Improved OR Composition of Sigma-Protocols

Michele Ciampi
Università di Salerno
ITALY
mciampi@unisa.it

Giuseppe Persiano
Università di Salerno
ITALY
pino.persiano@unisa.it

Alessandra Scafuro
Boston University and
Northeastern University
USA
scafuro@bu.edu

Luisa Siniscalchi
University of Salerno
ITALY
lsiniscalchi@unisa.it

Ivan Visconti
Università di Salerno
ITALY
visconti@unisa.it

Abstract

In [LS90] Lapidot and Shamir provide a 3-round witness-indistinguishable (WI) proof of knowledge for Graph Hamiltonicity (the LS proof) with a special property: the prover uses the statement to be proved *only* in the last round.

This property has been instrumental in constructing round-efficient protocols for various tasks [KO04, DPV04, YZ07, SV12]. In all such constructions, the WI proofs are used to prove the OR composition of statements that are specified at different stages of the main protocol. The special property of LS proofs is used precisely to allow a player of the main protocol to start a proof, even if only one of the statements of the OR relation is available, thus saving rounds of communication.

If, on the one hand, the usage of LS proofs saves rounds, on the other hand it necessarily requires NP-reductions to Graph Hamiltonicity. As such, even if each of the statements to be proved in the main protocol admits an efficient Σ-protocol (e.g., if the statement consists in proving knowledge of a committed value or of a secret key), the reduced round complexity is paid for with a loss of efficiency. Hence, round-efficient constructions that rely on LS proofs are typically inefficient.

A natural question is why one would go through the NP-reduction to use LS proof, instead of composing the Σ-protocols using the OR composition technique introduced by Cramer, Damgård and Schoenmakers (CDS) in [CDS94]. The answer is that the CDS technique requires *both* statements to be available at the beginning of the protocol. Due to this limitation, constructions that use the CDS technique have in some cases a worse round complexity than the ones based on LS proofs.

In this paper we introduce a new OR composition technique for Σ-protocols that needs only one statement to be fixed when the proof begins. This seemingly weaker property is sufficient to replace the use of LS proofs in many applications that do not need both theorems to be undefined when the proof starts. In fact, we show how the new OR composition technique can directly improve the round complexity of the efficient perfectly simulatable argument of Pass [Pas03] (from four to three rounds) and of efficient resettable WI arguments (from five to four rounds).

Our OR technique can not compose any arbitrary pair of Σ-protocols. Nevertheless, we provide a precise classification of the Σ-protocols that can be composed and show that all the widely used Σ-protocols can be composed with our technique.

Contents

1 Introduction

Witness indistinguishable (WI) proofs. WI[1] proofs are fundamental for the design of cryptographic protocols in particular when combined with a proof of knowledge (PoK) property. In a WIPoK the prover $\mathcal{P}$ proves knowledge of a witness certifying the veracity of a statement $x \in L$ to a verifier $\mathcal{V}$. WIPoKs can be used directly in some applications (e.g., in identification schemes) or can be a building block for stronger security notions (e.g., for zero-knowledge proofs using the FLS [FLS90] paradigm or for round-optimal secure computation [KO04]).

Round complexity of cryptographic protocols has always been extensively studied both for its practical relevance and for its natural and conceptual interest. Regarding WIPoKs, we know already from Blum's protocol [Blu86] that 3-round WIPoKs exist for all NP-languages under the sole assumptions that one-way permutations exist. This is a theoretical result based on reducing any NP-language to the language of Hamiltonian graphs. Under stronger cryptographic assumptions, 2-round WI proofs, called ZAPs, and non-interactive WI (NIWI) have been shown in [DN00, GOS06, BP15]. Both ZAPs and NIWI protocols however are *not* proofs of knowledge.

Since NP-reductions are extremely expensive, several practical interactive PoKs have been designed for languages that are used in real-world cryptographic protocols (e.g., for proving knowledge of a discrete logarithm (DLog)). The study of such ad-hoc protocols mainly concentrates on a standardized form of a 3-round PoK referred to as Σ-protocol [Dam10, Sch89].

Σ-protocols. A Σ-protocol for an NP-language L, with witness relation R_L, is a 3-round proof system jointly run by a prover $\mathcal{P}$ and a verifier $\mathcal{V}$ in which $\mathcal{P}$ proves knowledge of a witness w for $x \in L$. In a Σ-protocol the only message sent by $\mathcal{V}$ is a random string. Such proof systems have two very useful properties: special soundness, which is a strong form of proof of knowledge, and special honest-verifier zero knowledge (SHVZK), which guarantees a weak form of witness hiding. The latter property basically says the following: if one knows the challenge in advance, then by just knowing also the theorem, he can generate an accepting transcript without using the witness. This is formalized through the existence of a special simulator, called the SHVZK simulator that, on input *a theorem* x and a challenge c, will output (a, z) such that (a, c, z) is an accepting 3-message transcript for x. Blum's protocol for Graph Hamiltonicity is an example of a Σ-protocol. Other examples of Σ-protocols are Schnorr's protocol [Sch89] for the DLog and DDH problems.

The security provided by the SHVZK property is clearly insufficient as it gives no immediate guarantees against verifiers who deviates from the protocol. Despite of this, the success of Σ-protocols and their impact in various constructions [Lin15, CPSV15, LP15, CG15, GK15, ORV14, AOS13, SV12, OPV10, BPSV08, CV07, CDV06, Vis06, GMY06, CV05a, CV05b, DG03, BFGM01, PS96] is a fact. This is due to a breakthrough of Cramer et al. [CDS94] that upgrades security of Σ-protocols to WI.

OR composition of Σ-protocols. Let L be a language that admits a Σ-protocol Π_L. [CDS94] shows how to use Π_L and its properties to construct a new Σ-protocol, Π_L^{OR}, for proving the OR composition of theorems in L *avoiding* the NP-reduction. [CDS94] achieves this by crucially exploiting the honest-verifier zero-knowledge (HVZK[2]) property of Π_L. The rationale behind the transformation can be informally explained as follows. The prover has to prove a statement of the form $(x_0 \in L \vee x_1 \in L)$. The naïve idea of simply running Π_L twice in parallel would not work

[1] We will use WI to mean both "witness indistinguishability" and "witness indistinguishable".

[2] HVZK requires the existence of a simulator that by receiving in input the theorem gives in output an accepting triple (a, c, z). Clearly HVZK is implied by SHVZK.

because the prover knows only one of the witnesses, say w_b, and cannot compute two accepting transcripts without knowing w_{1-b}. However, due to the HVZK property, the prover can generate an accepting transcript for $x_{1-b} \in L$ even without knowing w_{1-b}, by running the HVZK simulator Sim associated with Π_L. Indeed, Sim "only" needs in input the theorem x_{1-b} and will output the entire transcript, challenge included. The trick is then to generate the challenges for the two executions of Π_L, in such a way that the prover can control the challenge of exactly one of them (but not both), and set it to the value generated by Sim. Note that, if running the algorithm of Sim is as efficient as running the algorithm of $\mathcal{P}$, then the composed protocol is efficient. We stress that this OR-composition technique preserves SHVZK and will refer to it as the CDS-OR technique.

A very interesting property of this transformation, besides the fact that it does not need NP-reduction, is that if Sim is a simulator for perfect HVZK then Π_L^{OR} satisfies the stronger notion of witness indistinguishability (this was shown in [CDS94]). In other words this technique strengthens the security of Σ-protocols by adding WI to SHVZK essentially for free[3]. This result was further extended by Garay et al. [GMY06] that noted that the CDS-OR technique can be used also for Σ-protocols that are computational HVZK. In this case the relation proved is slightly different, namely, starting with a relation $\mathcal{R}_L$ and instances x_0 and x_1, the resulting Π_L^{OR} protocol is computational WI for the relation $\mathcal{R}_L^{\mathsf{OR}} = \{((x_0, x_1), w) : ((x_0, w) \in \mathcal{R}_L \wedge (x_1 \in L))\mathsf{OR}((x_1, w) \in \mathcal{R}_L \wedge (x_0 \in L))\}$.

Input-delayed proofs. It is often the case in cryptographic protocols to have a preamble phase that has the purpose of establishing, at least in part, a statement to be proven with a WI proof. In such cases, since one of the statements is fully specified only when the preamble is completed, the WI proof can start only after the preamble ends. Hence, the overall round complexity of protocols that follow this paradigm amounts to the sum of the round complexity of the preamble and of the WI proof.

In [LS90], Lapidot and Shamir (and later on Feige et al. in [FLS90]) show a 3-round proof of knowledge for Hamiltonian Graphs which has the special property that a prover can compute the first round of the proof, *without* knowing the theorem to be proved (that is, the graph) but only needs to know its size (that is, the number of vertices). Such a 3-round protocol is a Σ-protocol (and thus satisfies the SHVZK property) and is a WI proof. We will refer to this protocol as the LS construction. Also, we will call *Input-delayed* a Σ-protocol where the prover computes the first message without knowledge of the statement to be proved.

This special input-delayed property directly improves the round complexity of all the cryptographic protocols that follow the paradigm described above. The reason is that now the WI proof can start even if the preamble that generates the statement is not completed yet. It is worthy to note that in many applications the preamble serves as a mean to generate some trapdoor theorem, that is used only by in the security proof. The "honest" theorem instead is typically known already at the beginning of the protocol. This technique has been used extensively and, most notably, it led to the celebrated FLS paradigm that upgrades any WI protocol into a zero-knowledge (ZK) protocol.

The input-delayed property of LS has been instrumental to provide round-efficient constructions from general assumptions, such as: 4-round (optimal) secure 2PC where only one player gets the output (5 rounds when both players get the output) [KO04]; 4-round (optimal) ZK-arguments for NP from OWF, 4-round resettable WI [YZ07, SV12], 4-round (optimal) resettable ZK for NP in the $\mathtt{BPK}$ model [YZ07, SV12].

Despite being so influential to achieve round efficiency for cryptographic protocols, the power

[3] We stress that if Sim is as efficient as the prover then running Π_L^{OR} corresponds to running Π_L twice.

of LS construction unfortunately vanishes as soon as practical constructions are desired. Indeed, similarly to Blum's protocol, the LS construction is crucially based on specific properties of Hamiltonian graphs. Thus, when used to prove more natural languages, which is the case of most of the applications using WI proofs, it requires to perform rather inefficient NP-reductions.

Efficient protocols and limits of the CDS-OR technique. A natural question is what happens if we want to avoid the NP-reduction and we try to use the CDS-OR technique to construct input-delayed WI proofs. A bit more specifically, we know that there exist Σ-protocols that are Input-delayed (Schnorr's protocol [Sch89] for DLog is such an example since the first message can be computed without knowing the instance, but only a group generator). Thus the question is what happens if we apply the CDS-OR technique to an Input-delayed Σ-protocol. Do we obtain a WI Σ-protocol that is input-delayed as well?

Unfortunately the answer is negative. The CDS-OR technique does *not* preserve the input-delayed property, even when composing two Σ-protocols that are both input-delayed. To see why, recall that the CDS-OR composition technique, for a language L, requires the prover to compute two accepting transcripts of Π_L, one of which is computed by running the HVZK simulator Sim. Recall that Sim needs in input the theorem to be proved. Hence, to prove knowledge of a theorem $(x_0 \vee x_1) \in L$, the prover, who knows one witness, say w_b, needs to know also x_{1-b} already at the first round to be able to run the simulator. Thus, in CDS-OR technique the prover can successfully complete the protocol if and only if *both*[4] the instances are specified already at the first round.

Because of this missing feature, the CDS-OR technique has limited power in allowing one to obtain round-efficient/optimal cryptographic protocols, compared to the round complexity that is obtained by using the generic (non-efficient) LS proof. As such, in some cases when focusing on efficient constructions, the *best* round-complexity that we can achieve using efficient Σ-protocols and avoiding NP-reductions need at least one additional round, therefore requiring at least 5-round if one wants to match the previously mentioned applications (e.g., 5-round resettable ZK for NP in the $\mathtt{BPK}$ model [YZ07, SV12] and 5-round resettable WI [YZ07, SV12]).

Additionally, we note that the CDS-OR technique is the bottleneck in the round-complexity of the 4-round straight-line black-box perfect ZK argument shown by Pass in [Pas03]. This ZK protocol uses quasi-polynomial time simulation and, potentially, it would only need three rounds as any Σ-protocol. The additional first round is required precisely to define the trapdoor theorem. Hence, the following natural question arises:

> *Given a language L with an input-delayed Σ-protocol Π_L, is it possible to design an efficient Witness Indistinguishable Σ-protocol Π^L_{OR} for proving knowledge of a witness certifying that $(x_0 \in L \vee x_1 \in L)$ that does* not *require knowledge of both x_0 and x_1 to play the first round?*

1.1 Our Contribution

In this paper we answer the above question positively for a large class of Σ-protocols that includes *all* Σ-protocols used in efficient constructions. Specifically, we propose a new OR composition technique for Σ-protocols that relaxes the need of having both instances fixed before the Σ-protocol starts. Our technique allows the composition of Σ-protocols for different languages and of Σ-protocols that only have computational HVZK. Our new OR composition allows to achieve improved round

[4] To see why, note that the WI property requires that the prover would be able to prove any of the two theorems, and thus potentially use the simulator on either x_0 or x_1.

complexity in previous efficient constructions based on CDS-OR technique. Namely, we describe the following applications of our new OR composition technique:

- Efficient 3-round straight-line perfect ZK with quasi-polynomial time simulation. The previous construction required four rounds [Pas03].
- Efficient 4-round rWI. Previous constructions required five rounds [YZ07, SV12].

Our new technique can also be used to replace LS towards obtaining efficient round-optimal resettable zero-knowledge arguments in the BPK model (using the constructions of [YZ07, SV12]), round-optimal secure two-party computation (using the construction of [KO04]) and 4-round non-malleable commitments (using the construction of [GRRV14]).

Finally, we provide a precise classification of the Σ-protocols that can be used in our new OR composition technique. In the following paragraphs we first provide a high-level description our OR composition technique, then we discuss the applications in more details.

1.2 Our Techniques

Overview. We start by defining the setting we are considering. Let L_0 and L_1 be any pair of languages admitting Σ-protocols Π_0 and Π_1. We want to construct a Σ-protocol Π_L^{OR} for the language $L = L_0 \vee L_1$. An instance of L is a pair (x_0, x_1) and we want only x_0 to be specified before Π_L^{OR} starts while x_1 is specified only upon the last round of the protocol[5]. We assume that Π_1 is an *Input-delayed* Σ-protocol and thus the first prover message of Π_1 can be computed without knowing x_1. As mentioned earlier this property is satisfied by the most popular Σ-protocols such as the ones for Discrete Log, Diffie-Hellman triples, and of course, LS itself.

Now, recall that the problem with the CDS-OR technique was that a prover needs to run Sim to compute the first round of the protocol, and this necessarily requires knowledge of *both* theorems before the protocol starts. We want instead that the prover uses only knowledge of x_0.

We solve this problem by introducing a new OR technique that does not require the prover to run Sim on x_1 already in the first round. Instead, our technique allows the prover to wait and take action only in the third round when x_1 is finally defined.

Our starting point is the well known fact that given any Σ-protocol there exists an instance-dependent trapdoor commitment (IDTC) scheme where the witness for the membership of the instance in the language can be used as a trapdoor to open a committed message as any desired message. Our next observation is that, instead of having the prover send the first round for protocol Π_1 in the clear, we can have him send a commitment to it, and such commitment can be computed using an instance-dependent trapdoor commitment based on Π_0 with respect to instance x_0. Recall that this is possible, as in our setting we assume that Π_1 is an Input-delayed Σ-protocol, so the prover can honestly compute the first message of Π_1 without knowing x_1. Therefore, the first round of our Π_L^{OR} protocol, is simply an IDTC of a honest Π_1's first round.

Later on, upon receiving the challenge c from the verifier, and after the theorem x_1 is defined, the prover computes the third round as follows. If she has received a witness for x_0, then she will run Sim on input (x_1, c) to compute an accepting transcript of Π_1 for x_1. Then, using the witness w_0 she will equivocate the commitment sent in the first round, according to the message output by Sim. Otherwise, if she has received a witness for x_1 then she does not need to equivocate: she

[5] Like LS, we will just need the size of x_1 to be known when Π_L^{OR} starts.

will honestly open the commitment, and honestly compute the third message of Π_1. Therefore, the third round of our Π_L^{OR} protocol, simply consists of an opening of the IDTC together with the third message of Π_1.

Now note that this idea works only if we have a special IDTC scheme that has the following strong trapdoor property: a sender can equivocate even a commitment that has been computed honestly. Unfortunately, this property is not satisfied in general by any trapdoor commitment based on Σ-protocols, but only for some. This would restrict the class of languages that we can use as L_0 in our technique. For example, this class would not contain Blum's protocol.

Our next contribution is the construction of IDTC schemes that satisfy this strong trapdoor property, for a large class of Σ-protocols. Towards this goal, we define the notion of a t-IDTC scheme which are IDTCs for which the ability to open a commitment in t ways implies knowledge of a witness for the instance associated with the commitment. Next, we construct 2-IDTC and 3-IDTC schemes based on two different classes of Σ-protocols, the union of which includes all the Σ-protocols that are commonly used in cryptographic protocols. Finally, we provide a general OR composition technique for any pair of languages L_0 and L_1 such that L_0 has a t-IDTC scheme and L_1 has an Input-delayed Σ-protocol.

***t*-Instance-Dependent Trapdoor Commitment.** For integer $t \geq 2$, a t-IDTC scheme for a polynomial-time relation $\mathcal{R}$ admitting Σ-protocol $\Pi_{\mathcal{R}}$ is a triple $(\mathsf{TCom}, \mathsf{TDec}, \mathsf{TFake})$ where TCom, TDec are the honest commitment/decommitment procedures and TFake is the equivocation procedure that, given a witness for an instance x, equivocates any commitment with respect to x computed by TCom. The crucial differences between a t-IDTC scheme and a regular trapdoor commitment scheme are: (a) the trapdoor property is strong in the sense that knowledge of the trapdoor (that is, the witness of the instance x) allows to equivocate even commitments that have been honestly computed; (b) the binding property is relaxed: in a t-IDTC scheme, the sender can open the same commitment in $t-1$ different ways, even without the trapdoor. This relaxation allows us to build an IDTC scheme from a wider class of Σ-protocols, which will cover all the Σ-protocols we know of.

Constructing a 2-IDTC scheme. A 2-IDTC scheme can be straight-forwardly constructed from any Σ-protocol Π_0 that has the following property: even if the first message a_0 was computed by the SHVZK simulator Sim, an accepting z_0 can be efficiently computed, for every challenge c_0, by using knowledge of the witness and of the randomness used by Sim to produce a_0. We call the Σ-protocols that satisfy this property, *chameleon* Σ-protocols, and we denote by $\mathsf{P}_{\mathsf{sim}}$ the special prover strategy that can answer any challenge even starting from a simulated a_0.

More precisely, given a *Chameleon* Σ-protocol Π_0 for a language L_0, one can construct a 2-IDTC scheme as follows. Let $x_0 \in L_0$. To commit to a message m, the sender runs $\mathsf{Sim}(m, x_0)$ and obtains a_0, z_0 and the randomness r_0 used for such computation. The commitment is the value a_0. The opening is the pair m, z_0. The commitment is accepted iff (x_0, a_0, m, z_0) is accepting. To equivocate a_0, as a message m', run the special prover algorithm $\mathsf{P}_{\mathsf{sim}}(x_0, a_0, r_0, w_0, m')$ and obtain an accepting z_0.

