

Improved OR-Composition of Sigma-Protocols

Michele Ciampi¹, Giuseppe Persiano², Alessandra Scafuro³, Luisa Siniscalchi¹,
and Ivan Visconti¹

¹ DIEM, University of Salerno, ITALY
{mciampi,lsiniscalchi,visconti}@unisa.it

² DISA-MIS, University of Salerno, ITALY
giuper@gmail.com

³ Boston University and Northeastern University, USA
scafuro@bu.edu

In [18] Cramer, Damgård and Schoenmakers (CDS) devise an OR-composition technique for Σ -protocols that allows to construct highly-efficient proofs for compound statements. Since then, such technique has found countless applications as building block for designing efficient protocols.

Unfortunately, the CDS OR-composition technique works only if *both* statements are fixed before the proof starts. This limitation restricts its usability in those protocols where the theorems to be proved are defined at different stages of the protocol, but, in order to save rounds of communication, the proof must start even if not all theorems are available. Many round-optimal protocols ([30, 21, 44, 41]) crucially need such property to achieve round-optimality, and, due to the inapplicability of CDS’s technique, are currently implemented using proof systems that requires expensive NP reductions, but that allow the proof to start even if no statement is defined (a.k.a., LS proofs from Lapidot-Shamir [31]).

In this paper we show an improved OR-composition technique for Σ -protocols, that requires only one statement to be fixed when the proof starts, while the other statement can be defined in the last round. This seemingly weaker property is sufficient for the applications, where typically one of the theorems is fixed before the proof starts. Concretely, we show how our new OR-composition technique can directly improve the round complexity of the efficient perfect quasi-polynomial time simulatable argument system of Pass [38] (from four to three rounds) and of efficient resettable WI arguments (from five to four rounds).

1 Introduction

Witness-indistinguishable (WI) proofs. WI¹ proofs are fundamental for the design of cryptographic protocols, particularly when they are also proofs of knowledge (PoK). 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 [24] paradigm or for round-optimal secure computation [30]).

¹ We will use WI to mean both “witness indistinguishability” and “witness indistinguishable”.

Round complexity of cryptographic protocols has been extensively studied both for its practical relevance and for its natural and conceptual interest. Regarding WIPoKs, we know from Blum’s protocol [5] that 3-round WIPoKs exist for all NP languages under the sole assumptions that one-way permutations exist. This result is obtained by designing a WIPoK for the language of Hamiltonian graphs and then by leveraging on the NP-completeness of the language of Hamiltonian graphs. Under stronger cryptographic assumptions, 2-round WI proofs, called ZAPs, and non-interactive WI (NIWI) proofs have been shown in [23, 28, 4]. Neither ZAPs nor NIWI proofs are PoKs.

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 [19, 42].

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). The latter property basically says the following: if the challenge is known in advance, then by just knowing also the theorem, it is possible to 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 and is indistinguishable from the transcript produced by the honest prover when the challenge is c . Blum’s protocol for Graph Hamiltonicity is an example of a Σ -protocol. Another popular example of Σ -protocols is Schnorr’s protocol [42] for proving knowledge of a discrete logarithm.

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 [32, 14, 15, 33, 12, 27, 37, 1, 41, 36, 6, 11, 9, 43, 25, 10, 22, 20, 2, 40] is a fact. This is due to a breakthrough of Cramer et al. [18] that adds WI to the security of Σ -protocol.

OR composition of Σ -protocols. Let L be a language that admits a Σ -protocol Π_L . In [18] it is shown 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 by crucially exploiting the honest-verifier zero-knowledge (HVZK²) property of Π_L . The rationale behind the transformation can be informally explained as follows. The prover wishes 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 because the prover knows only one of the witnesses, say w_b ,

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.

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} is WI (this was shown in [18]). This result was further extended by Garay et al. [25] 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

$$\mathcal{R}_L^{\text{OR}} = \{((x_0, x_1), w) : ((x_0, w) \in \mathcal{R}_L \wedge (x_1 \in L)) \vee ((x_1, w) \in \mathcal{R}_L \wedge (x_0 \in L))\}.$$

Input-delayed proofs. Often in cryptographic protocols there is 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 [31], Lapidot and Shamir (and later on Feige et al. in [24]) 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 LS. Also, we will call *input delayed* a Σ -protocol where the prover computes the first message without knowledge of the statement to be proved.

The 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 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 proof system into a zero-knowledge (ZK) proof system.

The input-delayed property of LS has been instrumental to provide round-efficient constructions from general assumptions, such as: 4-round (optimal) se-

cure 2PC where only one player gets the output (5 rounds when both players get the output) [30], 4-round resettable WI arguments [44, 41], 4-round (optimal) resettable ZK for NP in the BPK model [44, 41].

Despite being so influential to achieve round efficiency for cryptographic protocols, the power of LS unfortunately vanishes as soon as practical constructions are desired. Indeed, similarly to Blum’s protocol, LS 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 adaptive WI proofs. A bit more specifically, we know that there exist Σ -protocols that are input delayed. Schnorr’s protocol [42] 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, not even when used to compose two Σ -protocols that are both input delayed. To see why, recall that the CDS-OR composition technique when applied to Σ -protocol Π_L for language L requires the prover to compute two accepting transcripts, 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 witness for the compound theorem ($x_0 \in L \vee x_1 \in L$), 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 the CDS-OR technique the prover can successfully complete the protocol if and only if *both*³ 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 number of rounds obtained by using LS. 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 needs 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 BPK model [44, 41] and 5-round resettable WI [44, 41]) argument systems.

Additionally, we note that the CDS-OR technique is the bottleneck in the round-complexity of the 4-round straight-line perfect simulatable in quasi-polynomial time argument shown by Pass in [38]. This argument 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 Indis-*

³ 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 .

tinguishable Σ -protocol Π_{OR}^L 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 leads to improved round complexity in previous efficient constructions based on CDS-OR technique. Namely, we describe the following two results that we obtain by making use of our new OR-composition technique:

- Efficient 3-round straight-line perfect quasi-polynomial time simulatable argument system for a large class of useful languages. The previous construction required four rounds [38].
- Efficient 4-round rWI argument system. Previous constructions required five rounds [44, 41].

