Concurrent Zero-Knowledge is Easy in Practice

Ivan Damgård

Aarhus University, BRICS*

(preliminary version)

Abstract. We show that if any one-way function exists, then 3-round concurrent zero-knowledge arguments for all NP problems can be built in a model where a short auxiliary string with a prescribed distribution is available to the players. We also show that all known efficient identification schemes using specialized assumptions can be modified to work in this model with no essential loss of efficiency. We argue that the assumptions of the model will be satisfied in most practical scenarios where public key cryptography is used, in particular our construction works given any secure public key infrastructure. Finally, we point out that in a model with preprocessing (and no auxiliary string) proposed earlier, concurrent zero-knowledge for NP can be based on any one-way function.

1 Introduction

In a zero-knowledge protocol [17], a prover convinces a verifier that some statement is true, while the verifier learns nothing except the validity of the assertion. Apart from being interesting as theoretical objects, it is well-known that zero-knowledge protocols are extremely useful tools for practical problems, e.g., stand-alone for identification schemes, but perhaps even more as subprotocols in schemes for voting, electronic cash, etc.

Hence the applicability of the theory of zero-knowledge in real life is of extreme importance. One important aspect of this is composition of protocols, and the extent to which such composition preserves zero-knowledge. While sequential composition does preserve zero-knowledge, this is not always the case for parallel composition [16].

In [9] Dwork, Naor and Sahai pointed out that the strict synchronization usually assumed when composing zero-knowledge protocols is unrealistic in scenarios such as Internet based communication. Here, many instances of the same or different protocols may start at different times and may run with no fixed timing of messages. What is needed here is a stronger property known as *concurrent zero-knowledge*, i.e., even an arbitrary interleaving of several instances of zero-knowledge protocols is again zero-knowledge, even when the verifiers are all controlled by a single adversary, who may use information obtained from one protoccol to determine its behavior in another instance.

* Basic Research in Computer Science, Center of the Danish National Research Foundation

Unfortunately, standard constructions for zero-knowledge protocols fail to provide this property. This is because they are based on simulation by rewinding the verifier. In a concurrent setting, the simulator may be forced to rewind an exponential number of times. In fact, it seems that concurrent zero-knowledge cannot be provided at all in the usual model with as few rounds as ordinary zero-knowledge. Kilian, Petrank and Rackoff [14] show that only BPP languages have concurrent zero-knowledge proofs or arguments with 4 rounds or less, if black-box simulation is assumed [1].

Thus, a lot of research has gone into finding ways of getting around this problem. In [9], it was shown that given constraints on the time at which messages must arrive, concurrent zero-knowledge can be achieved for all of NP in a constant number of rounds. Subsequently it was shown that the need for timing constraints could be pushed into a preprocessing phase[10]. In [7] it was shown that the timing constraints in the preprocessing can be reduced to merely ensuring that all preprocessings are finished before the main proofs start. This comes at the price that the work needed in the preprocessing depends on the size and number of statements to be proved later. Finally, Richardson and Kilian [20] show that it is possible to do without timing constraints, at the expense of a non-constant number of rounds.

We note that a completely different approach is possible: one could go for a weaker property than zero-knowledge, one that would be preserved in a concurrent setting. One such possibility is the Witness-Hiding (WH) protocols of Feige and Shamir [12]. Most WH protocols are based on the standard paradigm of the prover proving knowledge of one of two "computationally independent" witnesses without revealing which one he knows. Such protocols are also WH when used concurrently, and can be used to construct secure identification systems. In [6], very efficient methods for building such protocols are developed. However, for more general use, e.g., as subrutines in multiparty computation or verifiable secret sharing protocols, WH is not always sufficient, one needs simulatability to prove the overall protocol secure.

2 Our Work

Our main objective is to show that concurrent zero-knowledge can often be obtained in a simple way using standard tools. We do not claim any major new techniques, in fact our solution is quite straightforward. Nevertheless, we believe it is useful to realize that concurrent zero-knowledge is very easy to achieve in many practical scenarios. We do not mean to suggest that this makes the large body of theoretical work on the subject less interesting or important, or that our solution can handle any practical scenario. Independently, Kilian and Petrank have made observations very similar to ours [13].