Constructing a 3-IDTC scheme. We now discuss a different committing strategy that works for Σ-protocols in which the simulated first message a_0 can only be continued for the one challenge specified by Sim, even if a witness is made available. Blum's protocol for Hamiltonicity is an example of a Σ-protocol with this property.

To commit to m, the sender sends a pair (a_0, a_0') where, with probability $1/2$, a_0 is obtained by running the $\mathsf{Sim}(x_0, m)$ while a'_0 is computed by running the prover of Π_0, and with probability $1/2$ the order is inverted. One can think of a commitment as composed of two threads: a *simulated* thread and a *honest* thread. To open the commitment, the prover sends m and z^*, and the verifier accepts the decommitment if m, z^* are accepting for one of the threads; namely, the verifier checks that either (a_0, m, z^*) or (a'_0, m, z^*) is accepting for $x_0 \in L_0$. To equivocate (a_0, a'_0) to a message m', the sender simply continues the thread of the honest prover, using m' as challenge and computes z^* using the witness. Clearly, a malicious sender can open in two different ways even when $x_0 \notin L$. Nevertheless, three openings allow the extraction of the witness for x_0.

When our OR technique is instantiated with a 3-IDTC scheme we have that the resulting protocol is still WI since no power is added to the verifier. However the protocol is *not* a Σ-protocol since the special-soundness property is not guaranteed. The reason is that, in a 3-IDTC scheme the sender can open the commitment in two different ways even without having the trapdoor; i.e., the witness for $x_0 \in L_0$. Therefore, for any challenge c sent by $\mathcal{V}$, the fact that the commitment of a_1 can be opened in two ways gives a malicious prover $\mathcal{P}^*$ two chances (a_1, c, z_1) and (a'_1, c, z'_1) to successfully complete the protocol for a false statement x_1. Nevertheless, this extra freedom does not hurt soundness as both openings, a_1 and a'_1 are fixed in advance, and thus when x_1 is not an instance of the language there exist only two challenges c' and c'' that would allow $\mathcal{P}^*$ to succeed. When the challenge space is large enough the success probability of $\mathcal{P}^*$ is therefore negligible.

The OR protocol derived from a 3-IDTC scheme is 3-special sound (i.e., answering to 3 challenges allows one to compute a witness efficiently), and is a proof of knowledge when the challenge length is large enough.

1.3 Discussion

What really matters. Our new OR composition technique works only when the theorem that has not been defined yet, x_1, admits an input-delayed Σ-protocol. We stress that this is not a limitation for the applications that we have in mind. In fact, in all *efficient* protocols that make use of input-delayed proofs that we are aware of, the preamble has always the purpose of generating the trapdoor theorem. In practical scenarios[6] L_1 usually corresponds to DLog or DDH. The fact that we can not have Blum's Σ-protocol for L_1 when L_1 is the language of Hamiltonian graphs, is therefore not relevant as the actual theorem of interest is x_0.

Comparison with CDS-OR Technique. Notice that even in the extremely simplified case of considering (a) two instances x_0, x_1 for the same language L, (b) L admits an *input-delayed* Σ-protocol Π_L which is also special HVZK; (c) Π_L is *chameleon*, and thus one can compute the first message using Sim and then continue with the prover to answer to arbitrary challenges, (d) the prover knows in advance the witness w and instance x_b for which she will be able to honestly complete the protocol; the CDS-OR technique still fails in obtaining a Σ-protocol (or a WIPoK) for the OR composition of instances of L if any one of the instances is not known when the protocol starts.

Beyond Schnorr's protocol: Maurer's result. The recent work of Maurer [Mau09, Mau15] showed that a protocol (referred to as *Main Protocol*) for proving knowledge of a pre-image of a group homomorphism is the abstraction of a large class of protocols. It represents a useful level of abstraction in proofs of knowledge since it unifies and generalizes a large number of protocols in

[6]These are the only scenarios of interest for our work since if practicality is not desired than one can just rely on the LS Σ-protocol and use NP-reductions.

the literature. Classic Σ-protocols as Schnorr's protocols and the protocol of Guillou-Quisquater are particular cases of this abstraction.

We will show that the *Main Protocol* is a Σ-protocol such that even when starting with a simulator that computes the first message a, it is possible to complete the protocol computing the third round z using a witness and answering to any challenge c. Namely, Maurer's protocol satisfies our definition of chameleon Σ-protocol.

What is included and what is out. As mentioned previously, L_1 can be any language that admits an *input-delayed* Σ-protocol. We now discuss which languages can be used to instantiate L_0 in our OR transform. For this purpose, we identify four classes of Σ-protocols and we prove that any Σ-protocol that falls in any of the first three classes can be used in our OR transform (by instantiating either a 2-IDTC or a 3-IDTC scheme).

Consequently, L_0 can be any language admitting a Σ-protocol that belongs in one of the first 3 classes listed below. We also identify a class of Σ-protocols that is not suitable for any of our techniques. Luckily, we have no example of natural Σ-protocols that fall in this class, and in order to prove the separation we had to construct a very contrived scheme. The four classes are listed below.

- *(Class 1)* Σ-protocols that are Chameleon and *do not* require the witness to compute the first round. This class of Σ-protocol can be used to construct both 2-IDTC and 3-IDTC schemes.
- *(Class 2)* Σ-protocols that are Chameleon and require the prover to use the witness already to compute the first round. This class of Σ-protocol can be used to construct a 2-IDTC scheme.
- *(Class 3)* Σ-protocols that are not Chameleon but do *not* require the prover to use the witness in the first round. This class of Σ-protocol can be used to construct a 3-IDTC scheme.
- *(Class 4)* Σ-protocols that are not Chameleon and require the witness to be used already in the first round. This class of Σ-protocol can not be used in our techniques.

The input-delayed features. We stress here that our techniques allow to start and complete an efficient OR composition of two Σ-protocols (with the discussed restrictions) provided that one instance is known and another one will be known later. Having a witness for the first or the second instance always allows $\mathcal{P}$ to convince $\mathcal{V}$. This contrasts with the CDS-OR technique where knowing a witness for x_0 would block $\mathcal{P}$ immediately since $\mathcal{P}$ would need immediately x_1 to continue, but x_1 will not be available until the third round.

1.4 Applications

Our new OR technique does not provide the full power of LS because it needs one theorem to be known before the protocol starts. However, as we show below, this seemingly weaker property suffices to improve the round-complexity of some of the previous constructions based on the CDS-OR technique. Such constructions aim to efficiently[7] transform a Σ-protocol for a relation $\mathcal{R}$ into a *round-efficient* argument with more appealing features.

[7]By efficiently we mean that no NP-reduction is needed and only a constant number of modular exponentiations are added. We do not discuss the practicality of the achieved constructions.

Efficient 3-round straight-line perfect ZK with quasi-polynomial time simulation. We achieve this result directly, using the construction of Pass [Pas03] and replacing the CDS-OR technique with our technique. As a result the first round of the verifier of [Pas03] can be postponed and played along with the second round of the original protocol, therefore reducing the round complexity from four to three rounds.

Efficient 4-round resettable WI arguments. It is well known [CGGM00] how to transform a Σ-protocol into a resettable WI protocol: the verifier commits to the challenge c using a perfectly hiding commitment scheme and sends it to the prover in the first round; the prover then computes its messages with randomness derived by applying a pseudo-random function (PRF) on the commitment received. Soundness follows directly from the soundness of the Σ-protocol due to the perfect hiding of the commitment. WI follows from the fact that the protocol is zero knowledge against a stand-alone verifier and thus concurrent WI. Then the use of the PRF and the fact that all messages of the verifier are committed in advance upgrades concurrent WI to resettable WI. This approach, however, generates a 5-round protocol.

Achieving the same result *efficiently*, namely, avoiding NP-reductions, in four rounds only is non-trivial. The reason is that if we attempt to replace the 2-round perfectly hiding commitment with a non-interactive commitment, we lose the unconditional soundness property, and then it is not clear how to argue about computational soundness. More specifically, black-box extraction of the witness is not possible (black-box extraction and resettable WI can not coexist) and the adversarial prover could try to maul the commitment of the verifier and adaptively generate the first round of the Σ-protocol. In fact, even allowing complexity-leveraging arguments (and thus, straight-line extraction), constructing a 4-round WI protocol that avoids NP-reductions and adds only a few modular exponentiations to the underlying Σ-protocol has remained so far an open problem.

We solve this problem by using our new OR technique. We have the verifier commit to the challenge in the first round, but then later, instead of sending the decommitment, she will directly send the challenge and prove that either the challenge is the correct opening of the commitment or she solved some hard puzzle (in our construction, computing the Discrete Log of a random group element chosen by the prover). The puzzle is sent by the prover in the second round and it will be solved by the reduction in super-polynomial time in the proof of soundness

This trick has been proposed in literature in various forms [Pas03, DPV04] and we are using the form used in [DPV04] where the puzzle is sent only in the second round. [DPV04] must use the LS transform and therefore needs NP-reduction. As explained earlier, going through LS *was* necessary as the CDS-OR transform can be applied only if both statements are fixed at the beginning.

Our new OR-transform solves precisely this problem, and it allows the verifier to start the proof before the puzzle is defined, and this proof can be done efficiently with no NP-reduction.

To use our technique of composing efficient Σ-protocol, we need to use a specific commitment scheme that allows to prove knowledge of the opening by means of an efficient Σ-protocol. We use the simulatable commitment of [MP03] and we modify it for our setting. [MP03]'s scheme allows the sender to give a 3-round public-coin HVZK proof of the message committed. In [MP03], the theorem proved consists in the commitment *and* the claimed opening which is known to the parties at the beginning of the proof. We cannot use such a proof system directly since, in our construction, we cannot reveal the challenge (the opening) before the prover has produced the first message. For this reason, we have to modify the proof system of [MP03] so that it retains its properties even if the committed message is revealed only in the third round.

Resettable WI follows from the CGGM transformation and the WI property of the proof generated by the prover. The groups used for the commitment of the challenge and for the puzzle sent by the prover, will be chosen appropriately so that the hardness of computing discrete logarithms are different and guarantee that our reductions work (i.e., we make use of complexity leveraging).

Further applications. Our new OR-composition technique can find various other applications. Indeed, wherever there is a round-efficient (but otherwise inefficient) construction based on the use of LS without a corresponding efficient construction with the *same* round complexity, then our technique constitutes a powerful tool towards achieving efficient and round-efficient constructions. For instance, the 4-round (optimal) resettable ZK argument systems in the BPK model provided in [YZ07, SV12], consists (roughly) of the parallel execution of a (resettable) WI protocol from the prover to the verifier, where the prover proves that either $x \in L$ or he knows the secret key associated to the public identity of the verifier, and a 3-round (resettably-sound) WI protocol from the verifier to the prover, where $\mathcal{V}$ proves knowledge of the secret key associate to its public key, or knowledge of the solution of a puzzle computed by the prover. When instantiated with efficient Σ-protocols, such construction requires 5-rounds, where the additional round, from the prover to the verifier, is used to send the puzzle necessary for the verifier to start a CDS's style proof. We observe that this setting closely resembles the setting of the 4-round resettable WI ($\mathsf{r}\mathcal{WI}$) protocol that we provide in this paper. As such, one could directly instantiate the proof provided by the prover of the BPK model, with our 4-round $\mathsf{r}\mathcal{WI}$ protocol, and have the verifier just prove knowledge of its secret keys, thus avoiding the need of the additional first round.

Similarly, our new OR transform could replace the use of LS in the 4-round non-malleable commitment scheme of [GRRV14], and in the round-optimal secure two-party computation protocol of [KO04].

1.5 Open Problems

Our OR transform relaxes the requirement of CDS-OR of having to know *all* theorems already at the beginning of the protocol, however it does not match the power of LS where *no* theorem is required for the protocol to start. An immediate open question is whether one can improve our OR transform so that the first round can be run without the knowledge of any theorem.

Perhaps a first step in this direction would be to answer a related relaxed question, which is to design an OR transform for a 1-out-of-N proof of knowledge that requires the knowledge of only k-out-of-N theorems instead of $N-1$ theorem that are required by our transform.

It would also be interesting to extend our technique in order to make it applicable to *all* Σ-protocols.

2 Definitions

In this section we set-up our notation and review some standard definitions and assumptions that will be used in the paper.

If A is a probabilistic algorithm then $A(x)$ denotes the probability distribution of the output of A when it receives x as input. By $A(x; R)$ instead we denote the output of A on input x when coin tosses R are used as randomness.

A *polynomial-time relation* $\mathcal{R}$ is a subset of $\{0,1\}^\star \times \{0,1\}^\star$ for which membership of (x, w) to $\mathcal{R}$ can be decided in time polynomial in $|x|$. We define the NP-language $L_\mathcal{R}$ as $L_\mathcal{R} = \{x | \exists w : (x, w) \in \mathcal{R}\}$. If $(x, w) \in \mathcal{R}$, we say that w is a *witness* for *instance* x. Following [GMY06], we define $\hat{L}_\mathcal{R}$ to be the input language that includes both $L_\mathcal{R}$ and all well formed instances that do not have a witness. More formally, $L_\mathcal{R} \subseteq \hat{L}_\mathcal{R}$ and membership in $\hat{L}_\mathcal{R}$ can be tested in polynomial time. We implicitly assume that the verifier of a protocol for relation $\mathcal{R}$ executes the protocol only if the common input x belongs to $\hat{L}_\mathcal{R}$ and rejects immediately common inputs not in $\hat{L}_\mathcal{R}$.

For two interactive machines A and B, we denote by $\langle A(\alpha), B(\beta)\rangle(\gamma)$ the output of B after running on private input β with A using private input α, both running on common input γ.

2.1 Number Theoretic Assumptions

We define *group generator* algorithms to be probabilistic polynomial-time algorithms that take as input security parameter 1^λ and output $(\mathcal{G}, q, g)$, where $\mathcal{G}$ is (the description of) a cyclic group of order q and g is a generator of $\mathcal{G}$. We assume that membership in $\mathcal{G}$ and its group operations can be performed in time polynomial in the length of q and that there is an efficient procedure to randomly select elements from $\mathcal{G}$. Moreover, with a slight abuse of notation, we will use $\mathcal{G}$ to denote the group and its description.

We consider the sub-exponential versions of the DLog and of the DDH assumptions that posit the hardness of the computation of discrete logarithms and of breaking the Decisional Diffie-Hellman assumption with respect to the group generator algorithm IG that, on input λ, randomly selects a λ-bit prime q such that $p = 2q + 1$ is also prime and outputs the order q group $\mathcal{G}$ of the quadratic residues modulo p along with a random generator g of $\mathcal{G}$. The strong versions of the two assumptions posit the hardness of the same problems even if p (and q) and generator g are chosen adversarially. More precisely:

Assumption 1 (DLog Assumption)**.** *There exists a constant c such that for every probabilistic algorithm A running in time 2^{λ^c} the following probability is a negligible function of λ*

$$\text{Prob}\left[\ (\mathcal{G}, q, g) \leftarrow \mathsf{IG}(1^\lambda); y \leftarrow \mathbb{Z}_q : A(g^y) = y\ \right].$$

Assumption 2 (Strong DLog Assumption [CD08])**.** *Consider a pair of probabilistic algorithms (A_0, A_1) such that A_0, on input 1^λ, outputs $(\mathcal{G}, q, g)$, where $\mathcal{G}$ is the group of the quadratic residues modulo $2q + 1$, q is a λ-bit prime and $g \in \mathcal{G}$, along with some auxiliary information* $\mathtt{aux}$. *There exists a constant c such that for any such pair (A_0, A_1) running in time 2^{c^λ} the following probability is a negligible function of λ:*

$$\text{Prob}\left[\ ((\mathcal{G}, q, g), \mathtt{aux}) \leftarrow A_0(1^\lambda); y \leftarrow \mathbb{Z}_q : A_1(g^y, \mathtt{aux}) = y\ \right].$$

We next introduce the DDH Assumption and the Strong DDH Assumption which imply the DLog Assumption and the Strong DLog Assumption, respectively.

Assumption 3 (DDH Assumption)**.** *There exists a constant c_{ddh} such that, for every probabilistic algorithm A running in time $2^{\lambda^{c_{ddh}}}$, the following is a negligible function of λ*

$$\Big|\text{Prob}\left[\ (\mathcal{G}, q, g) \leftarrow \mathsf{IG}(1^\lambda); x, y, z \leftarrow \mathbb{Z}_q : A((\mathcal{G}, q, g), g^x, g^y, g^z) = 1\ \right] -$$
$$\text{Prob}\left[\ (\mathcal{G}, q, g) \leftarrow \mathsf{IG}(1^\lambda); x, y, z \leftarrow \mathbb{Z}_q : A((\mathcal{G}, q, g), g^x, g^y, g^{xy}) = 1\ \right]\Big|.$$

Assumption 4 (Strong DDH Assumption)**.** *Consider a pair of probabilistic algorithms* (A_0, A_1) *such that* A_0*, on input* 1^λ*, outputs* $(\mathcal{G}, q, g)$*, where* $\mathcal{G}$ *is the group of the quadratic residues modulo* $2q+1$*,* q *is a* λ*-bit prime and* $g \in \mathcal{G}$*, along with some auxiliary information* $\mathtt{aux}$*. There exists a constant* c_{Sddh} *such that, for any such pair* (A_0, A_1) *running in time* $2^{\lambda^{c_{Sddh}}}$*, the following is a negligible function of* λ

$$\Big|\mathrm{Prob}\Big[\ ((\mathcal{G}, q, g), \mathtt{aux}) \leftarrow A_0(1^\lambda); x, y, z \leftarrow \mathbb{Z}_q : A_1((\mathcal{G}, q, g), g^x, g^y, g^z, \mathtt{aux}) = 1\ \Big] -$$
$$\mathrm{Prob}\Big[\ ((\mathcal{G}, q, g), \mathtt{aux}) \leftarrow A_0(1^\lambda); x, y, z \leftarrow \mathbb{Z}_q : A_1((\mathcal{G}, q, g), g^x, g^y, g^{xy}, \mathtt{aux}) = 1\ \Big]\Big|.$$

3 Σ-Protocols

We consider *3-move protocols* Π for a polynomial-time relation $\mathcal{R}$. Protocol Π is played by a prover $\mathcal{P}$ and a verifier $\mathcal{V}$ that receive a common input x. Prover $\mathcal{P}$ receives as an additional private input a witness w for x. The protocol Π has the following form:

1. $\mathcal{P}$ executes algorithm P_1 on input common input x, private input w and randomness R obtaining $a = \mathsf{P}_1(x, w; R)$ and sends a to $\mathcal{V}$.
2. $\mathcal{V}$, after receiving a from $\mathcal{P}$, chooses a random *challenge* $c \leftarrow \{0,1\}^l$ and sends c to $\mathcal{P}$.
3. $\mathcal{P}$ executes algorithm P_2 on input x, w, R, c and sends $z \leftarrow \mathsf{P}_2(x, w, R, c)$ to $\mathcal{V}$.
4. $\mathcal{V}$ executes and outputs $\mathsf{V}(x, a, c, z)$ (i.e., $\mathcal{V}$'s decision to accept $(b = 1)$ or reject $(b = 0)$).

