

# Signatures from Sequential-OR Proofs

Marc Fischlin, Patrick Harasser, and Christian Janson

Cryptoplexity, Technische Universität Darmstadt, Darmstadt, Germany  
{`marc.fischlin`, `patrick.harasser`, `christian.janson`}@cryptoplexity.de

**Abstract.** OR-proofs enable a prover to show that it knows the witness for one of many statements, or that one out of many statements is true. OR-proofs are a remarkably versatile tool, used to strengthen security properties, design group and ring signature schemes, and achieve tight security. The common technique to build OR-proofs is based on an approach introduced by Cramer, Damgård, and Schoenmakers (CRYPTO'94), where the prover splits the verifier's challenge into random shares and computes proofs for each statement in parallel.

In this work we study a different, less investigated OR-proof technique, highlighted by Abe, Ohkubo, and Suzuki (ASIACRYPT'02). The difference is that the prover now computes the individual proofs sequentially. We show that such sequential OR-proofs yield signature schemes which can be proved secure in the non-programmable random oracle model. We complement this positive result with a black-box impossibility proof, showing that the same is unlikely to be the case for signatures derived from traditional OR-proofs. We finally argue that sequential-OR signature schemes can be proved secure in the quantum random oracle model, albeit with very loose bounds and by programming the random oracle.

**Keywords.** Sequential-OR Proofs · Zero-knowledge · Signatures · Non-programmable random oracle model · Quantum random oracle model

## 1 Introduction

In a zero-knowledge  $\Sigma$ -protocol between a prover  $\mathsf{P}$  and a verifier  $\mathsf{V}$ , the prover holds a statement  $x$  and a witness  $w$  for  $x$ , and the verifier only  $x$ . Both parties engage in an interactive execution, resulting in an initial commitment `com` sent by the prover, a verifier random challenge `ch`, and a final response `resp` computed by the prover. With such a proof,  $\mathsf{P}$  shows to  $\mathsf{V}$  that  $x$  is true (in proof systems), or that it knows a witness  $w$  for  $x$  (in proofs of knowledge). At the same time, the zero-knowledge property guarantees that nothing beyond this fact is revealed.

### 1.1 OR-Proofs

Now assume that one has two interactive proof systems of the above form for two statements  $x_0$  and  $x_1$ , and a witness  $w_b$  for  $x_b$ ,  $b \in \{0, 1\}$ . The goal is to combine them into a single protocol which proves the logical OR of  $x_0$  and  $x_1$ ;

$P_{\text{par-OR}}(1^\lambda; (x_0, x_1), (b, w)):$	$P_{\text{par-OR}}(1^\lambda; (x_0, x_1), (b, w), (\text{com}_0, \text{com}_1), \text{ch}):$
11: $\text{com}_b \leftarrow_{\mathcal{S}} P_b(1^\lambda; x_b, w)$	21: $\text{ch}_b \leftarrow \text{ch} \oplus \text{ch}_{1-b}$
12: $\text{ch}_{1-b} \leftarrow_{\mathcal{S}} \{0, 1\}^{\ell(\lambda)}$	22: $\text{resp}_b \leftarrow_{\mathcal{S}} P_b(1^\lambda; x_b, w, \text{com}_b, \text{ch}_b)$
13: $(\text{com}_{1-b}, \text{resp}_{1-b}, \text{ch}_{1-b}) \leftarrow_{\mathcal{S}}$ $\quad \leftarrow_{\mathcal{S}} S_{1-b}(1^\lambda; x_{1-b}, \text{ch}_{1-b})$	23: <b>return</b> $(\text{ch}_0, \text{ch}_1, \text{resp}_0, \text{resp}_1)$
14: <b>return</b> $(\text{com}_0, \text{com}_1)$	

Fig. 1: Description of the prover algorithm  $P_{\text{par-OR}}$  from the parallel-OR construction by Cramer et al. [23] in the standard model. On the left, generation of the first message  $\text{com} = (\text{com}_0, \text{com}_1)$ . On the right, computation of the final response  $\text{resp} = (\text{ch}_0, \text{ch}_1, \text{resp}_0, \text{resp}_1)$  answering the verifier challenge  $\text{ch}$ .

that is, the prover should be able to convince a verifier that it holds a witness for one of the two statements, ideally without revealing which one. The first instantiation of such general OR-proofs, sometimes called CDS-OR proofs, was given by Cramer, Damgård, and Schoenmakers [23]. Their construction works under the assumption that the two protocols are special honest-verifier zero-knowledge, meaning that a simulator  $\mathcal{S}$ , given  $x$  and a random challenge  $\text{ch}$  at the outset, is able to generate a verifier view  $(\text{com}, \text{resp}, \text{ch})$  without knowing a witness for  $x$ , in such a way that this view is indistinguishable from a genuine interaction between the real prover and an honest verifier using the given challenge. The prover in the CDS-OR protocol from [23] is described in Figure 1. For reasons that will become apparent soon, we call such CDS-OR proofs also *parallel-OR* proofs.

An important observation is that the resulting protocol is witness indistinguishable, i.e., it does not reveal for which statement the prover holds a witness. Moreover, since the resulting protocol is again a  $\Sigma$ -protocol, one can apply the Fiat-Shamir transform [32] to it and obtain a non-interactive version or a signature scheme in the random oracle model. Also, the construction easily generalizes to the case of 1-out-of- $n$  proofs.

## 1.2 Applications of OR-Proofs

OR-proofs have turned out to be a very powerful tool in the design of efficient protocols. Early on they have been identified as a means to thwart man-in-the-middle attacks [22] and, similarly in spirit, to give designated-verifier proofs [43]. The idea in both cases is to have the verifier send its public key to the prover, who then shows that the statement  $x$  it originally wanted to prove is true or that it knows the verifier's secret key. This proof is still convincing for the verifier (who knows it is the only holder of its secret key), but not transferable to other parties. Garay et al. [38] apply the same idea to make zero-knowledge proofs simulation-sound and non-malleable, by putting a verification key into a common reference

string (CRS). The prover then shows that the original statement  $x$  is true or that it knows the secret to the verification key in the CRS.

The idea of giving a valid proof when knowing a witness for only one of several statements can also be used in the context of group signatures [19] and ring signatures [56]. Given a set of public keys  $x_1, \dots, x_n$ , where the signer knows only one witness  $w_i$  (their own secret key), an OR-proof allows to sign anonymously on behalf of the entire group, and witness indistinguishability implies that the identity of the signer remains hidden. This design strategy appears explicitly for example in the group signature scheme of Camenisch [13].

The OR-technique has also proved very useful in deriving tightly-secure schemes. This approach has appeared in several works in the literature [6, 39, 42]. The idea is to first derive tightly-secure signature schemes from the OR-combination of some  $\Sigma$ -protocols. These schemes are then used within higher-level solutions (like key exchange protocols), passing on the tight security guarantees to these protocols.

### 1.3 Non-Programmable Random Oracles

Another important feature of the OR-technique is that it facilitates the design of schemes in the non-programmable random oracle model. The general random oracle model comes with several remarkable technical properties, rooted in the formalization of the hash function as a truly random, oracle-based function. One of the most extraordinary consequences of this formalization is the programmability property of the random oracle, saying that one can adaptively choose the answers to random oracle queries made by the adversary. Indeed, the ability to change answers on the fly is a necessary feature of security proofs of some signature schemes [33, 35, 37, 61]. In practice, however, hash functions are not programmable and their values are fixed. Therefore, one would ideally prefer to forgo the programming of random oracle replies.

The fact that the OR-technique can be used to bypass the programmability issues with the random oracle model can already be observed in the early constructions of  $\Sigma$ -protocols, namely, the Okamoto variant [52] of the Schnorr signature scheme [57] and the Guillou-Quisquater variant [41] of the Fiat-Shamir signature protocol [32]. In these variants, based on number-theoretic specifics, one uses “embedded” OR-proofs which allow to simulate signatures without having to program the random oracle, as opposed to [32, 57] and explicitly carried out in [55]: One can then simply use the known witness to generate signatures.

Unfortunately, the security proofs of the signature schemes in [41, 52] still need programming at another step. Namely, in order to show that the adversary cannot forge signatures, one rewinds the execution and re-programs the random oracle in order to extract a witness (a technique called forking in [55]). This also comes with a loose security bound. Abdalla et al. [1] overcome the forking technique by considering passively-secure identification schemes, where the adversary is allowed to see transcripts of honest executions. Still, they program the random oracle when simulating signatures.

Later, Abdalla et al. [2] used the notion of lossy identification schemes to give non-forking security proofs for signatures derived via the Fiat-Shamir heuristic. Lossiness here roughly means that valid statements  $x$  are indistinguishable from so-called lossy ones, for which it is statistically impossible to find convincing proofs. This idea has later been adopted by lattice-based and LWE-based signature schemes such as [7,49] (in the classical random oracle model) or the TESLA signature scheme [4] (in the quantum random oracle model [10]). Still, all approaches program the random oracle in order to be able to simulate signatures.

#### 1.4 Sequential-OR Proofs

The above construction is the classical technique to combine  $\Sigma$ -protocols and prove OR-statements, but it is not the only possible solution. Indeed, there is at least one other way to prove the disjunction of two or more statements in the random oracle model, which in its spirit already appears in a work by Rivest, Shamir, and Tauman [56]. Here, we follow the exposition given by Abe, Ohkubo, and Suzuki [3] in the context of group signature schemes, and call this approach the *sequential-OR* technique.

