Verifiable Encryption, Group Encryption, and their Applications to Separable Group Signatures and Signature Sharing Schemes

Jan Camenisch
IBM Research
Zurich Research Laboratory
CH–8803 Rüschlikon
jca@zurich.ibm.com

Ivan Damgård
Department of Computer Science, BRICS*
University of Aarhus
DK-8000 Aarhus C, Denmark
ivan@daimi.au.dk

Abstract

We generalize and improve the security and efficiency of the verifiable encryption scheme of Asokan et al., such that it can rely on more general assumptions, and can be proven secure without assuming random oracles. We extend our basic protocol to a new primitive called verifiable group encryption. We show how our protocols can be applied to construct group signatures, identity escrow, and signature sharing schemes from a wide range of signature, identification, and encryption schemes already in use. In particular, we achieve perfect separability for all these applications, i.e., all participants can choose their signature and encryption schemes and the keys thereof independent of each other, even without having these applications in mind.

1 Introduction

A *verifiable encryption scheme* is in its basic form a two-party protocol between a prover P and a verifier V. Their common inputs are a public encryption key E, a public value x, and a binary relation $\mathcal{R}$. As a result of the protocol, V either rejects or obtains the encryption of some value w under E such that $(x, w) \in \mathcal{R}$ holds. For instance, $\mathcal{R}$ could be defined such that $(x, w) \in \mathcal{R}$ if and only if w is a signature on message x w.r.t. to some fixed public key. In other words, P claims to have given V the encryption of a valid signature on x.

The protocol should ensure that V accepts an encryption of an invalid w with only negligible probability. Moreover, V should learn nothing except the encryption of w and the fact that w is valid w.r.t. x. In particular, if the encryption scheme is semantically secure, the protocol should be zero-knowledge.

The encryption key E can belong to P, but typically belongs to a third party in which case the third party should not need to take part in the protocol, in other words, P does not need to know the secret key corresponding to E.

Verifiable encryption schemes are employed in many cryptographic protocols (although the term "verifiable encryption" is not always used). Examples are digital payment systems with revocable anonymity (e.g., [7, 23]), verifiable signature sharing (e.g., [24]), (publicly) verifiable secret sharing (e.g., [35]), escrow schemes [32, 37], or fair exchange of signatures [1, 2, 4]. However, only the schemes presented in [2, 28, 36] do not apply ad-hoc constructions using a

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

specific encryption scheme that suits the particular application; in fact, our protocols can be seen as a generalization of the protocols employed in [2, 28, 36].

The concept of verifiable encryption was introduced in [35] in the context of publicly verifiable secret sharing schemes, and in a more general form in [2], for the purpose of fair exchange of signatures.

The idea here is that if A and B wish to exchange their signatures on some message, they will first exchange verifiable encryptions of them, using as E the public key of some trusted third party. If this was successful, it will be safe for A to just reveal his signature to B. Even if B never answers, A can get B's signature by having the trusted party decrypt it. (Some care needs to be taken here to avoid that one party falsely accuses the other of cheating, see [2] for details).

Micali [29] also proposes the use of provable encryption of data for third parties to solve several variants of the fair exchange problem.

In this paper, we first show how to modify and generalize the verifiable encryption scheme from [2] to achieve the following:

- The relation $\mathcal{R}$ can be any relation possessing a three-move proof of knowledge which is an Arthur-Merlin game, i.e., as the second message, the verifier sends a random challenge. This proof should be honest-verifier zero-knowledge, and a cheating prover should be unable to answer more than one challenge correctly.
- It can be based on any public-key encryption scheme.
- The verifier needs to store only $O(\log k)$ encryptions of the underlying encryption scheme.
- In its interactive form, our scheme can be proved secure without relying on random oracles.

In comparison, the scheme from [2] only works for relations $\mathcal{R}$ containing pairs of form $(x, f^{-1}(x))$, where f is a one-way group homomorphism and the verifier is required store k encryptions of the underlying encryption scheme. Finally, the scheme from [2] is only secure in the random oracle model, even in its interactive form — in other words, it relies on random oracles for more than just standard removal of interaction *à la* Fiat-Shamir[1].

Our results are especially suited for a situation where a public-key infrastructure already exists, i.e., users already have (certified) public keys for encryption, signature, or identification schemes. However, we assume that these keys are not necessarily generated with other and more advanced primitives in mind, such as group signatures, identity escrow, fair contract signing, or blind signatures with revocable anonymity. We believe this is a very realistic scenario.

We provide examples of how our verifiable encryption schemes can be applied in this scenario to build such advanced primitives, where their security can be proved based only on security of the existing infrastructure: No special assumptions are needed on the existing encryption scheme. The signature scheme only needs to satisfy that one can prove knowledge of a signature on a given message by a so-called Σ-protocol. All standard signature schemes (RSA, El-Gamal, DSS) satisfy this. Any identification scheme that is a three-move honest-verifier proof of knowledge is also suitable.

All our solutions for these applications possess *perfect separability*, i.e., all participants can choose (or renew) their keys independently of each other or even use different kinds of encryption

[1] We are referring here to the proceedings version of [2]. Having been informed of our results, the authors of [2] later modified their protocols such that they do not rely on random oracles. These protocols appear in the journal version of their paper [3].

and signature schemes which is a requirement met for the first time for these applications. The term separability originates from [28] and was diversified in [9]. There exist weaker forms of separability, i.e., *partial* separability allows only a subset of the participants to choose their keys independently whereas *weak* separability requires that all participants use common system parameters, e.g., they must all use DSS and choose their keys from the same algebraic group (cf. [9]). Clearly, both these forms of separability are not sufficient if one accepts no special purpose key-setup procedure.

Furthermore, we introduce a new primitive called *verifiable group encryption* involving $n > 1$ third parties (called *proxies*) where only certain subsets of them can jointly decrypt the secret. This new primitive is an extension of our verifiable encryption scheme: we are given n public encryption keys $E_1, \ldots, E_n$ and the prover and the verifier agree on any monotone access structure over $\{1, \ldots, n\}$. Then, if V accepts, he is convinced that he has obtained an encryption of a valid secret which a subset of the proxies can decrypt if and only if that subset is contained in the access-structure. For example, one can decide that for some $t < n$, any subset of at least t players can decrypt, whereas less than t players cannot. This notion of group encryption should not be confused with the notion of threshold encryption [18]. In the latter case, a number of parties publish a *single* public key and the access structure is determined by these parties during the setup of the system. In contrast, a group encryption scheme allows to choose a (possibly different) access structure each time when (verifiably) encrypting a message. Another distinguishing feature is that a group encryption enables the proxies to choose their encryption keys independently of each other and allows them even to use different encryption schemes (perfect separability).

We also show how to get verifiable signature sharing from verifiable group encryption, yielding more general solutions for this problem than what were previously known and thereby solve an open problem raised in [11].

We believe that our verifiable encryption primitives facilitate the design of provably secure protocols for many applications such that their setup remains minimal and perfectly separable, i.e., protocols that are tailored for a public key infrastructure as described above. Previously, one had to resort to general zero-knowledge techniques for such solutions and thus accept a prohibitive loss of efficiency. Also, much more efficient schemes can sometimes be obtained by ad-hoc constructions, but this requires relying on particular properties of the encryption and/or signature scheme involved, and sometimes means that no proofs of security can be given, and always means that separability cannot be provided. For a comparison to some concrete previous schemes of this type, please refer to Section 5.2.