We call $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$ the algorithms *associated* with Π and l the challenge length such that, wlog, the challenge space $\{0,1\}^l$ is composed of 2^l different challenges.

The triple (a, c, z) of messages exchanged is called a *3-move transcript.* A 3-move transcript is *honest* if it corresponds to the messages computed following the honest algorithms of $\mathcal{P}$ and $\mathcal{V}$. A 3-move transcript (a, c, z) is *accepting* for x if and only if $\mathsf{V}(x, a, c, z) = 1$. Two accepting 3-move transcripts (a, c, z) and (a', c', z') for an instance x constitute a *collision* if $a = a'$ and $c \neq c'$.

Definition 1 (Σ-protocol [CDS94])**.** *A 3-move protocol* Π *with challenge length* l *is a* Σ-protocol *for a relation* $\mathcal{R}$ *if it enjoys the following properties:*

1. ***Completeness****. If* $(x, w) \in \mathcal{R}$ *then all honest 3-move transcripts for* (x, w) *are accepting.*
2. ***Special Soundness****. There exists an efficient algorithm* Extract *that, on input* x *and a collision for* x*, outputs a witness* w *such that* $(x, w) \in \mathcal{R}$*.*
3. **Special Honest-Verifier Zero Knowledge (SHVZK)**. *There exists a PPT simulator algorithm* Sim *that takes as input* $x \in L_{\mathcal{R}}$ *and* $c \in \{0,1\}^l$ *and outputs an accepting transcript for* x *where* c *is the challenge. Moreover, for all* l*-bit strings* c*, the distribution of the output of the simulator on input* (x, c) *is computationally indistinguishable from the distribution of the 3-move honest transcript obtained when* $\mathcal{V}$ *sends* c *as challenge and* $\mathcal{P}$ *runs on common input* x *and any private input* w *such that* $(x, w) \in \mathcal{R}$*. We say that a* Σ*-protocol is* Perfect *when the two distributions are identical.*

In the rest of the paper, we will call a 3-move protocol that enjoys Completeness, Special Soundness and Honest-Verifier Zero Knowledge (HVZK[8]) a $\tilde{\Sigma}$*-protocol*. The next theorem shows that SHVZK can be added to a 3-move protocol with HVZK without any significant penalty in terms of efficiency.

Theorem 1 ([Dam10]). *Suppose relation $\mathcal{R}$ admits a 3-move protocol Π' that is HVZK (resp., perfect HVZK). Then $\mathcal{R}$ admits a 3-move protocol Π that is SHVZK (resp., perfect SHVZK) and with the same efficiency.*

Proof. Let l be the challenge length of Π', let $(\mathsf{P}'_1, \mathsf{P}'_2, \mathsf{V}')$ be the algorithms associated with Π' and let Sim' be the simulator for Π'. Consider the following algorithms.

1. P_1, on input (x, w) and randomness R_1, parses R_1 as (r_1, c'') where $|c''| = l(x)$, computes $a' \leftarrow \mathsf{P}'_1(x, w; r_1)$, and outputs $a = (a', c'')$.
2. P_2, on input $(x, w) \in \mathcal{R}$, parses R_1 as (r_1, c''), c and randomness R_2, sets $c' = c \oplus c''$, computes $z' \leftarrow \mathsf{P}'_2(x, w, r_1, c'; R_2)$, and sends it to $\mathcal{V}$.
3. V, on input, x, $a = (a', c'')$, c and z' outputs the output of $\mathsf{V}'(x, a', c \oplus c'', z')$ to decide whether to accept or not.

Consider the following PPT simulator Sim that, on input an instance x and a challenge c, runs Sim' on input x and obtains (a', c', z'). Then Sim sets $c'' = c \oplus c'$ and $a = (a', c'')$ and outputs (a, c, z'). It is easy to see that if Sim' is a HVZK (resp. perfect HVZK) simulator for Π' then Sim is a SHVZK (resp. perfect SHVZK) simulator for Π. □

Definition 2 ([Dam10]). *Let $k : \{0,1\}^* \to [0,1]$ be a function. A protocol $(\mathcal{P}, \mathcal{V})$ is a* proof of knowledge *for the relation $\mathcal{R}$ with* knowledge error k *if the following properties are satisfied:*

- **Completeness** *If $\mathcal{P}$ and $\mathcal{V}$ follow the protocol on input x and private input w to $\mathcal{P}$ where $(x, w) \in \mathcal{R}$, then $\mathcal{V}$ always accepts.*
- **Knowledge soundness:** *There exists a constant $c > 0$ and a probabilistic oracle machine E, called the* extractor, *such that for every interactive prover $P^\star$ and every $x \in L_{\mathcal{R}}$, the machine E satisfies the following condition. Let $\epsilon(x)$ be the probability that $\mathcal{V}$ accepts on input x after interacting with $P^\star$. If $\epsilon(x) > k(x)$, then upon input x and oracle access to $P^\star$, the machine E outputs a string w such that $(x, w) \in \mathcal{R}$ within an expected number of steps bounded by*

$$\frac{|x|^c}{\epsilon(x) - k(x)}.$$

Theorem 2 ([Dam10]). *Let Π be a Σ-protocol for a relation $\mathcal{R}$ with challenge length l. Then Π is a proof of knowledge with knowledge error 2^{-l}.*

Definition 3 (Input-delayed Σ-protocol.). *A Σ-protocol $\Pi = (\mathcal{P}, \mathcal{V})$ with $\mathcal{P}$ running PPT algorithms $(\mathsf{P}_1, \mathsf{P}_2)$ is an* Input-delayed Σ*-protocol if P_1 takes as input only the length 1^λ of the common instance and P_2 takes as input the common instance x, the witness w, the randomness R_1 used by P_1 and the challenge c received from the verifier.*

[8]Recall that HVZK requires the existence of a simulator that generates a full transcript. This is a seemingly weaker requirement than SHVZK where the challenge is an input for the simulator.

In a *Chameleon* Σ*-protocol*, the prover can compute the first message by using the simulator and thus knowing only the input but not the witness. Once the challenge has been received, the prover can compute the last message (thus completing the interaction) by using the witness w (which is thus used only to compute the last message) and the coin tosses used by the simulator to compute the first message.

Definition 4 (Chameleon Σ-protocol.)**.** *A* Σ*-protocol* Π *for polynomial-time relation* $\mathcal{R}$ *is a* Chameleon Σ*-protocol if there exists an algorithm* $\mathsf{P}_{\mathsf{sim}}$ *satisfying the following property:*

Delayed Indistinguishability*: for all pairs of challenges* c_0 *and* c_1 *and for all* $(x, w) \in \mathcal{R}$*, the following two distributions* $\{R \leftarrow \{0,1\}^{|x|^d}; (a, z_0) \leftarrow \mathsf{Sim}(x, c_0; R);\ z_1 \leftarrow \mathsf{P}_{\mathsf{sim}}((x, c_0, a, R), w, c_1) : (x, a, c_1, z_1)\}$ *and* $\{(x, a, c_1, z_1) \leftarrow \mathsf{Sim}(x, c_1) : (x, a, c_1, z_1)\}$ *are indistinguishable, where* Sim *is the Special HVZK simulator and* d *is such that* Sim*, on input an* λ*-bit instance, uses at most* λ^d *random coin tosses. If the two distributions above are identical then we say that* Π *is a* Perfect *Chameleon* Σ*-protocol.*

We remark that a Chameleon Σ-protocol Π has two modes of operations: the standard mode is used when $\mathcal{P}$ has access to the witness from the first round; the *delayed* mode is used when $\mathcal{P}$ learns the witness only after the challenge has been received. In the delayed mode, $\mathcal{P}$ computes the first message by using the simulator Sim and then, once the challenge has been received and the witness become available, $\mathcal{P}$ computes the final message by executing algorithm $\mathsf{P}_{\mathsf{sim}}$. Moreover, observe that since Sim is a simulator for Π, it follows from the Delayed Indistinguishability property that, for all challenges c and $\tilde{c}$ and common inputs x, distribution

$$\{R \leftarrow \{0,1\}^{|x|^d}; (a, \tilde{z}) \leftarrow \mathsf{Sim}(x, \tilde{c}; R); z \leftarrow \mathsf{P}_{\mathsf{sim}}((x, \tilde{c}, a, R), w, c) : (a, c, z)\}$$

is indistinguishable from

$$\{R \leftarrow \{0,1\}^{|x|^d}; a \leftarrow \mathsf{P}_1(x, w; R); z \leftarrow \mathsf{P}_2((x, \tilde{c}, a, R), w, c) : (a, c, z)\}.$$

That is, the two modes of operations of Π are indistinguishable. This property make us able to claim that if Π is WI when a WI challenger interacts with an adversary using the algorithm $(\mathsf{P}_1, \mathsf{P}_2)$, is WI even when $(\mathsf{Sim}, \mathsf{P}_{\mathsf{sim}})$ are used. Finally, we observe that Chameleon Σ-protocols do exist and Schnorr's protocol [Sch89] is one example. When considering the associated algorithms to a Chameleon Σ-protocol, we will add $\mathsf{P}_{\mathsf{sim}}$.

3.1 Σ-Protocols and Witness Indistinguishability

Definition 5. *A* 3*-move protocol* $\Pi = (\mathcal{P}, \mathcal{V})$ *is* Witness Indistinguishable (WI) *for a relation* $\mathcal{R}$ *if, for all instances* x*, all pairs* (w, w') *of witnesses for* x *and all adversarial verifiers* $\mathcal{V}^\star$*, the distribution* $\langle \mathcal{P}(w), \mathcal{V}^\star \rangle(x)$ *is computationally indistinguishable from the distribution* $\langle \mathcal{P}(w'), \mathcal{V}^\star \rangle(x)$*.*

The notion of a *perfect* WI 3-move protocol is obtained by requiring the two distributions to be identical. We start by recalling the following result.

Theorem 3 ([CDS94])**.** *Every Perfect* $\tilde{\Sigma}$*-protocol*[9] *is Perfect WI.*

For completeness, in Appendix B we show a $\tilde{\Sigma}$-protocol that it is not WI.

[9] We remind the reader that we call a 3-move protocol that enjoys Completeness, Special Soundness and Honest-Verifier Zero Knowledge (HVZK) a $\tilde{\Sigma}$-protocol.

3.2 OR Composition of $\tilde{\Sigma}$-protocols: the CDS-OR Transform

In this section we describe the CDS-OR [CDS94] transform in details. Let Π be a $\tilde{\Sigma}$-protocol for polynomial-time relation $\mathcal{R}$ with challenge length l, associated algorithms $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$ and HVZK simulator Sim. The CDS-OR transform constructs a $\tilde{\Sigma}$-protocol Π_{OR} with associated algorithms $(\mathsf{P}_1^{\mathsf{OR}}, \mathsf{P}_2^{\mathsf{OR}}, \mathsf{V}_{\Sigma}^{\mathsf{OR}})$ for the relation

$$\mathcal{R}_{\mathsf{OR}} = \left\{((x_0, x_1), w) : \Big((x_0, w) \in \mathcal{R} \wedge x_1 \in \hat{L}_{\mathcal{R}}\Big) \,\mathsf{OR}\, \Big((x_1, w) \in \mathcal{R} \wedge x_0 \in \hat{L}_{\mathcal{R}}\Big)\right\}.$$

We describe Π_{OR} below.

Protocol 1. *CDS-OR Transform.*
Common input: (x_0, x_1).
$\mathcal{P}$'s private input: (b, w) *with* $b \in \{0,1\}$ *and* $(x_b, w) \in \mathcal{R}$.

$\mathsf{P}_1^{\mathsf{OR}}((x_0, x_1), (b, w); R_1)$**.** Set $a_b = \mathsf{P}_1(x_b, w; R_1)$. Compute $(a_{1-b}, c_{1-b}, z_{1-b}) \leftarrow \mathsf{Sim}(x_{1-b})$. Output (a_0, a_1).

$P_2^{\mathsf{OR}}((x_0, x_1), (b, w), c, R_1)$**.** Set $c_b = c \oplus c_{1-b}$. Compute $z_b \leftarrow \mathsf{P}_2(x_b, w, c_b, R_1)$. Output $((c_0, c_1), (z_0, z_1))$.

$\mathsf{V}_{\Sigma}^{\mathsf{OR}}((x_0, x_1), (a_0, a_1), c, ((c_0, c_1), (z_0, z_1)))$**.** $\mathsf{V}_{\Sigma}^{\mathsf{OR}}$ accepts if and only if $c = c_0 \oplus c_1$ and $\mathsf{V}(x_0, a_0, c_0, z_0) = 1$ and $\mathsf{V}(x_1, a_1, c_1, z_1) = 1$.

Theorem 4 ([CDS94, GMY06])**.** *If* Π *is a* $\tilde{\Sigma}$*-protocol for* $\mathcal{R}$ *then* Π_{OR} *is a* $\tilde{\Sigma}$*-protocol for* $\mathcal{R}_{\mathsf{OR}}$ *and is WI for relation*

$$\mathcal{R}'_{\mathsf{OR}} = \left\{((x_0, x_1), w) : \Big((x_0, w) \in \mathcal{R} \wedge x_1 \in L_{\mathcal{R}}\Big) \,\mathsf{OR}\, \Big((x_1, w) \in \mathcal{R} \wedge x_0 \in L_{\mathcal{R}}\Big)\right\}.$$

Moreover, if Π *is a Perfect* $\tilde{\Sigma}$*-protocol for* $\mathcal{R}$ *then* Π^{OR} *is WI for* $\mathcal{R}_{\mathsf{OR}}$*.*

It is possible to extend the above construction to handle two different relations $\mathcal{R}_0$ and $\mathcal{R}_1$ that admit $\tilde{\Sigma}$-protocols. Indeed by Theorem 14, we can assume, wlog, that $\mathcal{R}_0$ and $\mathcal{R}_1$ have $\tilde{\Sigma}$-protocols Π_0 and Π_1 with the same challenge length. Hence, the construction outlined above can be used to construct $\tilde{\Sigma}$-protocol $\Pi_{\mathsf{OR}}^{\mathcal{R}_0, \mathcal{R}_1}$ for relation

$$\mathcal{R}_{\mathsf{OR}} = \left\{((x_0, x_1), w) : \Big((x_0, w) \in \mathcal{R}_0 \wedge x_1 \in \hat{L}_{\mathcal{R}_1}\Big) \,\mathsf{OR}\, \Big((x_1, w) \in \mathcal{R}_1 \wedge x_0 \in \hat{L}_{\mathcal{R}_0}\Big)\right\}.$$

We have the following theorem.

Theorem 5. *If* Π_0 *and* Π_1 *are* $\tilde{\Sigma}$*-protocols for* $\mathcal{R}_0$ *and* $\mathcal{R}_1$*, respectively, then* $\Pi_{\mathsf{OR}}^{\mathcal{R}_0, \mathcal{R}_1}$ *is a* $\tilde{\Sigma}$*-protocol for relation* $\mathcal{R}_{\mathsf{OR}}$ *and is WI for relation*

$$\mathcal{R}'_{\mathsf{OR}} = \left\{((x_0, x_1), w) : \Big((x_0, w) \in \mathcal{R}_0 \wedge x_1 \in L_{\mathcal{R}_1}\Big) \,\mathsf{OR}\, \Big((x_1, w) \in \mathcal{R}_1 \wedge x_0 \in L_{\mathcal{R}_0}\Big)\right\}.$$

Moreover, if Π_0 *and* Π_1 *are Perfect* $\tilde{\Sigma}$*-protocols for* $\mathcal{R}_0$ *and* $\mathcal{R}_1$ *then* Π^{OR} *is WI for* $\mathcal{R}_{\mathsf{OR}}$*.*

We remark that if Π_0 and Π_1 are Σ-protocols then the CDS-OR transform yields a Σ-protocol for $\mathcal{R}_{\mathsf{OR}}$ and the equivalent of Theorem 5 (and of Theorem 4) holds.

4 t-Instance-Dependent Trapdoor Commitment Schemes

In this section, for integer $t \geq 2$, we define the notion of a t-Instance-Dependent Trapdoor Commitment scheme associated with a polynomial-time relation $\mathcal{R}$ and show constructions for $t = 2$ and $t = 3$.

Definition 6 (t-Instance-Dependent Trapdoor Commitment Scheme). *Let $t \geq 2$ be an integer and let $\mathcal{R}$ be a polynomial-time relation. A* t-Instance-Dependent Trapdoor Commitment *Scheme (a t-IDTC, in short) for $\mathcal{R}$ with message space M is a triple of PPT algorithms* $(\mathsf{TCom}, \mathsf{TDec}, \mathsf{TFake})$ *where* TCom *is the randomized* commitment *algorithm that takes as input security parameter 1^λ, an instance $x \in \hat{L}_{\mathcal{R}}$ and a message $m \in M$ and outputs* commitment $\mathtt{com}$, decommitment $\mathtt{dec}$, *and* auxiliary information $\mathtt{rand}$*;* TDec *is the* verification *algorithm that takes as input $(x, \mathtt{com}, \mathtt{dec}, m)$ and decides whether m is the decommitment of $\mathtt{com}$;* TFake *is the randomized* equivocation *algorithm that takes as input $(x, w) \in \mathcal{R}$, messages m_1 and m_2 in M, commitment $\mathtt{com}$ of m_1 with respect to instance x and associated auxiliary information $\mathtt{rand}$ and produces decommitment information $\mathtt{dec}_2$ such that TDec, on input $(x, \mathtt{com}, \mathtt{dec}_2, m_2)$, outputs 1.*

A t-Instance-Dependent Trapdoor Commitment scheme has the following properties:

- **Correctness***: for all $x \in \hat{L}_{\mathcal{R}}$, all $m \in M$, it holds that*

$$\text{Prob}\left[\ (\mathtt{com}, \mathtt{dec}, \mathtt{rand}) \leftarrow \mathsf{TCom}(1^\lambda, x, m) : \mathsf{TDec}(x, \mathtt{com}, \mathtt{dec}, m) = 1\ \right] = 1.$$

- **t-Special Extract***: there exists a PPT algorithm* $\mathsf{ExtractTCom}$ *that, on input x, commitment $\mathtt{com}$, pairs $(\mathtt{dec}_i, m_i)_{i=1}^t$ of openings and messages such that*
  - *for $1 \leq i < j \leq t$ we have that $m_i \neq m_j$;*
  - *$\mathsf{TDec}(x, \mathtt{com}, \mathtt{dec}_i, m_i) = 1$, for $i = 1, \ldots, t$;*

  outputs w such that $(x, w) \in \mathcal{R}$.

