In *Advances in Cryptology—CRYPTO '94*, Vol. 839 of *Lecture Notes in Computer Science*, Springer-Verlag, 1994. pp. 174-187.

Proofs of Partial Knowledge and Simplified Design of Witness Hiding Protocols

Ronald Cramer, CWI
Ivan Damgård, Aarhus University, Denmark
Berry Schoenmakers, CWI

Abstract. Suppose we are given a proof of knowledge $\mathcal{P}$ in which a prover demonstrates that he knows a solution to a given problem instance. Suppose also that we have a secret sharing scheme $\mathcal{S}$ on n participants. Then under certain assumptions on $\mathcal{P}$ and $\mathcal{S}$, we show how to transform $\mathcal{P}$ into a witness indistinguishable protocol, in which the prover demonstrates knowledge of the solution to some subset of n problem instances out of a collection of subsets defined by $\mathcal{S}$. For example, using a threshold scheme, the prover can show that he knows at least d out of n solutions without revealing which d instances are involved. If the instances are independently generated, we get a witness hiding protocol, even if $\mathcal{P}$ did not have this property. Our results can be used to efficiently implement general forms of group oriented identification and signatures. Our transformation produces a protocol with the same number of rounds as $\mathcal{P}$ and communication complexity n times that of $\mathcal{P}$. Our results use no unproven complexity assumptions.

1 Introduction

In this work[1], we assume that we are given an interactive proof where the prover P convinces the verifier V that P knows some secret. Typically, the secret is the preimage under some one-way function of a publicly known piece of information. Thus the secret could be for example a discrete log or an RSA root. Such a proof is called a proof of knowledge [5], and can be used in practice to design identification schemes or signature systems.

We assume in the following that the proof of knowledge has a special form in that the verifier only sends uniformly chosen bits. This is also known as a *public coin protocol*. For simplicity, we restrict ourselves to 3-round protocols, where the prover speaks first (generalization of our results to any number of rounds is possible). We also assume that the protocol is honest verifier zero-knowledge (HVZK), i.e. the protocol does not reveal anything (for example about the prover's secret) to the honest verifier, but it is not necessarily secure against a cheating verifier.

Numerous protocols are known to satisfy the conditions described above. Concrete examples are Schnorr's discrete log protocol [13] and Guillou-Quisquater's RSA root protocol [8]. None of these protocols are known to be zero-knowledge or even witness hiding. In general, a parallelization of a sequential zero-knowledge (ZK) proof [7] will often satisfy the conditions.

[1] Partly done while visiting Aarhus University.

The second ingredient we need is a secret sharing scheme, i.e. a scheme for distributing a secret among a set of participants such that some subsets of them are qualified to reconstruct the secret while other subsets have no information about it. The collection of qualified subsets is called the access structure. The secret sharing scheme has to satisfy some properties which will be made more precise below. Shamir's secret sharing scheme [14] has the properties we need.

Our main result uses a proof of knowledge $\mathcal{P}$, an access structure Γ for n participants, and a secret sharing scheme $\mathcal{S}$ for the access structure dual to Γ to build a new protocol, in which the prover shows that he knows solutions to a subset of n problem instances corresponding to a qualified set in the access structure of Γ (see Section 3 for details on access structures). The protocol is witness indistinguishable, i.e. the prover reveals no Shannon information about which qualified subset of solutions he knows. The new protocol has the same number of rounds as $\mathcal{P}$ and communication complexity roughly n times that of $\mathcal{P}$. We also show that for some access structures, the new protocol is in fact witness hiding (WH), i.e. even even a cheating verifier will not learn enough to be able to compute the prover's secret. Although WH is a weaker property than general ZK, it can replace ZK in many protocol constructions, including identification schemes.

Since a simple 1 out of 2 structure is enough for our result to produce a WH protocol, we obtain as a corollary a general method simplifying the design of WH protocols: first build a protocol $\mathcal{P}$ with properties as described above - for security against the verifier only the weak and therefore easy to obtain property of HVZK is needed. Then apply our result using a 1 out of 2 structure to get a WH protocol. This new protocol will have complexity equivalent to running $\mathcal{P}$ twice in parallel.

After surveying related work, we give in the following two sections more details on the protocols and the secret sharing schemes we consider. Section 4 then contains the main result and corollaries, and Section 5 describes a nice application of our results to group oriented identification and signatures.

1.1 Related Work

Our techniques are to some extent related to those of De Santis et al. [11]. The models are quite different, however: [11] considers non-interactive proofs of membership, while we consider interactive proofs of knowledge. Also, [11] considers variants of the quadratic residuosity problem, while we consider any problem that affords a protocol of the right form.

In some independent work, De Santis et al. [12] apply techniques similar to ours to proofs of membership in random self-reducible languages. This leads to perfect ZK proofs for monotone Boolean operations over such languages.

