable commitment scheme in Section 4.

Minimal complexity assumptions. Our protocol is based on the existence of perfectly secure commitment schemes and pseudo-random generators. Both can be constructed under the assumption of the existence of collision-intractable functions, and number-theoretic assumptions as the intractability of deciding quadratic residuosity modulo composite integers.

Any polynomial number of proofs. Our proposed protocol allows a number of proofs bounded by the size of the preprocessing. By combining our techniques with those in [13] for non-interactive zero-knowledge proofs, we can allow any polynomial (in the size of the preprocessing) number of proofs. This variation,

however, involves a general NP reduction, thus making the protocol much less efficient in terms of communication and computation complexity.

Extending to any public-coin honest-verifier zero-knowledge proof system. This extension is obtained by just observing that the transformation done for the challenge message in the 3-round protocol in the proof of Theorem 5 can be performed to all public-coin messages from the verifier. In particular, our results apply to all languages having an interactive proof system since they have a public-coin honest-verifier zero-knowledge proof system [2, 19].

6 Concurrent Zero-Knowledge Arguments with preprocessing

We now consider the problem of constructing concurrent perfect zero-knowledge arguments with preprocessing. Specifically, we show that the protocol in Section 5, when implemented using the perfectly equivocal commitment scheme in Section 4, allows to obtain concurrent perfect zero-knowledge arguments. Formally, we achieve the following

Theorem 8. *Let L be a language in NP. Assuming the existence of perfectly secure commitment schemes, there exists (constructively) a concurrent perfect zero-knowledge argument with preprocessing for L where the preprocessing and proof phase require 3 rounds of communication, and the soundness property holds under the intractability of computing discrete logarithms modulo a prime.*

The proof of the above theorem goes along the same lines as the proof of Theorem 5 and therefore we only sketch it, by pointing out the few differences. First of all, we reduce L to an NP-complete language, say, Hamiltonian Graphs. Then we can use the 3-round public-coin honest-verifier zero-knowledge proof system given in [5], call it (A,B). Starting from this protocol, we construct a protocol (P,V) as follows. Whenever, during the protocol (A,B), A is required to use a computationally equivocal commitment scheme, we will require A to run the perfectly equivocal commitment scheme in Section 4. The intuition here, also used in [14], is that the ability of decommitting each commitment to mes_i as any desired string allows to perform a perfect simulation of protocol (A,B) without need of a witness for the original input graph. We then have that the soundness holds under the same properties than the binding property of such commitment scheme, that is, the intractability of computing discrete logarithms modulo a prime. The concurrent zero-knowledge property is proved as for Theorem 5. We note that this technique also applies to languages having 3-round public-coin honest-verifier perfect zero-knowledge proof systems as those in [18, 17, 25, 6], in which case it returns a perfect zero-knowledge argument with preprocessing and with efficiency (in terms of communication and computational complexity) comparable to the starting protocol. Specifically, we do not need any NP reduction during the construction of the protocol, even if we want to run a number of proofs polynomial in the length of the preprocessing.

Acknowledgement. We thank Ran Canetti for useful discussions regarding equivocal commitments.

References

1. D. Beaver, *Adaptive Zero-Knowledge and Computational Equivocation*, Proc. of FOCS 96.
2. M. Ben-Or, J. Hastad, J. Kilian, O. Goldreich, S. Goldwasser, S. Micali, and P. Rogaway, *Everything Provable is Provable in Zero-Knowledge*, Proc. of CRYPTO 88.
3. T. Beth and Y. Desmedt, *Identification Tokens - or: Solving the Chess Grandmaster Problem*, Proc. of CRYPTO 90.
4. G. Brassard, D. Chaum, and C. Crepeau, *Minimum Disclosure Proofs of Knowledge*, Journal of Computer and System Sciences, vol. 37, n. 2, 1988, pp. 156–189.
5. M. Blum, *How to Prove a Theorem So No One Else Can Claim It*, Proc. of the International Congress of Mathematicians, California, 1986, pp. 1444–1451.
6. A. De Santis, G. Di Crescenzo, G. Persiano, and M. Yung, *On Monotone Formula Closure of SZK*, Proc. of FOCS 94.
7. A. De Santis, S. Micali, and G. Persiano, *Non-Interactive Zero-Knowledge Proof Systems with Preprocessing*, Proc. of CRYPTO 88.
8. G. Di Crescenzo, *A Note on Distributed Zero-Knowledge Proofs*, unpublished manuscript (submitted to PODC, January 1997).
9. G. Di Crescenzo, Y. Ishai, and R. Ostrovsky, *Non-Interactive and Non-Malleable Commitment*, Proc. of STOC 98.
10. D. Dolev, C. Dwork, and M. Naor, *Non-Malleable Cryptography*, Proc. of STOC 91.
11. C. Dwork, M. Naor, and A. Sahai, *Concurrent Zero-Knowledge*, Proc. of STOC 98.
12. C. Dwork and A. Sahai, *Concurrent Zero-Knowledge: Reducing the Need for Timing Constraints*, Proc. of CRYPTO 98.
13. U. Feige, D. Lapidot, and A. Shamir, *Multiple Non-Interactive Zero-Knowledge Proofs based on a Single Random String*, Proc. of FOCS 90.
14. U. Feige and A. Shamir, *Zero-Knowledge Proofs of Knowledge in Two Rounds*, Proc. of CRYPTO 89.
15. O. Goldreich, *Foundations of Cryptography: Fragments of a Book*, in Electronic Colloquium on Computational Complexity.
16. O. Goldreich and A. Kahan, *How to construct constant-round zero-knowledge proof systems for NP*, in Journal of Cryptology, vol. 9, n. 3, pp. 167–190.
17. 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, n. 1, 1991, pp. 691–729.
18. S. Goldwasser, S. Micali, and C. Rackoff, *The Knowledge Complexity of Interactive Proof-Systems*, SIAM Journal on Computing, vol. 18, n. 1, February 1989.
19. R. Impagliazzo and M. Yung, *Direct Minimum-Knowledge Computation*, Proc. of CRYPTO 87.
20. J. Kilian, E. Petrank, and C. Rackoff, *Lower Bounds for Zero-Knowledge on the Internet*, Proc. of FOCS 98.
21. M. Naor, *Bit Commitment from Pseudo-Randomness*, Proc. of CRYPTO 91.
22. M. Naor, R. Ostrovsky, R. Venkatesan, and M. Yung, *Perfect Zero-Knowledge Arguments under General Complexity Assumptions*, Proc. of CRYPTO 92.
23. T. Pedersen, *Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing*, Proc. of CRYPTO 91.
24. R. Richardson and J. Kilian, *On the Concurrent Composition of Zero-Knowledge Proofs*, Proc. of EUROCRYPT 99.
25. M. Tompa and H. Woll, *Random Self-Reducibility and Zero-Knowledge Interactive Proofs of Possession of Information*, Proc. of FOCS 87.

This article was processed using the L^AT_EX macro package with LLNCS style