- **Hiding (resp., Perfect Hiding)***: for every PPT (resp., unbounded) adversary $\mathcal{A}$ there exists a negligible function ν (resp., $\nu(\lambda) = 0$) such that, for all $x \in L_{\mathcal{R}}$ and all $m_0, m_1 \in M$, it holds that*

$$\text{Prob}\left[\ b \leftarrow \{0,1\}; (\mathtt{com}, \mathtt{dec}, \mathtt{rand}) \leftarrow \mathsf{TCom}(1^\lambda, x, m_b) : b = \mathcal{A}(\mathtt{com})\ \right] \leq \frac{1}{2} + \nu(\lambda).$$

- **Trapdoor***: the following two families of probability distributions are indistinguishable:*

$$\{(\mathtt{com}, \mathtt{dec}_1, \mathtt{rand}) \leftarrow \mathsf{TCom}(1^\lambda, x, m_1); \mathtt{dec}_2 \leftarrow \mathsf{TFake}(x, w, m_1, m_2, \mathtt{com}, \mathtt{rand}) : (x, \mathtt{com}, \mathtt{dec}_2, m_2)\}$$

  and

$$\{(\mathtt{com}, \mathtt{dec}_2, \mathtt{rand}) \leftarrow \mathsf{TCom}(1^\lambda, x, m_2) : (x, \mathtt{com}, \mathtt{dec}_2, m_2)\}$$

  over all families $\{(x, w, m_1, m_2)\}$ such that $(x, w) \in \mathcal{R}$ and $m_1, m_2 \in M$.

  The Perfect Trapdoor *property requires the two probability distributions to coincide for all (x, w, m_1, m_2) such that $(x, w) \in \mathcal{R}$ and $m_1, m_2 \in M$.*

Constructing a 2-IDTC scheme from a Chameleon Σ-protocol. Let $\Pi = (\mathcal{P}, \mathcal{V})$ with associated algorithms $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V}, \mathsf{P}_{\mathsf{sim}})$ be a Chameleon Σ-protocol for polynomial-time relation $\mathcal{R}$. Let l be the challenge length of protocol Π and let Sim be a SHVZK simulator associated to Π. We construct a t-IDTC scheme $(\mathsf{TCom}_\Pi, \mathsf{TDec}_\Pi, \mathsf{TFake}_\Pi)$ for $\mathcal{R}$ with messages space $M = \{0,1\}^l$ for $x \in \hat{L}_R$ as follows.

Protocol 2. 2-IDTC scheme from Chameleon Σ-protocol Π.

- $\mathsf{TCom}_\Pi(x, m_1)$*: On input* x *and* $m_1 \in M$*, pick randomness* R *and compute* $(a, z) \leftarrow \mathsf{Sim}(x, m_1; R)$*. Output* $\mathtt{com} = a$*,* $\mathtt{dec} = z$ *and* $\mathtt{rand} = R$*;*

- $\mathsf{TDec}_\Pi(m_1, x, \mathtt{com}, \mathtt{dec})$*: On input* $x, \mathtt{com}, \mathtt{dec}$ *and* m_1*, run* $b = \mathsf{V}(x, \mathtt{com}, m_1, \mathtt{dec})$ *and accept* m_1 *as the decommitted message iff* $b = 1$*.*

- TFake_Π*: On input* $(x, w) \in \mathcal{R}$*, messages* $m_1, m_2 \in M$*, commitment* $\mathtt{com}$ *for* m_2 *and* $\mathtt{rand}$ *for* $\mathtt{com}$*, output* $z = \mathsf{P}_{\mathsf{sim}}(x, (m_1, \mathtt{com}, \mathtt{rand}), w, m_2)$*.*

Theorem 6. *If* Π *is a Chameleon* Σ*-protocol for* $\mathcal{R}$ *then Protocol 2 is a* 2-IDTC *scheme for* $\mathcal{R}$*. Moreover, if* Π *is Perfect then so is Protocol 2.*

Proof. Correctness follows directly from the Completeness property of Π.
2-Special-Extract. Suppose $\mathtt{com}$ is a commitment with respect to instance x and let $\mathtt{dec}_1$ and $\mathtt{dec}_2$ be two openings of $\mathtt{com}$ as messages $m_1 \neq m_2$, respectively. Then, triplets $(\mathtt{com}, m_1, \mathtt{dec}_1)$ and $(\mathtt{com}, m_2, \mathtt{dec}_2)$ are accepting transcripts for Π on common input x with the same first round; that is, they constitute a collision for Π. Therefore, we define algorithm $\mathsf{ExtractTCom}$ to be the algorithm that runs algorithm $\mathsf{Extract}$ (that exists by the special soundness of Π) on input the collision. $\mathsf{ExtractTCom}$ returns the witness for x computed by $\mathsf{Extract}$.
(Perfect) Hiding. Hiding follows from the (perfect) SHVZK property of Π.
(Perfect) Trapdoorness. It follows from the Delayed Indistinguishability property of Π as well as the (perfect) Hiding property.

□

Constructing a 3-IDTC scheme. Let $\mathcal{R}$ be a polynomial-time relation as above admitting a Σ-protocol Π with associated algorithms $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$ in which the prover uses the witness only in the third round (we stress that this property is typically enjoyed by most Σ-protocols). Let l denote the challenge length of Π. We construct a 3-IDTC scheme for message space $M = \{0,1\}^l$ for $x \in \hat{L}_R$, as follows.

Protocol 3. 3-IDTC scheme.

- TCom_Π**:** *On input* x *and* $m_1 \in M$*, pick randomness* R *and compute* $(a_0, z) \leftarrow \mathsf{Sim}(x, m_1)$ *and* $a_1 \leftarrow \mathsf{P}_1(x; R)$*. Let* $\mathtt{com}_0 = a_0$ *and* $\mathtt{com}_1 = a_1$*. Output* $\mathtt{com} = (\mathtt{com}_b, \mathtt{com}_{1-b})$ *for a randomly selected bit* b*,* $\mathtt{dec} = z$ *and* $\mathtt{rand} = R$*.*

- TDec_Π**:** *On input* $x, \mathtt{com} = (\mathtt{com}_0, \mathtt{com}_1), \mathtt{dec}$ *and* m_1*, accept* m_1 *if and only if either* $\mathsf{V}(x, \mathtt{com}_0, m_1, \mathtt{dec}) = 1$ *or* $\mathsf{V}(x, \mathtt{com}_1, m_1, \mathtt{dec}) = 1$*.*

- TFake_Π**:** *On input* $(x, w) \in \mathcal{R}$*, messages* $m_1, m_2 \in M$*, commitment* $\mathtt{com}$ *for* m_1 *and* $\mathtt{rand}$ *for* $\mathtt{com}$*, output* $z \leftarrow \mathsf{P}_2(x, w, \mathtt{rand}, m_2)$*.*

Theorem 7. *If Π is a Σ-protocol, with the associated algorithms $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$, for $\mathcal{R}$ where P_1 can be executed without using the witness then Protocol 3 is a* 3-IDTC *scheme for $\mathcal{R}$. Moreover, if Π is Perfect then so is Protocol 3.*

Proof. Correctness follows from the completeness of Π.

3-*Special Extract.* It follows from the special soundness of Π. Assume that the committer generates 3 accepting openings $\mathtt{dec}_1$, $\mathtt{dec}_2$ and $\mathtt{dec}_3$, for distinct messages m_1, m_2 and m_3, for the same commitment $\mathtt{com}$ computed w.r.t. x. In this case, we have three accepting conversations for Π and therefore at least two of them must share the same first message, i.e., it is a collision. Thus we can run the extractor $\mathsf{Extract}$ for Π on the collision and obtain a witness for x.

Trapdoorness. It follows from the SHVZK property of Π. We prove this property via hybrid arguments.

The first hybrid, $\mathcal{H}_1$ is the real execution, where a honest prover commits to a message following the honest commitment and decommitment procedure, without using the trapdoor. More formally, in the hybrid $\mathcal{H}_1$ the prover performs the following steps:

- The prover on input x and $m_1, m_2 \in M$, selects random coin tosses R and computes $(a_0, z) \leftarrow \mathsf{Sim}(x, m_2)$, $a_1 \leftarrow \mathsf{P}_1(x; R)$. It picks $b \leftarrow \{0,1\}$ and sends $\mathtt{com} = (a_b, a_{1-b})$, $\mathtt{dec} = z$, m_2.

The second hybrid $\mathcal{H}_2$ is equal to $\mathcal{H}_1$ with the difference that a_0 is computed using the algorithm P_1 and z using P_2. Formally:

- On input x and $m_1, m_2 \in M$, the prover selects random coin tosses $R = (r_1, r_2)$ and computes $a_0 \leftarrow \mathsf{P}_1(x; r_1)$, $z \leftarrow \mathsf{P}_2(x, w, r_1, m_2)$ and $a_1 \leftarrow \mathsf{P}_1(x; r_2)$. It picks $b \leftarrow \{0,1\}$ and sends $\mathtt{com} = (a_b, a_{1-b})$, $\mathtt{dec} = z$, m_2.

Due to the SHVZK property of Π, $\mathcal{H}_1$ is indistinguishable from $\mathcal{H}_2$. Now we consider the hybrid $\mathcal{H}_3$ in which a_1 is computed using $\mathsf{Sim}(x, m_2)$. Formally:

- On input x and $m_1, m_2 \in M$, the prover selects random coin tosses R and sets $a_0 \leftarrow \mathsf{P}_1(x; R)$, $z \leftarrow \mathsf{P}_2(x, w, R, m_2)$ and $(a_1, \overline{z}) \leftarrow \mathsf{Sim}(x, m_1)$. It picks $b \leftarrow \{0,1\}$ and sends $\mathtt{com} = (a_b, a_{1-b})$, $\mathtt{dec} = z$, m_2.

Even in this case, we can claim that $\mathcal{H}_3$ is indistinguishable from $\mathcal{H}_2$ because of the SHVZK of Π. The proof ends with the observation that $\mathcal{H}_3$ is the experiment in which a sender commits to a message m_2 and opens to m_1 using the trapdoor.

If Π is a perfect SHVZK protocol, then the sequence of hybrids produces identical distributions. $\square$

5 Our New OR-Composition Technique

In this section we formally describe our new OR transform. Let $\mathcal{R}_0$ be a relation admitting a t-IDTC scheme, $I = (\mathsf{TCom}_{\Pi_0}, \mathsf{TDec}_{\Pi_0}, \mathsf{TFake}_{\Pi_0})$, with $t = 2$ or $t = 3$, and $\mathcal{R}_1$ a relation admitting an Input-delayed Σ-protocol Π_1 with associated algorithms $(\mathsf{P}_1^1, \mathsf{P}_2^1, \mathsf{V}^1)$ and simulator Sim^1. We show a Σ-protocol Π^{OR} for the OR relation:

$$\mathcal{R}_{\mathsf{OR}} = \{((x_0, x_1), w) : ((x_0, w) \in \mathcal{R}_0 \wedge x_1 \in \hat{L}_{\mathcal{R}_1}) \text{ OR } ((x_1, w) \in \mathcal{R}_1 \wedge x_0 \in \hat{L}_{\mathcal{R}_0})\}.$$

We denote by $(\mathsf{P}_1^{\mathsf{OR}}, \mathsf{P}_2^{\mathsf{OR}}, \mathsf{V}^{\mathsf{OR}})$ the algorithms associated with Π^{OR}. We assume that the initial common input is x_0. The other input x_1 and the witness w for (x_0, x_1) are made available to the prover only after the challenge has been received. We let $b \in \{0,1\}$ be such that $(x_b, w) \in \mathcal{R}_b$ and assume that the message space of the t-IDTC scheme I includes all possible first-round messages of Π_1. Note that for the constructions of the t-IDTC scheme we provide, the message space coincides with the set of challenges of the underlying Σ-protocol and, in Appendix B.1, we show that the challenge length of a Σ-protocol can be easily expanded/reduced.

We remind that prover algorithm $\mathsf{P}_2^{\mathsf{OR}}$ receives as further input the randomness $(R_1, \mathtt{rand}_1)$ used by $\mathsf{P}_1^{\mathsf{OR}}$ to produce the first-round message.

Protocol 4. Protocol Π_{OR} for $\mathcal{R}_{\mathsf{OR}}$.

Common input: $(x_0, 1^\lambda)$, *where λ is the length of the (still to come) instance of* $\hat{L}_{\mathcal{R}_1}$.

1. $\mathsf{P}_1^{\mathsf{OR}}(x_0, 1^\lambda)$. *Pick random R_1 and compute* $a_1 \leftarrow \mathsf{P}_1^1(1^\lambda; R_1)$. *Then commit to a_1 by running* $(\mathtt{com}, \mathtt{dec}_1, \mathtt{rand}_1) \leftarrow \mathsf{TCom}_{\Pi_0}(x_0, a_1)$. *Output* $\mathtt{com}$.
2. $\mathsf{P}_2^{\mathsf{OR}}((x_0, x_1), c, (w, b), (\mathtt{rand}_1, R_1))$ $(x_b, w) \in \mathcal{R}_b$.

   If $b = 1$, compute $z_1 \leftarrow \mathsf{P}_2^1(x_1, w, c, R_1)$ *and output* $(\mathtt{dec}_1, a_1, z_1)$.

   If $b = 0$, compute $(a_2, z_2) \leftarrow \mathsf{Sim}^1(x_1, c)$, $\mathtt{dec}_2 \leftarrow \mathsf{TFake}_{\Pi_0}(x_0, w, a_1, a_2, \mathtt{rand}_1)$ *and output* $(\mathtt{dec}_2, a_2, z_2)$.
3. V^{OR}, *on input* (x_0, x_1), $\mathtt{com}$, c, *and* $(\mathtt{dec}, a, z))$ *received from* Π^{OR}, *outputs* 1 *iff*

$$\mathsf{TDec}_{\Pi_0}(x_0, \mathtt{com}, a, \mathtt{dec}) = 1 \text{ and } \mathsf{V}^1(x_1, a, c, z) = 1;$$

Theorem 8. *If $\mathcal{R}_0$ admits a* 2-IDTC *(resp.,* 3-IDTC*) scheme and if $\mathcal{R}_1$ admits an Input-delayed Σ-protocol, then Π^{OR} is a Σ-protocol (resp., is a 3-round public-coin SHVZK PoK) for relation* $\mathcal{R}_{\mathsf{OR}}$.

Proof. Completeness follows by inspection. We next prove the properties of Protocol 4 when instantiated with a 2-IDTC and 3-IDTC schemes.

Proof for the construction based on the 2-IDTC scheme.

Special Soundness. It follows from the special soundness of the underlying Σ-protocol Π_1 and the 2-Special Extract of the 2-IDTC scheme. More formally, consider a collision $(\mathtt{com}, c, (z, a, \mathtt{dec}))$ and $(\mathtt{com}, c', (z', a', \mathtt{dec}'))$ for input (x_0, x_1). We observe that:

- if $a = a'$ then (a, c, z) and (a', c', z') is a collision for Π_1 for input x_1; then we can obtain a witness w_1 for x_1 by the Special Soundness property of Π_1;
- if $a \neq a'$, then $\mathtt{dec}$ and $\mathtt{dec}'$ are two openings of $\mathtt{com}$ with respect to x_0 for messages $a \neq a'$; then we can obtain a witness w_0 by the 2-Special Extract of the 2-IDTC scheme.

SHVZK property. Consider simulator $\mathsf{Sim}^{\mathsf{OR}}$ that, on input (x_0, x_1) and challenge c, sets $(a, c, z) \leftarrow \mathsf{Sim}_1(x_1, c)$ and $(\mathtt{com}, \mathtt{dec}) \leftarrow \mathsf{TCom}_{x_0}(a)$, and outputs $(\mathtt{com}, c, (\mathtt{dec}, a, z))$. Next, we show that the transcript generated by $\mathsf{Sim}^{\mathsf{OR}}$ is indistinguishable from the one generated by a honest prover.

Let us first consider the case in which the prover of Π^{OR} receives a witness for x_1. In this case, if we sample a random distribution $(\mathtt{com}, c, (\mathtt{dec}, a, z))$ of Π^{OR} on input (x_0, x_1) constrained to c being the challenge we have that (a, c, z) has the same distribution as in random conversation of Π_1 on input x_1 constrained to c being the challenge; moreover, $(\mathtt{com}, \mathtt{dec})$ is a pair of commitment and decommitment of a with respect to x_0. By the property of Sim_1, this distribution is indistinguishable from (a, c, z) computed as $\mathsf{Sim}_1(x_1, c)$ which is exactly as in the output $\mathsf{Sim}^{\mathsf{OR}}$.

Let us now consider the case in which the prover of Π^{OR} receives a witness for x_0. If we sample a random distribution $(\mathtt{com}, c, (\mathtt{dec}, a, z))$ of Π^{OR} on input (x_0, x_1) constrained to c being the challenge we have that (a, c, z) are distributed exactly as in the output of $\mathsf{Sim}^{\mathsf{OR}}$ (that is by running Sim_1 on input x_1 and c). In addition, in the output of $\mathsf{Sim}^{\mathsf{OR}}$, $(\mathtt{com}, \mathtt{dec})$ are commitment and decommitment of a whereas in the view of Π^{OR} they are computed by means of TFake algorithm. However, the two distributions are indistinguishable by the property of the Instance-Dependent Trapdoor Commitment.

Proof for the construction based on the 3-IDTC scheme.

3-*Special Soundness.* This property ensures that there exists an efficient algorithm that, given three accepting transcripts, (a, c_0, z_0), (a, c_1, z_1), (a, c_2, z_2) with $c_i \neq c_j$ for $1 \leq i < j \leq 3$, for the same common input, outputs a witness for x. Note that 3-Special Soundness implies the Proof of Knowledge property.

Consider three accepting transcripts for Π^{OR} and input (x_0, x_1):

$$(\mathtt{com}, c_1, (z_1, a_1, \mathtt{dec}_1)), \qquad (\mathtt{com}, c_2, (z_2, a_2, \mathtt{dec}_2))$$

and

$$(\mathtt{com}, c_3, (z_3, a_3, \mathtt{dec}_3)).$$

We observe that:

- if $a_i = a_j$ for some $i \neq j$ then (a_i, c_i, z_i) and (a_j, c_j, z_j) is a collision for Π_1 for input x_1; thus we can obtain a witness w_1 for x_1 by the Special Soundness property of Π_1;
- if $a_i \neq a_j$ for all $i \neq j$, then, $\mathtt{dec}_1$ and $\mathtt{dec}_2$ and $\mathtt{dec}_3$ are three openings of the same $\mathtt{com}$ with respect to x_0 for messages a_1, a_2 and a_3; then we can obtain a witness w_0 by the 3-Special Extract of the 3-IDTC scheme.

SHVZK property. Similar to the proof for the construction based on 2-IDTC. □

5.1 Witness Indistinguishability of Our Transform

In this section we prove that Π^{OR} is *adaptive WI*. Roughly speaking, adaptive means that, in the WI experiment, the adversary $\mathcal{A}$ does not have to choose *both* theorems x_0 and x_1 before the protocol starts. Rather, she can choose theorem x_1 and the witnesses w_0, w_1, adaptively, *after* seeing the first message of Π^{OR} played by the prover on input x_0. After x_1, w_0, w_1 have been selected by $\mathcal{A}$, the experiment randomly selects $b \leftarrow \{0, 1\}$ and the prover completes the protocol on input x_1 and w_b. The adversary wins the game if she can guess b with probability non-negligibly greater than 1/2. More formally, we prove that Π^{OR} is adaptive WI for

$$\mathcal{R}^c_{\mathsf{OR}} = \{((x_0, x_1), w) : ((x_0, w) \in \mathcal{R}_0 \wedge x_1 \in L_{\mathcal{R}_1}) \,\mathsf{OR}\, ((x_1, w) \in \mathcal{R}_1 \wedge x_0 \in L_{\mathcal{R}_0})\}$$

and that this can be strengthened to

$$\mathcal{R}^p_{\mathsf{OR}} = \Big\{((x_0,x_1),w) : \Big((x_0,w) \in \mathcal{R}_0 \wedge x_1 \in \hat{L}_{\mathcal{R}_1}\Big) \;\mathsf{OR}\; \Big((x_1,w) \in \mathcal{R}_1 \wedge x_0 \in \hat{L}_{\mathcal{R}_0}\Big)\Big\}$$

in case Π^{OR} is Perfect HVZK.