In [4], Feige and Shamir introduce the concepts of witness indistinguishable (WI) and witness hiding (WH) protocols and prove the existence of WH protocols for a large class of problems, including the ones we consider (Corollary 4.4). This was done using general zero-knowledge techniques and the assumption that one-way functions exist. Compared to [4], our result shows that if we

start from a proof of knowledge with properties as described above, WH protocols can be constructed much more efficiently and without using computational assumptions.

In [3], a transformation from HVZK proofs was given for protocols including the type we consider. That transformation produced ZK protocols, but on the other hand greatly increased the communication and round complexity so that, contrary to ours, the practical value of that transformation is quite limited. If the target is ZK, however, the increased round complexity seems to be unavoidable.

2 Proofs of Knowledge

Let a binary relation $R = \{(x, w)\}$ be given, for which membership can be tested in polynomial time. For any x, its *witness set* $w(x)$ is the set of w's, such that $(x, w) \in R$.

In the following, we assume that we are given a protocol $\mathcal{P}$, which is a *proof of knowledge* for R, i.e. there is a common input x (of length k bits) to prover P and verifier V and a private input w to P. The prover tries to convince the verifier that $w \in w(x)$. Refer to [5] or [4] for a formal definition.

In order for the constructions in the following to work, $\mathcal{P}$ needs to satisfy a few special properties.

First, we will assume that $\mathcal{P}$ is a three round public coin protocol (although the three round restriction can be removed). Conversations in the protocol will be ordered triples of the form

$$m_1, c, m_2$$

The second message in the protocol is a random bit string c chosen by the verifier. We refer to this as a challenge, and to the prover's final message as the answer.

We also assume that completeness holds for $\mathcal{P}$ with probability 1, i.e. if indeed $w \in w(x)$, then the verifier always accepts.

We assume that $\mathcal{P}$ satisfies knowledge soundness in the following sense: the length of c is such that the number of possible c-values is super-polynomial in k, and for any prover P^*, given two conversations between P^* and V, (m_1, c, m_2) and (m_1, c', m_2'), where $c \neq c'$, an element of $w(x)$ can be computed in polynomial time. We call this the *special soundness property*. It is easily seen to imply the standard soundness definition, which calls for the existence of a knowledge extractor, which can extract a witness in polynomial time from any prover that is successful with non-negligible probability.

Although special soundness is less general than the standard definition, all known proofs of knowledge have this property, or at least a variant where computation of the witness follows from some small number of correct answers. Assuming special soundness is therefore not a serious restriction.

Finally, we assume that $\mathcal{P}$ is *honest verifier zero-knowledge*: there is a simulator S that on input x produces conversations that are indistinguishable from real conversations with input x between the honest prover and the honest verifier. For simplicity we assume perfect indistinguishability in the following; generalization to other flavors of indistinguishability is easy. Most known honest verifier

zero-knowledge protocols in fact satisfy something stronger, viz. that there is a procedure that can take any c as input and produce a conversation indistinguishable from the space of all conversations between the honest prover and verifier in which c is the challenge. We call this *special honest verifier zero-knowledge*.

We will later need the concepts of *witness indistinguishable* (WI) and *witness hiding* (WH) protocols, which were introduced in [4]. Informally, a protocol is witness indistinguishable if conversations generated with the same x but different elements from $w(x)$ have indistinguishable distributions, i.e. even a cheating verifier cannot tell which witness the prover is using. If the problem instance x is generated with a certain probability distribution by a generator G which outputs pairs (x, w) with $w \in w(x)$, we can define the concept of *witness hiding*. A protocol is witness hiding over G, if it does not help even a cheating verifier to compute a witness for x with non-negligible probability when the x is generated by G. We refer to [4] for details.

With respect to the witness indistinguishable property, we can already now note the following:

Proposition 1. *Let $\mathcal{P}$ be a three round public coin proof of knowledge for relation R. If $\mathcal{P}$ is honest verifier zero-knowledge, then $\mathcal{P}$ is witness indistinguishable.*

Proof. We trivially have WI for conversations with the honest verifier: The use of any witness w leads to the distribution produced by the simulator. This implies that the distribution of m_2, given any fixed m_1 and c, is independent of w. The proposition then follows from noting that in conversations with a general verifier, the distribution of m_1, and hence of c, is independent of w.

In many concrete cases, this proposition is not interesting because there is only one witness, in which case WI is trivial and cannot imply anything. Nevertheless, Proposition 1 will be needed in the following for technical reasons.

2.1 An Example

As a concrete example of a protocol with the properties we need, we present Schnorr's protocol from [13] for proving knowledge of a discrete log in a group G of prime order q. Let $g \neq 1$, and let $x = g^w$ be the common input. P is given w as private input. In the language of the above section, the protocol is a proof of knowledge for the relation that consists of pairs $(\ (x, g, G),\ w)$ such that $x = g^w$ in G. Then the protocol works as follows:

1. The prover chooses z at random in $[0..q)$, and sends $a = g^z$ to V.
2. The verifier chooses c at random in $[0..q)$, and sends it to P.
3. P sends $r = (z + cw) \bmod q$ to V, and V checks that $g^r = a\, x^c$.