In this construction, the non-interactive prover computes the individual proofs sequentially, starting with the commitment  $\text{com}_b$  for the statement  $x_b$  for which it knows the witness  $w_b$ . Next it derives the challenge  $\text{ch}_{1-b}$  for the proof of  $x_{1-b}$  (with unknown witness) as the hash value of  $\text{com}_b$ . This in turn allows the OR-prover to simulate a view  $(\text{com}_{1-b}, \text{resp}_{1-b}, \text{ch}_{1-b})$  for  $x_{1-b}$  with this pre-determined challenge, as done in parallel-OR proofs. The simulated commitment  $\text{com}_{1-b}$  again yields the challenge  $\text{ch}_b$  for the first proof through the hash function, which the prover now can answer with a valid response  $\text{resp}_b$  since it knows the witness  $w_b$ . The details of the prover in the sequential-OR protocol from [3] are described in Figure 2.

Note that this technique generalizes to the 1-out-of- $n$  case (we provide all details in the full version [34]). In fact, Abe et al. [3] and follow-up works like [8,47], use this more general version of the sequential-OR technique to build group signature schemes, yet still programming the random oracle to fork and extract. The paradigm proposed by Abe et al. has also been applied in the area of cryptocurrencies, in particular Monero [58] and Mimblewimble [44, 54]. There, in order to prevent overflow attacks, it is necessary to prove that committed values fall within a specific range. One instance of such range proofs uses a special type of ring signature, called borromean ring signature [50], which is based on ideas presented in [3]. Observe that, in the aforementioned range proofs, borromean signatures have since been superseded by more efficient bulletproofs [12].

#### 1.5 Our Results

At first glance, the sequential-OR technique does not seem to give any significant advantage over the parallel version. Both protocols are based on the idea that one can easily give a proof for a statement for which the witness is known, and simulate the proof for the other statement where the challenge is known

$P_{\text{seq-OR}}(1^\lambda; (x_0, x_1), (b, w)):$ <hr style="border: 0.5px solid black;"/> 11: $\text{com}_b \leftarrow_{\S} P_b(1^\lambda; x_b, w)$ 12: $\text{ch}_{1-b} \leftarrow \mathcal{H}(b, x_0, x_1, \text{com}_b)$ 13: $(\text{com}_{1-b}, \text{resp}_{1-b}, \text{ch}_{1-b}) \leftarrow_{\S} S_{1-b}(1^\lambda; x_{1-b}, \text{ch}_{1-b})$ 14: $\text{ch}_b \leftarrow \mathcal{H}(1-b, x_0, x_1, \text{com}_{1-b})$ 15: $\text{resp}_b \leftarrow_{\S} P_b(1^\lambda; x_b, w, \text{com}_b, \text{ch}_b)$ 16: <b>return</b> $(\text{com}_0, \text{com}_1, \text{resp}_0, \text{resp}_1)$
--

Fig. 2: Description of the prover algorithm  $P_{\text{seq-OR}}$  from the sequential-OR construction by Abe et al. [3] in the random oracle model.

in advance. This, however, misses one important point if we combine these two approaches with the idea of lossy statements as in the work by Abdalla et al. [2]: We show that signatures derived from sequential-OR proofs are secure in the non-programmable random oracle model, whereas those originating from parallel-OR proofs do not seem to have a security proof in this model.

The signature scheme in the sequential-OR case is based on two valid statements  $x_0$  and  $x_1$  (the public keys), for which we know one of the two witnesses  $w_b$  (one of the secret keys). A signature for a message  $m$  is basically a sequential-OR proof, where  $m$  is included in the hash evaluations. In contrast to the proof in [3], which is based on forking, we can now reduce unforgeability to a decisional problem about the languages. This allows us to avoid rewinding and re-programming the random oracle.

The idea of our proof in the sequential-OR case can be illustrated by looking at the honest signer first. If one was able to observe the signer's random oracle queries, then their order reveals which witness the signer is using: The signer first queries the commitment  $\text{com}_b$  of the instance  $x_b$  for which it knows the witness  $w_b$ . We will use the same idea against the adversary, helping us to decide if some random input  $x_{1-b}$  is in the language or not. If  $x_{1-b}$  is not in the language, and thus does not have a witness, the special soundness of the  $\Sigma$ -protocol guarantees that the adversary will never make the first query about this part, since it will then not be able to answer the random challenge.<sup>1</sup> Hence, by merely observing the adversary's queries, we can decide membership of  $x_{1-b}$ . We use the other part  $x_b$  in the key and its witness  $w_b$  to simulate signatures without programming the random oracle. But we need to make sure that the adversary is not biased by our signatures. This follows from the witness indistinguishability of the proofs (against an adversary who cannot observe random oracle queries).

We next argue that it is in general hard to show that the parallel-OR technique of Cramer et al. [23] yields a secure signature scheme in the non-programmable random oracle model. Our result assumes a black-box reduction  $R$

<sup>1</sup> One can think of this as a very lossy mode.

transforming any (PPT or unbounded) adversary against the signature scheme into a solver of some hard problem, and makes a mild assumption about the zero-knowledge simulators of the languages (namely, that they work independently of how the statements  $x$  are generated). Remarkably, we do not make any stipulations about the reduction’s executions of the adversary instances: The reduction can run an arbitrary (bounded) number of instances of the adversary, and there are no restrictions on the inputs of these instances or their scheduling. However, the reduction  $R$  can only use the external random oracle.

Our approach is based on the meta-reduction technique [11, 40, 53]. That is, we start with an unbounded adversary  $A$ , who breaks the signature scheme easily with its super-polynomial power by computing a secret key and signing as the honest prover would. This means that the reduction  $R$  also solves the underlying problem when interacting with  $A$ . Afterwards, we show how to simulate  $A$  efficiently, resulting in an efficient algorithm solving the problem directly. This implies that there cannot exist such a reduction  $R$  in the first place.

The crucial difference between the sequential and the parallel version of the OR-technique is that in the latter case observing the random oracle queries of the adversary does *not* reveal which witness is being used. By the zero-knowledge property one cannot distinguish real and simulated sub-proofs in the parallel case. Indeed, our negative result relies exactly on this zero-knowledge property, taking advantage of the fact that the random oracle is external to the reduction.

## 1.6 Further Related Work

The issue of non-programmability of random oracles also appears in recent works related to Canetti’s universal composability (UC) framework [15]. In this model, random oracles can be cast as an ideal functionality  $\mathcal{F}_{\text{RO}}$ , and protocols can be developed in the hybrid setting where  $\mathcal{F}_{\text{RO}}$  is present. A technical consequence of this design choice is that the random oracle is programmable, and a compositional consequence is that one would need a fresh random oracle for each protocol instance. Therefore, the global random oracle model [18], based on ideas of global set-ups [16, 26], defines a random oracle functionality  $\mathcal{G}_{\text{sRO}}$  which can be used by all protocols, obliterating also the programmability of the random oracle in this model.

We stress, however, that protocols designed in the global random oracle model are not necessarily secure for non-programmable random oracles. The discrepancy lies in the distinction between the model and the security proof: In the global random oracle model, one may no longer be able to program the random oracle when devising a simulator *in the model*, but a reduction may still program the random oracle *in the security proof* showing that the simulator is good. Indeed, this can be observed in the security reductions in [14] proving that all signature schemes which have a stand-alone proof of unforgeability in the “isolated” random oracle model, including schemes with a security reduction via programming, remain secure in the strict global random oracle model  $\mathcal{G}_{\text{sRO}}$ .

The impossibility of proving the security of specific types of signatures derived via the Fiat-Shamir transform in the non-programmable random oracle

model has already been discussed in prior works, e.g., [33, 36]. These works usually make some restrictions on the reduction being ruled out (like key preservation or being single-instance), whereas we do not need any such condition. We remark here that our impossibility result for parallel-OR signatures does likely not follow in a general way from these results, since the same approach fails in the sequential-OR case.

In terms of OR-proofs, Ciampi et al. [20], based on an earlier approach by Lindell [46], *use* the OR-technique to build non-interactive zero-knowledge proofs from  $\Sigma$ -protocols in the non-programmable random oracle model. For technical reasons they also need a common reference string, which is used to form the OR-language. Note that this is orthogonal to our goal here, where we aim to *build* OR-proofs for two languages in the non-programmable random oracle model. In another work, Ciampi et al. [21] consider extensions of parallel-OR proofs where (some of) the languages are not specified yet when the execution starts. This includes the solution in the common reference string model in [20].

## 1.7 Extension to the Quantum Random Oracle Model

The results discussed so far are in the classical random oracle model. In terms of the quantum random oracle model (QROM), introduced by Boneh et al. [10], the situation regarding OR-proofs is less scrutinized. Our approach in the (classical) sequential-OR case is based on the observability of queries to the random oracle, a technique that usually does not carry over to the QROM because of superposition queries. In the parallel-OR case, we have seen that observability may not even help in the classical setting.

Fortunately, there have been two recent results regarding the security of Fiat-Shamir protocols in the QROM [27, 48], bypassing previous negative results concerning the Fiat-Shamir transform in this model [5, 24]. These works both yield a non-tight bound, but give an immediate solution for the parallel-OR case in the QROM. There, one first combines the two interactive proofs via the parallel-OR construction to get an interactive Fiat-Shamir proof, and then applies these techniques. We show in Section 6 that one can also prove security of signatures derived from the sequential-OR construction in the QROM via the measure-and-reprogram technique described in [27]. The price we pay is that we inherit the loose security bound from the solution in [27] and we, like all currently known constructions in the QROM, need to program the quantum random oracle.