2 Preliminaries

2.1 Σ-Protocols

A Σ-protocol [14, 16] for a boolean relation $\mathcal{R} \subseteq \{0,1\}^* \times \{0,1\}^*$ is a three-move honest-verifier zero-knowledge proof of knowledge for $\mathcal{R}$. That is, a string x is the common input to a prover P and a verifier V, and P demonstrates knowledge of a w such that $(x, w) \in \mathcal{R}$. We call w a *witness* for x, and the set of x's that have witnesses is called $L_{\mathcal{R}}$. Of course, carrying out a Σ-protocol for some relation R makes sense only if it is a hard relation for (at least) one of the two involved parties. A Σ-protocol can be defined by three (probabilistic) procedures σ_t, σ_s, and σ_v as follows. On input x and w, the prover uses σ_t to compute a so-called *commitment* t and some side-information r. The prover sends the verifier t as the first message. Then, the verifier sends back a random bit string c, called a *challenge*. The prover uses x, w, r, and c as

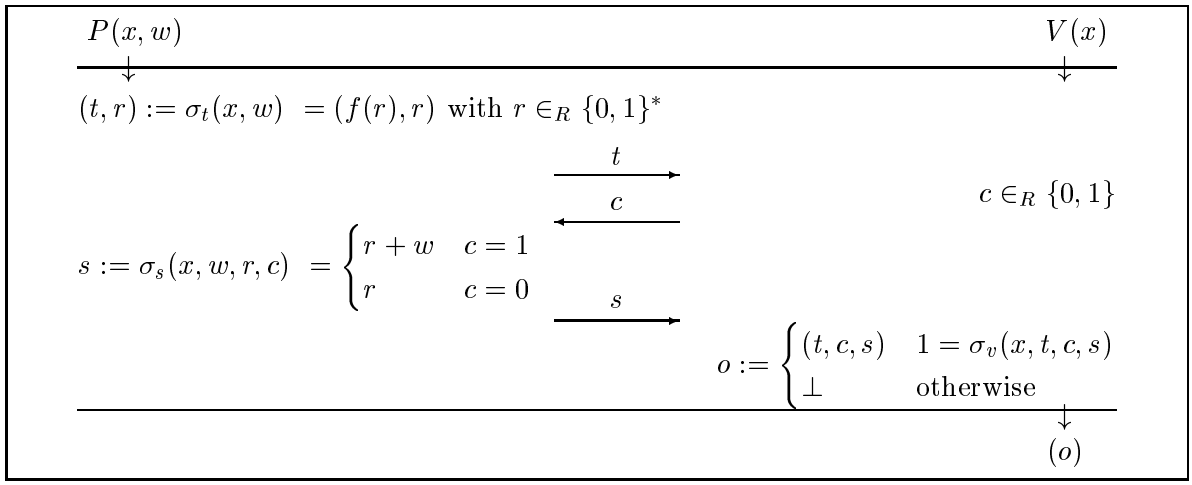

Figure 1: An example of a Σ-protocol: a proof of knowledge of a pre-image under a (one-way) group homomorphism $f(\cdot)$, i.e., $\mathcal{R} = \{(x, f^{-1}(x))\}$. The predicate $\sigma_v(x,t,c,s)$ equals 1 if and only if $f(s) = t \cdot x^c$ holds.

input to σ_s to compute the *response* s which he sends the verifier. Finally, σ_v is a predicate taking x, t, c, and s as input that V uses to check whether s is a valid response, i.e., V accepts if $\sigma_v(x,t,c,s) = 1$ holds. A triple (t,c,s) such that $\sigma_v(x,t,c,s) = 1$ is called an *accepting* triple for x.

We require that if P and V follow the protocol, V always accepts, whereas a cheating prover can answer at most one challenge correctly per commitment. More precisely, there is some polynomial-time procedure ρ that, given two accepting triples (t,c_1,s_1) and (t,c_2,s_2) with $c_1 \neq c_2$, computes w such that $(x,w) \in \mathcal{R}$.

We also require that a Σ-protocol is honest verifier perfect zero-knowledge in the particular sense that there is a simulator which, given input x and challenge c, computes a t and an s such that (t,c,s) is accepting w.r.t. x, and has a distribution equal to that of real conversations with the honest verifier where c occurs as challenge[2]. Figure 1 shows an example of a Σ-protocol.

Cramer et al. [16] show that different Σ-protocols can be composed to obtain Σ-protocols for statements such as *"I know a witness to $x_1 \in L_{\mathcal{R}_1}$ or a witness to $x_2 \in L_{\mathcal{R}_2}$"* while retaining efficiency. Note that the zero-knowledge property in particular implies that V does not learn whether P knows a witness to x_1 or to x_2.

More generally, we have the following lemma.

Lemma 1 (Composition of Σ-protocols [16]). *Given Σ-protocols for relations $\mathcal{R}_1, \ldots, \mathcal{R}_n$, then one can construct Σ-protocols for the relations*

- $\mathcal{R}_{\binom{n}{1}} = \left\{ \left((x_1,\ldots,x_n), w\right) \;\middle|\; \exists i \;:\; (x_i,w) \in \mathcal{R}_i \right\}$ *and*
- $\mathcal{R}_\Gamma = \left\{ \left((x_1,\ldots,x_n),(w_1,\ldots,w_n)\right) \;\middle|\; \exists S \in \Gamma \;:\forall i \in S:\; (x_i,w_i) \in \mathcal{R}_i \right\}$,

where Γ is a monotone access structure over $\{1,\ldots,n\}$.

The latter Σ-protocol allows P to demonstrate knowledge of witnesses to a qualified subset of the instances $x_1, \ldots, x_n$ without revealing which subset is involved. The construction of the protocol employs a secret sharing scheme for the access structure Γ^* which is the structure dual

[2]We require perfect zero-knowledge only for simplicity. Computational zero-knowledge could be handled too, however, known examples of Σ-protocols are typically perfect zero-knowledge.

to Γ. Such a scheme always exists since Γ is a monotone access structure. Furthermore, if Γ is a threshold access structure then so is Γ^* in which case there exist efficient secret sharing schemes for Γ^*, e.g., the one by Shamir [34]. We refer to [16] for further details.

Probably, the best-known special case of relations $\mathcal{R}$ with Σ-protocols are public-key identification schemes such as the ones by Feige, Fiat, and Shamir [21], by Guillou and Quisquater [27], or by Schnorr [33]. Furthermore, many proofs of knowledge of or about discrete logarithms found in literature are Σ-protocols. A protocol that is of interest in the context of this paper is a proof of knowledge of a pre-image under a group homomorphism (see Figure 1). It turns out that this protocol is very useful in practice because demonstrating (in zero-knowledge) that one knows a signature on given message reduces to demonstrating that one knows a pre-image under a group homomorphism. This is true for any of the standard signature schemes in use today (RSA, DSA, etc.) (see [2, 24]).