[1] Virtually all known zero-knowledge protocols are black-box simulatable

2.1 The Model

Our work starts from the following assumption: an auxiliary string with a prescribed distribution is available to the prover and verifier. Given this assumption we will see that concurrent zero-knowledge is easy to achieve in constant round with no timing constraints or preprocessing. Informally, zero-knowledge in such a setting means as usual that the verifiers entire view can be simulated efficiently, which here means its view of the interaction with the prover, as well as the auxiliary string. Soundness means that no polynomial time prover can cheat the verifier with non-negligible probability where the probability is taken over the choice of the auxiliary string as well as the coin tosses of the players. Formal definitions will appear in the final version of this paper.

Note that the standard non-interactive zero-knowledge model (where the auxiliary string is a uniformly chosen random string) [2] is a special case, and indeed by their very nature non-interactive zero-knowledge proofs do not require rewinding to simulate, and so are robust in a concurrent setting. It is even possible to do any polynomial number of non-interactive proofs based on the same globally shared random string [11].

However, there are still several reasons why non-interactive zero-knowledge proofs are not the answer to all our problems: they are in general much less efficient than interactive ones and - as far as we know - require stronger cryptographic assumptions (trapdoor one-way permutations as opposed to arbitrary one-way functions). We would like a solution allowing us to use standard efficient constructions of protocols securely in a concurrent setting, without significant loss of efficiency. Moreover, non-interactive proofs are always proofs of language membership, and cannot be proofs of knowledge for non-trivial problems. There are cases, for instance when ensuring so called input awareness in multiparty computation, where proofs of knowledge (rather than membership proofs) are required.

The model we use (with a general auxiliary string) was also used in [4] (for a different purpose). The rationale for allowing a general distribution of the reference string is of course that one may hope that this allows for more efficient protocols, for example a much shorter auxiliary string. The problem, on the other hand, may be that requiring a more powerful resource makes the model less realistic.

However, as we shall see, our protocols do in fact apply to a realistic situation, namely a public-key cryptography setting where users have public/private key pairs. In fact our prover and verifier do not need to have key pairs themselves, nevertheless, they will be able to prove and verify general NP statements in concurrent zero-knowledge by using the public key P_A of a third party A as auxiliary string. This will work, provided that

- The verifier believes that A's secret key is not known to the prover.
- The prover believes that P_A was generated using the proper key generation algorithm for the public-key system in use.

We stress that A does not need to take part in the protocols at all, nor does he need to be aware that his public key is being used this way, in particular keys for standard public key systems like RSA, El Gamal or DSS can be used directly.

Note that if we have a secure public key infrastructure where public keys are being certified by a certification authority (CA), then all our demands are already automatically satisfied because the CA can serve as player A in the above: in order for the infrastructure to be secure in the first place, each user needs to have an authentic copy of the CA's public key available, and one must of course trust that the CA generated a proper key and does not reveal its private key to anyone else.

So although our model does make stronger assumptions on the environment than the standard one, we believe that this can be reasonable: The problem of concurrent zero-knowledge arises from the need to apply zero-knowledge protocols in real situations. But then solutions to this problem should be also allowed to take advantage of resources that may exist in such scenarios.

2.2 The Results

Our first result is a general construction for protocols of a particular form. Assume we have a binary relation R, and a 3-move proof of knowledge for R, where the verifier sends a random challenge as the second message. Thus conversations in this protocol are of form (a, e, z), where the prover chooses a, z. We will assume that this protocol is honest verifier zero-knowledge in the sense that given e, one can efficiently compute a correctly distributed conversation where e is the challenge. Finally we assume that a cheating prover can answer only one of the possible challenges, or more precisely, from the common input x and any pair of accepting conversations $(a, e, z), (a, e', z')$ where $e \neq e'$, one can compute a witness of x, i.e. w such that $(x, w) \in R$. We call this a Σ-protocol. We have

Theorem 1. *Given any binary relation R and a Σ-protocol for R. If one-way functions exist, then there exists a computationally convincing and concurrent zero-knowledge 3-move proof of knowledge (with negligible knowledge error and no timing constraints) for R in the auxiliary string model.*

The construction behind this result can be applied in practice to the well known identification schemes of Schnorr and Guillou-Quisquater to yield concurrent zero-knowledge identification schemes in the auxiliary string model with negligible loss of efficiency compared to the original protocols (which were not even zero-knowledge in the usual sense!). The idea behind this result also immediately gives:

Theorem 2. *If one-way functions exist, there exist 3-move concurrent zero-knowledge interactive arguments in the auxiliary string model (with no timing constraints) for any NP problem.*

In both these results, the length of the auxiliary string is essentially the size of the computational problem the prover must solve in order to cheat. The length does not depend on the size or the number of statements proved.