## 2 Preliminaries

### 2.1 Basic Notation

We denote by  $\mathbb{N} = \mathbb{Z}_{\geq 0}$  the set of non-negative integers, and by  $\lambda \in \mathbb{N}$  the security parameter (often written in unary notation as  $1^\lambda$ ). A function  $\mu: \mathbb{N} \rightarrow \mathbb{R}$  is called *negligible* if, for every constant  $c \in \mathbb{R}_{>0}$ , there exists  $\lambda_c \in \mathbb{N}$  such that,

for every  $\lambda \in \mathbb{N}$  with  $\lambda \geq \lambda_c$ , we have  $\mu(\lambda) \leq \lambda^{-c}$ . For a random variable  $X$ , we write  $x \leftarrow_{\S} X$  to denote that  $x$  is a random variate of  $X$ . For a finite set  $S$  of size  $|S|$ , we use  $s \leftarrow_{\S} S$  as a shorthand for  $s \leftarrow_{\S} U_S$ , where  $U_S$  is a random variable uniformly distributed over  $S$ . The arrow  $\leftarrow$  will be used for assignment statements. We denote the length of a string  $x \in \{0, 1\}^*$  by  $|x|$ , and we write  $\varepsilon$  for the empty string. We consider an injective, efficiently computable, efficiently reversible, and length-increasing encoding function  $(\{0, 1\}^*)^* \rightarrow \{0, 1\}^*$ . This allows us to represent sequences of strings again as strings, and will be tacitly used throughout the paper.

In this work we use the computational model of probabilistic oracle Turing machines, also called algorithms. We assume that they are equipped with a separate security parameter tape containing the value  $1^\lambda$ . The running time of algorithms, which is intended to be bounded by the worst case, is a function of the security parameter input length only. A uniform algorithm is called *probabilistic polynomial-time (PPT)* if its running time is bounded by a polynomial, whereas a non-uniform algorithm is *PPT* if it corresponds to an infinite sequence of Turing machines, indexed by the security parameter  $\lambda$ , whose description sizes and running times are bounded by a polynomial in  $\lambda$ . Queries to the oracles always count as one operation each. For an algorithm  $A$ , we denote by  $A^{\mathcal{O}}(1^\lambda; x)$  the random variable representing the output of  $A$  when run on security parameter  $\lambda$  and input  $x \in \{0, 1\}^*$ , with access to oracles  $\mathcal{O} = (\mathcal{O}_1, \dots, \mathcal{O}_t)$ .

We use  $\perp$  as a special symbol to denote rejection or an error, and we assume that  $\perp \notin \{0, 1\}^*$ . Both inputs and outputs of algorithms can be  $\perp$ , and we convene that if any input to an algorithm is  $\perp$ , then its output is  $\perp$  as well. Double square brackets  $[\![\cdot]\!]$  enclosing boolean statements return the bit 1 if the statement is true, and 0 otherwise.

## 2.2 Random Oracle Model

Let  $\ell: \mathbb{N} \rightarrow \mathbb{N}$  be a polynomial-time computable function. For a security parameter  $\lambda \in \mathbb{N}$ , a *random oracle (RO)* [9, 17] is an oracle  $\mathcal{H}$  that implements a function randomly chosen from the space of all functions  $\{0, 1\}^* \rightarrow \{0, 1\}^{\ell(\lambda)}$ , to which all parties have access. In other words, it is an oracle that answers every query with a truly random response chosen from the range  $\{0, 1\}^{\ell(\lambda)}$ . For every repeated query the random oracle consistently returns the same output.

Constructions established and statements proved in the presence of a RO are said to hold in the *random oracle model (ROM)*. Throughout the paper, whenever a security game is set in the ROM, we assume that at the beginning of the experiment a random oracle is sampled uniformly from the aforementioned function space, and then used throughout the experiment. In this setting, it will sometimes be necessary to record queries to the random oracle  $\mathcal{H}$ , and we will do so via a set  $Q_{\mathcal{H}}$ : If  $(i, x) \in Q_{\mathcal{H}}$ , this means that the  $i$ -th query to  $\mathcal{H}$  was  $x$ .

We also define the “zero oracle” as a function  $\mathcal{Z}: \{0, 1\}^* \rightarrow \{0, 1\}^{\ell(\lambda)}$ , with  $\mathcal{Z}(x) = 0^{\ell(\lambda)}$  for all  $x \in \{0, 1\}^*$ . This allows us to state our definitions simultaneously in the standard model and in the ROM: Parties will be given access to a



generic oracle  $\mathcal{O}$ , and it is understood that  $\mathcal{O} := \mathcal{Z}$  if the definition is formulated in the standard model, and  $\mathcal{O} := \mathcal{H}$  if it is in the ROM.

The quantum analogue of the above is the so-called *quantum random oracle model* (QROM), introduced by Boneh et al. [10]. Here, a quantum algorithm may query the random oracle  $\mathcal{H}$  in superposition, i.e., submit superposition queries of the form  $\sum_x \alpha_x |x\rangle|0\rangle$  and obtain the output  $\sum_x \alpha_x |x\rangle|\mathcal{H}(x)\rangle$ . We refer to [51] for further background and conventions regarding quantum information.

### 2.3 Languages and Relations

A *language* is a subset  $L \subseteq \{0, 1\}^*$ . In this work, we assume that every language  $L$  is equipped with a uniform PPT algorithm  $G_L$  (called *instance generator*) which, on input  $(1^\lambda; b)$  with  $b \in \{0, 1\}$ , returns an element  $x \in L$  if  $b = 1$  (*yes-instance*), and an element  $x \notin L$  if  $b = 0$  (*no-instance*). Usually, the complexity of  $x$  is closely related to the security parameter  $\lambda$ , e.g.,  $|x| = \lambda$ , but we can allow for other (polynomial) dependencies as well.

A *binary relation* is a subset  $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$  which is *polynomially bounded*, i.e., there exists a polynomial  $p$  such that, for every  $(x, w) \in R$ , we have  $|w| \leq p(|x|)$ . If  $(x, w) \in R$ , we call  $x$  an *R-instance* and  $w$  an *R-witness* of  $x$ . For every  $x \in \{0, 1\}^*$ , we denote the set of all *R-witnesses* of  $x$  by  $W_R(x) := \{w \mid (x, w) \in R\}$  (if  $x$  is not an *R-instance*, then  $W_R(x) = \emptyset$ ). Note that every binary relation  $R$  defines a language  $L_R := \{x \mid \exists w : (x, w) \in R\}$ . Just like before for languages, we also assume that every binary relation  $R$  is equipped with a uniform PPT algorithm  $G_R$  (called *instance generator*) which, on input  $(1^\lambda; b)$  with  $b \in \{0, 1\}$ , returns a pair  $(x, w) \in R$  if  $b = 1$  (*yes-instance*), and an element  $x \notin L_R$  if  $b = 0$  (*no-instance*). Observe that from an instance generator  $G_R$  for a binary relation  $R$  we get an instance generator  $G_{L_R}$  for  $L_R$  by simply running  $G_R$  and returning the first component only if  $b = 1$ .

An  *$\mathcal{NP}$ -relation* is a binary relation that is *polynomial-time recognizable*, i.e.,  $R \in \mathcal{P}$ . Observe that if  $R$  is an  *$\mathcal{NP}$ -relation*, then  $L_R \in \mathcal{NP}$ , and vice-versa if  $L \in \mathcal{NP}$ , then the set  $R_L$  of all string pairs  $(x, w) \in \{0, 1\}^* \times \{0, 1\}^*$  with  $x \in L$  and  $w$  an  *$\mathcal{NP}$ -witness* for  $x$  (w.r.t. a fixed polynomial and Turing machine) is an  *$\mathcal{NP}$ -relation*. In this situation, we have of course  $L_{R_L} = L$  and  $R_{L_R} \supseteq R$ .

We next define the OR-combination of two relations and its instance generator. Here and in the following, we present all definitions and constructions with respect to the OR of two relations only, but all results extend to the more general 1-out-of- $n$  case. A *yes-instance* of the OR-relation is a pair of values  $(x_0, x_1)$ , each in its respective language, together with a witness  $w$  of either value. A *no-instance* of the OR-relation is again a pair of values, where at least one is not in the corresponding language, while the other may or may not belong to its language. The convention that a *yes-instance* has both inputs in their respective languages corresponds to the setting of group signature schemes, where all parties choose their public keys honestly; only in security reductions one may diverge from this. It is also in general necessary to ensure completeness of the OR-protocol, since the simulator for  $x_{1-b}$  is only guaranteed to output a valid transcript for *yes-instances*.

**Definition 1.** Let  $R_0$  and  $R_1$  be two binary relations. Define the OR-relation of  $R_0$  and  $R_1$  as the binary relation

$$R_{\text{OR}} := \{((x_0, x_1), (b, w)) \mid b \in \{0, 1\} \wedge (x_b, w) \in R_b \wedge x_{1-b} \in L_{R_{1-b}}\},$$

equipped with the instance generator  $G_{R_{\text{OR}}}$  defined in Figure 3. We denote the corresponding OR-language by  $L_{\text{OR}} := L_{R_{\text{OR}}}$ .

Observe that, for binary relations  $R_0$  and  $R_1$ , the relation  $R_{\text{OR}}$  is indeed a binary relation, and that  $L_{\text{OR}} = L_{R_0} \times L_{R_1}$ .