2.2 Probabilistic Encryption Schemes

A triple (G, E, D) of probabilistic polynomial-time algorithms is a polynomially secure public key encryption system (see for instance [26, 30]) if we have the following:

1. For every output $(E, D) \in G(1^k)$ and all messages $m \in \{0,1\}^k$ we have $D(E(m)) = m$.

2. For all probabilistic algorithms T and M, all polynomials $p(\cdot)$, and all sufficiently large k we have

$$Pr[T(1^k, E, m_0, m_1, \alpha) = m : (E, D) := G(1^k); (m_0, m_1) := M(E, 1^k); \\ m \in_R \{m_0, m_1\}; \alpha := E(m)] < \frac{1}{2} + \frac{1}{p(k)} .$$

That is, the adversary gets to see the public key, selects messages m_0 and m_1 in any way he wants, and gets an encryption of either m_0 or m_1. The system is secure if he cannot guess which is the case with probability essentially better than $1/2$.

For convenience, E denotes the public key as well as the actual encryption algorithm, and D the secret key as well as the decryption algorithm. Furthermore, as we often need to make the random choices in the encryption algorithm explicit, we will write sometimes $E(r, m)$ rather than $E(m)$, where r contains all coinflips to be made for encryption; hence E will in these cases denote a "deterministic" algorithm.

3 Verifiable Encryption

3.1 Definition of Verifiable Encryption

We give a definition of a secure verifiable encryption scheme for a relation $\mathcal{R}$ following [2].

Definition 1 (Secure Verifiable Encryption). *Let $\mathcal{R}$ be a binary relation and let $L_{\mathcal{R}} = \{x \mid \exists w : (x, w) \in \mathcal{R}\}$. A* secure verifiable encryption *scheme for a relation $\mathcal{R}$ consists of a two-party protocol (P, V) and a recovery algorithm R. Let $V_P(E, x, k)$ denote the output of V when interacting with P on input (E, x, k), where k is a security parameter. The following properties must hold:*

Completeness: *If P and V are honest then $V_P(E, x, k) \neq \perp$ for all $(E, D) \in G(1^k)$ and for all $x \in L_{\mathcal{R}}$.*

Validity: *For all polynomial time* $\tilde{P}$*, all positive polynomials* $p(\cdot)$*, all sufficiently large* k*, and all* $(E, D) = G(1^k)$ *we have*

$$Pr\left[\big(x, R(D, \alpha)\big) \notin \mathcal{R} \text{ and } \alpha \neq \perp \ : \ \alpha := V_{\tilde{P}}(E, x, k)\right] < \frac{1}{p(k)} \ .$$

Computational Zero-Knowledge: *For every* $\tilde{V}$ *there exists a expected polynomial-time simulator* $S_{\tilde{V}}$ *with black-box access to* $\tilde{V}$ *s.t. for all distinguishers* A*, all positive polynomials* p*, all* $x \in L_{\mathcal{R}}$*, and all sufficiently large* k*, we have*

$$Pr\big[A(E, x, \alpha_i) = i \ : \ (E, D) := G(1^k); \alpha_0 := S_{\tilde{V}}(E, x, k); \\ \alpha_1 := \tilde{V}_P(E, x, k); i \in_R \{0, 1\}\big] < \frac{1}{2} + \frac{1}{p(k)} \ .$$

The completeness property should need no discussion. Validity ensures that it almost never happens that an honest verifier accepts simultaneously with recovery failing to compute a witness for $x \in L_{\mathcal{R}}$. A consequence of this is that a verifiable encryption scheme is necessarily also a proof of knowledge of such a witness. Finally, the zero-knowledge condition assures the prover that no verifier can learn anything beyond the fact that $x \in L_{\mathcal{R}}$ and that a witness for x can be recovered from the output.

We note that computational zero-knowledge is the best we can achieve due to the requirement that a witness is (and should be) recoverable from the conversation. Also note that the encryption scheme must be semantically secure (cf. Section 2.2) in order for the zero-knowledge property to be satisfied.

Our zero-knowledge definition allows the simulator to rewind the verifier. In some applications it may be desirable to require that simulation can be done without rewinding, since this implies that the protocol is *concurrent* zero-knowledge [19], i.e., even an arbitrary interleaving of different instances of the protocol is simulatable. We note that one variant of our protocol in fact has this stronger property.

The verifiable encryption scheme for group homomorphisms described in [2] can be seen as a special case of our definition with respect to the relations $\mathcal{R}$ that are considered. In particular, the relation defined by $\{(x, w) | x = f(w)\}$, where f is a group homomorphism, is a subclass of the relations having a Σ-protocol.

3.2 A Verifiable Encryption Scheme

We first present a very simple but not very efficient scheme, and then move on with improvements.

Let (G, E, D) be a semantically secure crypto system. Let $(E, D) := G(1^k)$ be the public and secret key of a third party. Also, we are given a relation with a Σ-protocol defined by procedures σ_t, σ_s, and σ_v. Assume for simplicity that the verifier can choose only between 0 and 1 as challenges. The idea is now simply that given x, the prover will start a conversation in the Σ-protocol. Using his knowledge of a witness w he can compute answers to both $c = 0$ and $c = 1$, and he supplies encryptions under E of both of these. The verifier can now ask P to open one of these encryptions to check if it contains a valid answer. If this is true for both encryptions, decrypting the other one allows to recover w (due to the properties of the Σ-protocol), whereas if at least one of them contains garbage, the prover will be caught with probability 1/2. Concretely, we have the following.

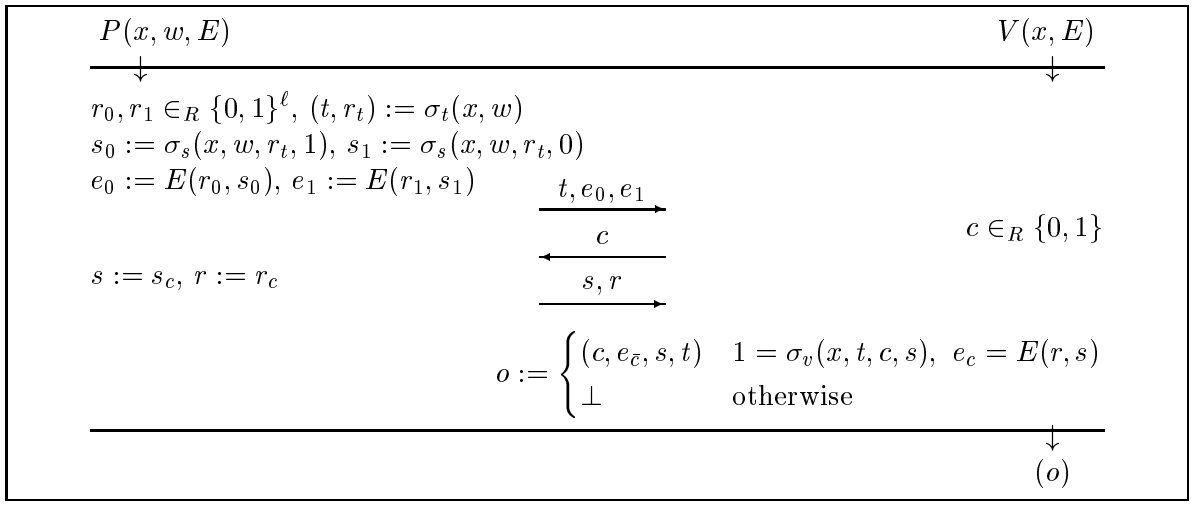

Figure 2: One round of the basic verifiable encryption scheme. Recovery takes place by decrypting $e_{\bar{c}}$ and using the soundness property of the Σ-protocol to compute w. The parameter ℓ is defined by the encryption algorithm E.