Our final result is an observation concerning the preprocessing model of Dwork and Sahai [10] (where there is no auxiliary string). It was shown in [10] that prover and verifier can do a once-and-for-all preprocessing (where timing constraints are applied), and then do any number of interactive arguments for any NP problem in concurrent zero-knowledge (with no timing constraints) in 4 rounds. This was shown under the assumption that one-way trapdoor permutations exist. Below, we observe the following:

Theorem 3. *If any one-way functions exists, then any NP problem has a 3-round concurrent zero-knowledge argument in the preprocessing model of Dwork and Sahai.*

We note that our preprocessing is once-and-for-all, like the one in [10]: once the preprocessing is done, the prover and verifier can execute any polynomial number of proofs securely, and the complexity of the preprocessing does not depend on the number or size of the statements proved.

3 The Protocols

3.1 Trapdoor Commitments Schemes

In a commitment scheme, a committer C can commit himself to a secret s chosen from some finite set by sending a *commitment* to a reciever R. The receiver should be unable to find s from the commitment, yet C chould be able to later open the commitment and convince R about the original choice of s.

A *trapdoor commitment scheme* is a special case that can be loosely described as follows: first a public key pk is chosen based on a security parameter value k, usually by R, and is sent to C. There is a fixed function *commit* that C can use to compute a commitment c to s by choosing some random input r, and setting $c = commit(s, r, pk)$. Opening takes place by revealing s, r to R, who can then check that $commit(r, s, pk)$ is the value he received originally. We require the following:

Hiding: For correctly chosen pk, uniform r, r' and any s, s', the distributions of $commit(s, r, pk), commit(s', r', pk)$ are polynomially indistinguishable.

Binding: For any polynomially bounded C, the probability that C on input pk computes s, r, s', r' such that $commit(s, r, pk) = commit(s', r', pk)$ and $s \neq s'$ is negligible.

Trapdoor Property: The algorithm for generating pk also outputs a string t, the trapdoor. There is an efficient algorithm which on input t, pk outputs a commitment c, and then on input any s produces r such that $c = commit(s, r, pk)$. The distribution of c is poynomially indistinguishable from that of commitments computed in the usual way.

In other words, the commitment scheme is binding if you know only pk, but given the trapdoor, you can cheat arbitrarily.

From the results in Shamir et al.[15], it follows that existence of any one-way function f implies the existence of a trapdoor commitment scheme, where the public key is simply $f(y)$, where y is chosen uniformly in the input domain of f, and y is the trapdoor. Based on standard intractability assumptions such as hardness of discrete log or RSA root extraction, very efficient trapdoor commitment schemes can be built, see e.g. [5].

3.2 A construction for Σ-protocols

In what follows, we will assume that we have a relation R and a Σ-protocol $\mathcal{P}$ for R. Also, we will be in the auxiliary string model, where the auxiliary string will be the public key pk of a trapdoor commitment scheme. Our protocol in the auxiliary string model gets as common input for prover and verifier x, while the prover gets as private input w, such that $(x, w) \in R$. For simplicity, we assume that the commitment scheme allows to commit in one commitment to any string a, that may occur as the first message in $\mathcal{P}$ (in case of a bit commitment scheme, we could just commit bit by bit). The protocol then proceeds as follows:

1. On input x, w, the prover computes a using the prover's algorithm from $\mathcal{P}$, chooses r at random and sends $c = \mathit{commit}(a, r, pk)$ to the verifier.
2. The verifier chooses e at random and sends it to the prover.
3. The prover computes z, the answer to challenge e in $\mathcal{P}$ and sends z, a, r to the verifier.
4. The verifier accepts iff it would have accepted on x, a, e, z in $\mathcal{P}$, and if $c = \mathit{commit}(a, r, pk)$.

It is straightforward to show that this protocol has the desired properties. First, a simulator for the protocol given an arbitrary verifier V^*:

1. Generate pk with known trapdoor t and give x, pk to V^*.
2. Send a commitment c computed according to the trapdoor property to V^* and get e back.
3. Run the honest verifier simulator on input e to get an accepting conversation (a, e, z) in the original protocol. Use the trapdoor to compute r such that $c = \mathit{commit}(a, r, pk)$. Send z, a, r to V^*.