We now recall two hardness notions a binary relation  $R$  may satisfy. Intuitively,  $R$  is *decisionally hard* if no PPT distinguisher can decide if it is given an  $R$ -instance or a no-instance. It is *computationally hard* if no PPT adversary can efficiently compute an  $R$ -witness  $w$  for a given  $R$ -instance  $x$ .

**Definition 2.** Let  $R$  be a binary relation. We say that  $R$  is:

1. Decisionally Hard if, for every PPT distinguisher  $D$ , there exists a negligible function  $\mu: \mathbb{N} \rightarrow \mathbb{R}$  such that, for every  $\lambda \in \mathbb{N}$  and every  $z \in \{0, 1\}^*$ ,

$$\left| \Pr \left[ \mathbf{Exp}_{D,R}^{\text{DHR},0}(\lambda, z) = 1 \right] - \Pr \left[ \mathbf{Exp}_{D,R}^{\text{DHR},1}(\lambda, z) = 1 \right] \right| \leq \mu(\lambda),$$

where  $\mathbf{Exp}_{D,R}^{\text{DHR},0}(\lambda, z)$  and  $\mathbf{Exp}_{D,R}^{\text{DHR},1}(\lambda, z)$  are defined in Figure 3.

2. Computationally Hard if, for every PPT algorithm  $A$ , there exists a negligible function  $\mu: \mathbb{N} \rightarrow \mathbb{R}$  such that, for every  $\lambda \in \mathbb{N}$  and every  $z \in \{0, 1\}^*$ ,

$$\Pr \left[ \mathbf{Exp}_{A,R}^{\text{CHR}}(\lambda, z) = 1 \right] \leq \mu(\lambda),$$

where  $\mathbf{Exp}_{A,R}^{\text{CHR}}(\lambda, z)$  is defined in Figure 3.

It is readily verified that two binary relations  $R_0$  and  $R_1$  are computationally hard if and only if  $R_{\text{OR}}$  is computationally hard. Furthermore, if an  $\mathcal{NP}$ -relation  $R$  is decisionally hard, it is also computationally hard.

## 2.4 Interactive Protocols

An *interactive protocol*  $\Pi$  between two parties, called *prover* and *verifier*, is a pair of uniform algorithms  $\Pi = (P, V)$ . We write  $P^{\mathcal{O}}(1^\lambda; x, w) \leftrightarrow V^{\mathcal{O}}(1^\lambda; x, z)$  to denote the interaction between  $P$  and  $V$  on security parameter  $\lambda$ , common input  $x$ , respective auxiliary inputs  $w$  and  $z$ , and with access to oracle  $\mathcal{O}$ .

Algorithms  $P$  and  $V$  compute the next-message function of the corresponding party. In more detail,  $P^{\mathcal{O}}(1^\lambda; \beta_i, \text{stp}_P)$  is the random variable which returns the prover's next message  $\alpha_{i+1}$  and its updated state  $\text{stp}_P$ , both computed on input the security parameter  $\lambda$ , the last incoming message  $\beta_i$ , and the current state  $\text{stp}_P$ . Here we assume that  $\text{stp}_P$  contains all the information necessary for  $P$  to perform its computation, including at least the common input, its auxiliary input, and the messages exchanged thus far. Similar considerations hold for  $V$ .

$\mathbf{G}_{R_{\text{OR}}}(1^\lambda; b):$	$\mathbf{Exp}_{\mathcal{D}, R}^{\text{DHR}, b}(\lambda, z):$
11: <b>if</b> $b = 0$ <b>then</b> 12: $b', b'' \leftarrow_{\S} \{0, 1\}$ 13: $x_{b'} \leftarrow_{\S} \mathbf{G}_{L_{b'}}(1^\lambda; 0)$ 14: $x_{1-b'} \leftarrow_{\S} \mathbf{G}_{L_{R_{1-b'}}}(1^\lambda; b'')$ 15: <b>return</b> $(x_0, x_1)$ 16: <b>else</b> 17: $b' \leftarrow_{\S} \{0, 1\}$ 18: $(x_0, w_0) \leftarrow_{\S} \mathbf{G}_{R_0}(1^\lambda; 1)$ 19: $(x_1, w_1) \leftarrow_{\S} \mathbf{G}_{R_1}(1^\lambda; 1)$ 20: <b>return</b> $((x_0, x_1), (b', w_{b'}))$	31: $x \leftarrow_{\S} \mathbf{G}_R(1^\lambda; 0)$ 32: <b>if</b> $b = 1$ <b>then</b> 33: $(x, w) \leftarrow_{\S} \mathbf{G}_R(1^\lambda; 1)$ 34: $b' \leftarrow_{\S} \mathcal{D}^{\mathcal{O}}(1^\lambda; x, z)$ 35: <b>return</b> $b'$  $\mathbf{Exp}_{\mathcal{A}, R}^{\text{CHR}}(\lambda, z):$ 41: $(x, w) \leftarrow_{\S} \mathbf{G}_R(1^\lambda; 1)$ 42: $w^* \leftarrow_{\S} \mathcal{A}^{\mathcal{O}}(1^\lambda; x, z)$ 43: <b>return</b> $\llbracket (x, w^*) \in R \rrbracket$

Fig. 3: Definition of the instance generator  $\mathbf{G}_{R_{\text{OR}}}$  of the relation  $R_{\text{OR}}$ , and of the experiments  $\mathbf{Exp}_{\mathcal{D}, R}^{\text{DHR}, b}(\lambda, z)$  and  $\mathbf{Exp}_{\mathcal{A}, R}^{\text{CHR}}(\lambda, z)$  from Definition 2. Recall that  $\mathcal{O}$  is either a random oracle or the trivial all-zero oracle.

We write  $\text{trans}[\mathbf{P}^{\mathcal{O}}(1^\lambda; x, w) \rightleftharpoons \mathbf{V}^{\mathcal{O}}(1^\lambda; x, z)] = (A_1, B_1, \dots, A_t, B_t, A_{t+1})$  for the *transcript* of the interaction between  $\mathbf{P}$  and  $\mathbf{V}$ . This is the random variable which returns a sequence of messages  $(\alpha_1, \beta_1, \dots, \alpha_t, \beta_t, \alpha_{t+1})$ , where  $(\alpha_{i+1}, \text{stp}) \leftarrow_{\S} \mathbf{P}^{\mathcal{O}}(1^\lambda; \beta_i, \text{stp})$  and  $(\beta_j, \text{stv}) \leftarrow_{\S} \mathbf{V}^{\mathcal{O}}(1^\lambda; \alpha_j, \text{stv})$  for every  $0 \leq i \leq t$  and  $1 \leq j \leq t$ . Here we assume that  $\text{stp}$ ,  $\text{stv}$  and  $\beta_0$  are initialized to  $\text{stp} \leftarrow (x, w)$ ,  $\text{stv} \leftarrow (x, z)$  and  $\beta_0 \leftarrow \varepsilon$ . The *view* of  $\mathbf{V}$  in the interaction with  $\mathbf{P}$ , denoted  $\text{view}_{\mathbf{V}}[\mathbf{P}^{\mathcal{O}}(1^\lambda; x, w) \rightleftharpoons \mathbf{V}^{\mathcal{O}}(1^\lambda; x, z)]$ , is the random variable  $(A_1, A_2, \dots, A_t, A_{t+1}, R_{\mathbf{V}})$ , where  $R_{\mathbf{V}}$  is the random variable representing  $\mathbf{V}$ 's random coins.

The interaction between prover and verifier terminates with  $\mathbf{V}$  computing a decision  $v \leftarrow_{\S} \mathbf{V}^{\mathcal{O}}(1^\lambda; \alpha_{t+1}, \text{stv})$ , where  $v \in \{0, 1\}$ , on whether to accept or reject the transcript. This is also called  $\mathbf{V}$ 's *local output*, and the corresponding random variable will be denoted by  $\text{out}_{\mathbf{V}}[\mathbf{P}^{\mathcal{O}}(1^\lambda; x, w) \rightleftharpoons \mathbf{V}^{\mathcal{O}}(1^\lambda; x, z)]$ .

We say that a protocol  $\Pi = (\mathbf{P}, \mathbf{V})$  is *efficient* if  $\mathbf{V}$  is a PPT algorithm. For a binary relation  $R$ , we say that  $\Pi$  has an *efficient prover w.r.t.  $R$*  if  $\mathbf{P}$  is a PPT algorithm and, on security parameter  $\lambda$ , it receives common and auxiliary inputs  $x$  and  $w$  such that  $(x, w) \leftarrow_{\S} \mathbf{G}_R(1^\lambda; 1)$ . Note that we will only consider protocols which are efficient, have an efficient prover w.r.t. a specified binary relation  $R$ , and where the honest verifier is independent of its auxiliary input (we can therefore assume  $z = \varepsilon$  in this case). We call these *protocols w.r.t.  $R$* .

We call  $\Pi$  *public-coin (PC)* if all the messages the honest verifier sends to  $\mathbf{P}$  consist of disjoint segments of its random tape, and if  $\mathbf{V}$ 's local output is computed as a deterministic function of the common input and the transcript, that is  $v \leftarrow \mathbf{V}^{\mathcal{O}}(1^\lambda; x, \alpha_1, \beta_1, \dots, \alpha_t, \beta_t, \alpha_{t+1})$ . In this situation we say that a transcript is *accepting for  $x$*  if  $v = 1$ .