Theorem 2. *Let $\mathcal{R}$ be a relation that has a Σ-protocol. The protocol depicted in Figure 2 is a secure verifiable encryption scheme for $\mathcal{R}$ when sequentially repeated k times.*

Proof. Completeness: This can easily be verified.

Validity: The algorithm R will fail, only if in every iteration at least one of e_0, e_1 contains invalid protocol responses. This means that V will accept with probability at most $1/2$ in each iteration, and so it accepts all iterations with probability at most 2^{-k}.

Computational Zero-Knowledge: We can build a simulator for one iteration using standard resetting techniques: we first choose a random bit c', our guess at what V's challenge will be. Then we simulate a conversation (t, c', s) in the Σ-protocol. We encrypt s to get $e_{c'}$, and encrypt something random to get $e_{1-c'}$. Then we send t, e_0, e_1 to $\tilde{V}$, and get c back. If $c' = c$, we open e_c as the honest prover would do. Otherwise, we rewind and try again. By honest verifier zero-knowledge of the Σ-protocol and semantic security of the encryption, the probability that $c' \neq c$ is at most $1/2 + 1/p(k)$ for some polynomial p and the output is computationally indistinguishable from a real conversation. It also follows that the expected run-time of the simulator is $O(kT(k))$, where $T(k)$ is the time needed for one attempt to simulate one iteration. □

The main drawback of our basic scheme is that the verifier must store an encryption and a conversation (accepting triple) in the Σ-protocol for *each* repetition and that it needs to be repeated $\Theta(k)$ times sequentially. In Figure 3 an improved scheme is depicted, that allows to store much less encryptions and triples. The idea is that in the basic step, the prover will supply a valid triple where the challenge is 0, but where the prover's response is encrypted. The verifier can then ask the prover either to open the encryption or to supply a valid answer to challenge 1. The point is that the verifier only needs to remember the unopened encryptions and the related triples.

Furthermore, the protocol is made constant-round by relying on a commitment scheme, i.e., a function *Com* that takes as input the string α to commit to and an additional input string β. Then a player can commit by sending $T := Com(\alpha, \beta)$, where β is randomly chosen (we write just $Com(\alpha)$ in the following). The commitment scheme must possess the following three

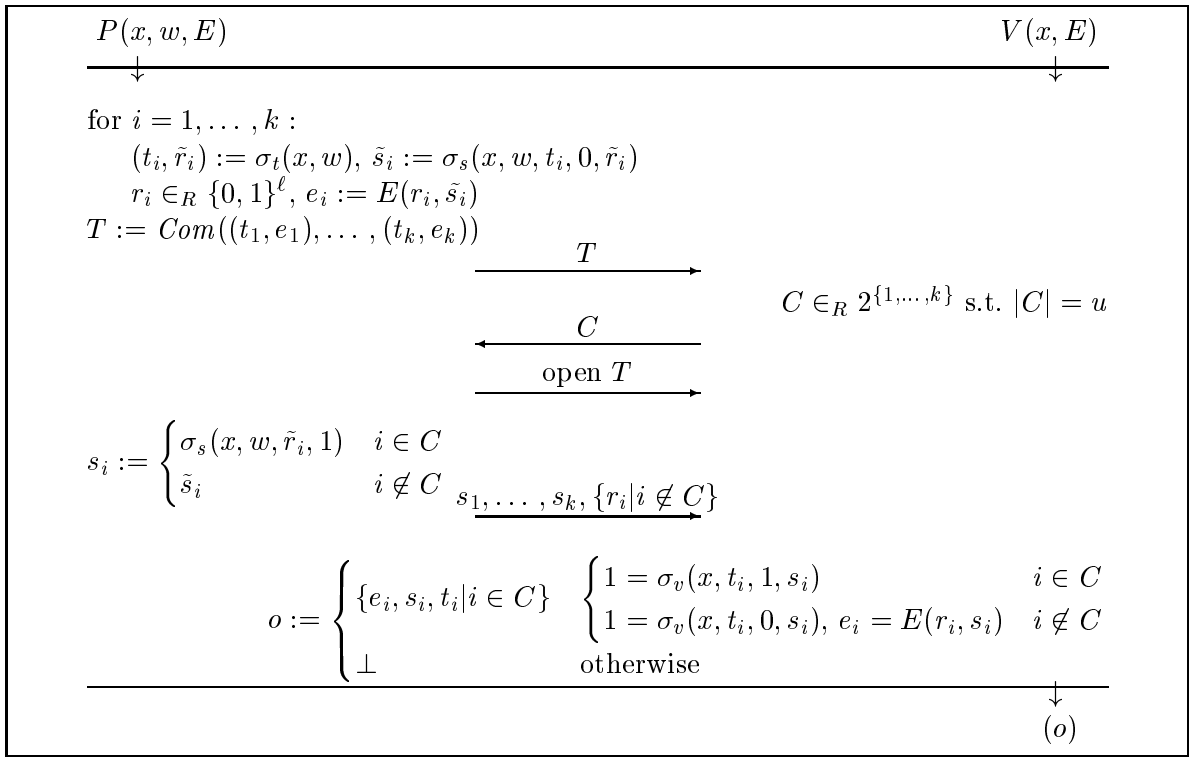

Figure 3: An improved verifiable encryption scheme using a commitment scheme. Reconstruction is straightforward by decrypting the e_i's in the output. The integers k, u, and ℓ are security parameters and $|C|$ denotes the cardinality of the set C.

properties:

Hiding: The distributions of commitments to different α's must be computationally indistinguishable.

Binding: It should be computationally hard to open a commitment in two different ways, i.e., to find $\alpha \neq \alpha'$ and β, β' such that $Com(\alpha, \beta) = Com(\alpha', \beta')$.

Trapdoor: There is a piece of trapdoor information the knowledge of which allow to open a commitment in an arbitrary way.

There are numerous efficient constructions known of such schemes (see for instance [15]), but in fact our assumptions in this scenario are already sufficient to ensure their existence: $\mathcal{R}$ must be a hard relation, in order for our scenario to make sense and it is known that a Σ-protocol for a hard relation implies the existence of a secure commitment scheme with these three properties [17].

We require a once-and-for-all set-up phase for the commitment schemes where the verifier generates the *Com*-function together with the trapdoor information, sends the function to the prover, and proves knowledge of the trapdoor. For simplicity, we do not include this set-up phase explicitly in the protocol description; nevertheless, this set-up phase needs to be considered in the proof of security.

Theorem 3. *Let $\mathcal{R}$ be a relation possessing a Σ-protocol. The protocol depicted in Figure 3, when using a secure commitment scheme Com, is a secure verifiable encryption scheme for $\mathcal{R}$ for any u such that* $\log k < u < k/2$.

Proof. Completeness: This can easily be verified.

Validity: Consider an arbitrary prover P^*. Assume P^* convinces V about an incorrect encryption with probability larger than a polynomial fraction $1/p(k)$ for infinitely many k (and some x). We then show that P^* can be used to break the commitment scheme with non-negligible probability, contradicting the assumption.

Fix a k and x such that P^* cheats V with probability larger than $1/p(k)$. Fix a random set of coin-flips R for P^*. Say R is good if, when using R as random input, P^* cheats V with probability at least $1/2p(k)$. Then the probability that R is good is also at least $1/2p(k)$. So assume in the following that we have a good R. Since all inputs to P^* are now fixed, any conversation starts with a fixed commitment T.