Completeness trivially holds with probability 1. Correct answers to two different c-values give two equations $r_1 = z + wc_1 \bmod q$ and $r_2 = z + wc_2 \bmod q$ so we find that $w = (r_1 - r_2)/(c_1 - c_2) \bmod q$. So special soundness holds also. Finally, note that by choosing c and r at random, we can make a simulated conversation $(g^r x^{-c}, c, r)$ between the honest verifier and prover. Since c can be chosen freely, we even get special honest verifier zero-knowledge.

3 Secret Sharing

A secret sharing scheme is a method by which a secret s can be distributed among n participants, by giving a *share* to each participant. The shares are computed in such a way that some subsets of participants can, by pooling their shares, reconstruct s. These subsets are called *qualified* sets. Participants forming a non-qualified set should be able to obtain no information whatsoever about s. Such a secret sharing scheme is called *perfect*.

The collection of qualified sets is called the *access structure* for the secret sharing scheme. Clearly if participants in some set can reconstruct s, so can any superset, and therefore in order for the scheme to make sense, it must be the case that if A is a qualified set, then any set containing A is also qualified. An access structure with this property is called *monotone*.

A special case of monotone access structures is structures containing all subsets larger than some threshold value. Such structures are called *threshold structures*.

Any monotone access structure has a natural dual structure. This concept was first defined in [15].

Definition 2. *Let Γ be an access structure containing subsets of a set M. If $A \subseteq M$, then $\bar{A}$ denotes the complement of A in M. Now Γ^*,* the dual access structure *is defined as follows:*

$$A \in \Gamma^* \Leftrightarrow \bar{A} \notin \Gamma.$$

The next propositions follow directly from the definition.

Proposition 3. *The dual Γ^* of a monotone access structure is monotone as well, and satisfies*

$$(\Gamma^*)^* = \Gamma.$$

Furthermore, if Γ is a threshold structure, then so is Γ^*.

Proposition 4. *Let Γ be monotone. A set is qualified in Γ exactly when it has a non-empty intersection with every qualified set in Γ^*.*

In the next section, we will assume we are given a protocol of the form described in Section 2. For each input length k we will assume we are given a monotone access structure $\Gamma(k)$ on n participants, where $n = n(k)$ is a polynomially bounded function of k. Thus we have a *family of access structures*

$\{\Gamma(k) |\ k = 1, 2, \ldots\}$ We can then build a new protocol for proving statements on n problem instances provided we have a perfect secret sharing scheme $\mathcal{S}(k)$ for $\Gamma(k)^*$ satisfying certain requirements to be defined below.

Let $D(s)$ denote the joint probability distribution of all shares resulting from distributing the secret s. For any set A of participants, $D_A(s)$ denotes the restriction of $D(s)$ to shares in A. As $\mathcal{S}(k)$ is perfect, $D_A(s)$ is independent from s for any non-qualified set A. So we will write D_A instead of $D_A(s)$, whenever A is non-qualified. The requirements then are:

1. All shares generated in $\mathcal{S}(k)$ have length polynomially related to k.
2. Distribution and reconstruction of a secret can be done in time polynomial in k.
3. Given secret s and a full set of n shares, one can test in time polynomial in k that the shares are all consistent with s, i.e. that all qualified sets of shares determine s as the secret.
4. Given any secret s, a set of shares for participants in a non-qualified set A (distributed according to D_A) can always be completed to a full set of shares distributed according to $D(s)$ and consistent with s. This completion process can be done in time polynomial in k.
5. For any non-qualified set A, the probability distribution D_A is such that shares for the participants in A are independent and uniformly chosen.

Definition 5. *A perfect secret sharing scheme satisfying requirements 1–4 is called* semi-mooth. *If, in addition, requirement 5 is satisfied it is called* smooth.

It is natural to ask if for any family of monotone access structures there is a family of smooth secret sharing schemes. This question is easy to answer in case of threshold structures. In that case it is clear that Shamir's secret sharing scheme [14] can be used. This scheme is even *ideal*, i.e. the shares are of the same length as the secret. Given d or more shares, the secret s can be found, whereas with $d-1$ or fewer shares, s is completely unknown.

The following alternative to Shamir's scheme (which is also ideal) can lead to more efficient protocols than Shamir's when used in our construction (Theorem 8) with a threshold structure where $d < n/2$.