The adaptive WI experiment, $\mathsf{ExpWI}^{\delta}_{\mathcal{A}}(x_0, \mathsf{aux})$ with $\delta \in \{c, p\}$, is parameterized by PPT adversary $\mathcal{A}$ and has two inputs: x_0 and auxiliary information z for $\mathcal{A}$.

$\mathsf{ExpWI}^{\delta}_{\mathcal{A}}(x_0, \mathsf{aux})$:

1. $a = \mathsf{P}^{\mathsf{OR}}_1(x_0; R_1)$, for random coin tosses R_1;
2. $\mathcal{A}(x_0, a, \mathsf{aux})$ outputs $((x_1, w_0, w_1), c, \mathtt{state})$ such that $((x_0, x_1), w_0), ((x_0, x_1), w_1) \in \mathcal{R}^{\delta}_{\mathsf{OR}}$;
3. $b \leftarrow \{0, 1\}$;
4. $z \leftarrow \mathsf{P}^{\mathsf{OR}}_2((x_0, x_1), w_b, R_1, c)$;
5. $b' \leftarrow \mathcal{A}(x_0, a, z, \mathtt{state})$;
6. If $b = b'$ then output 1 else output 0.

We stress that when Π^{OR} is only Computational HVZK (that is, $\delta = c$) then it is adaptive WI when restricted to inputs $x_0 \in L_{\mathcal{R}_0}$ and $x_1 \in L_{\mathcal{R}_1}$; this restriction is not necessary for the perfect case. We notice that this is exactly the same as the OR composition result of [CDS94] (the proof for the perfect case appears in [CDS94] and the one for the computational case in [GMY06]).

We set

$$\mathsf{Adv}^{\delta}_{\mathcal{A}}(x_0, \mathsf{aux}) = \left|\text{Prob}\left[\ \mathsf{ExpWI}^{\delta}_{\mathcal{A}}(x_0, \mathsf{aux}) = 1\ \right] - \frac{1}{2}\right|$$

and we have the following theorems.

Theorem 9. *For every PPT $\mathcal{A}$ with auxiliary information* aux*, there exists a negligible function ν such that for all $x_0 \in L_{\mathcal{R}_0}$*

$$\mathsf{Adv}^{c}_{\mathcal{A}}(x_0, \mathsf{aux}) \le \nu(|x_0|).$$

Proof. Let $\mathcal{A}$ be an adversary with non-negligible $\mathsf{Adv}^{c}_{\mathcal{A}}$. Without loss of generality, we can assume that $\mathcal{A}$ outputs w_0 and w_1 such that $(x_0, w_0) \in \mathcal{R}_0$ and $(x_1, w_1) \in \mathcal{R}_1$. Suppose in fact that both w_0 and w_1 output by $\mathcal{A}$ are witnesses for x_0. Then, by definition of $\mathcal{R}^{c}_{\mathsf{OR}}$, $x_1 \in L_{\mathcal{R}_1}$ and thus there exists a witness w_2 for x_1. The output of $\mathcal{V}^{\star}$ interacting with $\mathcal{P}((x_0, x_1), w_2)$ would necessarily be distinguishable from the output of the interaction with either $\mathcal{P}((x_0, x_1), w_0)$ or $\mathcal{P}((x_0, x_1), w_1)$. Then one could use w_0 and w_2 or w_1 and w_2 as witnesses.

Let us then consider the following hybrids. The first hybrid experiment $\mathcal{H}_1(x_0, \mathsf{aux})$ is the original experiment $\mathsf{ExpWI}^{c}_{\mathcal{A}}(x_0, \mathsf{aux})$ in which $b = 1$ (and thus $\mathcal{P}$ uses witness w_1). That is,

- In Step 1 of $\mathsf{ExpWI}^{c}_{\mathcal{A}}(x_0, \mathsf{aux})$, the following steps are executed:
  1. $a = \mathsf{P}^1_1(1^{\lambda}; R_1)$, for random coin tosses R_1;
  2. $(\mathtt{com}, \mathtt{dec}, \mathtt{rand}) \leftarrow \mathsf{TCom}_{\Pi_0}(x_0, a)$ and outputs $\mathtt{com}$.
- In Step 4 of $\mathsf{ExpWI}^{c}_{\mathcal{A}}(x_0, \mathsf{aux})$, the following steps are executed:

1. set $a' = a$;
2. $z \leftarrow \mathsf{P}_2^1(x_1, w_1, c, R_1)$;
3. set $\mathtt{dec}' = \mathtt{dec}$;
4. output $(\mathtt{dec}', a', z)$.

The second hybrid experiment $\mathcal{H}_2(x_0, \mathsf{aux})$ differs from $\mathcal{H}_1(x_0, \mathsf{aux})$ in the way a' and $\mathtt{dec}'$ are computed. More specifically,

- Step 1 of $\mathsf{ExpWI}^c_{\mathcal{A}}(x_0, \mathsf{aux})$ stays the same.
  1. $a = \mathsf{P}_1^1(1^\lambda; R_1)$, for random coin tosses R_1;
  2. $(\mathtt{com}, \mathtt{dec}, \mathtt{rand}) \leftarrow \mathsf{TCom}_{\Pi_0}(x_0, a)$ and outputs $\mathtt{com}$.
- In Step 4 of $\mathsf{ExpWI}^c_{\mathcal{A}}(x_0, \mathsf{aux})$, the following steps are executed:
  1. $a' = \mathsf{P}_1^1(1^\lambda; R_1')$, for random coin tosses R_1';
  2. $z \leftarrow \mathsf{P}_2^1(x_1, w_1, c, R_1')$;
  3. $\mathtt{dec}' \leftarrow \mathsf{TFake}_{\Pi_0}(x_0, w_0, a, a', \mathtt{com}, \mathtt{rand})$;
  4. $(\mathtt{dec}', a', z)$.

The Trapdoor property of the Instance-Dependent Trapdoor Commitment based on Π_0 guarantees that $\mathcal{H}_1(x_0, \mathsf{aux})$ and $\mathcal{H}_2(x_0, \mathsf{aux})$ are indistinguishable.

The third hybrid experiment $\mathcal{H}_3(x_0, \mathsf{aux})$ differs from $\mathcal{H}_2(x_0, \mathsf{aux})$ in the way a' and z are computed. More specifically,

- Step 1 of $\mathsf{ExpWI}^c_{\mathcal{A}}(x_0, \mathsf{aux})$ stays the same.
  1. $a = \mathsf{P}_1^1(1^\lambda; R_1)$, for random coin tosses R_1;
  2. $(\mathtt{com}, \mathtt{dec}, \mathtt{rand}) \leftarrow \mathsf{TCom}_{\Pi_0}(x_0, a)$ and outputs $\mathtt{com}$.
- In Step 4 of $\mathsf{ExpWI}^c_{\mathcal{A}}(x_0, \mathsf{aux})$, the following steps are executed:
  1. $(a', z) \leftarrow \mathsf{Sim}^1(x_1, c)$;
  2. $\mathtt{dec}' \leftarrow \mathsf{TFake}_{\Pi_0}(x_0, w_0, a, a', \mathtt{com}, \mathtt{rand})$;
  3. $(\mathtt{dec}', a', z)$.

By the SHVZK of Π^1 we have that $\mathcal{H}_2(x_0, \mathsf{aux})$ and $\mathcal{H}_3(x_0, \mathsf{aux})$ are indistinguishable. The proof ends with the observation that $\mathcal{H}_3(x_0, \mathsf{aux})$ is exactly the experiment $\mathsf{ExpWI}^c_{\mathcal{A}}(x_0, \mathsf{aux})$ when $b = 0$. □

Theorem 10. *If Π_0 and Π_1 are Perfect then so is Π^{OR} and for all PPT $\mathcal{A}$ with auxiliary information* aux *and for all $x_0 \in L_{\mathcal{R}_0}$ we have that*

$$\mathsf{Adv}^p_{\mathcal{A}}(x_0, \mathsf{aux}) = 0.$$

Proof. We prove this theorem considering the following three cases:

1. $(x_0, w_0) \in \mathcal{R}_0$ and $(x_1, w_1) \in \mathcal{R}_1$;

2. $(x_0, w_0) \in \mathcal{R}_0$ and $(x_0, w_1) \in \mathcal{R}_0$;
3. $(x_1, w_0) \in \mathcal{R}_1$ and $(x_1, w_1) \in \mathcal{R}_1$.

For each case we present a sequence of hybrids and prove that pairs of consecutive hybrids are perfectly indistinguishable.

Case 1. In this case, the proof follows closely the one of Case 1 of Theorem 9 with the difference that we have that $\mathcal{H}_1(x_0, \mathsf{aux}) \equiv \mathcal{H}_2(x_0, \mathsf{aux}) \equiv \mathcal{H}_3(x_0, \mathsf{aux})$, for all $x_0 \in L_{\mathcal{R}_0}$.

Case 2. The first hybrid experiment $\mathcal{H}_1(x_0, \mathsf{aux})$ is the original experiment $\mathsf{ExpWI}^{p}_{\mathcal{A}}(x_0, \mathsf{aux})$ in which $b = 1$ (and thus $\mathcal{P}$ uses witness w_1). That is,

- Step 1 of $\mathsf{ExpWI}^{p}_{\mathcal{A}}(x_0, \mathsf{aux})$ stays the same.
  1. $a = \mathsf{P}^1_1(1^\lambda; R_1)$, for random coin tosses R_1;
  2. $(\mathtt{com}, \mathtt{dec}, \mathtt{rand}) \leftarrow \mathsf{TCom}_{\Pi_0}(x_0, a)$ and outputs $\mathtt{com}$.
- In Step 4 of $\mathsf{ExpWI}^{p}_{\mathcal{A}}(x_0, \mathsf{aux})$, the following steps are executed:
  1. $(a', z) = \mathsf{Sim}^1(x_1, c)$;
  2. $\mathtt{dec}' \leftarrow \mathsf{TFake}_{\Pi_0}(x_0, w_1, a, a', \mathtt{com}, \mathtt{rand})$;
  3. $(\mathtt{dec}', a', z)$.

The second hybrid experiment $\mathcal{H}_2(x_0, \mathsf{aux})$ differs from $\mathcal{H}_1(x_0, \mathsf{aux})$ in the way TFake is executed (namely, using as input w_0 instead of w_1). More specifically,

- Step 1 of $\mathsf{ExpWI}^{p}_{\mathcal{A}}(x_0, \mathsf{aux})$ stays the same.
  1. $a = \mathsf{P}^1_1(1^\lambda; R_1)$, for random coin tosses R_1;
  2. $(\mathtt{com}, \mathtt{dec}, \mathtt{rand}) \leftarrow \mathsf{TCom}_{\Pi_0}(x_0, a)$ and outputs $\mathtt{com}$.
- In Step 4 of $\mathsf{ExpWI}^{p}_{\mathcal{A}}(x_0, \mathsf{aux})$, the following steps are executed:
  1. $(a', z) = \mathsf{Sim}^1(x_1, c)$;
  2. $\mathtt{dec}' \leftarrow \mathsf{TFake}_{\Pi_0}(x_0, w_0, a, a', \mathtt{com}, \mathtt{rand})$;
  3. $(\mathtt{dec}', a', z)$.

We observe that the WI property of Π_0, the protocol used to construct the IDTC, implies that $\mathcal{H}_1(x_0, \mathsf{aux}) \equiv \mathcal{H}_2(x_0, \mathsf{aux})$. The proof ends with the observation that $\mathcal{H}_2(x_0, \mathsf{aux})$ is exactly the experiment $\mathsf{ExpWI}^{p}_{\mathcal{A}}(x_0, \mathsf{aux})$ when $b = 0$.

Case 3. The first hybrid experiment $\mathcal{H}_1(x_0, \mathsf{aux})$ is again the original experiment $\mathsf{ExpWI}^p_{\mathcal{A}}(x_0, \mathsf{aux})$ in which $b = 1$ (and thus $\mathcal{P}$ uses witness w_1). The second hybrid experiment $\mathcal{H}_2(x_0, \mathsf{aux})$ differs from $\mathcal{H}_1(x_0, \mathsf{aux})$ in the way that z is computed (using as input w_1 instead of w_0 when P_2 is executed). More specifically,

- In Step 1 of $\mathsf{ExpWI}^p_{\mathcal{A}}(x_0, \mathsf{aux})$, the following steps are executed:
  1. $a = \mathsf{P}^1_1(1^\lambda; R_1)$, for random coin tosses R_1;
  2. $(\mathtt{com}, \mathtt{dec}, \mathtt{rand}) \leftarrow \mathsf{TCom}_{\Pi_0}(x_0, a)$ and outputs $\mathtt{com}$.
- In Step 4 of $\mathsf{ExpWI}^p_{\mathcal{A}}(x_0, \mathsf{aux})$, the following steps are executed:
  1. $z \leftarrow \mathsf{P}^1_2(x_1, w_0, c, R_1)$;
  2. output $(\mathtt{dec}, a, z)$

From the $\mathcal{WI}$ property of Π_1 we can easily claim that $\mathcal{H}_1(x_0, \mathsf{aux}) \equiv \mathcal{H}_2(x_0, \mathsf{aux})$. The proof ends with the observation that $\mathcal{H}_2(x_0, \mathsf{aux})$ is exactly the experiment $\mathsf{ExpWI}^p_{\mathcal{A}}(x_0, \mathsf{aux})$ when $b = 0$. □

6 Applications

In this section, we describe the application of our new OR composition technique for constructing a 3-round straight-line perfect ZK argument with quasi-polynomial time simulation and a 4-round resettable WI argument. The latter is based on the concept of Simulatable Commitment++ (see Appendix A), a stronger version of the Simulatable Commitment of [MP03].

6.1 A 3-Round Efficient Perfectly Simulatable Argument System

In [Pas03], Pass introduced relaxed notions of zero knowledge and knowledge extraction in which the simulator and the extractor are allowed to run in quasi-polynomial time. Allowing the simulator to run in quasi-polynomial time typically dispenses with the need of rewinding the verifier; that is, the simulator is *straight-line*. In [Pas03], Pass first describes the following 2-round perfect ZK argument for any language L: the verifier $\mathcal{V}$ sends a value $Y = f(y)$ for a randomly chosen y where f is a sub-exponentially hard OWF and the first round of a ZAP protocol. The prover $\mathcal{P}$ then sends a commitment to $(y'|w')$ and uses the second round of the ZAP to prove that either $y' = f^{-1}(y)$ or w' is a witness for $x \in L$. If language L admits a Σ-protocol Π_L then the above construction can be implemented as an efficient 4-round argument with quasi-polynomial time simulation: the function f is concretely instantiated to be an exponentiation in a group in which the Discrete Log problem is hard and the ZAP is replaced with the CDS-OR composition of Π_L and Schnorr's Σ-protocol for the Discrete Log.

Note that Schnorr's Σ-protocol is Input-delayed and thus we can use it as Σ-protocol Π_1 in our OR transform in conjunction with any Σ-protocol Π_0 that belongs to one of the first 3 Classes described in Section 1.1.

6.1.1 Preliminary Definitions

We start by providing some useful definitions.

Simulation in quasi-polynomial time. Since the verifier in an interactive argument is often modeled as a PPT machine, the classical zero-knowledge definition requires that the simulator runs also in (expected) polynomial time. In [Pas03], the simulator is allowed to run in time $\lambda^{\mathtt{poly}(\mathtt{log}(\lambda))}$. Loosely speaking, we say that an interactive argument is $\lambda^{\mathtt{poly}(\mathtt{log}(\lambda))}$-perfectly simulatable if for any adversarial verifier there exists a simulator running in time $\lambda^{\mathtt{poly}(\mathtt{log}(\lambda))}$, where λ is the size of the statement being proved, whose output is identically distributed to the output of the adversarial verifier.

Definition 7 (One-way functions for sub-exponential circuits. [Pas03])**.** *A function $f : \{0,1\}^* \rightarrow \{0,1\}^*$ is called one-way for sub-exponential circuits if there exists a constant α such that the following two condition holds:*

- *there exist a deterministic polynomial-time algorithm that on input y outputs $f(y)$;*
- *for every probabilistic algorithm $\mathcal{A}$ with running time bounded by 2^{λ^α}, all sufficiently large λ's, and every auxiliary input $z \in \{0,1\}^{\mathtt{poly}(\lambda)}$*

$$\text{Prob}\left[\, x \stackrel{R}{\leftarrow} \{0,1\}^* : \mathcal{A}(f(x), z) \in f^{-1}(f(x)) \,\right] < \frac{1}{\mathtt{poly}(2^{\lambda^\alpha})}.$$

Now we define straight-line $T(\lambda)$-perfectly simulatable interactive arguments.

Definition 8 (straight-line $T(\lambda)$ simulatability, Def. 31 of [Pas04])**.** *Let $T(\lambda)$ be a class of functions that is closed under composition with any polynomial. We say that an interactive argument (proof) $(\mathcal{P}, \mathcal{V})$ for the language $L \in NP$, with the witness relation $\mathcal{R}_L$, is* straight-line $T(\lambda)$-simulatable *if for every PPT machine $\mathcal{V}^\star$ there exists a probabilistic simulator S with running time bounded by $T(\lambda)$ such that the following two ensembles are computationally indistinguishable (when the distinguish gap is a function in $\lambda = |x|$)*

- $\{(\langle \mathcal{P}(y), \mathcal{V}^\star(z)\rangle(x))\}_{z\in\{0,1\}^*, x\in L}$ *for arbitrary $y \in \mathcal{R}_L(x)$*
- $\{(\langle S, \mathcal{V}^\star(z)\rangle(x))\}_{z\in\{0,1\}^*, x\in L}$

We note that the above definition is very restrictive. In fact, the simulator is supposed to act as a cheating prover, with its only advantage being the possibility of running in time $T(\lambda)$, instead of in polynomial time. Trivially, it do not exist a straight-line $T(\lambda)$-simulatable proof for non-trivial languages (this should be contrasted with straight-line simulatable interactive arguments, which instead do exist).

The following theorem shows the importance of straight-line $\lambda^{\mathtt{poly}(\mathtt{log}(\lambda))}$-perfect simulatability by connecting it to concurrent composition of arguments.

Theorem 11. *If the interactive argument $\Pi = (\mathcal{P}, \mathcal{V})$ is straight-line $\lambda^{\mathtt{poly}(\mathtt{log}(\lambda))}$-simulatable then it is also straight-line concurrent $\lambda^{\mathtt{poly}(\mathtt{log}(\lambda))}$-simulatable.*

6.1.2 The Protocol

For any NP-language L we consider the Σ-protocol Π for the relation R_L and use our OR-composition technique to compose Π_{Dlog} and Π_L to obtain a new Σ-protocol $\Pi^{\mathsf{OR}} = (\mathcal{P}^{\mathsf{OR}}, \mathcal{V}^{\mathsf{OR}})$ for the relation

$$\mathcal{R}_{\mathsf{OR}} = \Big\{((x_L, x_{Dlog}), w) : \Big((x_L, w) \in \mathcal{R}_L \wedge x_{Dlog} \in \hat{L}_{\mathcal{R}_{Dlog}}\Big) \,\mathsf{OR}\, \Big((x_{Dlog}, w) \in \mathcal{R}_{Dlog} \wedge x_L \in \hat{L}_{\mathcal{R}_L}\Big)\Big\}$$

with challenge length $l = \lambda$ and associated algorithm $\mathsf{P}_1^{\mathsf{OR}}$, $\mathsf{P}_2^{\mathsf{OR}}$ and V^{OR}.