The number of different subsets I that can be chosen by V is $\binom{k}{u}$. Using Sterlings bounds on $k!$,$u!$, and $(k-u)!$, we have

$$\binom{k}{u} = \frac{k!}{u!(k-u)!} > \frac{1}{q(k)}(\beta-1)^{k/\beta}\Big(\frac{\beta}{\beta-1}\Big)^{k} \ ,$$

where $q(k)$ is some polynomial and we have set $u = k/\beta$. For $\log k < u < k/2$ this is clearly super-polynomial. It follows that we can by rewinding P^* and sending it random values of C efficiently find conversations $(T, C, v), (T, C', v')$, where $C \neq C'$, and where P^* has cheated V. If v and v' involve the same opening of T, then the same set of triples from the underlying Σ-protocol are involved in both cases. Then since $C \neq C'$ there is an index i such that v contains t_i, e_i, s_i, where $(t_i, 1, s_i)$ is an accepting triple, and where v' contains s_i', the decryption of e_i, where $(t_i, 0, s_i')$ is also accepting. But by soundness of the Σ-protocol this implies that the recovery algorithm does manage to find a correct w, contradicting the fact that P^* cheats V. Thus it must be the case that opening of T contained in v is different from the one occurring in v', and we have broken the commitment scheme.

Computational Zero-Knowledge: Note that if we know the subset C that $\tilde{V}$ will choose, it is easy to simulate: for all $i \notin C$, that is, where we have to open the encryption e_i, we prepare an accepting triple $(t_i, 0, s_i)$ and encrypt s_i to get e_i. For the other i's, we prepare an accepting triple $(t_i, 1, s_i)$ and encrypt something random in e_i. This will clearly allow us to satisfy the verifier's requests, and the only difference between this and a real conversation is in the contents of the unopened encryptions.

Hence the strategy of a simulator is to use the proof of knowledge of the trapdoor that V^* gives in the set-up phase of the commitment scheme. The simulator attempts to extract the trapdoor by rewinding V^* in this phase, and in parallel (for technical reasons) runs an exhaustive search for the trapdoor. This ensures that the simulator will in expected polynomial time extract the trapdoor. Then any commitment can be opened to reveal any value desired, and the rest of the simulation is trivial, by the above. □

There is another variant of our protocol which requires 4 messages, but has the advantage of being secure against an unbounded prover. Here, V commits to his choice of C in advance, P sends $(t_1, e_1), ..., (t_k, e_k)$, V opens the commitment, and P responds to the C revealed as above. This can be proven zero-knowledge using techniques from [25].

Remark. Note that our simulator in the proof of the zero-knowledge property uses rewinding, but only in the set-up phase for the commitment scheme. Since this phase can be done once and for all in a preprocessing, it is feasible to ensure that this one phase is not interleaved with other protocols. This will imply that the entire protocol is concurrent zero-knowledge (see also [20]).

In practice one is often interested in a particular error probability ϵ. There will then be many values of k and u with $1/\binom{k}{u} \leq \epsilon$, and one should then of course choose whatever particular k and u fit the application best.

It is easy to make a non-interactive variant of this construction using the Fiat-Shamir heuristic: the prover computes $(t_1, e_1), \ldots, (t_k, e_k)$, determines the challenge C from $h((t_1, e_1), \ldots, (t_k, e_k))$, where h is some suitable (hash) function, and finally appends valid responses $(s_1, r_1), \ldots, (s_k, r_k)$. All this can be verified as before by V. It is straightforward to show that this is secure in the random oracle model, replacing calls to h by calls to the oracle.

4 Verifiable Group Encryption

In our basic primitive, the prover and the verifier have to trust the third party to behave as expected. That is, the prover must trust him/her to recover the encrypted witness when appropriate only, whereas the verifier relies on the third party to decrypt the witness when required. To achieve higher security against fraudulent third parties, we extend our basic primitive to verifiable *group* encryption, which on its own is a useful concept in many cases. Here, the witness gets encrypted for n third parties, called *proxies*, such that only designated subsets of them can jointly recover the witness. Although it superficially looks more complicated, it turns out to be trivial to implement using our basic verifiable encryption scheme, as we shall see.

Informally, verifiable group encryption takes place in a similar model as ordinary verifiable encryption, that is, P and V interact on common input x, where P knows w such that $(x, w) \in \mathcal{R}$. As before, it is instructive to think of w as being a signature on x w.r.t. some fixed public key. Now, however, n public encryption keys $E_1, \ldots, E_n$ are involved, and a monotone access structure Γ over $\{1, \ldots, n\}$ is agreed upon by P and V. Then an honest V obtains from P encryptions $E_1(w_1), \ldots, E_n(w_n)$ such that a valid w can be reconstructed from a subset A of the w_i's, if $A \in \Gamma$, whereas a set $A \notin \Gamma$ gives no information. Finally, if honest proxies forming a set $A \in \Gamma$ decide to reconstruct w, they can do so successfully, even if dishonest proxies also participate.

This notion of "group encryption" should not be confused with the notion of threshold encryption [18]. In the latter case, a number of parties publishes a *single* public key and the access structure is determined by these parties during the setup of the system. In contrast, a group encryption scheme allows to choose a (possibly different) access structure each time when (verifiably) encrypting a message. Another distinguishing feature is that a group encryption enables the proxies to choose their encryption keys independent of each other and allows them even to use different encryption schemes (perfect separability). We finally note that one could also use any threshold encryption scheme on its own to achieve higher security against fraudulent third parties for the basic verifiable encryption scheme. However, such a solution will not offer perfect separability and not fit the framework of an existing public key infrastructure we are targeting.

4.1 Realization

A verifiable group encryption scheme can be realized using a secret sharing scheme for the chosen access structure Γ. We just need to observe that our verifiable encryption scheme from before works based on *any* public key encryption scheme. In particular, we will execute it using the following encryption scheme: given input s, use the given secret sharing scheme with Γ to get shares $s_1, \ldots, s_n$ of s, and then let the output ciphertext be $E_1(s_1), \ldots, E_n(s_n)$. Clearly,