$\mathbf{Exp}_{V^*, D, \Pi}^{\text{CWI}, b}(\lambda, x, w, w', z, z')$ :	$\mathbf{Exp}_{D, \Pi}^{\text{SCZK}, b}(\lambda, x, w, z)$ :
11: $y \leftarrow w$	21: $(\text{ch}, \text{st}_D) \leftarrow_{\S} D_0^{\mathcal{O}}(1^\lambda; x, z)$
12: <b>if</b> $b = 1$ <b>then</b>	22: $\text{st}_P \leftarrow (x, w)$
13: $y \leftarrow w'$	23: $(\text{com}, \text{st}_P) \leftarrow_{\S} P^{\mathcal{O}}(1^\lambda; \text{st}_P)$
14: $v^* \leftarrow_{\S} \text{out}_{V^*} [P^{\mathcal{O}}(1^\lambda; x, y) \Leftrightarrow V^{*\mathcal{O}}(1^\lambda; x, z)]$	24: $(\text{resp}, \text{st}_P) \leftarrow_{\S} P^{\mathcal{O}}(1^\lambda; \text{ch}, \text{st}_P)$
15: $d \leftarrow_{\S} D^{\mathcal{O}}(1^\lambda; x, z, z', v^*)$	25: $v \leftarrow (\text{com}, \text{resp}, \text{ch})$
16: <b>return</b> $d$	26: <b>if</b> $b = 1$ <b>then</b>
	27: $v \leftarrow_{\S} S^{\mathcal{O}}(1^\lambda; x, \text{ch})$
	28: $d \leftarrow_{\S} D_1^{\mathcal{O}}(1^\lambda; x, z, v, \text{st}_D)$
	29: <b>return</b> $d$

Fig. 4: Definition of the experiments  $\mathbf{Exp}_{V^*, D, \Pi}^{\text{CWI}, b}(\lambda, x, w, w', z, z')$  and  $\mathbf{Exp}_{D, \Pi}^{\text{SCZK}, b}(\lambda, x, w, z)$  from Definitions 3 and 4.

We now recall the notion of computational witness indistinguishability [31], which is the property of general interactive protocols that is most relevant to our work. Intuitively, this notion captures the idea that protocol runs for a fixed  $R$ -instance but different witnesses should be indistinguishable. For the sake of completeness, we state the precise definitions of the completeness, soundness, honest-verifier zero-knowledge (HVCZK), and computational witness hiding (CWH) properties in the full version [34].

**Definition 3.** *Let  $R$  be a binary relation, and let  $\Pi = (P, V)$  be a protocol w.r.t.  $R$ . We say that  $\Pi$  is Computationally Witness Indistinguishable (CWI) if, for every uniform PPT algorithm  $V^*$  and every PPT distinguisher  $D$ , there exists a negligible function  $\mu: \mathbb{N} \rightarrow \mathbb{R}$  such that, for every  $\lambda \in \mathbb{N}$ , every  $x \leftarrow_{\S} G_{LR}(1^\lambda; 1)$ , every  $w, w' \in W_R(x)$ , and every  $z, z' \in \{0, 1\}^*$ ,*

$$\left| \Pr \left[ \mathbf{Exp}_{V^*, D, \Pi}^{\text{CWI}, 0}(\lambda, x, w, w', z, z') = 1 \right] - \Pr \left[ \mathbf{Exp}_{V^*, D, \Pi}^{\text{CWI}, 1}(\lambda, x, w, w', z, z') = 1 \right] \right| \leq \mu(\lambda),$$

where  $\mathbf{Exp}_{V^*, D, \Pi}^{\text{CWI}, b}(\lambda, x, w, w', z, z')$  is defined in Figure 4.

Note that we will later require a stronger version of CWI, which we term multi-query computational witness indistinguishability (mqCWI) and define formally in the full version [34]. This is basically an oracle extension of ordinary CWI, where the distinguisher can query arbitrarily many protocol executions before guessing which witness was used to generate them. One can prove via a simple hybrid argument that CWI and mqCWI are equivalent, albeit with a polynomial loss in the distinguishing advantage.

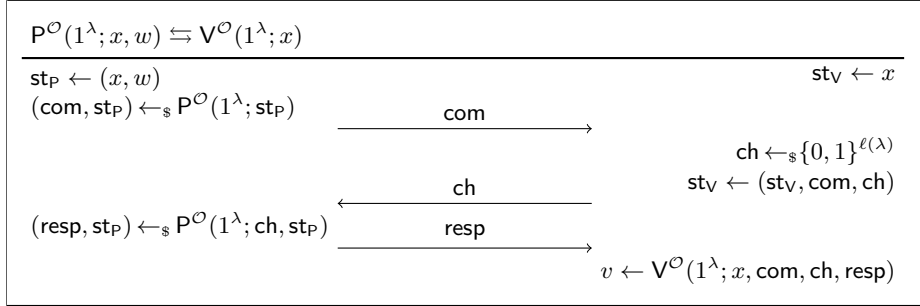


Fig. 5: Representation of a 3PC protocol w.r.t. a binary relation  $R$ .

## 2.5 3PC-Protocols and $\Sigma$ -Protocols

Let  $R$  be a binary relation. We will be mainly interested in so-called *3PC-protocols w.r.t.  $R$* , i.e., protocols w.r.t.  $R$  which are public-coin, and where the two parties exchange exactly three messages. We also assume that, on security parameter  $\lambda$ , the only message sent by the verifier to the prover has fixed length  $\ell(\lambda)$ , for a function  $\ell: \mathbb{N} \rightarrow \mathbb{N}$  called the *length function* associated to the protocol. A graphical representation of such a protocol is given in Figure 5.

In this particular context, we call the three messages exchanged between prover and verifier the *commitment*, the *challenge*, and the *response*, and denote them by  $com := \alpha_1$ ,  $ch := \beta_1$ , and  $resp := \alpha_2$ , respectively. Furthermore, we say that two accepting transcripts  $(com, ch, resp)$  and  $(com', ch', resp')$  for an element  $x$  constitute a *transcript collision for  $x$*  if  $com = com'$  and  $ch \neq ch'$ .

We now recall the critical notion of special computational zero-knowledge. Intuitively, it means that there exists a simulator which, for any maliciously chosen challenge given in advance, is able to create an authentic-looking transcript.

**Definition 4.** *Let  $R$  be a binary relation, and let  $\Pi = (P, V)$  be a 3PC protocol w.r.t.  $R$ . We say that  $\Pi$  is Special Computational Zero-Knowledge (SCZK), if there exists a uniform PPT algorithm  $S$  with the following property: For every two-stage PPT distinguisher  $D = (D_0, D_1)$ , there exists a negligible function  $\mu: \mathbb{N} \rightarrow \mathbb{R}$  such that, for every  $\lambda \in \mathbb{N}$ , every  $(x, w) \leftarrow_{\mathcal{S}} G_R(1^\lambda; 1)$ , and every  $z \in \{0, 1\}^*$ ,*

$$\left| \Pr \left[ \mathbf{Exp}_{D, \Pi}^{\text{SCZK}, 0}(\lambda, x, w, z) = 1 \right] - \Pr \left[ \mathbf{Exp}_{D, \Pi}^{\text{SCZK}, 1}(\lambda, x, w, z) = 1 \right] \right| \leq \mu(\lambda),$$

where  $\mathbf{Exp}_{D, \Pi}^{\text{SCZK}, b}(\lambda, x, w, z)$  is defined in Figure 4.

The definitions of other properties of 3PC protocols, like optimal and special soundness, are included in the full version [34]. Roughly, optimal soundness says that for every  $x \notin L$  and every commitment, there is at most one challenge which can lead to a valid response. Special soundness says that for  $x \in L$ , any



both inputs  $x_0$  and  $x_1$  with the corresponding challenge share. If the prover knows a witness  $w$  for  $x_b$ , it can use the HVCZK-simulator  $S_{1-b}$  of  $\Pi_{1-b}$  to generate a simulated view  $(\text{com}_{1-b}, \text{resp}_{1-b}, \text{ch}_{1-b})$  for  $x_{1-b}$ , and then compute a genuine transcript  $(\text{com}_b, \text{ch}_b, \text{resp}_b)$  for  $x_b$  using the witness  $w$  it knows.

Observe that the same idea works with minor changes if  $\Pi_0$  and  $\Pi_1$  are both SCZK w.r.t.  $R_0$  and  $R_1$  (instead of HVCZK). The only difference is that  $\text{P}_{\text{par-OR}}$  must now sample a random challenge  $\text{ch}_{1-b}$  before running the SCZK-simulator  $S_{1-b}$  in the first step. The main properties of  $\text{par-OR}[\Pi_0, \Pi_1, S_0, S_1]$  are summarized in the following.

**Theorem 6.** *Let  $R_0$  and  $R_1$  be binary relations, and let  $\Pi_0$  and  $\Pi_1$  be two 3PC HVCZK protocols w.r.t.  $R_0$  and  $R_1$ , such that the length functions satisfy  $\ell_0 = \ell_1 =: \ell$ . Consider the protocol  $\Pi = \text{par-OR}[\Pi_0, \Pi_1, S_0, S_1]$ . Then:*

1.  $\Pi$  is a 3PC CWI HVCZK protocol w.r.t.  $R_{\text{OR}}$ .
2. If  $\Pi_0$  and  $\Pi_1$  are complete, then  $\Pi$  is also complete.
3. If  $R_0$  and  $R_1$  are  $\mathcal{NP}$ -relations and  $R_{\text{OR}}$  is computationally hard, then  $\Pi$  is CWH.