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 [44, 41]), round-optimal secure two-party computation (using the construction of [30]) and 4-round non-malleable commitments (using the construction of [26]).

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 protocol⁴. 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 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

⁴ Like LS, we will just need the size of x_1 to be known when Π_L^{OR} starts.

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-composition 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, as in [20]. 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 will honestly open the commitment, and honestly compute the third message of Π_1 . Therefore, the third round of Π_L^{OR} , 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 Σ -protocols 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 scheme. 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 $(\text{TCom}, \text{TDec}, \text{TFake})$ 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 that have been used in literature.

Constructing a 2-IDTC scheme. A 2-IDTC scheme can be directly 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 P_{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 $\text{Sim}(x_0, m; r_0)$ and obtains a_0, z_0 . 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 $P_{\text{sim}}((x_0, m, r_0), w_0, m')$ 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 challenge specified by Sim, even if a witness is made available. Blum's protocol for Hamiltonicity is an example of such Σ -protocol.

To commit to m , the sender sends a pair (a_0, a'_0) where, with probability $1/2$, a_0 is obtained by running $\text{Sim}(x_0, m)$ while a'_0 is computed by running the prover of Π_0 , and with probability $1/2$ the above 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-composition 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 (i.e., 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 is long enough the success probability of \mathcal{P}^* is therefore negligible.

Our construction when starting from a 3-IDTC scheme is 3-special sound (i.e., answering to 3 challenges allows one to compute a witness efficiently), and therefore it is a proof of knowledge when the challenge is long enough.

1.3 Discussion

What really matters. Our new OR-composition technique works only when the theorem that has not been defined yet (i.e., 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 scenarios⁵ 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 language of interest is L_0 .

Comparison with the CDS-OR technique. We remark that even in the extremely simplified case where:

1. the two instances x_0 and x_1 are for the same language L ,
2. L admits an *input-delayed* Σ -protocol Π_L which is also special HVZK,
3. 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,
4. 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 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. The works of Cramer [16], Cramer and Damgård [17], and Maurer [35, 34] showed that a protocol (referred to as the *Pre-Image Protocol*) for proving knowledge of a pre-image of a group homomorphism unifies and generalizes a large number of protocols in the literature. Classic Σ -protocols, such as Schnorr's protocol [42] and the Guillou-Quisquater protocol [29], are particular cases of this abstraction. We show that the *Pre-Image Protocol* is a chameleon Σ -protocol and can thus be used in our construction.

What is in and what is out. As mentioned previously, the Σ -protocol for L_1 can be any *input-delayed* Σ -protocol. We now discuss which Σ -protocols 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

⁵ 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.

first three classes can be used in our OR transform (by instantiating either a 2-IDTC or a 3-IDTC scheme).

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 -protocols 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 -protocols 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 -protocols 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 -protocols 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-composition 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 efficiently⁶ transform a Σ -protocol for a relation \mathcal{R} into a *round-efficient* argument with more appealing features.

Efficient 3-round straight-line perfect quasi-polynomial time simulatable argument system. We achieve this result directly, using the construction of Pass [38] and replacing the CDS-OR technique with our technique. As a result the first round of the verifier of [38] can be postponed, therefore reducing the round complexity from four to three rounds. Our construction works for all languages admitting a perfect chameleon Σ -protocol.

⁶ 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 resulting constructions.

Efficient 4-round resettable WI arguments. It is well known [8] 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 only four rounds 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 argument system 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-composition 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 [38, 21] and we are using the form used in [21] where the puzzle is sent only in the second round. [21] 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 without NP reductions.

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 computationally efficient and round-efficient constructions. For instance, the 4-round (optimal) resettable ZK argument systems in the BPK model provided in [44, 41], 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 (resettable-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 proof using the CDS-OR technique. We observe that this setting closely resembles the setting of the 4-round resettable WI (rWI) 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 rWI protocol, and have the verifier just prove knowledge of its secret keys, thus avoiding the need of the additional first round.

Our OR-composition technique could also be useful in replacing the use of LS in the 4-round non-malleable commitment scheme of [26], and in the round-optimal secure two-party computation protocol of [30].

1.5 Open Problems

Our OR-composition technique relaxes the requirement of CDS-OR of requiring knowledge of *all* instances already at the beginning of the protocol. However still our result 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 proving (still preserving WI) knowledge of 1 out of n theorems and that requires knowledge of (at least some) theorems only after the second round. 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 give some useful definitions. More definitions can be found in the full version.

We denote the security parameter by λ . 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} (or, simply, a *relation*) is a subset of $\{0, 1\}^* \times \{0, 1\}^*$ 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 $L_{\mathcal{R}} = \{x \mid \exists w : (x, w) \in \mathcal{R}\}$. If $(x, w) \in \mathcal{R}$, we say that w is a *witness* for *instance* x . Following [25], 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 $\hat{L}_{\mathcal{R}}$.

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.

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 . \mathcal{P} receives as an additional private input a witness w for x and the security parameter 1^λ in unary. The protocol Π has the following form:

1. \mathcal{P} runs algorithm P_1 on common input x , private input w , security parameter 1^λ and randomness R obtaining $a = P_1(x, w, 1^\lambda; R)$ 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} runs algorithm P_2 on input x, w, R, c and sends $z \leftarrow P_2(x, w, R, c)$ to \mathcal{V} .
4. \mathcal{V} outputs $V(x, a, c, z)$ (i.e., \mathcal{V} 's decision to accept ($b = 1$) or reject ($b = 0$)).

We call (P_1, P_2, 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 a, z correspond to the messages computed running the honest algorithms, respectively, of P_1 and P_2 , and c is a random string, in $\{0, 1\}^l$. A 3-move transcript (a, c, z) is *accepting* for x if and only if $V(x, a, c, z) =$

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 [18]). 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 $x \in L_{\mathcal{R}}$, security parameter 1^λ 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 Π is Perfect when the two distributions are identical.