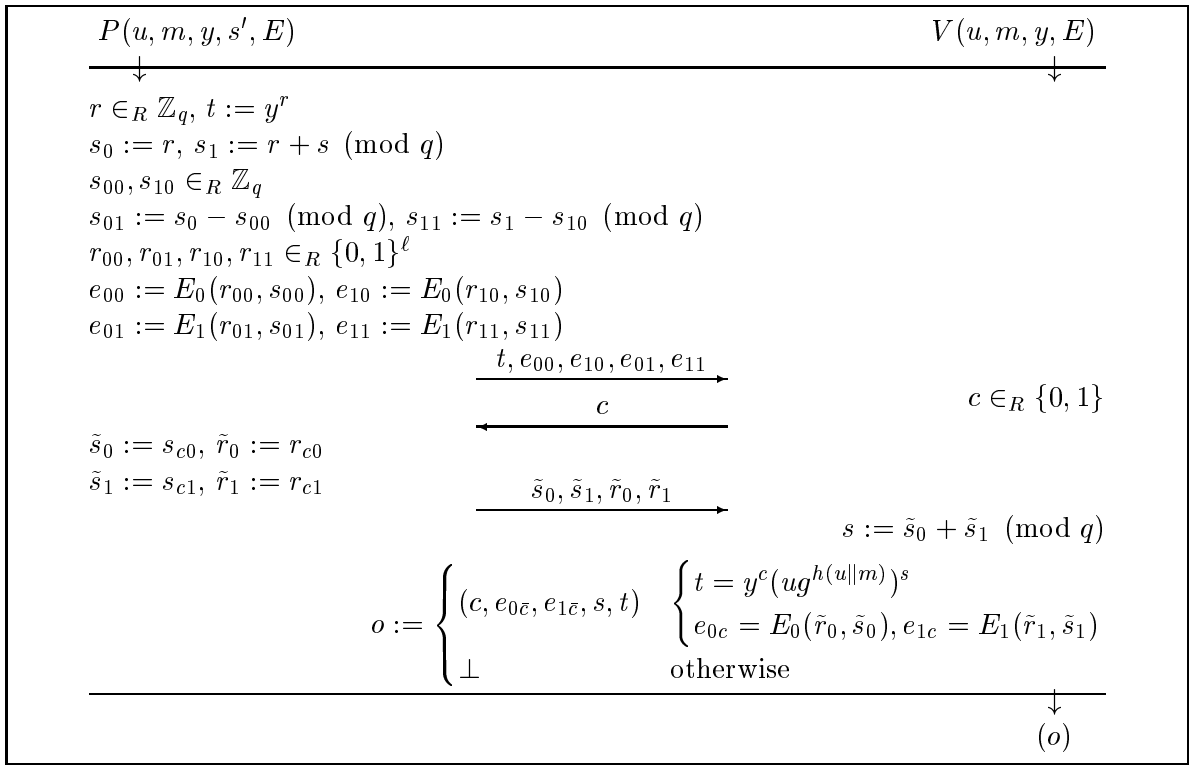

Figure 4: An example of verifiable group encryption for $\Gamma = \{\{0,1\}\}$. For simplicity only one round is described (cf. previous section). Schnorr signature on m is defined to be (c', s') such that $c' = h(u\|m)$ holds, where $u = g^{c'} y^{s'}$ and y is the public key of signer. The signature can be verifiable encrypted by providing u and m and using the relation $\{(u/g^{h(u\|m)}, w) | w = \log_y(u/g^{h(u\|m)}) \pmod q\}$. Recovery takes place by computing $(-1)^c(D_0(e_{0\bar{c}}) + D_1(e_{1\bar{c}}) - s) \pmod q$.

one can compute s from a correctly formed ciphertext, if one can decrypt a subset of the s_i's corresponding to a set in Γ.

From the construction it is clear that, due to the properties of secret sharing schemes, the proxies can reconstruct the witness when given $E_1(s_1), \ldots, E_n(s_n)$ and Γ by decrypting, pooling the shares, reconstructing s, and then using the properties of the basic verifiable encryption protocol. This is possible even with the participation of malicious proxies as long as the honest proxies form a set $A \in \Gamma$, however, they might not be able to do this efficiently in the presence of malicious proxies who provide incorrect values for the s_i's. If the encryption scheme used allows proxy i to prove that s_i is indeed the value encrypted in $E_i(s_i)$ directly, the problem is trivial to solve. If this is not the case, we can modify the encryption scheme and encrypt the random choices used for encryption of the shares as well for the respective proxies, i.e., our encryption scheme would output $\big(E_1(r_1, s_1), E_1(\tilde{r}_1, r_1)\big), \ldots, \big(E_n(r_n, s_n), E_n(\tilde{r}_n, r_n)\big)$. Now, proxy i can prove correct decryption by providing r_i and s_i.

In case the scheme is made non-interactive using the Fiat-Shamir heuristic [22], the access structure Γ should also be included in the hash-function.

4.2 Verifiable Group Encryption vs. Verifiable Signature Sharing

The concept of *verifiable signature sharing* (VΣS) [11, 24] involves a *signature receiver* who distributes shares of a signature on a public message to a set of proxies such that all proxies can verify that this has been done correctly, and a qualified set of proxies can always reconstruct the signature, even in the presence of malicious proxies. Trivially, a verifiable group encryption scheme can be used to implement VΣS: we simply execute verifiable group encryption of a signature on the given message, such that the signature receiver plays the role of the prover, and the proxies together play the role of the verifier. If the interactive variant is used, this is done by having proxies generate challenges for the prover by collective coin-flipping. With the non-interactive variant, no special precautions are needed, and the scheme becomes even publicly verifiable. The possibility to use any encryption scheme for the individual proxies solves an open problem raised in [11]. Moreover, a (publicly) verifiable secret sharing scheme, e.g., see [35], can be obtained along the same lines.

5 Application to Group Signatures and Identity Escrow

A *group signature scheme* [6, 9, 10, 12, 13, 31] allows a member of a group of users to sign a message on the group's behalf. The scheme protects the privacy of signers in that the verifier should not be able to find the identity of the signer. However, to handle special cases where the scheme is misused by some user, there is a *revocation manager* who can indeed find the identity of the signer from the signature. In some schemes, there is also a *membership manager* who takes care of the key set-up and enrollment of users in the group. No collusion of users (even including the membership manager) should be able to forge signatures such that it is not possible for the revocation manager to reveal the identity of the signer. Furthermore, no collusion of users (even including both managers) should be able to forge signatures such that upon revocation they seem to originate from another user. Moreover, we want to minimize the involvement of parties in the protocols. This means that managers should not be involved in creating a signature, and ideally also that the revocation manager is not involved in establishing the group, and is in fact completely inactive until revocation of anonymity is needed.

The interactive equivalent of a group signature scheme is a group identification scheme with revocable anonymity (also called identity escrow scheme [28]). Here, the goal is for a group member to anonymously identify himself as a member — rather than being able to sign a message. Except for this, the security properties of group signatures carry over directly.

For both kind of schemes, it is of course desirable that they possess perfect separability, i.e., that they can be implemented based on keys that users and managers already have established, even if those keys were not intended to be used in these schemes.

In the following, we show how to use a verifiable encryption scheme to design a separable group identification scheme with revocable anonymity. This scheme can be proved secure, assuming only security of the encryption, signature, and identification schemes involved. We then modify this to a group signature scheme using the Fiat-Shamir heuristic [22]. This scheme is secure assuming in addition that the heuristic is valid for the protocol and hash function involved. It can be proved secure with no additional assumptions in the random oracle model. For a formal model of group signatures and identity escrow we refer to [10, 28].

5.1 Realizations

We describe the basic idea for the case where the users' public keys are of some signature schemes and then extend the scheme such that the users' public keys can also be of some interactive identification schemes.

We assume we are given public keys $P_1, \ldots, P_n$ of secure signature schemes for n players, and that for each signature scheme employed, there is a Σ-protocol for the relation $\{(x, w) |\ w$ *is a valid signature on message* $x\}$. As mentioned earlier, this is true of any signature scheme for which reduction functions to a group-homomorphism exists [2]. To prove our scheme's security formally, we need that the signature schemes are secure against chosen message attacks. Finally, we assume that a revocation manager has been selected, who has published a public key E of a semantically secure encryption scheme. By Lemma 1 and Theorem 3 we get a group identification scheme with revocable anonymity as follows:

The group's public key consists just of the membership manager's signature (certificate) on the tuple $(P_1, \ldots, P_n)$. To authenticate himself as a group member, the prover computes the signature on message x (randomly chosen by the verifier) with respect to his/her public key. Then the prover and the verifier carry out the verifiable encryption protocol for the relation $\{(x, w) | w$ *is a valid signature on* x *with respect to one of the public keys* $P_1, \ldots, P_n\}$, where encryption public key used is E, the one of the revocation manager. We can do this because, by Lemma 1, an efficient Σ-protocol for this relation can be derived from the Σ-protocols we assumed exist for each single signature scheme.