Let f be a sub-exponentially hard one-way function implemented using DLog as described before, with the only change that for some constant α, f is one-way w.r.t circuits of size 2^{λ^α}. Let $L \in NP$ and $k = \frac{1}{\alpha} + 1$. Our 3-round straight-line ZK argument for $x \in L$ is the following.

Protocol 5. A 3-round straight-line ZK argument.
Common input: *An instance x of a language $L \in NP$ with witness relation R_L, and 1^λ as security parameter.*
Private input: *Prover has w as a private input, s.t. $(x, w) \in \mathcal{R}_L$.*
Round 1. $\mathcal{P} \to \mathcal{V}$:

1. *On input a randomness R_1, $\mathcal{P}$ uniformly chooses (p, q, g) where $p = 2q + 1$ is a safe prime and g is a generator of a group $\mathcal{G}_q$ of size q. We remark that (p, q, g) are parameters selected so that the function $f(y) = g^y$ is one-way function for some constant α w.r.t circuits of size 2^{λ^α}.*
2. *$\mathcal{P}$ computes $a \leftarrow \mathsf{P}_1^{\mathsf{OR}}((x, (p, q, g)); R_1)$.*
3. *$\mathcal{P}$ sends (p, q, g) and a to $\mathcal{V}$.*

Round 2. $\mathcal{V} \to \mathcal{P}$:

1. *$\mathcal{V}$ chooses $y \leftarrow \mathbb{Z}_q$ and computes $Y = g^y$.*
2. *$\mathcal{V}$ chooses $c \leftarrow \{0, 1\}^l$.*
3. *$\mathcal{V}$ sends c and Y to $\mathcal{P}$.*

Round 3. $\mathcal{P} \to \mathcal{V}$:

1. *$\mathcal{P}$ computes $z \leftarrow \mathsf{P}_2((x, ((p, q, g), Y)), w, c, R_1)$.*
2. *$\mathcal{P}$ sends z to $\mathcal{V}$.*
3. *$\mathcal{V}$ accepts if and only if $\mathsf{V}^{\mathsf{OR}}((x, ((p, q, g), Y)), a, c, z) = 1$.*

We remark that we are using the same assumption of [CD08] that allows the adversary of DLog to generate the DLog parameters while the challenger selects the random element of the group.

Theorem 12. *If Π^{OR} is a Σ-protocol for OR composition of L and L_{Dlog}, then Protocol 5 is a 3-round straight-line $\lambda^{O(\log^k \lambda)}$-perfectly simulatable argument of knowledge.*

Proof. Completeness follows directly from the completeness of Π^{OR}.

Soundness/knowledge extraction. We show that Π is an argument of knowledge; this directly implies soundness. The claim follows from the fact that the argument system Π^{OR} used is a proof of knowledge, for a large enough challenge l, with knowledge error 2^{-l}, and from the fact that a PPT adversary only finds a pre-image to Y (for f) with negligible probability. More formally, we construct a polynomial-time extractor E for every polynomial-time $\mathcal{P}^\star$ for protocol Π. E internally incorporates $\mathcal{P}^\star$ and each time Π^{OR} proves a new theorem it proceeds as follows. E invokes the extractor E^{OR} for Π^{OR}. E outputs whatever E^{OR} outputs. By the proof knowledge property of Π^{OR}, the output of E will either be a witness w for the statement proved, or the pre-image of Y. If E outputs w, we are done. Otherwise, if it outputs y with non-negligible probability, then we can construct a reduction that breaks the DLog assumption (still in the form proposed by [CD08]).

Zero knowledge. Consider a straight-line simulator Sim that computes the first round as the honest prover. This is possible because Π^{OR} does not need any witness to computes the first round. After the simulator receives Y it checks if that indeed has a pre-image. Sim thereafter performs an exhaustive search to find a pre-image y of a value Y for the function f. To perform this task Sim tries all possible values $y' \in \{0,1\}^{\log^k \lambda}$ and checks if $f(y') = Y$. This thus takes time $\mathtt{poly}(2^{\log^k \lambda})$, since the time it takes to evaluate the function f is a polynomial in λ. After having found a value y such that $f(y) = Y$, Sim uses y as witness to complete the execution of Π^{OR} (instead of using a real witness for x, as the honest prover would do). Clearly the running time of Sim is bounded by $\lambda^{O(\log^k \lambda)}$. We proceed to show that the output of the simulator is identically distributed to the output of any adversarial verifier in a real execution with an onest prover. Note that the only difference between a real execution and a simulated execution is in the choice of the witness used in the last stage of the protocol. Therefore, from the adaptive WI property of Π^{OR} we have that the output of the simulated execution is identically distributed to the output of the real execution. $\square$

6.2 Efficient Resettable WI Argument System

In this section, we show how to transform any *any* Σ-protocol into a resettable WI ($\mathsf{r}\mathcal{WI}$) argument system, adding only one extra round.

Resettable witness indistinguishability was introduced in [CGGM00]. Very roughly, a resetting verifier is a PPT adversary that is able to interact with the prover polynomially many times with possibly distinct inputs forcing the prover to execute the protocol using the same randomness several times. Namely, we can think of a prover under a reset as being equipped with a vector of inputs $\bar{x}$ and a vector of random tapes $\vec{\omega}$. The malicious verifier can adaptively start a new interaction with the prover by specifying the input x_i and the randomness ω_j to be used in the interaction. Moreover, the malicious verifier has complete control over the schedule of the messages. A concurrent verifier is a restricted form of resetting verifier that cannot start two interactions with the same randomness. A formal definition is found in [CGGM00].

An informal description. In our construction, we use Simulatable Commitment++ [MP03]. Very informally, Simulatable Commitment++ is a regular commitment scheme that is equipped with a Σ[10]-protocol, where the committer proves the knowledge of the message committed, without revealing the opening information (we will describe this primitive in details in Appendix A).

Our proposed $\mathsf{r}\mathcal{WI}$ argument system $\Pi^{\mathrm{WI}} = (\mathcal{P}, \mathcal{V})$ for $\mathcal{R}$ with associated NP-language L consists of the interleaved execution of two Σ-protocols: Π_L, in which $\mathcal{P}$ acts a prover, and Π^{OR}, in which $\mathcal{V}$ acts as a prover. Π_L is the Σ-protocol obtained for relation $\mathcal{R}$ and we denote by $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$ its associated algorithms. Π^{OR} is the Σ-protocol obtained from the OR composition of two Σ-protocols: Π^{simcom}, which is the Σ-protocol associated to Simulatable Commitment++ which is a Chameleon Σ-protocol (as we show in Appendix A) and Π^{dlog}, which is Schnorr's Σ-protocol for proving knowledge of a DLog, and is input-delayed. We denote by $(\mathsf{P}_1^{\mathsf{OR}}, \mathsf{P}_2^{\mathsf{OR}}, \mathsf{V}^{\mathsf{OR}})$ the algorithms associated with Π^{OR}.

In the first round, $\mathcal{V}$ randomly selects the challenge c for Π_L and commits to it by using the Simulatable Commitment++, thus obtaining $\mathtt{com}$. The pair $(\mathtt{com}, c)$ defines the first theorem for Π^{OR}, although $\mathcal{V}$ only sends $\mathtt{com}$ to $\mathcal{P}$. In the same round, $\mathcal{V}$ starts Σ-protocol Π^{OR} in which she

[10] More precisely, the Σ-protocol associated to Simulatable Commitment++ does not enjoy special soundness but only optimal soundness.

proves that she knows the message committed to by $\mathtt{com}$ or that she knows the discrete log of a random element Y selected by $\mathcal{P}$ and sent to $\mathcal{V}$ as part of the second round of Π^{WI}. We make two remarks here. First of all, c is not revealed to $\mathcal{P}$ (that acts as a verifier) when Π^{OR} starts but $\mathcal{P}$ is only given commitment $\mathtt{com}$. This is so because c cannot be revealed before $\mathcal{P}$ has sent the first message a of Π and this happens only at the second round of Π^{WI}. This does not constitute a problem since in the Σ-protocol associated with a Simulatable Commitment++ the verifier ($\mathcal{P}$ in this case) learns the message committed to by $\mathtt{com}$ only at the last round. We also stress that Y is not known to $\mathcal{V}$ (that acts as a prover) when Π^{OR} starts. For this reason, we need to use our OR-composition technique that we have presented in Section 5 since it allows the Σ-protocol to be started even if the prover (in this case $\mathcal{V}$) does not have the entire input. Execution of Π starts at the second round of Π^{WI} and $\mathcal{P}$ also sends at the same round a random element Y in a prime order group in which discrete log is considered difficult whose description is sent by $\mathcal{V}$ as part of the first round.

Let us proceed more formally and denote by l^{OR} the challenge length of Π^{OR} challenge length of Π. Furthermore, we denote by x the common input and by w the witness received by $\mathcal{P}$. Security parameters λ and λ' will be determined as part of the proof.

We denote by $(\mathsf{Com}, \mathsf{Dec})$ the commitment and decommitment procedures of Simulatable Commitment++ (the procedures are formally described in Appendix A.3). We denote by Π^{simcom} the Σ-protocol associated to Simulatable Commitment++.

Π^{simcom} is a Chameleon Σ-protocol (as we prove in Appendix A) and admits a 2-IDTC scheme.

Finally, let Π^{dlog} be a Input-delayed Σ-protocol for the discrete log polynomial-time relation

$$\mathsf{DLOG} = \{((\mathcal{G}, q, g, Y), y) : g^y = Y\}$$

over groups $\mathcal{G}$ of prime-order q. Schnorr's Σ-protocol [Sch89].

The security of our proposed $\mathsf{r}\mathcal{WI}$ $\Pi^{\mathrm{WI}} = (\mathcal{P}, \mathcal{V})$ for $\mathcal{R}$ is based on the well-known DDH assumption and on a strengthening of its (see Assumption 4).

Protocol 6. *(4-round Resettable WI.)*
Public input: $x \in L$.
Private input to $\mathcal{P}$: *a witness* w *for* $\mathcal{R}(x, w) = 1$.
Round 1*: from* $\mathcal{V}$ *to* $\mathcal{P}$*.*

1.1 Picking the parameters for the second input to Π^{OR}*.*

1.1.1 $\mathcal{V}$ *randomly selects* $(\mathcal{G}, q, g) \leftarrow \mathsf{IG}(1^\lambda)$ *and sends it to* $\mathcal{P}$*;*

1.2 First round of Π^{OR} *with* $\mathcal{V}$ *acting as a prover.*

1.2.1 $\mathcal{V}$ *randomly selects* $c \leftarrow \{0,1\}^l$ *and computes* $(\mathtt{com}, \mathtt{dec}) = \mathsf{Com}(1^{\lambda'}, c)$*;*

1.2.2 $\mathcal{V}$ *sets* $x_0^{\mathsf{OR}} = (\mathtt{com}, c)$*,* $w_0^{\mathsf{OR}} = \mathtt{dec}$ *and sends* $\mathtt{com}$ *to* $\mathcal{P}$*;*

1.2.3 $\mathcal{V}$ *randomly selects coin tosses* R_1*, computes* $a^{\mathsf{OR}} = \mathsf{P}_1^{\mathsf{OR}}(x_0^{\mathsf{OR}}; R_1)$ *and sends it to* $\mathcal{P}$*;*

Round 2*: from* $\mathcal{P}$ *to* $\mathcal{V}$*.*

2.1 First round of Π*.*

2.1.1 $\mathcal{P}$ *reads* ω *from its random tape, computes* $\bar{R} = F_\omega(\mathtt{com})$ *and parses it as* $\bar{R} = (R_2, R_3)$*;*

2.1.2 $\mathcal{P}$ sets $a = \mathsf{P}_1(x, w; R_2)$ and sends it to $\mathcal{V}$;

2.2 Completing the input for Π^{OR} and sending the challenge for Π^{OR}.

2.2.1 $\mathcal{P}$ selects $y \leftarrow \mathbb{Z}_q$, compute $Y = g^y$ and send it to $\mathcal{V}$;

2.2.3 $\mathcal{P}$ picks $c^{\mathsf{OR}} \leftarrow \{0,1\}^{l^{\mathsf{OR}}}$ and sends it to $\mathcal{V}$;

Round 3*: from $\mathcal{V}$ to $\mathcal{P}$.*

3.1 Last round of Π^{OR} and sending the challenge for Π.

$\mathcal{V}$ sends c to $\mathcal{P}$;

$\mathcal{V}$ sets $x_1^{\mathsf{OR}} = (q, g, Y)$ and computes and sends $z^{\mathsf{OR}} = \mathsf{P}_2^{\mathsf{OR}}((x_0^{\mathsf{OR}}, x_1^{\mathsf{OR}}), w_0^{\mathsf{OR}}, R_1)$;

Round 4*: from $\mathcal{P}$ to $\mathcal{V}$.*

4.1 Last round of Π.

4.1.1 If $\mathsf{V}^{\mathsf{OR}}((x_0^{\mathsf{OR}}, x_1^{\mathsf{OR}}), a^{\mathsf{OR}}, c^{\mathsf{OR}}, z^{\mathsf{OR}}) = 0$, $\mathcal{P}$ aborts.

4.1.2 $\mathcal{P}$ computes $z \leftarrow \mathsf{P}_2(x, w, c, R_2; R_3)$ and sends it to $\mathcal{V}$;

$\mathcal{V}$'s decision*: $\mathcal{V}$ accepts if and only if $\mathsf{V}(x, a, c, z) = 1$.*

Completeness is straightforward.

Proof of Soundness. Suppose that there exists a prover $\mathcal{P}^\star$ that violates the soundness of Π^{WI}. We consider the following three indistinguishable hybrid experiments $\mathcal{H}_0, \mathcal{H}_1$, and $\mathcal{H}_2$.

The first hybrid experiment $\mathcal{H}_0(x)$ takes as input $x \notin L$ and consists simply of $\mathcal{P}^\star$ interacting with the honest verifier $\mathcal{V}$. In the second experiment, $\mathcal{H}_1(x)$, $\mathcal{P}^\star$ interacts with a verifier $\mathcal{V}_1$ that, upon receiving, $(\mathcal{G}_2, q_2, g_2, Y_2)$ at Step 2.2.2 computes the discrete log of Y_2 in group $\mathcal{G}_2$ with respect to g_2 and uses it as a witness at Step 3.1. $\mathcal{H}_0(x)$ and $\mathcal{H}_1(x)$ are indistinguishable because of the perfect adaptive WI of Π^{OR}.

In hybrid $\mathcal{H}_2(x)$, $\mathcal{P}^\star$ interacts with $\mathcal{V}_2$ that differs from $\mathcal{V}_1$ because it selects c and c_1 at Step 1.2.1, commits to c_1 and sends c at Step 3.1.

We point out a subtlety here. Recall that in the instantiation of Π^{OR} that we are using the first message a^{OR} is a commitment of the first message a of Schnorr's Σ-protocol for Y. Such commitment is computed using the 2-IDTC scheme derived from Π^{simcom}. Now, recall that such commitment is computed on the instance $(\mathtt{com}, c)$, which is known to $\mathcal{V}$ and defined already the first round, even though the prover will see c only in the last round of Π^{OR}.

Now, in $\mathcal{H}_2(x)$ we have that $\mathtt{com}$ is a commitment of c_1 but $\mathcal{V}$ sends $c \neq c_1$ in the last round. This seems to be problematic. Nevertheless, we stress that this is indeed no problem because the simulator of Simulatable Commitment++ (this simulator consists of combining the simulator for simulatable commitments of [MP03] and the one for equality of two DLogs, see App. A for details) works also for bad commitments and thus $\mathcal{V}_2$ can complete $\mathcal{H}_2(x)$ even without knowing the message committed to by $\mathtt{com}$. Therefore, we can conclude that the hybrids $\mathcal{H}_1(x)$ and $\mathcal{H}_2(x)$ are indistinguishable by the hiding of the Simulatable Commitment++ (which holds under the DDH assumption) provided that λ' and λ are chosen in such a way that the commitment scheme (whose security depends on λ') is still secure against adversary with the running time of $\mathcal{V}_2$ (that depends on λ).

Now we observe that if $\mathcal{P}^\star$ breaks the soundness of Π^{WI} then it succeeds in successfully completing an interaction with $\mathcal{V}_2$ with the same (up to negligible) probability. In this interaction however the first message $\mathcal{V}_2$ is independent from the actual challenge that will be received and thus, by using an argument similar to [GK96] (also used in [CGGM00] for the 5-round rWI for NP), we conclude that $\mathcal{P}^\star$ breaks the soundness of Π. Contradiction. Notice that here we are using the existence of algorithm Cont that makes the first message of the verifier independent from the actual challenge for protocol Π that $\mathcal{P}^\star$ will receive at the third round.

Proof of rWI. The idea of the proof for rWI is very simple. We consider the prover $\mathcal{P}_1$ that, when instructed to use randomness with index j and input with index i and receives first message $\mathtt{msg}$, checks first if a tuple $(j, i, \mathtt{msg}, R)$ has been stored in a previous step. If such a tuple is found then R is used as source of randomness; otherwise, a fresh R is selected and tuple $(j, i, \mathtt{msg}, R)$ is stored. Clearly, by pseudo-randomness, $\mathcal{P}_1$ is indistinguishable from the honest prover $\mathcal{P}$.

Now we observe that the resetting verifier $\mathcal{V}^\star$ is performing an attack on $\mathcal{P}_1$ in which two distinct interactions use the same randomness iff they share j, the input and the first message. More precisely, we say that interactions t_1 and t_2 between $\mathcal{P}_1$ and $\mathcal{V}^\star$ are a *collision* if they share the input, the randomness used by $\mathcal{P}_1$ and the first message and $\mathcal{V}^\star$ opens the commitment in the first messages in two different ways. If $\mathcal{V}^\star$ has a non-negligible probability of producing a conversation then we can break the Strong DLog Assumption (see Definition 2). Consider algorithm DL that receives as input $(\bar{x}, \bar{w}^0, \bar{w}^1)$ for which $\mathcal{V}^\star$ distinguishes $\mathcal{H}_2(\bar{x}, \bar{w}^0)$ from $\mathcal{H}_2(\bar{x}, \bar{w}^1)$. DL interacts with $\mathcal{V}^\star$ and at the start of the interaction guesses two interactions $t_1 < t_2$ (in the hope they constitute a collision). In all interactions other than t_1 and t_2, DL runs just like $\mathcal{P}_1$. When DL receives the discrete log parameters from $\mathcal{V}^\star$ as part of the first message of interaction t_1, it forwards them to the challenger of the discrete log and receives Y (and the task is two compute the discrete log of Y) and uses it as part of the second message of interaction t_1. When interaction t_2 is activated DL checks if the first message is the same as the first message of interaction t_1. If they are (and this happens with non-negligible) probability DL continues and sends the same second message, including Y. Notice that $\mathcal{V}^\star$ expects to receive the same message since it thinks it is interacting with $\mathcal{P}_1$ that is using the same randomness. Otherwise, DL aborts. Then DL rewinds $\mathcal{V}^\star$ and sends a different challenge in each of the two interactions. In at least one of them $\mathcal{V}^\star$ has opened the commitment to a different message than the one committed to by the commitment in the first message. This means that $\mathcal{V}^\star$ has used knowledge of the discrete log of Y to complete the interaction and thus DL can extract it.