This simulation works based on the hiding and trapdoor properties of the commitment scheme, and does not require rewinding of V^*, hence the protocol is also concurrent zero-knowledge.

For the knowledge soundness (or validity), we give only a very loose analysis in this preliminary version. Note that we are using the definition of Bellare and Goldreich, in the version modified for computationally convincing proofs of knowledge[3]. We will aim for a knowledge error $\kappa(x)$ satisfying that the binding property of the commitments can be broken in time $\mathit{poly}(|x|)/\kappa(x)$ with only negligible probability. Depending on the intractability assumption we are willing to make in the binding condition, we can set $\kappa(x)$ to be some concrete negligible function. The properties of a Σ-protocol are preserved under parallel composition, so we can assume without loss of generality that the length k of the challenge e is such that $2^{-k} < \kappa(x)$.

From any prover convincing the verifier with probability $p(x) > \kappa(x)$, we can extract, using rewinding, convincing answers to two different challenges e, e', in time proportional to $1/p(x)$ since $p(x) > 2^{-k}$. Note that this is less obvious than it may seem: the prover may be probabilistic, but we still have to fix his random tape once we start rewinding. And there is no gurantee that the prover has success probability $p(x)$ for all choices of random tapes, indeed $p(x)$ is the average over all such choices. However, a strategy for probing the prover can be devised that circumvents this problem. Details will appear in the final version of the paper.

Once we are successful, we get commitment c, conversations $(a, e, z), (a', e', z')$ that are accepting in the original protocol, and finally values r, r' such that $c = \mathit{commit}(a, r, pk) = \mathit{commit}(a', r', pk)$. This breaks the binding property of the commitment scheme if $a \neq a'$, but this possibility can be ignored because we have used time less than a polynomial times $1/\kappa(x)$ to reach this situation. On the other hand, if $a = a'$, a witness for the common input x can be computed by assumption on the original protocol.

This protocol and the result from [15] above on existence of trapdoor commitments now implies Theorem 1. As for Theorem 2, we just need to observe that the standard zero-knowledge interactive protocols for NP complete problems [18, 1] can in fact be based on any commitment scheme. They are usually described as sequential iterations of a basic 3-move protocol. However, in our model we will use a trapdoor commitment scheme, and do the iterations in parallel: it is then trivial that the protocols can be straight line simulated if the simulator knows the trapdoor. And soundness for a poly-time bounded prover follows by a standard rewinding argument. A more careful analysis of the error probability and the way it depends on the intractability assumption we make can be obtained using the definitions from [8].

This same idea applies easily to the preprocessing model (with no auxiliary string) of Dwork and Sahai [10]: here, the prover and verifier are allowed to do a preprocessing, where timing constraints are used in order to ensure concurrent zero-knowledge. After this, the goal is to be able to do any number of interactive arguments in concurrent zero-knowledge, without timing constraints. In [10], it is shown how to achieve this based on existence of one-way trapdoor permutations. However, an idea similar to the above will allow us to do it based on any one-way function (and a smaller number of rounds): In the preprocessing, the verifier chooses an instance of the trapdoor commitment scheme from [15] and sends the public key to the prover. The verifier then proves knowledge of the trapdoor. After this, any number of interactive arguments for NP problems can be carried out in constant round and concurrent zero-knowledge. We will use the parallel version of [18] or [1] based on the commitment scheme we established in the preprocessing. Simulation can be done by extracting the trapdoor from the verifier's proof of knowledge (here, rewinding is allowed because of the timing constraints) and then simulate the main proofs straight-line.

4 Implementation in Practice

In our arguments for practicality of our model, we claimed that the public key of a third party can be used as auxiliary string. Given the construction above, this amounts to claiming that the public key of any public-key crypto-system or signature scheme can also be used without modification as the public key of a trapdoor commitment scheme.

We can assume that the public key was generated using some known key generation algorithm (recall that we originally assumed about the third party that he generates his keys properly and does not give away the private key). Clearly, the function mapping the random bits consumed by this algorithm to the resulting public key must be one-way. Otherwise, the system could be broken by reconstructing the random input and running the algorithm to obtain the private key. Thus, from a theoretical point of view, we can always think of the public key as the image of a random input under a one-way function and apply the commitment scheme from [15].