This derived Σ-protocol proves that w is a valid signature w.r.t. one of the public keys, but yields no information about which one is involved. This and the zero-knowledge property of the verifiable encryption scheme implies anonymity for the prover. Furthermore, if some coalition of users could impersonate another user, then they could also forge this user's signatures. Finally, due to the validity of the verifiable encryption, the anonymity can be revoked just by decryption. With respect to the revocation manager's ability to prove correct decryption of the witness, the remarks of the previous section applies: either the underlying encryption scheme allows this or, if not, the verifiable encryption scheme is modified as described there.

Clearly, applying the Fiat-Shamir heuristic [22] yields a group *signature* scheme from this construction: the message x which was chosen by the verifier before, will now be the message to be signed. And we hash x and the prover's first message in the verifiable encryption protocol to get a challenge. This can be proved secure in the random oracle model. In summary, we have argued:

Theorem 4. *Given any set of secure signature schemes with an associated Σ-protocol each and any secure public key encryption scheme, a secure separable group identification scheme with revocable anonymity and a separable group signature scheme secure in the random oracle model can be constructed. The complexities of the schemes are linear in the group's size, the security parameter, and in the complexity of the signature and encryption schemes.*

What if we are not given a number of signature public keys, but instead public keys of interactive identification protocols, which we assume to be Σ-protocol (such as Fiat-Feige-Shamir or Schnorr)? Can we still build a group identification scheme? We could try using the identification protocol directly together with the 1-out-of-n method in the verifiable encryption scheme. But this would allow the revocation manager to compute the secret key of the prover and hence the scheme could be used only once.

In the following we provide a method to exploit the Σ-protocol of any identification scheme such that we can nevertheless build a group identification scheme from it. The idea is to

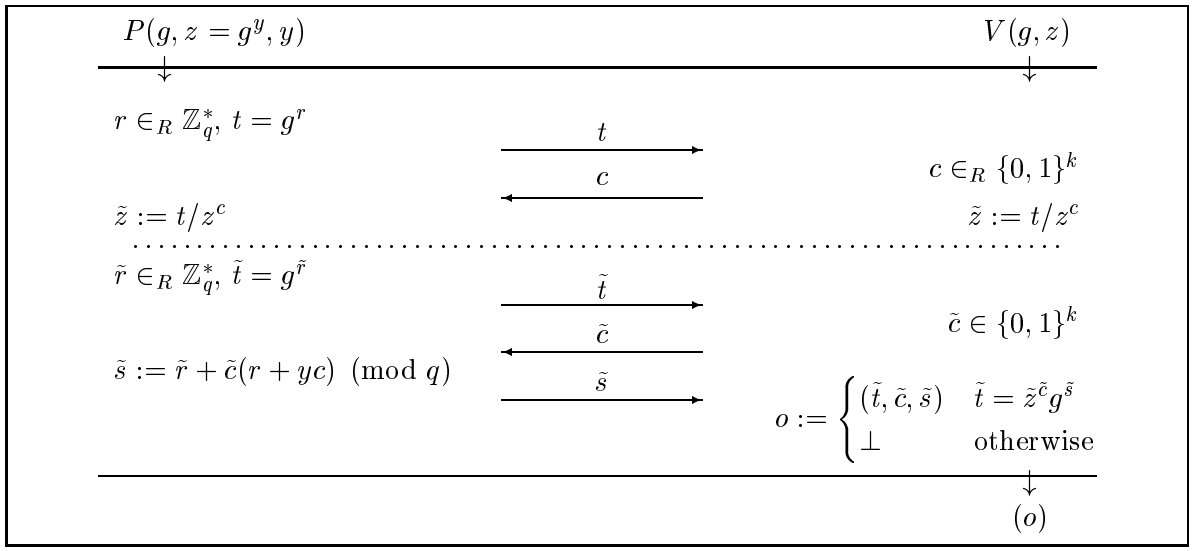

Figure 5: Transforming Schnorr's identification protocol. The prover's public key is $(z, G = \langle g \rangle)$ where $|G| = q$ and $y = \log_g z$ is the prover's secret key. In the protocol steps above the dotted line, the prover and the verifier jointly construct the instance $(\tilde{z}, G = \langle g \rangle)$ of the new relation. The protocol steps below the dotted line constitute the well-known protocol for proving the knowledge of $\tilde{z}$, i.e., the witness in the new relation.

only prove the knowledge of the response in an instance of the identification protocol rather than providing it. Let $\mathcal{R}_i$ be the relation of some identification scheme with the Σ-protocol $(\sigma_t, \sigma_s, \sigma_v)$ and let x be the instance of $\mathcal{R}_i$ consistent with the public key P_i. Define the relation $\tilde{\mathcal{R}}_i = \{(s, (x, t, c)) | 1 = \sigma_v(x, t, c, s), x \in L_{\mathcal{R}_i}\}$. In other words, t and c were the first two messages in the identification protocol and the third message is a witness to $\tilde{x} = (x, t, c)$ for $\tilde{\mathcal{R}}_i$. Thus, if the prover and the verifier construct $\tilde{x}$ together as the first two steps of the identification protocol, and then the prover proves the knowledge of a witness to $\tilde{x}$ for $\tilde{\mathcal{R}}_i$, then the verifier will be convinced, that the prover also knows a witness to x for $\mathcal{R}_i$. For most known identification protocol there exist efficient Σ-protocols for this latter relation. Figure 5 shows as an example such a transformation of the Schnorr identification protocol.

This transformation allows us now to construct a group identification scheme with revocable anonymity as follows. The prover and the verifier carry out the first two steps of the identification protocols for every public key P_i to get instances for the new relations $\tilde{\mathcal{R}}_i$ and then carry out the Σ-protocol for proving knowledge of a witness to (at least) one of these new instances (using Lemma 1). The rest of the construction is the same as in case we start from signature schemes. It is clear that we can also mix identification schemes and signature schemes.

5.2 Extensions and Related Work

Extensions to generalized group identification schemes with revocable anonymity or generalized group signature schemes [6] are straight forward. In this case the prover is actually a set of group members who demonstrate that they form a qualified set w.r.t. some access structure, but without revealing which set is present. Thus they have to do joint computations and some precautions must be taken. We refer to [6] for details.

To protect group members from fraudulent behavior of the revocation manager, the ability of revocation can be distributed to several parties by using our verifiable group encryption

instead of ordinary verifiable encryption (cf. Section 4).

Comparing our group signature and identity escrow schemes with previous proposals, we find that ours are the only ones that provide perfect separability. The identity escrow scheme in [28] provides only partial separability, i.e., separability w.r.t. the revocation manager; the group signature scheme in [9] provides only weak separability, i.e., group members can only use RSA or discrete logarithm based keys of a given size and the membership manager can only use the RSA signature scheme together with a particular algebraic padding function.

The schemes proposed in [9, 10, 28] have the property that the group's public key does not depend on the size of the group. These scheme are an order of magnitude more efficient than ours. However, our schemes have some advantages over them. Namely, in order to exclude group members, the membership manager has to interact with all remaining group members in these schemes whereas in our solution the membership manager just publishes a new group public key (our scheme does not require the group members to participate in the setup at all). Furthermore, only the kind of schemes as ours allows to make it publicly known who the group members are (note signatures are still anonymous) while this is not possible for the other type of scheme. Finally, generalized group signature schemes can only be realized this way. Therefore, we believe that schemes where the group's public key reflects the group's size and structure are indeed the only solution for certain applications.