Furthermore, if both  $\Pi_0$  and  $\Pi_1$  are SCZK, then  $\Pi$  is SCZK.

The proof of the above can be found in a slightly different syntactical version in [25], whereas the particular proof of the CWH property can be found in [59]. Note that one can build a secure signature scheme  $\text{sFS}[\Pi, \mathcal{H}]$  in the ROM from the protocol  $\Pi$  applying the Fiat-Shamir transform, which we discuss in more detail in the full version [34].

## 4 Sequential-OR Proofs

In this section, we discuss an alternative OR-proof technique which we call *sequential-OR*. This technique was first used in the context of group signature schemes by Abe et al. [3]. On a high level, in the sequential-OR variant the prover derives two sub-proofs, where data from one proof is used to derive the challenge for the other one.

### 4.1 Protocol

Similarly to Section 3, we denote by  $R_0$  and  $R_1$  two binary relations, and consider two 3PC SCZK protocols  $\Pi_0 = (\text{P}_0, \text{V}_0)$  and  $\Pi_1 = (\text{P}_1, \text{V}_1)$  w.r.t.  $R_0$  and  $R_1$  and simulators  $S_0$  and  $S_1$ , such that the two length functions  $\ell_0 = \ell_1 =: \ell$  coincide. Furthermore, let  $\mathcal{H}$  be a random oracle. The sequential-OR construction enables one to merge the two protocols  $\Pi_0$  and  $\Pi_1$  into a non-interactive protocol  $\text{seq-OR}[\Pi_0, \Pi_1, S_0, S_1, \mathcal{H}] = (\text{P}_{\text{seq-OR}}, \text{V}_{\text{seq-OR}})$  w.r.t. the binary relation  $R_{\text{OR}}$ . The formal details of the protocol are summarized in Figure 7.

The key idea of the construction is to compute the challenge for the instance the prover indeed does know the witness of, based on the commitment for which it does not know the witness (derived via the SCZK-simulator). In more detail,

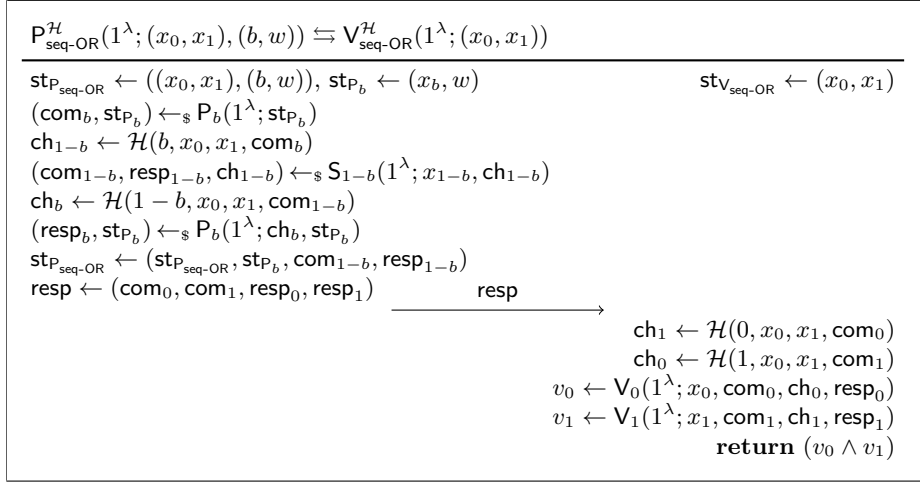


Fig. 7: Details of the sequential-OR construction by Abe et al. [3].

on input the security parameter  $\lambda \in \mathbb{N}$ , consider a yes-instance for the OR-relation  $((x_0, x_1), (b, w)) \leftarrow_{\S} G_{R_{\text{OR}}}(1^\lambda; 1)$ . The protocol  $\text{seq-OR}[\Pi_0, \Pi_1, S_0, S_1, \mathcal{H}]$  starts with the prover  $P_{\text{seq-OR}}$  and verifier  $V_{\text{seq-OR}}$  receiving  $(x_0, x_1)$  as common input. Additionally,  $P_{\text{seq-OR}}$  receives the witness  $(b, w)$  as auxiliary input. The protocol then proceeds in the following way:

1.  $P_{\text{seq-OR}}$  sets  $\text{st}_{P_b} \leftarrow (x_b, w)$  and computes  $(\text{com}_b, \text{st}_{P_b}) \leftarrow_{\S} P_b(1^\lambda; \text{st}_{P_b})$ . It then computes the challenge  $\text{ch}_{1-b}$  evaluating the random oracle  $\mathcal{H}$  on the common input  $(x_0, x_1)$  and the previously generated commitment  $\text{com}_b$ . It also includes the bit  $b$  from the witness for domain separation. Next, it runs the SCZK-simulator  $S_{1-b}$  to obtain a simulated view  $(\text{com}_{1-b}, \text{resp}_{1-b}, \text{ch}_{1-b}) \leftarrow_{\S} S_{1-b}(1^\lambda; x_{1-b}, \text{ch}_{1-b})$ . It then obtains the challenge  $\text{ch}_b$  for the first proof by evaluating  $\mathcal{H}$  on the common input  $(x_0, x_1)$ , the commitment  $\text{com}_{1-b}$  from the simulator, and the bit  $1-b$ . Finally,  $P_{\text{seq-OR}}$  computes  $(\text{resp}_b, \text{st}_{P_b}) \leftarrow_{\S} P_b(1^\lambda; \text{ch}_b, \text{st}_{P_b})$  using the witness for  $x_b$ , and sends  $\text{resp} \leftarrow (\text{com}_0, \text{com}_1, \text{resp}_0, \text{resp}_1)$  to  $V_{\text{seq-OR}}$ .
2.  $V_{\text{seq-OR}}$  first re-computes both challenge values using the random oracle  $\mathcal{H}$ . It then accepts the proof if and only if *both* transcripts verify correctly, i.e.,  $V_0(1^\lambda; x_0, \text{com}_0, \text{ch}_0, \text{resp}_0) = 1$  and  $V_1(1^\lambda; x_1, \text{com}_1, \text{ch}_1, \text{resp}_1) = 1$ .

In the following theorem, we establish the main properties of the protocol  $\text{seq-OR}[\Pi_0, \Pi_1, S_0, S_1, \mathcal{H}]$ .

**Theorem 7.** *Let  $R_0$  and  $R_1$  be binary relations, and let  $\Pi_0$  and  $\Pi_1$  be two 3PC SCZK protocols w.r.t.  $R_0$  and  $R_1$ , such that the length functions satisfy  $\ell_0 = \ell_1 =: \ell$ . Consider the protocol  $\Pi = \text{seq-OR}[\Pi_0, \Pi_1, S_0, S_1, \mathcal{H}]$ . Then the following holds in the ROM:*

1.  $\Pi$  is a 1-move CWI protocol w.r.t.  $R_{\text{OR}}$ .



KGen( $1^\lambda$ ):	Sign $^{\mathcal{H}}(1^\lambda; m, vk, sk)$ :
11: $((x_0, x_1), (b, w)) \leftarrow_{\mathcal{S}} \mathbf{G}_{R_{\text{OR}}}(1^\lambda; 1)$	21: <b>parse</b> $vk = (x_0, x_1)$
12: $vk \leftarrow (x_0, x_1)$	22: <b>parse</b> $sk = (b, w)$
13: $sk \leftarrow (b, w)$	23: $st_{P_b} \leftarrow (x_b, w)$
14: <b>return</b> $(vk, sk)$	24: $(com_b, st_{P_b}) \leftarrow_{\mathcal{S}} \mathbf{P}_b(1^\lambda; st_{P_b})$
<hr/>	25: $ch_{1-b} \leftarrow \mathcal{H}(b, vk, com_b, m)$
Verify $^{\mathcal{H}}(1^\lambda; m, \sigma, vk)$ :	26: $(com_{1-b}, resp_{1-b}, ch_{1-b}) \leftarrow_{\mathcal{S}}$
41: <b>parse</b> $\sigma = (com_0, com_1, resp_0, resp_1)$	$\leftarrow_{\mathcal{S}} \mathbf{S}_{1-b}(1^\lambda; x_{1-b}, ch_{1-b})$
42: $ch_1 \leftarrow \mathcal{H}(0, vk, com_0, m)$	27: $ch_b \leftarrow \mathcal{H}(1-b, vk, com_{1-b}, m)$
43: $ch_0 \leftarrow \mathcal{H}(1, vk, com_1, m)$	28: $(resp_b, st_{P_b}) \leftarrow \mathbf{P}_b(1^\lambda; ch_b, st_{P_b})$
44: $v_0 \leftarrow \mathbf{V}_0(1^\lambda; x_0, com_0, ch_0, resp_0)$	29: $\sigma \leftarrow (com_0, com_1, resp_0, resp_1)$
45: $v_1 \leftarrow \mathbf{V}_1(1^\lambda; x_1, com_1, ch_1, resp_1)$	30: <b>return</b> $\sigma$
46: <b>return</b> $(v_0 \wedge v_1)$	

Fig. 8: Description of the signature scheme  $\Gamma = (\text{KGen}, \text{Sign}, \text{Verify})$  obtained from the protocol  $\text{seq-OR}[\Pi_0, \Pi_1, \mathcal{S}_0, \mathcal{S}_1, \mathcal{H}]$  by appending the message  $m$  being signed to all random oracle queries.