Again $s \in GF(q)$ is the secret, but the i-th share now is a number $c_i \in GF(q)$, $1 \leq i \leq n$, such that $B\boldsymbol{c} = s\boldsymbol{e_1}$. Here, B is a $n-d+1$ by n matrix over $GF(q)$, $\boldsymbol{c} = (c_1, \ldots, c_n)$, and $\boldsymbol{e_1} = (1, 0, \ldots, 0)$ is a vector of length $n-d+1$. Matrix B should be such that any $n-d+1$ columns are linearly independent (which implies that the rank of B is equal to $n-d+1$). An appropriate choice for B is therefore the first $n-d+1$ rows of a Vandermonde matrix over $GF(q)$, say:

$$B = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & 2 & \cdots & n \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 2^{n-d} & \cdots & n^{n-d} \end{pmatrix}.$$

The secret s can be recovered from any d shares as follows. Since $B\boldsymbol{c} = s\boldsymbol{e_1}$, it follows that $s = \sum_{i=1}^{n} c_i$. Furthermore, when d entries of $\boldsymbol{c}$ are known, the

remaining $n-d$ entries follow uniquely from the equation $B'\boldsymbol{c}=\mathbf{0}$, where B' is the matrix B with the first row removed and $\mathbf{0}$ denotes a vector of $n-d$ zeros. This is true because B' is a $n-d$ by n matrix for which any $n-d$ columns are linearly independent. In case less than d shares are known, the remaining shares can be chosen such that any secret is matched.

For families of access structures other than threshold ones, the answer to the question on existence of smooth secret sharing schemes depends on whether the parameter n is a constant, or is allowed to increase polynomially as a function of k.

In case n is a constant, there exists a smooth secret sharing scheme for any monotone access structure. For any minimal qualified set A, we do the following: choose $s_1,\ldots,s_{|A|}$ at random under the condition that $s_1\oplus\cdots\oplus s_{|A|}=s$, and give one s_i to each participant in A. This scheme was first proposed in [9].

It is easy to check that this scheme is smooth. In particular, the size of shares and the work needed in this scheme is linear in k, but the constant involved depends of course on n and on the access structure. However, the number of possible subsets is exponential in n, so for non-constant n this scheme will not necessarily be smooth.

For non-constant n, it is an open question whether there are secret sharing schemes of the kind we need for any sequence of access structures. Benaloh and Leichter [1] have proposed secret sharing schemes for more general access structures defined by monotone formulae, i.e. Boolean formulae containing only AND and OR operators.

Consider a monotone formula F with n variables. Any subset A of n participants corresponds in a natural way to a set of values of the n variables by assigning a variable to each participant and let each variable be 1 if the corresponding participant is in A and 0 otherwise. We let $F(A)$ be the bit resulting from evaluating F on inputs corresponding to A. Then we can define an access structure Γ_F by

$$A\in\Gamma_F \Leftrightarrow F(A)=1$$

We let F^* denote the *dual formula*, which results from replacing in F all AND operators by OR's and vice versa. It is not hard to show the following proposition.

Proposition 6. *If F is monotone then Γ_F is also monotone. Conversely, for any monotone access structure Γ, there is a monotone formula F, such that $\Gamma=\Gamma_F$. We have that $(\Gamma_F)^*=\Gamma_{F^*}$.*

In [1], a generic method is given that, based on any monotone formula F, builds a perfect secret sharing scheme for the access structure Γ_F. The formula F may contain general threshold operators, in addition to simple AND and OR operators. For a polynomial size formula, it can be shown that the secret sharing scheme from [1] satisfies all of the above requirements except possibly requirement 5. This leads to:

Proposition 7. *Let $\{\Gamma(k)\}$ be a family of access structures such that $\Gamma(k)=\Gamma_{F_k}$ for a family of polynomial size monotone formula $\{F_k\}$. Then there exists a family of semi-smooth secret sharing schemes for $\{\Gamma(k)\}$.*

A final comment before we go on to the main result is that we will need to distribute secrets of length $t = t(k)$ bits, where t is polynomially bounded in k. This does not impose any restrictions on $\mathcal{S}(k)$ because any secret sharing scheme can distribute secrets of any length by running an appropriate number of copies of the scheme in parallel. We therefore assume that $\mathcal{S}(k)$ always distributes secrets of length t. Note that, if n is constant as a function of k, only one access structure and secret sharing scheme are involved.

4 Main Result

The next theorem describes the construction of a proof of knowledge from a basic proof of knowledge $\mathcal{P}$ for a relation R and a family of secret sharing schemes. In the constructed proof of knowledge both prover and verifier are probabilistic polynomial time machines, using the prover and verifier of $\mathcal{P}$, respectively, as subroutines.

For the statement of the result we need some notation. Let $\Gamma = \{\Gamma(k)\}$ be a family of access structures on $n(k)$ participants. Then R_Γ is a relation defined by the following condition: $((x_1, ..., x_m), (w_1, ..., w_m)) \in R_\Gamma$ iff all x_i's are of the same length, say, k bits, $m = n(k)$, and the set of indices i for which $(x_i, w_i) \in R$ corresponds to a qualified set in $\Gamma(k)$. In a proof of knowledge for relation R_Γ the prover thus proves to know witnesses to a set of the x_i's corresponding to a qualified set in $\Gamma(k)$.