This will not be a practical solution. But fortunately, standard public key systems used in real life allow much more efficient implementations. Any system based on discrete logarithms in a prime order group, such as DSS, many El Gamal variants, and Cramer-Shoup has as part of the public key some group element of form g^x where x is private and g is public, and where g has prime order q . This is precisely the public key needed for the trapdoor commitment scheme of Pedersen [19], which allows commitment to a string of length $\log q$ in one commitment.

If we have an RSA public key with modulus n, we can always construct from this a public key for the RSA based trapdoor commitment scheme described in [5]. We define q to be the least prime such that $q > n$ (this can easily be computed by both prover and verifier). We then fix some number b in Z_n^*, this could be for instance be a string representing the name of the verifier. The intractability assumption for the commitment scheme then is that the prover will not be able to extract a q'th root mod n of b (such a root always exists by choice of q). Also this scheme allows comitment to $\log q$ bits in one commitment.

Note that when executing a proof of the kind we constructed, it is always enough in practice for the prover to make only one commitment: he can always hash the string the wants to commit to using a standard collision intractable hash function and commit to the hash value. This means that well known efficient protocols can be executed in this model with no significant loss of efficiency.

References

1. Brassard, Chaum and Crépeau: *Minimum disclosure proofs of knowledge*, JCSS vol. 37, pp.156-189, 1988.
2. Blum, De Santis, Micali and Persiano: *Non-Interactive Zero-Knowledge*, SIAM J.Computing, vol.20, 1991.
3. Bellare and Goldreich: *Defining proofs of knowledge*, Proc. of Crypto 92, Springer Verlag LNCS series nr. 740.

4. Brandt, Landrock, Damgaård and Pedersen: *Zero-knowledge authentication scheme with secret key exchange*, J.Cryptology, vol.11, pp.147-160, 1998.
5. Cramer and Damgård: *Zero-Knowledge proofs for finite field arithmetic*, Proc. of Crypto 98, Springer Verlag LNCS series nr.1462.
6. Cramer, Damgaård and Schoenmakers: *Proofs of partial knowlegde*, Proc. of Crypto 94, Springer Verlag LNCS series nr. 839.
7. Di Crescenzo and Ostrovsky: *Concurrent Zero-Knowledge: Avoiding Impossibility with Pre-Processing*, Proc. of Crypto 99, to appear.
8. Damgård and Pfitzmann: *Squential Iteration of Interactive Arguments*, proc. of ICALP 98, Springer Verlag LNCS series.
9. Dwork, Naor and Sahai: *Concurrent Zero-Knowledge*, Proc. of STOC 98.
10. Dwork and Sahai: *Concurrent Zero-Knowledge. Reducing the need for timing constraints*, Proc. of Crypto 98, Springer Verlag LNCS series, nr.1462.
11. Feige, Lapidot and Shamir: *Multiple non-interactive zero knowledge proofs based on a single random string*, Proc. of FOCS 90.
12. Feige and Shamir: *Witness indistinguishability and witness hiding protocols*, Proc. of STOC 90.
13. Joe Kilian: Private communication.
14. Kilian, Petrank and Rackoff: *Lower Bounds for Zero-Knowledge on the Internet*, Proc. of FOCS 98.
15. Feige and Shamir: *Zero-knowledge proofs of knowledge in two rounds*, Proc. of Crypto 89, Springer Verlag LNCS series nr. 435.
16. Goldrecish and Krawczyk: *On the composition of zero-knowledge proof systems*, SIAM J.Computing, vol.25, pp.169-192, 1996.
17. Goldwasser, Micali and Rackoff: *The knowledge complexity of interactive proof systems*, SIAM J.Computing, vol. 18, pp.186-208, 1989.
18. Goldreich, Micali and Wigderson: *Proofs that yield nothing but their validity or all languages in NP have zero-knowledge interactive proofs systems*, JACM, vol 38, pp.691-729, 1991.
19. Pedersen: *Non-interactive and information theoretic secure verifiable secret sharing*, proc. of Crypto 91, Springer Verlag LNCS series, nr. 576.
20. Richardson and Kilian: *On the Concurrent Composition of Zero-Knowledge Proofs*, to appear in Proceedings of EuroCrypt 99.