2. If  $\Pi_0$  and  $\Pi_1$  are complete, then  $\Pi$  is also complete.
3. If  $R_0$  and  $R_1$  are  $\mathcal{NP}$ -relations and  $R_{\text{OR}}$  is computationally hard, then  $\Pi$  is CWH.

A detailed proof of Theorem 7 as well as an extension of the above technique to the more general 1-out-of- $n$  case can be found in the full version [34].

## 4.2 Sequential-OR Signatures

We now show how one can use the sequential-OR proof technique (see Figure 7) to build a secure signature scheme  $\Gamma = (\text{KGen}, \text{Sign}, \text{Verify})$  in the *non-programmable* ROM. On a high level, the signer runs a normal execution of the protocol  $\text{seq-OR}[\Pi_0, \Pi_1, \mathcal{S}_0, \mathcal{S}_1, \mathcal{H}]$ , but always includes the message  $m$  being signed when it queries the random oracle to obtain the challenges. Signatures in this scheme consist of the commitments and responses generated during the protocol execution, and verification can be achieved by re-computing the challenges (again, including the message) and checking whether the two transcripts verify. The formal details of the scheme can be found in Figure 8, and we provide a detailed description in the following.

The signature scheme's key generation algorithm runs the instance generator  $((x_0, x_1), (b, w)) \leftarrow_{\mathcal{S}} \mathbf{G}_{R_{\text{OR}}}(1^\lambda; 1)$  of the relation  $R_{\text{OR}}$ , which returns an  $R_{\text{OR}}$ -instance  $(x_0, x_1)$  and a witness  $w$  for  $x_b$ . The pair  $(x_0, x_1)$  then constitutes the public verification key, and  $(b, w)$  is set to be the secret signing key.

Signing a message  $m$  starts with running  $P_b$  on the instance  $x_b$  with the corresponding known witness (contained in the signing key), which results in a commitment  $\text{com}_b$ . The next step is to compute the challenge  $\text{ch}_{1-b}$  for the instance the prover does not know the witness for, and this is done querying the random oracle  $\mathcal{H}$ , as done before. The only difference is that now the message  $m$  is appended to the oracle's input. Next, the signer runs the SCZK-simulator of  $\Pi_{1-b}$  on the instance  $x_{1-b}$  and this challenge, generating a simulated view  $(\text{com}_{1-b}, \text{resp}_{1-b}, \text{ch}_{1-b})$ . To complete the signature, it is still necessary to derive the missing response  $\text{resp}_b$ . In order to do so, first the random oracle is invoked to output  $\text{ch}_b$  from  $\text{com}_{1-b}$  (again, the message  $m$  is appended to its argument), and on input this challenge the prover computes the response  $\text{resp}_b$ . Finally, the signature is  $(\text{com}_0, \text{com}_1, \text{resp}_0, \text{resp}_1)$ .

The verification algorithm checks whether the signature is valid for the given message. The signature is parsed in its components, and the algorithm queries the random oracle twice (including the message) to obtain the challenges  $\text{ch}_0$  and  $\text{ch}_1$ , as computed by the signing algorithm. It then verifies whether  $(\text{com}_0, \text{ch}_0, \text{resp}_0)$  and  $(\text{com}_1, \text{ch}_1, \text{resp}_1)$  are accepting transcripts for  $x_0$  and  $x_1$ , respectively. If both transcripts verify correctly then the verification algorithm accepts the signature, and rejects otherwise.

**Theorem 8.** *Let  $R_0$  and  $R_1$  be decisional hard relations, and let  $\Pi_0$  and  $\Pi_1$  be two 3PC optimally sound SCZK protocols w.r.t.  $R_0$  and  $R_1$ , such that the length functions satisfy  $\ell_0 = \ell_1 =: \ell$ . Consider the signature scheme  $\Gamma$  obtained from the protocol  $\Pi = \text{seq-OR}[\Pi_0, \Pi_1, S_0, S_1, \mathcal{H}]$  as depicted in Figure 8. Then  $\Gamma$  is an UF-CMA-secure signature scheme in the non-programmable random oracle model. More precisely, for any PPT adversary  $A$  against the UF-CMA-security of  $\Gamma$  making at most  $q_{\mathcal{H}}$  queries to the random oracle  $\mathcal{H}$ , there exist PPT algorithms  $C, V^*, D_0$  and  $D_1$  such that*

$$\begin{aligned} \text{Adv}_{A, \Gamma}^{\text{UF-CMA}}(\lambda) &\leq \text{Adv}_{V^*, C, \Pi}^{\text{mqCWI}}(\lambda) + \text{Adv}_{D_0, R_0}^{\text{DHR}}(\lambda) + \text{Adv}_{D_1, R_1}^{\text{DHR}}(\lambda) \\ &\quad + 2 \cdot (q_{\mathcal{H}}(\lambda) + 2)^2 \cdot 2^{-\ell(\lambda)}. \end{aligned}$$

In particular, for a perfectly witness indistinguishable proof system, where  $\text{Adv}_{V^*, C, \Pi}^{\text{mqCWI}}(\lambda) \leq q_s(\lambda) \cdot \text{Adv}_{V^*, C, \Pi}^{\text{CWI}}(\lambda) = 0$  (here and in the following,  $q_s$  denotes the number of queries the adversary makes to the signature oracle), the bound becomes tightly related to the underlying decisional problem. This holds for example if we have a perfect zero-knowledge simulator. We remark that our proof also works if the relations are not optimally sound but instead  $c$ -optimally sound, i.e., for every  $x \notin L_R$  and every commitment, there is a small set of at most  $c$  challenges for which a valid response can be found. In this case we get the term  $c(\lambda) \cdot 2^{-\ell(\lambda)}$  in place of  $2^{-\ell(\lambda)}$  in the above bound.

The complete proof of Theorem 8 can be found in the full version [34], but we still give a proof sketch here. We show that the obtained signature scheme  $\Gamma$  is secure via a sequence of game hops. The general approach is based on the following idea:

1. Assume that we have an adversary  $A$  which creates a forgery  $(\text{com}_0^*, \text{com}_1^*, \text{resp}_0^*, \text{resp}_1^*)$  for message  $m^*$ . We can modify  $A$  into

an adversary  $B$  which will always query both  $(0, x_0, x_1, \text{com}_0^*, m^*)$  and  $(1, x_0, x_1, \text{com}_1^*, m^*)$  to the random oracle when computing the forgery, simply by making the two additional queries if necessary.

2. Since the adversary is oblivious about which witness  $w_b$  is being used to create signatures,  $B$  will submit the query  $(1-b, x_0, x_1, \text{com}_{1-b}^*, m^*)$  first, before making any query about  $(b, x_0, x_1, \text{com}_b^*, m^*)$ , with probability roughly  $1/2$ , and will still succeed with non-negligible probability.
3. If we next replace  $x_{1-b}$  with a no-instance (which is indistinguishable for  $B$  because  $R_{1-b}$  is decisionally hard) we obtain the contradiction that  $B$ 's advantage must be negligible now, because finding a forgery when querying  $\text{com}_{1-b}^*$  first should be hard by the optimal soundness property of  $\Pi_{1-b}$ , since  $x_{1-b}$  is a no-instance.

In more detail, in the first step we transition from the classical unforgeability game  $G_0$  for the signature scheme  $\Gamma$  to a game  $G_1$  where the adversary is additionally required to query both  $(0, x_0, x_1, \text{com}_0^*, m^*)$  and  $(1, x_0, x_1, \text{com}_1^*, m^*)$  to the random oracle. It is always possible to make this simplifying assumption: Indeed, given any adversary  $A$  against the UF-CMA-security of  $\Gamma$ , we can modify it into an adversary  $B$  which works exactly like  $A$ , but whose last two operations before returning the forgery as computed by  $A$  (or aborting) are the two required oracle queries, in the order given above. It is clear that  $B$  is a PPT algorithm, that it makes at most  $q_{\mathcal{H}} + 2$  random oracle queries, and that the probabilities of adversaries  $A$  winning game  $G_0$  and  $B$  winning game  $G_1$  are the same.

We remark that it was also possible, albeit a bit lengthy, to prove that a successful adversary  $A$  against  $G_0$  would already make both oracle queries with overwhelming probability, so that one could replace this first step with a more cumbersome security proof ruling out adversaries that do not make both queries. We choose here not to do so, because it would make the proof much longer and worsen the overall bound on the advantage of  $A$ .

Next, we define a game  $G_2$  which is the same as game  $G_1$  with the change that the adversary is required to query  $(1-b, x_0, x_1, \text{com}_{1-b}^*, m^*)$  to the random oracle *before* submitting any query of the form  $(b, x_0, x_1, \text{com}_b^*, m^*)$ . By witness indistinguishability this should happen with roughly the same probability as the other case (with the opposite order), because from the adversary's perspective the signatures do not reveal which witness  $w_b$  is used by the signer. Indeed, we show that the difference between both games is (up to the factor  $\frac{1}{2}$ ) negligibly close. This is shown by building a distinguisher against a multi-query extension of the CWI property (see the full version [34] for its definition), and proving that the difference coincides with the distinguishing advantage of this distinguisher in the mqCWI experiment. As a result, the winning probability of  $B$  in game  $G_1$  is approximately twice its winning probability in game  $G_2$ .