Theorem 8. *Let $\mathcal{P}$ be a three round public coin, honest verifier zero-knowledge proof of knowledge for relation R, which satisfies the special soundness property. Let $\Gamma = \{\Gamma(k)\}$ be a family of monotone access structures and let $\{\mathcal{S}(k)\}$ be a family of smooth secret sharing schemes such that the access structure of $\mathcal{S}(k)$ is $\Gamma(k)^*$. Then there exists a three round public coin, witness indistinghuisable proof of knowledge for relation R_Γ.*

Proof. To improve readability we drop in the following the dependency on k from the notation, and write $\mathcal{S} = \mathcal{S}(k)$, $\Gamma = \Gamma(k)$ and $n = n(k)$. We will distribute secrets of length t in $\mathcal{S}$. If the length of any share resulting from this is larger than t, we will replace $\mathcal{P}$ by a number of parallel executions of $\mathcal{P}$ to make sure that a challenge is at least as long as any share.[2] Note that this does not violate the honest verifier zero-knowledge nor the special soundness property. A basic idea in the following will be to interpret a challenge as a share. If challenges are longer than shares, we will simply take the first appropriate number of bits of the challenge to be the corresponding share. If c is a challenge, $share(c)$ will denote the corresponding share.

The following now describes the new protocol, in which $A \in \Gamma$ denotes the set of indices i for which P knows a witness for x_i:

[2] For some secret sharing schemes, there is a lower bound on the length of shares in terms of n. For Shamir's scheme, the length of shares is at least $\log_2(n+1)$. If t is smaller than this bound, we can again replace $\mathcal{P}$ by a number of parallel executions.

1. For each $i \in \overline{A}$, P runs simulator S on input x_i to produce conversations (m_1^i, c_i, m_2^i). For each $i \in A$, P determines m_1^i as what the prover in $\mathcal{P}$ would send as m_1 given a witness for input x_i. P then sends the values $m_1^i,\ i = 1, \ldots, n$ to V.
2. V chooses a t-bit string s at random and sends it to P.
3. Consider the set of shares $\{share(c_i) | i \in \overline{A}\}$ that correspond to the c_i from the simulation in Step 1. As $\overline{A}$ is non-qualified in Γ^*, requirement 4 guarantees that P can complete these shares to a full set of shares consistent with s. P then forms challenges c_i for indices $i \in A$, such that $share(c_i)$ equals the share produced in the completion process. This is done by simply copying the bits of the shares and padding with random bits if necessary. In Step 1, S has produced a final message m_2^i in $\mathcal{P}$ for $i \in \overline{A}$. For $i \in A$, P knows a witness for x_i, and can therefore find a valid m_2^i for m_1^i and c_i by running the prover's algorithm from $\mathcal{P}$. Finally, P sends the set of messages c_i, m_2^i, $i = 1, \ldots, n$ to V.
4. V checks that all conversations (m_1^i, c_i, m_2^i) now produced would lead to acceptance by the verifier in $\mathcal{P}$, and that the shares $share(c_i)$ are consistent with secret s. He accepts if and only if these checks are satisfied.

It is clear from the assumptions on $\mathcal{S}$ that P and V need only poly-time and access to the prover and verifier of $\mathcal{P}$. It therefore remains to be seen that the protocol is a proof of knowledge and that it is witness indistinguishable.

Completeness is trivially seen to hold by inspection of the protocol. For *soundness*, assume that some prover P^* for a given first message $\{m_1^i |\ i = 1, \ldots, n\}$ can answer correctly a non-negligible fraction of the possible choices of s. This means that by rewinding P^*, we can efficiently get correct answers to two different values, say s and s'.[3] Let the shares of s and s' sent in the protocol be $share(c_i)$ and $share(c_i')$, $i = 1, \ldots, n$, respectively. Then for every qualified set $B \in \Gamma^*$, there must be an $i \in B$, such that $share(c_i) \neq share(c_i')$ since otherwise it would follow that $s = s'$. But then we also have that $c_i \neq c_i'$ and so by assumption on $\mathcal{P}$, we can compute a witness for x_i. So P^* knows a witness in every qualified set of Γ^*. On account of Proposition 4 the set of witnesses we thus extract is a qualified set in the access structure Γ.

As for *witness indistinguishability*, we have to show that the distribution of the conversation is independent of which qualified set $A \in \Gamma$ the prover uses. First observe that the distribution of each m_1^i depends only on x_i and equals the distribution of the prover's first message in an execution of $\mathcal{P}$ with x_i as input. This follows from Proposition 1, using that $\mathcal{P}$ is honest verifier zero-knowledge. In particular, the joint distribution of the m_1^i's, and hence the verifier's choice of s, is independent of A.