Not to overburden the descriptions of protocols and simulators, we will omit the specification of the security parameter when it is clear from the context.

In the rest of the paper, we will call a 3-move protocol that enjoys Completeness, Special Soundness and Honest-Verifier Zero Knowledge (HVZK⁷) 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 ([19]). 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 has the same efficiency.

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

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

⁷ 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.

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 .

We will use the definition of proof of knowledge given in [3, 19].

Theorem 2 ([19]). *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 2 (Input-delayed Σ -protocol). *A Σ -protocol $\Pi = (\mathcal{P}, \mathcal{V})$ with \mathcal{P} running PPT algorithms (P_1, P_2) is an input-delayed Σ -protocol if P_1 takes as input only the length 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.*

Definition 3 (Witness-delayed Σ -protocol). *A Σ -protocol $\Pi = (\mathcal{P}, \mathcal{V})$ for a relation \mathcal{R} with associated algorithms (P_1, P_2, V) is a witness-delayed Σ -protocol if P_1 takes as input only the common instance x .*

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 SHVZK simulator Sim and an algorithm P_{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*

$$\left\{ R \leftarrow \{0, 1\}^{|x|^d}; (a, z_0) \leftarrow \text{Sim}(x, c_0; R); z_1 \leftarrow P_{\text{sim}}((x, c_0, R), w, c_1) : \right. \\ \left. (x, a, c_1, z_1) \right\}$$

and

$$\left\{ (a, z_1) \leftarrow \text{Sim}(x, c_1) : (x, a, c_1, z_1) \right\}$$

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 delayed indistinguishability is perfect, and Π is a Perfect Chameleon Σ -protocol.

We remark that a chameleon Σ -protocol Π has two modes of operations: the standard mode when \mathcal{P} runs P_1 and P_2 , and a *delayed* mode when \mathcal{P} uses Sim and P_{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

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

is indistinguishable from

$$\{R \leftarrow \{0, 1\}^{|x|^d}; a \leftarrow P_1(x, w; R); z \leftarrow P_2(x, w, R, c) : (a, c, z)\}.$$

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 (P_1, P_2) , then Π is WI even when the pair $(\text{Sim}, P_{\text{sim}})$ is used. Finally, we observe that Chameleon Σ -protocols do exist and Schnorr's protocol [42] is one example. When considering the algorithms associated to a Chameleon Σ -protocol, we will add P_{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 every malicious verifier \mathcal{V}^* , there exists a negligible function ν such that for all x, w, w' such that $(x, w) \in \mathcal{R}_L$ and $(x, w') \in \mathcal{R}_L$

$$\left| \text{Prob} [\langle \mathcal{P}(w, 1^\lambda), \mathcal{V}^* \rangle(x) = 1] - \text{Prob} [\langle \mathcal{P}(w', 1^\lambda), \mathcal{V}^* \rangle(x) = 1] \right| \leq \nu(\lambda).$$

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 ([18]). Every Perfect $\tilde{\Sigma}$ -protocol⁸ is Perfect WI.

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

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

In this section we describe the CDS-OR [18] transform in details. Let Π be a $\tilde{\Sigma}$ -protocol for polynomial-time relation \mathcal{R} with challenge length l , associated algorithms (P_1, P_2, V) and HVZK simulator Sim . The CDS-OR transform constructs a $\tilde{\Sigma}$ -protocol Π_{OR} with associated algorithms $(P_1^{\text{OR}}, P_2^{\text{OR}}, V_{\Sigma}^{\text{OR}})$ for the relation

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

⁸ 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.

Protocol 1 *CDS-OR Transform.*Common input: (x_0, x_1) . \mathcal{P} 's private input: (b, w) with $b \in \{0, 1\}$ and $(x_b, w) \in \mathcal{R}$. $P_1^{\text{OR}}((x_0, x_1), (b, w); R_1)$. Set $a_b = P_1(x_b, w; R_1)$. Compute $(a_{1-b}, c_{1-b}, z_{1-b}) \leftarrow \text{Sim}(x_{1-b})$. Output (a_0, a_1) . $P_2^{\text{OR}}((x_0, x_1), (b, w), c, R_1)$. Set $c_b = c \oplus c_{1-b}$. Compute $z_b \leftarrow P_2(x_b, w, c_b, R_1)$. Output $((c_0, c_1), (z_0, z_1))$. $V_{\Sigma}^{\text{OR}}((x_0, x_1), (a_0, a_1), c, ((c_0, c_1), (z_0, z_1)))$. V_{Σ}^{OR} accepts if and only if $c = c_0 \oplus c_1$ and $V(x_0, a_0, c_0, z_0) = 1$ and $V(x_1, a_1, c_1, z_1) = 1$.**Theorem 4 ([18, 25]).** *If Π is a $\tilde{\Sigma}$ -protocol for \mathcal{R} then Π_{OR} is a $\tilde{\Sigma}$ -protocol for \mathcal{R}_{OR} and is WI for relation*

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

Moreover, if Π is a Perfect $\tilde{\Sigma}$ -protocol for \mathcal{R} then Π^{OR} is WI for \mathcal{R}_{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 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 (details are available in the full version). Hence, the construction outlined above can be used to construct $\tilde{\Sigma}$ -protocol $\Pi_{\text{OR}}^{\mathcal{R}_0, \mathcal{R}_1}$ for relation

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

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

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

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}_{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}_{OR} and Theorem 5 and Theorem 4 still hold.