Finally, let us consider the case in which $\mathcal{V}^\star$ produces a collision in his resetting attack only with negligible probability. This means that $\mathcal{V}^\star$ is conducting a successful concurrent WI attack on the argument system. Standard arguments show that this contradicts the WI of Π[11].

7 Acknowledgments

We thank Berry Schoenmakers for various useful discussions on Σ-protocols.

[11] We are implicitly using the fact that Σ-protocol Π is WI. This is certainly true if Π is perfect [CDS94]. If Π is only computational then we consider self OR composition of Π which by [GMY06] is WI.

References

[AOS13] Masayuki Abe, Tatsuaki Okamoto, and Koutarou Suzuki. Message recovery signature schemes from sigma-protocols. *IEICE Transactions*, 96-A(1):92–100, 2013.

[BFGM01] Mihir Bellare, Marc Fischlin, Shafi Goldwasser, and Silvio Micali. Identification protocols secure against reset attacks. In *EUROCRYPT*, volume 2045 of *Lecture Notes in Computer Science*, pages 495–511. Springer, 2001.

[Blu86] Manuel Blum. How to prove a theorem so no one else can claim it. In *International Congress of Mathematicians*, page 1444, 1986.

[BP15] Nir Bitansky and Omer Paneth. Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, pages 401–427, 2015.

[BPSV08] Carlo Blundo, Giuseppe Persiano, Ahmad-Reza Sadeghi, and Ivan Visconti. Improved security notions and protocols for non-transferable identification. In *Computer Security - ESORICS 2008, 13th European Symposium on Research in Computer Security, Málaga, Spain, October 6-8, 2008. Proceedings*, volume 5283 of *Lecture Notes in Computer Science*, pages 364–378, 2008.

[CD08] Ran Canetti and Ronny Ramzi Dakdouk. Extractable perfectly one-way functions. In *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part II - Track B: Logic, Semantics, and Theory of Programming & Track C: Security and Cryptography Foundations*, pages 449–460, 2008.

[CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In YvoG. Desmedt, editor, *Advances in Cryptology — CRYPTO '94*, volume 839 of *Lecture Notes in Computer Science*, pages 174–187. Springer Berlin Heidelberg, 1994.

[CDV06] Dario Catalano, Yevgeniy Dodis, and Ivan Visconti. Mercurial commitments: Minimal assumptions and efficient constructions. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 120–144. Springer, 2006.

[CG15] Pyrros Chaidos and Jens Groth. Making sigma-protocols non-interactive without random oracles. In *Public-Key Cryptography - PKC 2015 - 18th IACR International Conference on Practice and Theory in Public-Key Cryptography, Gaithersburg, MD, USA, March 30 - April 1, 2015, Proceedings*, pages 650–670, 2015.

[CGGM00] Ran Canetti, Oded Goldreich, Shafi Goldwasser, and Silvio Micali. Resettable zero-knowledge (extended abstract). In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 235–244, 2000.

[CPSV15] Michele Ciampi, Giuseppe Persiano, Luisa Siniscalchi, and Ivan Visconti. A transform for nizk almost as efficient and general as the fiat-shamir transform without programmable random oracles. *IACR Cryptology ePrint Archive*, 770, 2015.

[CV05a] Dario Catalano and Ivan Visconti. Hybrid trapdoor commitments and their applications. In *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings*, pages 298–310, 2005.

[CV05b] Giovanni Di Crescenzo and Ivan Visconti. Concurrent zero knowledge in the public-key model. In *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings*, volume 3580 of *Lecture Notes in Computer Science*, pages 816–827. Springer, 2005.

[CV07] Dario Catalano and Ivan Visconti. Hybrid commitments and their applications to zero-knowledge proof systems. *Theor. Comput. Sci.*, 374(1-3):229–260, 2007.

[Dam10] Ivan Damgård. On Σ-protocol. http://www.cs.au.dk/~ivan/Sigma.pdf, 2010.

[DG03] Ivan Damgård and Jens Groth. Non-interactive and reusable non-malleable commitment schemes. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, June 9-11, 2003, San Diego, CA, USA*, pages 426–437, 2003.

[DN00] Cynthia Dwork and Moni Naor. Zaps and their applications. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 283–293, 2000.

[DPV04] Giovanni Di Crescenzo, Giuseppe Persiano, and Ivan Visconti. Constant-round resettable zero knowledge with concurrent soundness in the bare public-key model. In Matthew K. Franklin, editor, *Advances in Cryptology - CRYPTO 2004, 24th Annual International CryptologyConference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, volume 3152 of *Lecture Notes in Computer Science*, pages 237–253. Springer, 2004.

[FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume I*, pages 308–317. IEEE Computer Society, 1990.

[GK96] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *J. Cryptology*, 9(3):167–190, 1996.

[GK15] Jens Groth and Markulf Kohlweiss. One-out-of-many proofs: Or how to leak a secret and spend a coin. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, pages 253–280, 2015.

[GMY06] Juan A. Garay, Philip MacKenzie, and Ke Yang. Strengthening zero-knowledge protocols using signatures. *Journal of Cryptology*, 19(2):169–209, 2006.

[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for NP. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, pages 339–358, 2006.

[GRRV14] Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 41–50. IEEE Computer Society, 2014.

[KO04] Jonathan Katz and Rafail Ostrovsky. Round-optimal secure two-party computation. In *Advances in Cryptology - CRYPTO 2004, 24th Annual International Cryptology-Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, pages 335–354, 2004.

[Lin15] Yehuda Lindell. An efficient transform from Sigma protocols to NIZK with a CRS and non-programmable random oracle. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part I*, pages 93–109, 2015.

[LP15] Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. *J. Cryptology*, 28(2):312–350, 2015.

[LS90] Dror Lapidot and Adi Shamir. Publicly verifiable non-interactive zero-knowledge proofs. In *Advances in Cryptology - CRYPTO*, 1990.

[Mau09] Ueli M. Maurer. Unifying zero-knowledge proofs of knowledge. In *Progress in Cryptology - AFRICACRYPT 2009, Second International Conference on Cryptology in Africa, Gammarth, Tunisia, June 21-25, 2009. Proceedings*, pages 272–286, 2009.

[Mau15] Ueli Maurer. Zero-knowledge proofs of knowledge for group homomorphisms. *Designs, Codes and Cryptography*, pages 1–14, 2015.

[MP03] Daniele Micciancio and Erez Petrank. Simulatable commitments and efficient concurrent zero-knowledge. In *Advances in Cryptology - EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4-8, 2003, Proceedings*, pages 140–159, 2003.

[OPV10] Rafail Ostrovsky, Omkant Pandey, and Ivan Visconti. Efficiency preserving transformations for concurrent non-malleable zero knowledge. In *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, pages 535–552, 2010.

[ORV14] Rafail Ostrovsky, Vanishree Rao, and Ivan Visconti. On selective-opening attacks against encryption schemes. In *Security and Cryptography for Networks - 9th International Conference, SCN 2014, Amalfi, Italy, September 3-5, 2014. Proceedings*, volume 8642 of *Lecture Notes in Computer Science*, pages 578–597. Springer, 2014.

[Pas03] Rafael Pass. Simulation in quasi-polynomial time, and its application to protocol composition. In Eli Biham, editor, *Advances in Cryptology - EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4-8, 2003, Proceedings*, volume 2656 of *Lecture Notes in Computer Science*, pages 160–176. Springer, 2003.

[Pas04] Rafael Pass. Alternative variants of zero-knowledge proofs. Master's thesis, Kungliga Tekniska Högskolan, 2004. Licentiate Thesis Stockholm, Sweden.

[Ped91] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Advances in Cryptology - CRYPTO '91, 11th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1991, Proceedings*, pages 129–140, 1991.

[PS96] David Pointcheval and Jacques Stern. Security proofs for signature schemes. In *Advances in Cryptology - EUROCRYPT '96, International Conference on the Theory and Application of Cryptographic Techniques, Saragossa, Spain, May 12-16, 1996, Proceeding*, pages 387–398, 1996.

[Sch89] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In Gilles Brassard, editor, *Advances in Cryptology — CRYPTO' 89 Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 239–252. Springer New York, 1989.

[SV12] Alessandra Scafuro and Ivan Visconti. On round-optimal zero knowledge in the bare public-key model. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 153–171, 2012.

[Vis06] Ivan Visconti. Efficient zero knowledge on the internet. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II*, pages 22–33, 2006.

[YZ07] Moti Yung and Yunlei Zhao. Generic and practical resettable zero-knowledge in the bare public-key model. In *Advances in Cryptology - EUROCRYPT 2007, 26th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Barcelona, Spain, May 20-24, 2007, Proceedings*, pages 129–147, 2007.

A Commitments and Simulatable Commitments

Here, we review the notion of a *Simulatable Commitment* (introduced in [MP03]). As we shall see, this notion does not suffice for our application to the construction of efficient 4-round resettable WI arguments and we will introduce the stronger notion of Simulatable Commitment++.

A.1 Commitment Schemes

A non-interactive commitment scheme, over a message space M, consists of two algorithms $(\mathsf{Com}, \mathsf{Dec})$ with the following syntax:

1. Com is a probabilistic algorithm that receives as input secure parameter 1^λ and message $m \in M$, and outputs a pair $(\mathtt{com}, \mathtt{dec})$ consisting of a *commitment* $\mathtt{com}$ of m and of an *opening* $\mathtt{dec}$ of $\mathtt{com}$ as m. With a slight abuse of notation, sometime we will simply say $\mathsf{Com}(m)$;

2. For a message $m \in M$, commitment $\mathtt{com}$ and opening $\mathtt{dec}$ $\mathsf{Dec}(\mathtt{com}, \mathtt{dec}, m)$ accepts or rejects $\mathtt{dec}$ as an opening of $\mathtt{com}$ as m.

We require a commitment scheme to enjoy the following properties.

1. Correctness: for all $m \in M$:

$$\text{Prob}\Big[\ (\mathtt{com}, \mathtt{dec}) \leftarrow \mathsf{Com}(1^\lambda, m) : \mathsf{Dec}(\mathtt{com}, \mathtt{dec}, m) = 1\ \Big] = 1.$$

2. Binding: for every PPT adversary $\mathcal{A}$ there exist a negligible function ν such that, for all $m_0, m_1 \in M$ s.t. $m_0 \neq m_1$ it holds that:

$$\text{Prob}\Big[\ (\mathtt{com}, \mathtt{dec}, \mathtt{dec}') \leftarrow \mathcal{A}(1^\lambda) : \mathsf{Dec}(\mathtt{com}, \mathtt{dec}, m_0) = 1 \wedge \mathsf{Dec}(\mathtt{com}, \mathtt{dec}', m_1) = 1\ \Big] \leq \nu(\lambda).$$

   The *Perfect* Binding property is obtained by requiring that the above holds for any adversary (and not just for PPT adversaries) with $\nu(\lambda) = 0$.

3. Hiding: for every PPT adversary $\mathcal{A}$ there exist a negligible function ν such that, for all $m_0, m_1 \in M$, it holds that:

$$\text{Prob}\Big[\ b \leftarrow \{0,1\}; (\mathtt{com}, \mathtt{dec}) \leftarrow \mathsf{Com}(1^\lambda, m_b) : b = \mathcal{A}(1^\lambda, \mathtt{com})\ \Big] = \frac{1}{2} + \nu(\lambda)$$

   The Perfect Hiding property is obtained by requiring that the above holds for any adversary (and not just for PPT adversaries) with $\nu(\lambda) = 0$.

In this paper we will use perfectly binding commitment schemes that are hiding with respect to sub-exponential adversaries. The formal definition is obtained by quantifying over all adversaries $\mathcal{A}$ running in time $2^{c\lambda}$, for some constant c.

A.2 Simulatable Commitment Schemes

A Simulatable Commitment scheme is a regular (non-interactive) commitment scheme $(\mathsf{Com}, \mathsf{Dec})$ with the usual Binding and Hiding properties which comes equipped with a special 3-move protocol Π^{simcom} with *Perfect Completeness*, *Optimal Soundness*, *Honest-Verifier Zero Knowledge* by which the committer can prove that a string $\mathtt{com}$ is the commitment of a message m. That is,

1. *Perfect Completeness:* for all pairs $(\mathtt{com}, m)$ such that $\mathtt{com}$ is a commitment of m, all 3-move transcripts (a, c, z) of Π^{simcom} with common input $(\mathtt{com}, m)$ are accepting.

2. *Optimal Soundness:* if $\mathtt{com}$ is not a commitment of m, then for every a there exists at most one c for which there exists z such (a, c, z) is accepting for $(\mathtt{com}, m)$.

3. *Honest-Verifier Zero Knowledge:* there exists a simulator Sim such that, for all pairs $(\mathtt{com}, m)$ such that $\mathtt{com}$ is a commitment of m outputs a 3-move transcript that is indistinguishable from a 3-move transcript of an interaction between prover (that has as private input the randomness used to produce $\mathtt{com}$) and verifier.

It is not difficult to observe that the hiding property of the commitment implies that the output of the simulator Sim is indistinguishable from real conversations also for *bad* commitments. More precisely, the output of the following two experiments are indistinguishable for every pair m and m' of messages:

1. compute a random commitment of $\mathtt{com}$ of m, run Sim on input $(\mathtt{com}, m)$ to obtain (a, c, z) and output $(\mathtt{com}, m, a, c, z)$.

2. compute a random commitment of $\mathtt{com}$ of m, run Sim on input $(\mathtt{com}, m')$ to obtain (a, c, z) and output $(\mathtt{com}, m', a, c, z)$.

We extend the notion of a Simulatable Commitment in the following way.

Definition 9. *A* Simulatable Commitment++ *is a Simulatable Commitment* $(\mathsf{Com}, \mathsf{Dec}, \Pi^{\mathsf{simcom}})$ *with the following extra properties*

1. Π^{simcom} *is a Σ-protocol for relation* $\{((\mathtt{com}, 1^\lambda, m), r) | \mathsf{Com}(1^\lambda, m; r) = \mathtt{com}\}$*;*

2. Delayed Input to Verifier*: at the first round the verifiers receives only commitment* $\mathtt{com}$ *and message m is disclosed only at the third round;*

A.3 Construction of a Simulatable Commitment++

We next describe our proposed Simulatable Commitment++. The commitment scheme is exactly the same as the one of [MP03] whereas the associated proof system is a bit more complex as it has to deal with the fact that the message m is revealed only at the last round. Potentially, this has impact on the soundness property since the prover might be able to reveal different messages depending on the challenge received.

Let IG be a group generator for which the DDH assumption holds (see Assumption 3). The commitment algorithm Com works as follows: on input security parameter 1^λ and message $m \in \mathbb{Z}_q$, the committer randomly selects $(\mathcal{G}, q, g) \leftarrow \mathsf{IG}(1^\lambda)$, picks random $h \in \mathcal{G}$ and $r \in \mathbb{Z}_q$ and outputs $\mathtt{com} = (\mathcal{G}, g, h, g^r, h^{r+m})$ and $\mathtt{dec} = r$. The Dec algorithm, on input $\mathtt{com} = (q, g, h, G, H), \mathtt{dec} = r$

and m simply checks that $G = g^r$ and that $H = h^{r+m}$. Clearly, $(\mathsf{Com}, \mathsf{Dec})$ is perfectly binding and computationally hiding, under the DDH assumption. Let us now describe the Σ-protocol Π^{simcom} with associated algorithms $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$ for the polynomial-time relation

$$\mathsf{ELG} = \{((\mathcal{G}, g, h, G, H, m), r) | G = g^r \text{ and } H = h^{r+m}\}.$$

The common input consists of $\mathtt{com} = (\mathcal{G}, g, h, g^r, h^{r+m})$ and m (to be disclosed only at the last round) and the prover has private input r such that $G = g^r$ and $H = h^{r+m}$.

1. Algorithm P_1, on input $\mathtt{com} = (\mathcal{G}, g, h, G, H)$, message m and randomness r, randomly selects $t, t' \in \mathbb{Z}_q$ and computes $A = g^t, B = h^t$ and $A' = g^{t'}, B' = h^{t'}$.

   P_1 sends (A, B, A', B') and commitment $\mathtt{com}$ to $\mathcal{V}$.

2. Algorithm P_2, on input $(\mathcal{G}, g, h, G, H)$, m and r and challenge c, sets $z = t + rc$ and $z' = t' + tc$. P_2 sends (m, z, z') to $\mathcal{V}$

3. Algorithm V, on input $\mathtt{com} = (\mathcal{G}, g, h, G, H)$ and m, c and z checks that

$$g^z = A \cdot G^c \text{ and } h^z = B \cdot \left(\frac{H}{h^m}\right)^c \text{ and } g^{z'} = A' \cdot A^c \text{ and } h^{z'} = B' \cdot B^c$$

   and accepts if all checks are successful.

Special Soundness. Let us now consider a collision for Π^{simcom}; that is, a pair of conversations

$$((A, B, A', B', \mathtt{com}), c_0, (z_0, z_0', m_0)) \text{ and } ((A, B, A', B', \mathtt{com}), c_1, (z_1, z_1', m_1))$$

with the same first message. We let a, b, a', b', r_1 and r_2 be such that $A = g^a$, $B = h^b$, $A' = g^{a'}$, $B' = h^{b'}$, $G = g^{r_1}$ and $H = h^{r_2}$. The accepting conditions for the two conversations can be written as follows:

$$\begin{array}{ll} z_0 = a + c_0 \cdot r_1 & z_1 = a + c_1 \cdot r_1 \\ z_0 = b + c_1 \cdot (r_2 - m_0) & z_1 = b + c_2 \cdot (r_2 - m_1) \\ z_0' = a' + c_0 \cdot a & z_1' = a' + c_1 \cdot a \\ z_0' = b' + c_0 \cdot b & z_1' = b' + c_1 \cdot b \end{array} \qquad (\star)$$

By subtracting the equations in pairs we obtain the following relations

$$\begin{array}{lcl lcl} a - b & = & c_0 \cdot (r_2 - r_1 - m_0) & a - b & = & c_1 \cdot (r_2 - r_1 - m_1) \\ c_0 \cdot (a - b) & = & b' - a' & c_1 \cdot (a - b) & = & b' - a' \end{array}$$

Since $c_0 \neq c_1$, the two equations in the second row imply $a = b$ and $a' = b'$. Therefore, the two equations in the first row give $m_0 = r_2 - r_1 = m_1$ and thus the two conservation must open the commitment to the same message. At the light of this, the first row of equations from $(\star)$ is a system of equations in two unknowns which can be used to compute a and r_1 which constitutes a witness for the the ELG relation.