Since the set $\{share(c_i)\}$ is constructed by completing a set of uniformly distributed shares in a non-qualified set of $\mathcal{S}$, the joint distribution of the $share(c_i)$'s is simply $D(s)$. Since the c_i's are constructed from the shares by possibly padding with random bits, the joint distribution of the c_i's is independent of A. Fi-

[3] There are 2^t possible s-values which is super-polynomial in k, whence any polynomial fraction of these contain at least 2 values for all large enough k.

nally, Proposition 1 implies that the distribution of each m_2^i depends only on x_i, m_1^i and c_i, and is therefore also independent of A.

Remark. If the secret sharing schemes are ideal, the communication complexity of the protocol in Theorem 8 is at most t bits plus n times that of $\mathcal{P}$. Note that instead of taking several instances of the same proof of knowledge, it is also possible to combine different proofs of knowledge. In this way, one may for instance prove knowledge of either a discrete log or an RSA root without revealing which.

Theorem 9. *As Theorem 8, but with* $\mathcal{P}$ special *honest verifier zero-knowledge and* $\mathcal{S}(k)$ semi-*smooth.*

Proof. In this case the protocol from Theorem 8 is changed as follows. In Step 1, the prover uses $\mathcal{S}$ to distribute an arbitrary secret, and discards all shares in A. The remaining shares are distributed according to $D_{\overline{A}}$. He then runs the special simulator on the corresponding challenges. Note that the completion process can still be performed on account of requirement 4, and as before, the honest prover can counter any challenge s by the verifier. Soundness is proven in the same way as before. Therefore, the modified scheme still constitutes a proof of knowledge for relation R_Γ.

As for witness indistinguishability, we only have to note that the distribution of any m_1^i generated by the (special) simulator is the same for any particular challenge value c_i used, because m_1^i in a real execution of $\mathcal{P}$ is independent of the challenge. Therefore the joint distribution of the m_1^i's is the same as in the case of Theorem 8. The rest of the proof is therefore the same as for Theorem 8.

The witness indistinguishable property of the protocol from Theorem 8 leads us to a generalization of Theorem 4.3 of [4]. To state the result, we need to introduce the concept of an *invulnerable generator* G for a relation R. Such generators were first introduced in [6] and later used in slightly modified form in [4]. Such a generator is a probabilistic polynomial time algorithm which outputs a pair $(x, w) \in R$. The generator is invulnerable if no probabilistic polynomial time enemy given only x can compute an element in $w(x)$ with non-negligible probability, taken over the coin flips of both G and the enemy.

Thus, asserting the existence of an invulnerable generator for a relation is a way of stating that it is feasible to generate hard, solved instances of the underlying computational problem.

For any generator G, we let G^n denote the generator that produces an n-tuple of pairs in R by running G independently n times in parallel. We will also need some notation for access structures: for a monotone access structure Γ, we let the sets in Γ correspond to subsets of the index set $N = \{1, ..., n\}$. Now let the set $I_\Gamma \subseteq N$ be defined by: $i \in I_\Gamma$ iff i is contained in every qualified set in Γ. It is easy to see by monotonicity of Γ that $i \in I_\Gamma$ precisely if $N \setminus \{i\}$ is not qualified (using Proposition 4).

Theorem 10. *Let $\mathcal{P}$ be a witness indistinguishable proof of knowledge for the relation R_Γ, where $\Gamma = \{\Gamma(k)\}$ is a family of monotone access structures on $n(k)$ participants, and R is a binary relation. If for all k, $\Gamma(k)$ contains at least two different minimal qualified sets, and there is an invulnerable generator G for R, then $\mathcal{P}$ is witness hiding over $G^{n(k)}$.*

Proof. We follow the line of reasoning from Thm. 4.3 of [4]. Suppose we are given an probabilistic polynomial time enemy $\mathcal{A}$ that has non-negligible probability of computing a witness, using the honest prover in the scheme from Theorem 8 as a subroutine. We show that $\mathcal{A}$ can be compiled into an algorithm that solves with non-negligible probability random instances x generated by G, thus contradicting the invulnerability of the generator (see [4]).

From the assumption on $\Gamma(k) = \Gamma$ (at least two minimal qualified sets) it follows that $N \setminus I_\Gamma$ must contain at least two elements, and that I_Γ is not qualified.

Our compilation now works as follows:

1. Recall that our input is a problem instance x generated by G. We now form an n tuple of instances $(x_1, ..., x_n)$ as follows: choose at random $j \in N$, and let $x_j = x$. For all other indices i, run G to produce a solved instance x_i and save the witness w_i.
2. Give $x_1, ..., x_n$ as input to $\mathcal{A}$. When $\mathcal{A}$ needs to interact with the prover, we simply simulate the prover's algorithm in $\mathcal{P}$. If $j \notin I_\Gamma$, this can be done, since then $N \setminus \{j\}$ is qualified and we know witnesses of all instances except x_j. If $j \in I_\Gamma$, we fail and stop.
3. If $\mathcal{A}$ is successful, it outputs a witness for the relation R_Γ which by definition is a set of witnesses $\{w_i\}$ corresponding to a qualified set A in Γ. If $j \in A$, we have success and can output w_j. Else output something random.