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 (a t -IDTC, in short) scheme for \mathcal{R} with message space M is a triple of PPT algorithms (TCom, TDec, TFake) where TCom

is the randomized commitment algorithm that takes as input security parameter 1^λ , an instance $x \in \hat{L}_{\mathcal{R}}$ (with $|x| = \text{poly}(\lambda)$) and a message $m \in M$ and outputs commitment com , decommitment dec , and auxiliary information rand ; TDec is the verification algorithm that takes as input $(x, \text{com}, \text{dec}, m)$ and decides whether m is the decommitment of 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 com of m_1 with respect to instance x and associated auxiliary information rand and produces decommitment information dec_2 such that TDec, on input $(x, \text{com}, \text{dec}_2, m_2)$, outputs 1.

A t -Instance-Dependent Trapdoor Commitment enjoys:

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

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

- **t -Special Extract:** there exists an efficient algorithm ExtractTCom that, on input x , commitment com , pairs $(\text{dec}_i, m_i)_{i=1}^t$ of openings and messages such that
 - for $1 \leq i < j \leq t$ we have that $m_i \neq m_j$;
 - $\text{TDec}(x, \text{com}, \text{dec}_i, m_i) = 1$, for $i = 1, \dots, t$;
 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., $\nu(\cdot) = 0$) such that, for all $x \in L_{\mathcal{R}}$ and all $m_0, m_1 \in M$, it holds that

$$\text{Prob}[b \leftarrow \{0, 1\}; (\text{com}, \text{dec}, \text{rand}) \leftarrow \text{TCom}(1^\lambda, x, m_b) : b = \mathcal{A}(x, \text{com}, m_0, m_1)] \leq \frac{1}{2} + \nu(\lambda).$$

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

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

and $\{(\text{com}, \text{dec}_2, \text{rand}) \leftarrow \text{TCom}(1^\lambda, x, m_2) : (\text{com}, \text{dec}_2)\}$ 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 trapdooriness 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 $(P_1, P_2, V, P_{\text{sim}})$ be a Chameleon Σ -protocol for polynomial-time relation \mathcal{R} with a security parameter 1^λ . Let l be the challenge length of Π and let Sim be a SHVZK simulator associated to Π . We construct a t -IDTC scheme $(\text{TCom}_\Pi, \text{TDec}_\Pi, \text{TFake}_\Pi)$ 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 .

- $\text{TCom}_\Pi(1^\lambda, x, m_1)$: On input x and $m_1 \in M$, pick randomness R and compute $(a, z) \leftarrow \text{Sim}(x, m_1; R)$. Output $\text{com} = a$, $\text{dec} = z$ and $\text{rand} = R$;
- $\text{TDec}_\Pi(x, \text{com}, \text{dec}, m_1)$: On input $x, \text{com}, \text{dec}$ and m_1 , run $b = \text{V}(x, \text{com}, m_1, \text{dec})$ and accept m_1 as the decommitted message iff $b = 1$.
- TFake_Π : On input $(x, w) \in \mathcal{R}$, messages $m_1, m_2 \in M$, for m_2 and rand for com , output $z = \text{P}_{\text{sim}}((x, m_1, \text{rand}), w, m_2)$.

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 com is a commitment with respect to instance x and let dec_1 and dec_2 be two openings of com as messages $m_1 \neq m_2$, respectively. Then, triplets $(\text{com}, m_1, \text{dec}_1)$ and $(\text{com}, m_2, \text{dec}_2)$ 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 ExtractTCom to be the algorithm that runs algorithm Extract (that exists by the special soundness of Π) on input the collision. ExtractTCom returns the witness for x computed by Extract .

(Perfect) Trapdooriness. It follows from the Perfect 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 witness-delayed Σ -protocol Π with associated algorithms $(\text{P}_1, \text{P}_2, \text{V})$ and security parameter 1^λ . Let l denote the challenge length of Π . We construct a 3-IDTC scheme for message space $M = \{0, 1\}^l$ for $x \in L_R$, as follows.

Protocol 3 3-IDTC scheme.

- TCom_Π : On input $1^\lambda, x$ and $m_1 \in M$, pick randomness R and compute $(a_0, z) \leftarrow \text{Sim}(x, m_1)$ and $a_1 \leftarrow \text{P}_1(x; R)$. Let $\text{com}_0 = a_0$ and $\text{com}_1 = a_1$. Output $\text{com} = (\text{com}_b, \text{com}_{1-b})$ for a randomly selected bit b , $\text{dec} = z$ and $\text{rand} = R$.
- TDec_Π : On input $x, \text{com} = (\text{com}_0, \text{com}_1), \text{dec}$ and m_1 , accept m_1 if and only if either $\text{V}(x, \text{com}_0, m_1, \text{dec}) = 1$ or $\text{V}(x, \text{com}_1, m_1, \text{dec}) = 1$.
- TFake_Π : On input $(x, w) \in \mathcal{R}$, messages $m_1, m_2 \in M$, commitment com for m_1 and rand for com , output $z \leftarrow \text{P}_2(x, w, \text{rand}, m_2)$.

Theorem 7. *If Π is a witness-delayed Σ -protocol for \mathcal{R} , with the associated algorithms $(\text{P}_1, \text{P}_2, \text{V})$, 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 $\text{dec}_1, \text{dec}_2$ and dec_3 , for distinct messages m_1, m_2 and m_3 , for the same commitment com computed w.r.t. x . In this case, we have three accepting transcript 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 Extract for Π on the collision and obtain a witness for x .

Trapdooriness. 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:

- On input x and $m_1, m_2 \in M$, the prover selects random coin tosses R and computes $(a_0, z) \leftarrow \text{Sim}(x, m_2)$, $a_1 \leftarrow P_1(x; R)$. It picks $b \leftarrow \{0, 1\}$ and sends $\text{com} = (a_b, a_{1-b})$, $\text{dec} = z, m_2$.

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 $a_0 \leftarrow P_1(x; r_1)$, $z \leftarrow P_2(x, w, r_1, m_2)$ and $a_1 \leftarrow P_1(x; r_2)$. It picks $b \leftarrow \{0, 1\}$ and sends $\text{com} = (a_b, a_{1-b})$, $\text{dec} = z, m_2$.