Special Honest-Verifier Zero Knowledge. The simulator Sim of Π^{simcom}, on input $\mathtt{com} = (\mathcal{G}, g, h, G, H)$, message m and challenge c works as follow:

1. randomly select $z, z' \in \mathbb{Z}_q$;
2. set $A = \dfrac{g^z}{G^c}$ and $B = \dfrac{h^z}{(H/h^m)^c}$;
3. set $A' = \dfrac{g^{z'}}{A^c}$ and $B' = \dfrac{h^{z'}}{B^c}$;
4. output $((A, B, A', B'), (z, z'))$;

Notice that A, A', B, and B' are uniformly distributed in $\mathcal{G}$ and that the three-move transcript is accepting. This is exactly the distribution of a honest transcript with challenge c. Therefore we can claim that the output of Sim is perfectly indistinguishable from a real transcript.

Π^{simcom} is a Chameleon Σ-protocol. We describe PPT algorithm $\mathsf{P}_{\mathsf{sim}}$ that satisfies the Delayed Indistinguishability requirement. Let (A, B, A', B') be the first message and (z_0, z'_0) the last message computed by Sim on input $\mathtt{com} = (\mathcal{G}, g, h, G, H), m$ and challenge c_0. Upon receiving a new challenge c_1 and randomness r used to compute the commitment (that is, r is such that $G = g^r$ and $H = h^{r+m}$), $\mathsf{P}_{\mathsf{sim}}$ proceeds as follows. First of all, $\mathsf{P}_{\mathsf{sim}}$ sets $t = z_0 + r \cdot c_0$ and $t' = z'_0 + t \cdot c$. Observe that t and t' are such that $A = g^t$ and $B = h^t$ and $A' = g^{t'}$ and $B = h^{t'}$; in other words, $\mathsf{P}_{\mathsf{sim}}$ has recovered the random coin tosses of P_1 that would generate (A, B, A', B') as a first message. Therefore from now on, $\mathsf{P}_{\mathsf{sim}}$ can just execute the code of P_2 to complete the conversation for challenge c_1. Specifically, $\mathsf{P}_{\mathsf{sim}}$ sets $z_1 = t + rc_1$ and $z'_1 = t' + tc_1$.

An extra property. For some applications, we slightly modify the three-move protocol Π^{simcom} associated with the Simulatable Commitment++ described above by allowing the verifier to send a two-part challenge (c, c'). The verifier expects to receive z and z' each satisfying two of the acceptance criteria for c and c', respectively. It is not difficult to see that the modified protocol is not special sound but only sound (and this suffices for the applications). Moreover, the so modified protocol admits an efficient algorithm Cont that, on input $((\mathcal{G}, g, h, G, H, m), r) \in \mathsf{ELG}$, first round message (A, B, A', B'), and message m', outputs challenges (c, c') and third round message (z, z') that constitute an accepting transcript for m'. In other words, for any commitment for which a witness is given and for any first round message, Cont can produce, for any message m', an accepting transcript opening the commitment as m'.

B More About Σ-Protocols

Theorem 13. *For every relation $\mathcal{R}$ such that $L_{\mathcal{R}} \notin \mathtt{BPP}$ there exist Σ-protocols that are not WI.*

Proof. Let $\Pi' = (\mathcal{P}', \mathcal{V}')$ be a Σ-protocol for the relation $\mathcal{R}$ with challenge length l and let $(\mathsf{P}'_1, \mathsf{P}'_2, \mathsf{V}')$ be the triple of PPT algorithms associated to Π'. We use these algorithms to describe a Σ-protocol Π with associated algorithms $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$ that is not WI. Consider $(x, w) \in \mathcal{R}$.

1. P_1 on input (x, w) and randomness R_1 parses it as (r_1, c_p) where c_p is an $l(x)$-bit string, computes $a' \leftarrow \mathsf{P}'_1(x, w; r_1)$, and outputs $a = (a', c_p)$.

2. P_2, on input (x, w), R_1, a challenge c computes $z' \leftarrow \mathsf{P}'_2(x, w, r_1, c, r_1)$ and if $c = c_p$ then it also sets $z = w$ otherwise it sets $z = z'$; finally it outputs z.

3. V, on input x, $a = a', c_p$, c and z, makes the following steps: in case c is different from c_p it outputs $\mathsf{V}'(x, a', c, z)$ otherwise it output 1 iff $(x, z) \in \mathcal{R}$.

We now check that Π is a Σ-protocol.

- **Completeness**: The completeness of Π follows from the completeness of Π' except when c is equal to c_p. In this case $\mathcal{P}$ has a witness and sends it to $\mathcal{V}$ that still accepts.
- **Special Soundness**: $\mathsf{Extract}$ on input a collision $(a = (a', c_p), c_1, z_1)$ $(a = (a', c_p), c_2, z_2)$ works as follows:
  - if c_1 and c_2 are different from c_p then it runs the extractor $\mathsf{Extract}'$ of Π' on input x and a collision (a', c_1, z_1) (a', c_2, z_2) returning its output.
  - if c_1 is equal to c_p, it outputs z_1 while instead if c_2 is equal to c_p it outputs z_2.
- **SHVZK** $\mathsf{Sim}(x, c)$ of Π works as follows:
  - computes $(a', z') \leftarrow \mathsf{Sim}'(x, c)$, where $\mathsf{Sim}'(x, c)$ is the simulator of Π';
  - picks $c_p \leftarrow \{0, 1\}^l$.
  - if c_p is equal to c then it aborts, otherwise it outputs $(a = (a', c_p), z')$.

We prove that Π is computational SHVZK, namely: for any l-bit string c, the transcript given in output by $\mathsf{Sim}(x, c)$ is computationally indistinguishable from a honest transcript where the challenge is c and $\mathcal{P}$ runs on common input x and private input w such that $(x, w) \in \mathcal{R}$.

Suppose there exists a distinguisher $\mathcal{A}$ for the SHVZK of Π, then we can show a distinguisher $\mathcal{A}'$ for the SHVZK of Π'.

$\mathcal{A}'$ runs $\mathcal{A}$ that outputs a pair (x, w) and a challenge c. $\mathcal{A}'$ then asks the challenger of SHVZK to produce a transcript (either honest or simulated) for instance x, witness w and challenge c. $\mathcal{A}'$ obtains from the challenger a pair (a', z') such that (a', c, z') is an accepting transcript. $\mathcal{A}'$ picks randomly an l-bit string c', sets $a = a'|c'$, $z = z'$, feeds (a, z) to $\mathcal{A}$, and outputs what $\mathcal{A}$ outputs.

We note that the success probability of $\mathcal{A}'$ is statistically close to the one of $\mathcal{A}$ since the probability that c is equal to c' is negligible and this case is the only deviation among the two distributions.

We finally note that Π' is not WI since an adversarial verifier $\mathcal{V}^\star$ can obtain a witness by just sending a challenge c that is equal to c_p. As a consequence $\mathcal{V}^\star$ can get and output a witness for $x \in L$ during an execution with $\mathcal{P}$. Clearly no PPT simulator can produce the same output unless $L \in \mathsf{BPP}$. □

B.1 Challenge Length of Σ-Protocols

In this section we show how one can reduce or stretch the size of the challenge in a Σ-protocol and in a $\tilde{\Sigma}$-protocol.

Challenge-length amplification. The challenge of a Σ-protocol can be extended through parallel repetition.

Lemma 1. *[CDS94] Let Π be a Σ-protocol (resp. $\tilde{\Sigma}$-protocol) for relation $\mathcal{R}$ and challenge length l. Running Π k-times in parallel for the same instance x corresponds to running Σ-protocol (resp. $\tilde{\Sigma}$-protocol) for $\mathcal{R}$ with challenge length $k \cdot l$.*

Challenge-length reduction.

Lemma 2. *Given a Σ-protocol of challenge length l for the relation $\mathcal{R}$, is possible to construct a Σ-protocol, for the same relation $\mathcal{R}$ with challenge length l' where $l' < l$ [CDS94].*

We now show that lemma 2 it is true even when we consider a $\tilde{\Sigma}$-protocol. One possibility to obtain this result is to convert the $\tilde{\Sigma}$-protocol to a Σ-protocol, and then use Lemma 2. We show how to obtain the same result without first converting the $\tilde{\Sigma}$-protocol to a Σ-protocol.

Lemma 3. *For any $\tilde{\Sigma}$-protocol $\Pi = (\mathcal{P}, \mathcal{V})$, for a relation $\mathcal{R}$ with challenge length l, simulator Sim, and the associated triple $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$, there exists a $\tilde{\Sigma}$-protocol $\Pi' = (\mathcal{P}', \mathcal{V}')$, for the same relation $\mathcal{R}$, with challenge length l', where $l' < l$ and with the same efficiency.*

Proof. We show Π' by presenting the associated triple $(\mathsf{P}'_1, \mathsf{P}'_2, \mathsf{V}')$ of efficient PPT algorithms.

1. P'_1 on input (x, w) and randomness R_1 computes and outputs $a \leftarrow \mathsf{P}_1(x, w; R_1)$.
2. P'_2 on input (x, w), $c \in \{0,1\}^{l'}$, R_1 and randomness R_2, parses R_2 as (pad, R'_2) where pad is an $(l - l')$-bit string, sets $c' = c|pad$, computes $z \leftarrow \mathsf{P}_2(x, w, R_1, c'; R'_2)$ and outputs $z' = (z, pad)$.
3. V' on input $x, a, z' = (z, pad)$ and c, outputs the output of $\mathsf{V}(a, c|pad, z)$.

Completeness follows directly from the completeness of Π.

HVZK We can consider the simulator Sim', that on input x runs as follows:

- picks $c' \leftarrow \{0,1\}^l$;
- runs $(a, z) \leftarrow \mathsf{Sim}(x, c')$;
- sets pad equal to the last $l - l'$ bits of c', and sets c equal to the fist l' bits of c';
- outputs $(a, c, (z, pad))$.

Special soundness follows directly from the special soundness of Π. □

From Lemma 1, 2 and 3, we can claim the following theorem.

Theorem 14. *Suppose that relation $\mathcal{R}$ has a Σ-protocol ($\tilde{\Sigma}$-protocol) Π. Then, for any challenge length l, $\mathcal{R}$ admits a Σ-protocol ($\tilde{\Sigma}$-protocol) Π' with challenge length l'. If $l' \leq l$ than Π' is almost as efficient as Π. Otherwise the communication and computation complexities of Π' are l'/l times the ones of Π.*

C Maurer's Result and Its Extension

Let $(\mathcal{G}, \star)$ and $(\mathcal{H}, \oplus)$ be two groups whose operations are efficiently computable, and let $f : \mathcal{G} \to \mathcal{H}$ be a one-way homomorphism from $\mathcal{G}$ to $\mathcal{H}$. That is, $f(x \star y) = f(x) \oplus f(y))$.

The main protocol Π for relation $\mathcal{R} = \{(x, w) : x = f(w)\}$ with associated algorithm $(\mathsf{P}_1, \mathsf{P}_2, \mathsf{V}))$ and with challenge length l is described below:

- Common input: (description of) $\mathcal{G}$ and $\mathcal{H}$ and $x \in \mathcal{H}$;
- Prover's private input: w such that $x = f(w)$.
- Algorithm P_1.

  On input $(x, w) \in \mathcal{R}$ and random coin tosses R_1, P_1 picks $k \leftarrow \mathcal{G}$, sets $a \leftarrow f(k)$ and outputs a.
- Algorithm P_2.

  On input $(x, w) \in \mathcal{R}$, k, and challenge c, P_2 computes and outputs $z = k \star w^c$.
- Algorithm V.

  On input (x, a, c, z), V outputs 1 iff $f(z) = a \oplus x^c$.

The simulator Sim of Π on input instance x and challenge c works as follows:

- randomly pick $z \leftarrow \mathcal{G}$;
- compute $a = f(z) \oplus x^{-c}$;
- return (a,z).

Theorem 3 of [Mau15] describes the condition under which The Main Protocol is a Σ-protocol. Specifically, for integer y, $u \in \mathcal{G}$ and $(x, w) \in \mathcal{R}$ we have:

- $\gcd(c_1 - c_2, y) = 1$, for all challenges $c_1 \neq c_2 \in \{0, 1\}^l$;
- $f(u) = x^y$.

C.1 Main Protocol is a Chameleon Σ-Protocol

Theorem 15. *The main protocol of [Mau15] is a Chameleon Σ-protocol.*

Proof. We describe algorithm $\mathsf{P}_{\mathsf{sim}}$. Let $(a, \widetilde{z})$ be the output of Sim on input x and challenge $\widetilde{c}$. PPT algorithm $\mathsf{P}_{\mathsf{sim}}$ on input $x, \widetilde{c}$ and the witness w for x and challenge c, computes and outputs $z = \widetilde{z} \star w^{-\widetilde{c}} \star w^c$.

The triple (a, c, z) is an accepting conversation because the test (f$(z) = a \oplus x^c$) of V is successful. Indeed we have

$$a \oplus x^c = f(\widetilde{z}) \oplus x^{-\widetilde{c}} \oplus x^c = f(\widetilde{z}) \oplus f(w)^{-\widetilde{c}} \oplus f(w)^c = f(\widetilde{z} \star w^{-\widetilde{c}} \star w^c) = f(z).$$

We prove now the property of indistinguishability for Chameleon Σ-protocols. We prove that for all pairs of challenges $\widetilde{c}$ and c and for all $(x, w) \in \mathcal{R}$ the two following distributions are indistinguishable:

- (a, c, z), where $a = f(z) \oplus x^{-c}$;
- (a, c, z), where $a = f(\widetilde{z}) \oplus x^{-\widetilde{c}}$ $z = \widetilde{z} \star w^{-\widetilde{c}} \star w^{c}$.

From the SHVZK property follows that the first distribution is perfect indistinguishable from a real transcript $a = f(k)$, c, $z = k \star w^c$ (where k is a random element of $\mathcal{G}$) produced by Π.

We note that $\widetilde{z}$ is a random element of $\mathcal{G}$, so $\widetilde{z} \star w^{-\widetilde{c}}$ is also a random element of $\mathcal{G}$, for this reason we can set $\bar{k} = \widetilde{z} \star w^{-\widetilde{c}}$, and obtain that the second distribution $a = f(\bar{k})$, c, $z = \bar{k} \star w^c$ is identical distributed from a transcript given in output by Π. $\square$

D Classification of Σ-Protocols

In this section, we provide examples for the four classes of Σ-protocols mentioned in Section 1. Table 1 summarizes the classes of Σ-protocols that can be used to construct either a 2-IDTC scheme or a 3-IDTC scheme, or both, and the class of Σ-protocols that cannot be used to instantiate any of our IDTC schemes.

	(Class 1)	*(Class 2)*	*(Class 3)*	*(Class 4)*
2-IDTC	Yes	Yes	No	No
3-IDTC	Yes	No	Yes	No

Table 1: Class of Σ-protocols and their compatibility with IDTC scheme.

Examples of Class 1 and Class 3 Σ-protocols. Schnorr's Σ-protocol for DLog is an example of Class 1 Σ-protocol. Indeed, its first round consists of the prover sending a random group element. Moreover, even if this value is computed by a simulator, knowledge of the witness and of the randomness used by the simulator suffices for the prover to answer any challenge. Blum's Σ-protocol for Hamiltonian graphs instead requires the prover to know the graph to compute the first round. As such, it belongs to Class 3.

An example of a Class 2 Σ-protocol. Recall that Class 2 is the class of Σ-protocols that are Chameleon and require the prover to use the witness to compute the first message. We construct a Class 2 Σ-protocol $\Pi = (\mathcal{P}, \mathcal{V})$ for the relation $\mathsf{DLOG} = \{((\mathcal{G}, q, g, Y), y) : g^y = Y\}$ by using Pedersen's commitment scheme [Ped91] as a 2-IDTC scheme. The TCom algorithm of the 2-IDTC scheme for DLOG based on Pedersen's commitment takes as input the description of a cyclic group $\mathcal{G}$ of order q, a generator g of $\mathcal{G}$ and an element $Y \in \mathcal{G}$. To commit to m, TCom selects $r \leftarrow \mathbb{Z}_q$ uniformly at random and returns $\mathtt{com} = g^r \cdot Y^m, \mathtt{dec} = r, \mathtt{rand} = r$. The decommitment algorithm TDec is straightforward and the trapdoor algorithm TFake, knowing the discrete log of Y, can open the commitment $\mathtt{com}$ as any message m'.

We are now ready to describe our proposed Class 2 Σ-protocol for DLOG with common input $(\mathcal{G}, q, g, Y)$. In the first round, $\mathcal{P}$ computes $(\mathtt{com}_p, \mathtt{dec}_p^0, \mathtt{rand}) \leftarrow \mathsf{TCom}(\mathcal{G}, g, Y, 0)$ and then uses TFake to compute an opening $\mathtt{dec}_p^1$ of $\mathtt{com}_p$ as 1. Finally, $\mathcal{P}$ computes $(\mathtt{com}_0, \mathtt{dec}_0, \mathtt{rand}_0) \leftarrow \mathsf{TCom}(\mathcal{G}, g, Y, \mathtt{dec}_p^0)$ and $(\mathtt{com}_1, \mathtt{dec}_1, \mathtt{rand}_1) \leftarrow \mathsf{TCom}(\mathcal{G}, g, Y, \mathtt{dec}_p^1)$ and sends $(\mathtt{com}_p, \mathtt{com}_0, \mathtt{com}_1)$ to $\mathcal{V}$ that replies with a one bit challenge b. $\mathcal{P}$ answers by sending $\mathtt{dec}_p^b$ and $\mathtt{dec}_b$. The simulator Sim for the Special HVZK receives an instance $(\mathcal{G}, q, g, Y)$ and a bit b and computes $(\mathtt{com}_p, \mathtt{dec}_p^b, \mathtt{rand}) \leftarrow$

$\mathsf{TCom}((\mathcal{G}, q, g, Y), b)$. Commitment $\mathtt{dec}_p^b$ is committed twice obtaining $\mathtt{com}_0$ and $\mathtt{com}_1$ and only $\mathtt{com}_b$ is opened. Notice that, since Pedersen's commitment is perfectly hiding, then the simulation of Sim is perfect. Clearly, $\mathcal{P}$ needs the witness of Y for computing the first round. Moreover, the proposed protocol is Chameleon since Sim commits twice to the same opening of the commitment $\mathtt{com}_p$ but then $\mathcal{P}$, once the discrete log of Y becomes available, can computed $\mathtt{dec}_p^{1-b}$ and then open $\mathtt{com}_{1-b}$ as $\mathtt{dec}_p^{1-b}$.

Examples of Class 4 Σ-protocols. As a first example of a Class 4 Σ-protocol we consider the protocol obtained from the Class 2 Σ-protocol described in the previous paragraph in which $\mathtt{com}_0$ and $\mathtt{com}_1$ are committed by using a new randomly chosen element Y' of $\mathcal{G}$. The perfect simulator Sim is essentially the same as in the previous paragraph but now we observe that after Y' has been chosen by Sim the prover has no way of continuing the protocol (as he does not have the discrete log of Y'). Thus the new Σ-protocol is not Chameleon anymore.

A second example is obtained by using a (non-interactive) commitment scheme that is perfectly binding e computationally hiding instead of a second instance of Pedersen's commitment scheme. For example, the ElGamal encryption scheme can be used to construct such a commitment scheme. In this case, the Σ-protocol obtained is only computational HVZK.