We now show that this compilation finds a witness for x with non-negligible probability. It is sufficient to show that we find a witness with non-negligible probability given that $j \notin I_\Gamma$ since this happens with probability at least $2/n$. Now note that the joint distribution of the x_i's we give to $\mathcal{A}$ is the same as in an ordinary interaction with the prover. Therefore $\mathcal{A}$ is successful with non-negligible probability by assumption. We therefore only have to bound the probability that j is in A, the set of witnesses we get from $\mathcal{A}$. Since I_Γ is not qualified, A must contain at least one index not in I_Γ. By witness indistinguishability, $\mathcal{A}$ has no information about which j in $N \setminus I_\Gamma$ we have chosen, and so the probability that $j \in A$ is at least $1/|N \setminus I_\Gamma|$. Hence if $\mathcal{A}$ has success probability ϵ, our success probability given that $j \notin I_\Gamma$ is at least ϵ/n, which is non-negligible.

Note that an access structure has at least two minimal qualified sets exactly when the corresponding minimal CNF-formula contains at least one OR-operator.

Note also that this result only shows that an enemy cannot compute a complete qualified set of witnesses. It does not rule out that the protocol could help him to compute a small, non-qualified set. Ideally, we would like to prove that

the enemy cannot compute even a single witness. With a stronger assumption on the access structure, this can be done:

Corollary 11. *Let $\mathcal{P}$ be a witness indistinguishable proof of knowledge for the relation R_Γ, where $\Gamma = \{\Gamma(k)\}$ is a family of monotone access structures on $n(k)$ participants, and R is a binary relation. Suppose that for all k the set $I_{\Gamma(k)}$ is empty. Suppose finally that there is an invulnerable generator G for R, and that inputs for $\mathcal{P}$ are generated by $G^{n(k)}$. Then no probabilistic polynomial time enemy interacting with the honest prover can with non-negligible probability compute a witness for any of the x_i in the input to the protocol.*

Proof. Since $I_{\Gamma(k)}$ is non qualified, there are at least two minimal sets, and therefore the proof is the same as for Theorem 10, except that it follows from the assumption that the index j is always chosen among all indices. Hence if the enemy outputs at least one correct witness, there is a non-negligible probability of at least $1/n$ that this is the witness we are looking for.

A certain special case of Theorem 8 is interesting in its own right:

Corollary 12. *Let $\mathcal{P}$ be a three round public coin, honest verifier zero-knowledge proof of knowledge for relation R, which satisfies the special soundness property. Then for any n, d there is a protocol with the same round complexity as $\mathcal{P}$ in which the prover shows that he knows d out of n witnesses without revealing which d witnesses are known.*

Proof. Use Theorem 8 with, for example, Shamir's secret sharing scheme for $\mathcal{S}$ and a threshold value of $n - d + 1$.

Corollary 13. *Consider the protocol guaranteed by Corollary 12, let $n = 2$ and $d = 1$, i.e. the prover proves that he knows at least 1 out of 2 solutions. For any generator G generating pairs in R, this protocol is witness hiding over G^2.*

Proof. Since protocols constructed from Theorem 8 are always witness indistinguishable, we can use Theorem 4.2 of Feige and Shamir[4].

Note that for this corollary, we do not need the assumption that G is invulnerable, as in Theorem 10.

To build the protocol of Corollary 13, we need a 2 out of 2 threshold scheme. Such a scheme can be implemented by choosing random shares c_1, c_2 such that $c_1 \oplus c_2$ equals the secret. Therefore, in the simple case of Corollary 13, the protocol constructed by Theorem 8 simply becomes a game where the verifier chooses a random s, and the prover shows that he can answer correctly a pair of challenges c_1, c_2, such that $s = c_1 \oplus c_2$. In the prover's final message, he only has to send c_1 because the verifier can then compute c_2 himself. Hence the communication complexity of the new protocol is exactly twice that of $\mathcal{P}$, whence the new protocol is just as practical as $\mathcal{P}$.