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 $\text{Sim}(x, m_2)$. Formally:

- On input x and $m_1, m_2 \in M$, the prover selects random coin tosses R and computes $a_0 \leftarrow P_1(x; R)$, $z \leftarrow P_2(x, w, R, m_2)$ and $(a_1, \bar{z}) \leftarrow \text{Sim}(x, m_1)$. It picks $b \leftarrow \{0, 1\}$ and sends $\text{com} = (a_b, a_{1-b})$, $\text{dec} = z, m_2$.

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_1 and opens to m_2 using the trapdoor.

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

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, $I = (\text{TCom}_{\Pi_0}, \text{TDec}_{\Pi_0}, \text{TFake}_{\Pi_0})$, with $t = 2$ or $t = 3$, and \mathcal{R}_1 a relation admitting an input-delayed Σ -protocol Π_1 with associated algorithms (P_1^1, P_2^1, V^1) and simulator Sim^1 . We show a Σ -protocol Π^{OR} for the OR relation:

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

We denote by $(P_1^{\text{OR}}, P_2^{\text{OR}}, V^{\text{OR}})$ 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 the full version we show that the challenge length of a Σ -protocol can be easily expanded/reduced.

We remind that prover algorithm P_2^{OR} receives as further input the randomness (R_1, \mathbf{rand}_1) used by P_1^{OR} to produce the first-round message.

Protocol 4 Protocol Π^{OR} for \mathcal{R}_{OR} .

Common input: $(x_0, 1^\lambda)$, where 1^λ is the security parameter.

1. $P_1^{\text{OR}}(x_0, 1^\lambda)$. Pick random R_1 and compute $a_1 \leftarrow P_1^1(1^\lambda; R_1)$. Then commit to a_1 by running $(\mathbf{com}, \mathbf{dec}_1, \mathbf{rand}_1) \leftarrow \text{TCom}_{\Pi_0}(1^\lambda, x_0, a_1)$. Output \mathbf{com} .
2. $P_2^{\text{OR}}((x_0, x_1), c, (w, b), (\mathbf{rand}_1, R_1))$ (with $(x_b, w) \in \mathcal{R}_b$).
If $b = 1$, compute $z_1 \leftarrow P_2^1(x_1, w, R_1, c)$ and output $(\mathbf{dec}_1, a_1, z_1)$.
If $b = 0$, compute $(a_2, z_2) \leftarrow \text{Sim}^1(x_1, c)$, $\mathbf{dec}_2 \leftarrow \text{TFake}_{\Pi_0}(x_0, w, a_1, a_2, \mathbf{com}, \mathbf{rand}_1)$ and output $(\mathbf{dec}_2, a_2, z_2)$.
3. V^{OR} , on input (x_0, x_1) , \mathbf{com} , c , and (\mathbf{dec}, a, z) received from Π^{OR} , outputs 1 iff

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

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}_{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 $(\mathbf{com}, c, (\mathbf{dec}, a, z))$ and $(\mathbf{com}, c', (\mathbf{dec}', a', z'))$ 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 \mathbf{dec} and \mathbf{dec}' are two openings of \mathbf{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 Sim^{OR} that, on input (x_0, x_1) and challenge c , sets $(a, c, z) \leftarrow \text{Sim}_1(x_1, c)$ and $(\mathbf{com}, \mathbf{dec}) \leftarrow \text{TCom}_{x_0}(a)$, and outputs $(\mathbf{com}, c, (\mathbf{dec}, a, z))$. Next, we show that the transcript generated by Sim^{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 $(\mathbf{com}, c, (\mathbf{dec}, a, z))$ 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 transcript of Π_1 on input x_1 constrained to c being the challenge; moreover, $(\mathbf{com}, \mathbf{dec})$ 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 $\text{Sim}_1(x_1, c)$ which is exactly as in the output Sim^{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 $(\text{com}, c, (\text{dec}, a, z))$ 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 Sim^{OR} (that is by running Sim_1 on input x_1 and c). In addition, in the output of Sim^{OR} , (com, 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 trapdooriness 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 .

Consider three accepting transcripts for Π^{OR} and input (x_0, x_1) : $(\text{com}, c_1, (\text{dec}_1, a_1, z_1))$, $(\text{com}, c_2, (\text{dec}_2, a_2, z_2))$ and $(\text{com}, c_3, (\text{dec}_3, a_3, z_3))$.

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, dec_1 and dec_2 and dec_3 are three openings of the same com with respect to x_0 for messages a_1 , a_2 and a_3 ; then we can obtain a witness w_0 for x_0 by the 3-Special Extract of the 3-IDTC scheme.

We stress that having a long enough challenge, 3-special soundness implies the proof of knowledge property.

SHVZK property. This is similar to the proof for the construction based on 2-IDTC.

5.1 Witness Indistinguishability of Our Transform

In this section we discuss the *adaptive WI* property of Π^{OR} . Roughly speaking, adaptive WI means that in the WI experiment the adversary \mathcal{A} is not forced to choose *both* theorems x_0 and x_1 at the onset of the experiment. Rather, she can choose theorem x_1 and 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 $b \leftarrow \{0, 1\}$. The prover then receives x_1 and w_b and proceeds to complete the protocol. The adversary wins the game if she can guess b with probability non-negligibly greater than $1/2$. More formally, we consider adaptive WI for polynomial-time relation

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

and for the weaker relation

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

The adaptive WI experiment, $\text{ExpWI}_{\mathcal{A}}^{\delta}(x_0, \lambda, \text{aux})$ with $\delta \in \{c, p\}$, is parameterized by PPT adversary \mathcal{A} and has three inputs: instance x_0 , security parameter λ , and auxiliary information aux for \mathcal{A} .

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

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

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

Definition 7. Π^{OR} is Adaptive Witness Indistinguishable (*resp.*, Adaptive Perfect Witness Indistinguishable) if for every adversary \mathcal{A} there exists a negligible function ν such that for all aux and x_0 it holds that $\text{Adv}_{\mathcal{A}}^c(x_0, \lambda, \text{aux}) \leq \nu(\lambda)$ (*resp.*, $\text{Adv}_{\mathcal{A}}^p(x_0, \lambda, \text{aux}) = 0$).