Finally, our verifiable encryption scheme can be used as a subroutine in the schemes [9, 10, 28] allowing the revocation manager to independently choose *any* encryption scheme.

It remains an open problem to find efficient schemes providing perfect separability and where the group's public key does not depend on the group's size.

Acknowledgments

The authors are grateful to Christian Cachin and Anna Lysyanskaya for their comments on earlier versions of this paper.

References

[1] G. Ateniese Efficient Verifiable Encryption (and Fair Exchange) of Digital Signatures, In *6th ACM CCS*, pp. 138–146, 1999.

[2] N. Asokan, V. Shoup, and M. Waidner. Optimistic fair exchange of digital signatures. In *EUROCRYPT '98*, vol. 1403 of *LNCS*, pp. 591–606, 1998.

[3] N. Asokan, V. Shoup, and M. Waidner. Optimistic fair exchange of digital signatures. *IEEE Journal on Selected Areas in Communications*, 18(4):591–610, Apr. 2000.

[4] F. Bao. An Efficient Verifiable Encryption Scheme for the Encryption of Discrete Logarithms, In *CARDIS '98* vol. 1820 of *LNCS*, 2000.

[5] J. Camenisch and I. Damgård. Verifiable encryption and applications to group signatures and signature sharing. Technical Report RS-98-32, BRICS, Department of Computer Science, University of Aarhus, Dec. 1998.

[6] J. Camenisch. Efficient and generalized group signatures. In *EUROCRYPT '97*, vol. 1233 of *LNCS*, pp. 465–479, 1997.

[7] J. Camenisch, U. Maurer, and M. Stadler. Digital payment systems with passive anonymity-revoking trustees. In *Computer Security — ESORICS 96*, vol. 1146 of *LNCS*, pp. 33–43. Springer Verlag, 1996.

[8] J. Camenisch and M. Michels. A group signature scheme with improved efficiency. In *ASIACRYPT '98*, vol. 1514 of *LNCS*, pp. 160–174. Springer Verlag, 1998.

[9] J. Camenisch and M. Michels. Separability and Efficiency for Generic Group Signature Schemes In M. Wiener, *CRYPTO '99*, vol. 1666 of *LNCS*, 1998.

[10] J. Camenisch and M. Stadler. Efficient group signature schemes for large groups. In *CRYPTO '97*, vol. 1296 of *LNCS*, pp. 410–424, 1997.

[11] D. Catalano and R. Gennaro. New efficient and secure protocols for verifiable signature sharing and other applications. In *CRYPTO '98*, vol. 1642 of *LNCS*, pp. 105–120, Berlin, 1998. Springer Verlag.

[12] D. Chaum and E. van Heyst. Group signatures. In *EUROCRYPT '91*, vol. 547 of *LNCS*, pp. 257–265. Springer-Verlag, 1991.

[13] L. Chen and T. P. Pedersen. New group signature schemes. In *EUROCRYPT '94*, vol. 950 of *LNCS*, pp. 171–181, 1995.

[14] R. Cramer. *Modular Design of Secure yet Practical Cryptographic Protocol*. PhD thesis, University of Amsterdam, 1997.

[15] R. Cramer and I. Damgård. Zero-knowledge proof for finite field arithmetic, or: Can zero-knowledge be for free? In *CRYPTO '98*, vol. 1642 of *LNCS*, 1998.

[16] R. Cramer, I. Damgård, and B. Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *CRYPTO '94*, vol. 839 of *LNCS*, pp. 174–187. Springer Verlag, 1994.

[17] I. B. Damgård. On the existence of bit commitment schemes and zero-knowledge proofs. In *CRYPTO '89*, vol. 435 of *LNCS*, pp. 17–27, 1990.

[18] Y. Desmedt and Y. Frankel. Threshold cryptography. In *CRYPTO '89*, vol. 435 of *LNCS*, pp. 307–315. Springer-Verlag, 1990.

[19] C. Dwork, M. Naor, and A. Sahai. Concurrent zero knowledge. In *Proc. 30th Annual ACM Symposium on Theory of Computing (STOC)*, 1998.

[20] C. Dwork and A. Sahai. Concurrent zero-knowledge: Reducing the need for timing constraints. In *CRYPTO '98*, vol. 1642 of *LNCS*, pp. 105–120, Berlin, 1998. Springer Verlag.

[21] U. Feige, A. Fiat, and A. Shamir. Zero-knowledge proofs of identity. *Journal of Cryptology*, 1:77–94, 1988.

[22] A. Fiat and A. Shamir. How to prove yourself: Practical solution to identification and signature problems. In *CRYPTO '86*, vol. 263 of *LNCS*, pp. 186–194, 1987.

[23] Y. Frankel, Y. Tsiounis, and M. Yung. "Indirect discourse proofs:" Achieving efficient fair off-line e-cash. In *ASIACRYPT '96*, vol. 1163 of *LNCS*, 1996.

[24] M. Franklin and M. Reiter. Verifiable signature sharing. In *EUROCRYPT '95*, vol. 921 of *LNCS*, pp. 50–63. Springer Verlag, 1995.

[25] O. Goldreich and A. Kahan. How to construct constant-round zero-knowledge proof systems for NP. *Journal of Cryptology*, 9(3):167–190, 1996.

[26] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, Apr. 1984.

[27] L. C. Guillou and J.-J. Quisquater. A practical zero-knowledge protocol fitted to security microprocessor minimizing both transmission and memory. In *EUROCRYPT '88*, vol. 330 of *LNCS*, pp. 123–128, 1988.

[28] J. Kilian and E. Petrank. Identity escrow. In *CRYPTO '98*, vol. 1642 of *LNCS*, pp. 169–185, Berlin, 1998. Springer Verlag.

[29] S. Micali. Efficient certificate revocation and certified e-mail with transparent post offices. Presentation at the 1997 RSA Security Conference.

[30] S. Micali, C. Rackoff, and B. Sloan. The notion of security for probabilistic cryptosystems. *SIAM Journal on Computing*, 17(2):412–426, April 1988.

[31] H. Petersen. How to convert any digital signature scheme into a group signature scheme. In *Security Protocols Workshop*, Paris, 1997.

[32] G. Poupard and J. Stern, Fair Encryption of RSA Keys. In *EUROCRYPT 2000*, LNCS, pp. 173–190. Springer Verlag, 2000.

[33] C. P. Schnorr. Efficient signature generation for smart cards. *Journal of Cryptology*, 4(3):239–252, 1991.

[34] A. Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, Nov. 1979.

[35] M. Stadler. Publicly verifiable secret sharing. In *EUROCRYPT '96*, vol. 1070 of *LNCS*, pp. 191–199. Springer Verlag, 1996.

[36] M. Stadler, J.-M. Piveteau, and J. Camenisch. Fair blind signatures. In *EUROCRYPT '95*, vol. 921 of *LNCS*, pp. 209–219, 1995.

[37] A. Young and M. Yung. Auto-Recoverable Auto-Certifiable Cryptosystems. In *EUROCRYPT '98*, vol. 1403 of *LNCS*, pp. 17–31, 1998.