Corollary 14. *Let $\{\Gamma(k) = \Gamma_{F_k}\}$ be a family of monotone access structure on $n(k)$ participants defined by a polynomial size family of formulas $\{F_k\}$, and let $\mathcal{P}$ be a three round public coin,* special *honest verifier zero-knowledge proof of knowledge for relation R, which satisfies the special soundness property. Then there exists a witness-indistinguishable proof of knowledge for relation $R_{\Gamma(k)}$. Let $M(k)$ be the maximal number of occurrences of a variable in $F(k)$. Then the communication complexity of the new protocol is at most $nM(k)$ times that of $\mathcal{P}$ plus t bits.*

Proof. By Proposition 3, $\Gamma(k)^* = \Gamma_{F_k^*}$, and since the size of F_k^* is the same as that of F_k, we can use the secret sharing scheme guaranteed by Proposition 7 when we do the construction of Theorem 8. The statement on the communication complexity follows from the fact that the shares of the secret sharing scheme constructed in [1] from $F(k)$ have maximal size $tM(k)$ bits, so that we have to use $M(k)$ parallel executions of $\mathcal{P}$ in the construction of Theorem 8.

5 Application to Identification and Signatures

Suppose we have n users, for example employees of a company, such that the i-th user has a public key x_i and secret key $w_i \in w(x_i)$. Suppose also that certain subsets of users are qualified in the sense that they are allowed to initiate certain actions, sign letters on behalf of the company, etc. This defines an access structure on the set of users. Theorem 8 now gives a way in which a subset of users can collaborate to identify themselves as a qualified subset, without revealing anything else about their identities. This makes good sense, if they are to assume responsibility on behalf of the company, rather than personally.

This also extends to digital signatures, since by using a hash function, any three round proof of knowledge as the one produced by Theorem 8 can be turned into a signature scheme by computing the challenge as a hash value of the message to be signed and the prover's first message (this technique was introduced in [5]). By this method, a signature can be computed which will show that a qualified subset was present, without revealing which subset was involved. This is a generalization of the results from e.g. [10] and also of the group signature concept, introduced by Chaum and Van Heyst [2]. One aspect of group signatures which is missing here, however, is that it is not possible later to "open" signatures to discover the identities of users involved.

Note also that our method allows participants to form groups completely freely, using the same keys in all groups. For example, two participants who normally use Schnorr signatures individually can go together and form a "1 out of 2" signature without changing their keys or the basic algorithms in which they are used.

6 Open Problems

Two obvious open problems remain. First, can Theorem 8 be proved assuming ordinary soundness of $\mathcal{P}$, and not special soundness? And secondly, can it be generalized to other types of protocols than public coin protocols?

Acknowledgement We thank Douglas Stinson for helping us with information about results on secret sharing schemes, and Matthew Franklin for useful discussions and comments on the presentation.

References

1. J. Benaloh and J. Leichter: *Generalized Secret Sharing and Monotone Functions*, Proc. of Crypto 88, Springer Verlag LNCS series, 25–35.
2. D. Chaum and E. van Heyst: *Group Signatures*, Proc. of EuroCrypt 91, Springer Verlag LNCS series.
3. I. Damgård: *Interactive Hashing can Simplify Zero-Knowledge Protocol Design Without Complexity Assumptions*, Proc. of Crypto 93, Springer Verlag LNCS series.
4. U. Feige and A. Shamir: *Witness Indistinguishable and Witness Hiding Protocols*, Proc. of STOC 90.
5. U. Feige, A. Fiat and A. Shamir: *Zero-Knowledge Proofs of Identity*, Journal of Cryptology 1 (1988) 77–94.
6. M. Abadi, E. Allender, A. Broder, J. Feigenbaum and L. Hemachandra: *On Generating Solved Instances of Computational Problems*, Proc. of Crypto 88, Springer Verlag LNCS series.
7. S. Goldwasser, S. Micali and C. Rackoff: *The Knowledge Complexity of Interactive Proof Systems*, SIAM Journal on Computing 18 (1989) 186–208.
8. L. Guillou and J.-J. Quisquater: *A Practical Zero-Knowledge Protocol fitted to Security Microprocessor Minimizing both Transmission and Memory*, Proc. of EuroCrypt 88, Springer Verlag LNCS series.
9. M. Ito, A. Saito, and T. Nishizeki: *Secret Sharing Scheme Realizing any Access Structure*, Proc. Glob.Com. (1987).
10. T.Pedersen: *A Threshold Cryptosystem without a Trusted Third Party*, Proc. of EuroCrypt 91.
11. A. De Santis, G. Di Crescenzo and G. Persiano: *Secret Sharing and Perfect Zero-Knowledge*, Proc. of Crypto 93, Springer Verlag LNCS series.
12. A. De Santis, G. Persiano, M. Yung: *Formulae over Random Self-Reducible Languages: The Extended Power of Perfect Zero-Knowledge*, manuscript.
13. C.P. Schnorr: *Efficient Signature Generation by Smart Cards*, Journal of Cryptology 4 (1991) 161–174.
14. A. Shamir: *How to Share a Secret*, Communications of the ACM 22 (1979) 612–613.
15. G.J. Simmons, W.A. Jackson and K. Martin: *The Geometry of Shared Secret Schemes*, Bulletin of the Institute of Combinatorics and its Applications 1 (1991) 71–88.