Next, in Theorem 9, we prove the Adaptive Perfect WI of Π^{OR} when both Π_0 and Π_1 are perfect SHVZK. When one of Π_0 and Π_1 is not perfect, we would like to prove that Π^{OR} is Adaptive WI. In Theorem 10 we prove a weaker form of Adaptive WI in which the adversary is restricted in his choice of witnesses (w_0, w_1) for relation $\mathcal{R}_{\text{OR}}^c$. We leave open the problem of an OR-composition technique that gives Adaptive WI when the Σ -protocol composed are not both perfect SHVZK.

Theorem 9. *If Π_0 and Π_1 are perfect SHVZK then Π^{OR} is Adaptive Perfect Witness Indistinguishable.*

Proof. The proof considers the following three cases:

- Case 1.** $(x_0, w_0) \in \mathcal{R}_0$ and $(x_1, w_1) \in \mathcal{R}_1$;
- Case 2.** $(x_0, w_0) \in \mathcal{R}_0$ and $(x_0, w_1) \in \mathcal{R}_0$;
- Case 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. The first hybrid experiment $\mathcal{H}_1(x_0, \lambda, \text{aux})$ is the original experiment $\text{ExpWI}_{\mathcal{A}}^p(x_0, \lambda, \text{aux})$ in which $b = 1$ (and thus \mathcal{P} uses witness w_1). That is,

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

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

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

The trapdooriness of the instance-dependent trapdoor commitment scheme based on Π_0 guarantees that $\mathcal{H}_1(x_0, \lambda, \mathbf{aux})$ and $\mathcal{H}_2(x_0, \lambda, \mathbf{aux})$ are perfectly indistinguishable for all λ .

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

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

By the perfect SHVZK of Π_1 , we have that $\mathcal{H}_2(x_0, \lambda, \mathbf{aux})$ and $\mathcal{H}_3(x_0, \lambda, \mathbf{aux})$ are perfectly indistinguishable for all λ . The proof ends with the observation that $\mathcal{H}_3(x_0, \lambda, \mathbf{aux})$ is exactly experiment $\mathsf{ExpWI}_{\mathcal{A}}^p(x_0, \lambda, \mathbf{aux})$ when $b = 0$.

Case 2. The first hybrid experiment $\mathcal{H}_1(x_0, \lambda, \mathbf{aux})$ is again the original experiment $\mathsf{ExpWI}_{\mathcal{A}}^p(x_0, \lambda, \mathbf{aux})$ in which $b = 1$ (and thus \mathcal{P} uses witness w_1). The second hybrid experiment $\mathcal{H}_2(x_0, \lambda, \mathbf{aux})$ differs from $\mathcal{H}_1(x_0, \lambda, \mathbf{aux})$ in the way TFake is executed (namely, using as input w_0 instead of w_1). More specifically,

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

The trapdooriness of the instance-dependent trapdoor commitment scheme based on Π_0 implies that $\mathcal{H}_1(x_0, \lambda, \mathbf{aux})$ is perfectly indistinguishable from $\mathcal{H}_2(x_0, \lambda, \mathbf{aux})$ for all λ . The proof ends with the observation that $\mathcal{H}_2(x_0, \lambda, \mathbf{aux})$ is exactly experiment $\mathsf{ExpWI}_{\mathcal{A}}^p(x_0, \lambda, \mathbf{aux})$ when $b = 0$.

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

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

The Perfect WI property of Π_1 implies that $\mathcal{H}_1(x_0, \lambda, \mathbf{aux})$ is perfectly indistinguishable from $\mathcal{H}_2(x_0, \lambda, \mathbf{aux})$. The proof ends with the observation that $\mathcal{H}_2(x_0, \lambda, \mathbf{aux})$ is exactly the experiment $\text{ExpWI}_{\mathcal{A}}^p(x_0, \lambda, \mathbf{aux})$ when $b = 0$.

Next we consider the computational case in which one of Π_0 and Π_1 is not Perfect SHVZK (but they are still both SHVZK).

Theorem 10. *If Π_0 and Π_1 are SHVZK then Π^{OR} is Adaptive Witness Indistinguishable with respect to adversaries that output (x_1, w_0, w_1) such that at least one of w_0 and w_1 is a witness for $x_1 \in L_{\mathcal{R}_1}$.*

Proof. We prove this theorem by considering the following two cases:

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

Case 1. In this case the proof follows closely the one of Case 1 of Theorem 9, with the difference that hybrids here are only computationally indistinguishable.

Case 2. In this case we show that there exists \mathcal{A}' for Case 1 that has the same success probability of \mathcal{A} . Suppose indeed that both w_0 and w_1 are witnesses for x_1 and that \mathcal{A} breaks the adaptive WI property of Π^{OR} . Then, by definition of $\mathcal{R}_{\text{OR}}^c$ and by Definition 7, there exists \mathcal{A}' that has in his description a witness w_2 for x_0 . Indeed, the output of \mathcal{A} 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 $\mathcal{P}((x_0, x_1), w_1)$. Therefore \mathcal{A}' would contradict Case 1 and thus there exists no successful \mathcal{A} for Case 2.

6 Applications

In this section, we describe the application of our new OR-composition technique for constructing a 3-round straight-line perfect quasi-polynomial time simulatable argument system. In the full version we also show an efficient 4-round resettable WI argument system and an efficient 4-round resettable zero knowledge with concurrent soundness argument system in the BPK model.

A 3-round efficient perfect quasi-polynomial time simulatable argument system. In [38], 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 [38], 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 Chameleon Σ -protocol Π_0 . One drawback of reducing to 3 rounds the result of [38] is that we can use only a perfect Σ -protocolsince the goal is to obtain perfect WI in 3 rounds.

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 [38], the simulator is allowed to run in time $\lambda^{\text{poly}(\log(\lambda))}$. Loosely speaking, we say that an interactive argument is $\lambda^{\text{poly}(\log(\lambda))}$ -perfectly simulatable if for any adversarial verifier there exists a simulator running in time $\lambda^{\text{poly}(\log(\lambda))}$, where λ is the size of the statement being proved, whose output is identically distributed to the output of the adversarial verifier.

Definition 8 (One-way functions for sub-exponential circuits. [38]). *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 $z \in \{0, 1\}^{\text{poly}(\lambda)}$*

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

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

For our result we consider a one-way functions for sub-exponential circuits that is also one-to-one.

Definition 9 (straight-line $T(\lambda)$ simulatability, Def. 31 of [39]). *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 $L \in \text{NP}$,*

with the witness relation \mathcal{R}_L , is straight-line $T(\lambda)$ -simulatable if for every PPT machine \mathcal{V}^* 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 $\lambda = |x|$)

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

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).

For any NP-language L we consider the perfect chameleon Σ -protocol Π_L for the relation R_L . Also we consider the Schnorr Σ -protocol Π_{DLOG} the following relation $\text{DLOG} = \{((\mathcal{G}, q, g, Y), y) : g^y = Y\}$ with the associated NP-language L_{DLOG} , over groups \mathcal{G} of prime-order q , and use our OR-composition technique to obtain a new Σ -protocol $\Pi^{\text{OR}} = (\mathcal{P}^{\text{OR}}, \mathcal{V}^{\text{OR}})$ for the relation

$$\mathcal{R}_{\text{OR}} = \left\{ ((x_L, x_{\text{DLOG}}), w) : \left((x_L, w) \in \mathcal{R}_L \wedge x_{\text{DLOG}} \in \hat{L}_{\text{DLOG}} \right) \text{OR} \right. \\ \left. \left((x_{\text{DLOG}}, w) \in \text{DLOG} \wedge x_L \in \hat{L}_{\mathcal{R}_L} \right) \right\}$$

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

Let f be a sub-exponentially hard one-to-one 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 $L \in \text{NP}$ and $k = \frac{1}{\alpha} + 1$. Our 3-round straight-line quasi-polynomial time simulatable argument system for $x \in L$ is the following.

Protocol 5 A 3-round straight-line quasi-polynomial time simulatable argument system.

Common input: An instance x of a language $L \in \text{NP}$ with witness relation R_L with a perfect chameleon Σ -protocol, and 1^λ as security parameter.

Private input: \mathcal{P} has w as a private input, s.t. $(x, w) \in \mathcal{R}_L$.

Round 1. $\mathcal{P} \rightarrow \mathcal{V}$:

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 a one-to-one one-way function for some constant α w.r.t circuits of size 2^{λ^α} .
2. \mathcal{P} computes $a \leftarrow \mathcal{P}_1^{\text{OR}}((x, 1^{\lambda^\alpha}); R_1)$.
3. \mathcal{P} sends (p, q, g) and a to \mathcal{V} .

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

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. $\mathcal{P} \rightarrow \mathcal{V}$:

1. \mathcal{P} computes $z \leftarrow P_2^{\text{OR}}((x, ((p, q, g), Y)), w, c, R_1)$.
2. \mathcal{P} sends z to \mathcal{V} .
3. \mathcal{V} accepts if and only if $V^{\text{OR}}((x, ((p, q, g), Y)), a, c, z) = 1$.

We remark that we are using the same assumption of [7] that allows the adversary of DLog to generate the DLog parameters while the challenger selects the random element of the group.

Theorem 11. *If Π^{OR} is a perfect Σ -protocol for OR composition of \mathcal{R}_L and DLOG, then Protocol 5 is a 3-round straight-line perfectly $\lambda^{O(\log^k \lambda)}$ -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 when the challenge is long enough. 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}^* for protocol Π . E internally incorporates \mathcal{P}^* 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 [7]).

Quasi-polynomial time perfect simulation. 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 compute the first round. After the simulator receives Y it checks that Y 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 $\text{poly}(2^{\log^k \lambda})$, 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 honest 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.

7 Acknowledgments

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

The work of the third author was supported by the MACS project under NSF Frontier grant CNS-1414119 and by NSF grant 1012798. This work was done in part while the third author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467.

For the full version of this work see [13].

References

1. Abe, M., Okamoto, T., Suzuki, K.: Message recovery signature schemes from sigma-protocols. *IEICE Transactions* 96-A(1), 92–100 (2013)
2. Bellare, M., Fischlin, M., Goldwasser, S., Micali, S.: Identification protocols secure against reset attacks. In: *EUROCRYPT '01*. LNCS, vol. 2045, pp. 495–511. Springer (2001)
3. Bellare, M., Goldreich, O.: On defining proofs of knowledge. In: *CRYPTO '92*. pp. 390–420 (1992)
4. Bitansky, N., Paneth, O.: Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In: *TCC*. pp. 401–427 (2015)
5. Blum, M.: How to prove a theorem so no one else can claim it. In: *International Congress of Mathematicians*. p. 1444 (1986)
6. Blundo, C., Persiano, G., Sadeghi, A., Visconti, I.: 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*. LNCS, vol. 5283, pp. 364–378 (2008)
7. Canetti, R., Dakdouk, R.R.: Extractable perfectly one-way functions. In: *ICALP*. pp. 449–460 (2008)
8. Canetti, R., Goldreich, O., Goldwasser, S., Micali, S.: Resettable zero-knowledge (extended abstract). In: *STOC*. pp. 235–244 (2000)
9. Catalano, D., Dodis, Y., Visconti, I.: 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*. LNCS, vol. 3876, pp. 120–144. Springer (2006)
10. Catalano, D., Visconti, I.: Hybrid trapdoor commitments and their applications. In: *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005*. pp. 298–310 (2005)
11. Catalano, D., Visconti, I.: Hybrid commitments and their applications to zero-knowledge proof systems. *Theor. Comput. Sci.* 374(1-3), 229–260 (2007)
12. Chaidos, P., Groth, J.: Making sigma-protocols non-interactive without random oracles. In: *PKC 2015*. pp. 650–670 (2015)

13. Ciampi, M., Persiano, G., Scafuro, A., Siniscalchi, L., Visconti, I.: Improved OR composition of Sigma-protocols. IACR Cryptology ePrint Archive 2015, 810 (2015), <http://eprint.iacr.org/2015/810>
14. Ciampi, M., Persiano, G., Siniscalchi, L., Visconti, I.: A transform for NIZK almost as efficient and general as the Fiat-Shamir transform without programmable random oracles. IACR Cryptology ePrint Archive 770 (2015), <http://eprint.iacr.org/2015/770>
15. Ciampi, M., Persiano, G., Siniscalchi, L., Visconti, I.: A transform for NIZK almost as efficient and general as the Fiat-Shamir transform without programmable random oracles. In: Theory of Cryptography - 13th Theory of Cryptography Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, (2016)
16. Cramer, R.: Modular design of secure yet practical cryptographic protocols. Ph.D. thesis, University of Amsterdam (1996)
17. Cramer, R., Damgård, I.: Zero-knowledge proofs for finite field arithmetic; or: Can zero-knowledge be for free? In: Krawczyk, H. (ed.) CRYPTO '98. LNCS, vol. 1462, pp. 424–441. Springer (1998)
18. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs of partial knowledge and simplified design of witness hiding protocols. In: Desmedt, Y. (ed.) Advances in Cryptology — CRYPTO '94, LNCS, vol. 839, pp. 174–187. Springer Berlin Heidelberg (1994)
19. Damgård, I.: On Σ -protocol. <http://www.cs.au.dk/~ivan/Sigma.pdf> (2010)
20. Damgård, I., Groth, J.: Non-interactive and reusable non-malleable commitment schemes. In: STOC '03. pp. 426–437 (2003)
21. Di Crescenzo, G., Persiano, G., Visconti, I.: Constant-round resettable zero knowledge with concurrent soundness in the bare public-key model. In: Franklin, M.K. (ed.) CRYPTO 2004, 24th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 2004. LNCS, vol. 3152, pp. 237–253. Springer (2004)
22. Di Crescenzo, G., Visconti, I.: Concurrent zero knowledge in the public-key model. In: Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005. LNCS, vol. 3580, pp. 816–827. Springer (2005)
23. Dwork, C., Naor, M.: Zaps and their applications. In: FOCS 2000. pp. 283–293 (2000)
24. Feige, U., Lapidot, D., Shamir, A.: Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In: FOCS '90. pp. 308–317. IEEE Computer Society (1990)
25. Garay, J.A., MacKenzie, P., Yang, K.: Strengthening zero-knowledge protocols using signatures. *Journal of Cryptology* 19(2), 169–209 (2006)
26. Goyal, V., Richelson, S., Rosen, A., Vald, M.: An algebraic approach to non-malleability. In: 55th FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014. pp. 41–50. IEEE Computer Society (2014)
27. Groth, J., Kohlweiss, M.: One-out-of-many proofs: Or how to leak a secret and spend a coin. In: EUROCRYPT 2015. pp. 253–280 (2015)
28. Groth, J., Ostrovsky, R., Sahai, A.: Perfect non-interactive zero knowledge for NP. In: EUROCRYPT 2006. pp. 339–358 (2006)
29. Guillou, L.C., Quisquater, J.: A practical zero-knowledge protocol fitted to security microprocessor minimizing both transmission and memory. In: Günther, C.G. (ed.) EUROCRYPT '88. LNCS, vol. 330, pp. 123–128. Springer (1988)
30. Katz, J., Ostrovsky, R.: Round-optimal secure two-party computation. In: CRYPTO 2004. pp. 335–354 (2004)

31. Lapidot, D., Shamir, A.: Publicly verifiable non-interactive zero-knowledge proofs. In: CRYPTO (1990)
32. Lindell, Y.: An efficient transform from Sigma protocols to NIZK with a CRS and non-programmable random oracle. In: TCC 2015. pp. 93–109 (2015)
33. Lindell, Y., Pinkas, B.: An efficient protocol for secure two-party computation in the presence of malicious adversaries. *J. Cryptology* 28(2), 312–350 (2015)
34. Maurer, U.: Zero-knowledge proofs of knowledge for group homomorphisms. *Designs, Codes and Cryptography* pp. 1–14 (2015)
35. Maurer, U.M.: Unifying zero-knowledge proofs of knowledge. In: AFRICACRYPT 2009. pp. 272–286 (2009)
36. Ostrovsky, R., Pandey, O., Visconti, I.: 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. pp. 535–552 (2010)
37. Ostrovsky, R., Rao, V., Visconti, I.: 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. LNCS, vol. 8642, pp. 578–597. Springer (2014)
38. Pass, R.: Simulation in quasi-polynomial time, and its application to protocol composition. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 160–176. Springer (2003)
39. Pass, R.: Alternative Variants of Zero-Knowledge Proofs. Master’s thesis, Kungliga Tekniska Högskolan (2004), licentiate Thesis Stockholm, Sweden
40. Pointcheval, D., Stern, J.: Security proofs for signature schemes. In: EUROCRYPT ’96. pp. 387–398 (1996)
41. Scafuro, A., Visconti, I.: On round-optimal zero knowledge in the bare public-key model. In: EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings. pp. 153–171 (2012)
42. Schnorr, C.: Efficient identification and signatures for smart cards. In: Brassard, G. (ed.) CRYPTO ’89, LNCS, vol. 435, pp. 239–252. Springer New York (1989)
43. Visconti, I.: Efficient zero knowledge on the internet. In: Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Part II. pp. 22–33 (2006)
44. Yung, M., Zhao, Y.: Generic and practical resettable zero-knowledge in the bare public-key model. In: EUROCRYPT 2007. pp. 129–147 (2007)