Finally, we move to a game  $G_3$  which is identical to  $G_2$  with the difference that the  $(1-b)$ -th instance is switched to a no-instance. Since the relations are decisionally hard, we can build another distinguisher playing the DHR experiment, showing that the winning probabilities are again roughly the same.

To conclude the proof we argue that the probability of the adversary winning game  $G_3$  can be bounded using the fact that  $\Pi_{1-b}$  is optimally sound. Indeed, by the winning condition in the game, the adversary needs to provide the commitment  $\text{com}_{1-b}^*$  early on. By the fact that the  $(1-b)$ -th instance is a no-instance, we know that for every such commitment there exists at most one challenge (derived querying  $\mathcal{H}$  on  $\text{com}_b^*$  later in the game) for which there exists a response such that the transcript for  $x_{1-b}$  verifies correctly. Since the adversary must ask  $\text{com}_{1-b}^*$  in one of the random oracle queries, there are at most  $q_{\mathcal{H}} + 2$  commitments  $\text{com}_{1-b}^*$  it can check. For every such commitment it can try at most  $q_{\mathcal{H}} + 2$  other oracle queries to find the matching challenge, so that we can bound  $B$ 's winning probability in  $G_3$  by  $(q_{\mathcal{H}}(\lambda) + 2)^2 \cdot 2^{-\ell(\lambda)+1}$ .

### 4.3 Example: Post-Quantum Ring Signatures

We discuss here briefly that our sequential-OR technique can be applied to build lattice-based ring signatures. We exemplify this for the case of the Dilithium signature scheme [28]. We stress that our solution may be less efficient than optimized lattice-based constructions such as [30] (but which, again, relies on programming the random oracle and yields a loose reduction). Our aim is to demonstrate that one can use the sequential-OR approach in principle to immediately obtain a solution with security guarantees in the non-programmable classical ROM (with tight security relative to the underlying lattice problem) and also in the QROM (with loose security at this point).

We briefly recall the Dilithium signature scheme [29]. The scheme works over a ring  $R_q = \mathbb{Z}_q[X]/(X^n + 1)$ . The public key consists of (a size-reduced version of)  $t = As_1 + s_2$ , where the matrix  $A \in R_q^{k \times \ell}$  and the vectors  $s_1, s_2$  become part of the secret key. The signature  $\sigma = (z, h, c)$  of a message  $m$  consists of a short response value  $z = y + cs_1$ , where  $y$  is chosen randomly and  $c = H(\mu, w_1)$  is a (deterministically post-processed) hash value of a salted hash  $\mu$  of the message  $m$  and the commitment of  $w = Ay$  in form of its higher-order bits  $w_1$ . The value  $h$  is a hint required for verification. When generating a signature, the process may not always create a sufficiently short value  $z$ , in which case the generation is started from scratch.

The security proof of Dilithium [45] is based on the presumably hard problem to distinguish genuine public keys  $(A, As_1 + s_2)$  from  $(A, t)$  for random  $t$ . As such we have our required decisional hard relation. Optimal soundness, in the sense that for random public keys there exists at most one challenge for which one can find a valid answer for a given commitment, has been also shown to hold with overwhelming probability in [45]. The zero-knowledge property in [45] reveals, by inspecting the construction of the simulator, that the construction is special zero-knowledge with perfectly indistinguishable distribution. The witness indistinguishability of the sequential-OR protocol hence follows from Theorem 7.

We can now apply Theorem 8 to conclude that the sequential-OR version is a secure signature scheme (in the non-programmable random oracle model). Note that it is irrelevant for us how many trials the signature generation takes, since we are merely interested in the point in time when we actually observe the

right random oracle queries. With Theorem 10 we can also conclude that the protocol is secure in the quantum random oracle model.

## 5 Impossibility of parallel-OR Signatures in the Non-Programmable Random Oracle Model

In this section we show that it may be hard to prove the unforgeability of the parallel-OR signature scheme  $\Gamma = \text{sFS}[\text{par-OR}[\Pi_0, \Pi_1, \mathcal{S}_0, \mathcal{S}_1], \mathcal{H}]$  in the non-programmable ROM (all formal details about the definition of  $\Gamma$  can be found in the full version [34]). On a high level, this means that we must rule out the existence of an efficient reduction  $R$  which has access to a random oracle but is not allowed to program it, and which transforms any (bounded or unbounded) successful adversary  $A$  against the unforgeability of  $\Gamma$  into an algorithm  $C$  solving some problem  $G$  assumed to be hard with non-negligible advantage.

Our proof will proceed in two steps. First, assuming by contradiction that such a reduction  $R$  indeed does exist, we will construct a specific unbounded adversary  $A$  which breaks the unforgeability of  $\Gamma$  with overwhelming probability. By the properties of  $R$ , this means that the *unbounded* algorithm  $C$  resulting from the interaction between  $R$  and  $A$  must successfully break instances of  $G$  in the non-programmable ROM with non-negligible probability. Then, we show how to efficiently simulate to  $R$  its interaction with  $A$ , thereby yielding an *efficient* algorithm  $B$  against  $G$  in the standard model with roughly the same advantage as  $C$ . This is impossible by the hardness of  $G$ , which means that  $R$  cannot exist.

In the following paragraphs we discuss which kinds of reductions  $R$  we are able to rule out, define what types of problems  $G$  the algorithms  $B$  and  $C$  play against, and discuss a pointwise version of zero-knowledge which the base protocols must satisfy for our result to work. We then come to the main result of this section.

*Reduction.* The efficient reductions  $R$  we consider have oracle access to the random oracle  $\mathcal{H}$ , as well as a (bounded) number of adversary instances  $A_i$  which themselves have oracle access to  $\mathcal{H}$ . The latter guarantees that the reduction cannot program the random oracle for the adversarial instances, but we stress that  $R$  gets to see all the queries made by any instance  $A_i$ . We let each adversarial instance be run on the same security parameter  $\lambda$  as the reduction itself.

Recall that, in the first step of our proof, the adversary  $A$  is unbounded. Therefore, we can assume that  $A$  incorporates a truly random function which it uses if random bits are required. With this common derandomization technique, we can make some simplifying assumptions about the reduction: Without loss of generality,  $R$  runs the instances of the adversary in sequential order, starting with  $A_1$ . It also never revisits any of the previous instances  $A_1, \dots, A_i$  once it switches to the next instance  $A_{i+1}$  by inputting a verification key  $\text{vk}_{i+1}$ . Furthermore, we can disallow any resets of the adversarial instances: The reduction can simply re-run the next instance up to the desired reset point and then diverge from there on.

*Games.* The hard problems that algorithms  $B$  and  $C$  are trying to solve are non-interactive (“oracle-free”) problems, like distinguishing between different inputs. Formally, we consider games of the form  $G = (I, V, \alpha)$  consisting of an instance generation algorithm  $I$  and a verification algorithm  $V$ , where  $(\text{inst}, \text{st}) \leftarrow_{\$} I(1^\lambda)$  generates a challenge  $\text{inst}$  of the game and some state information  $\text{st}$ . On input a potential solution  $\text{sol}$  computed by some algorithm, the deterministic algorithm  $V(1^\lambda; \text{inst}, \text{sol}, \text{st})$  returns 0 or 1, depending on whether  $\text{sol}$  is a valid solution of  $\text{inst}$ . The constant  $\alpha$  allows to measure the advantage of an algorithm trying to win the game over some trivial guessing strategy (e.g.,  $\alpha = \frac{1}{2}$  for distinguishing games). We say that an algorithm  $B$  has advantage  $\epsilon$  winning the game  $G = (I, V, \alpha)$  if

$$\Pr [V(1^\lambda; \text{inst}, \text{sol}, \text{st}) = 1 : (\text{inst}, \text{st}) \leftarrow_{\$} I(1^\lambda), \text{sol} \leftarrow_{\$} B(1^\lambda; \text{inst})] \geq \alpha + \epsilon(\lambda).$$

For us, the canonical problem to reduce security of the signature scheme to would be the distinguishing game against the hard instance generator for the underlying language. However, our theorem holds more generally for other problems.

*The all-powerful Adversary.* In our setting, the reduction  $R^{\mathcal{H}, A_1^{\mathcal{H}}, A_2^{\mathcal{H}}, \dots}(1^\lambda; \text{inst})$  has black-box access to a successful adversary  $A$  against  $\Gamma$ , receives as input some instance  $\text{inst}$  of a game  $G = (I, V, \alpha)$ , and is supposed to output a valid solution  $\text{sol}$ , winning the game with non-negligible advantage  $\epsilon$ , while interacting with  $A$ . Recall that  $R$  must be able to convert any (efficient or unbounded) adversary  $A$  into a solver for  $G$ ; in particular, this must be the case for the following all-powerful forger  $A$ , which we will consider throughout the proof:

1. Upon receiving a verification key  $\text{vk} = (x_0, x_1)$  as input, the adversary first queries its signing oracle for a signature on the message  $m_{\text{vk}} = \text{vk}$ .
2. When receiving the signature  $\sigma$ , adversary  $A$  verifies the signature and aborts if this check fails.
3. Else, adversary  $A$  uses its power to compute the lexicographic first witness  $w$  of  $x_0$  (if it exists), or of  $x_1$  (if it exists, and no witness for  $x_0$  has been found). If no witness can be found, then  $A$  aborts. Otherwise, let  $b \in \{0, 1\}$  be such that  $A$  has found a witness for  $x_b$ .
4. Adversary  $A$  picks a random  $\lambda$ -bit message  $m^*$  and runs the signing algorithm with secret key  $\text{sk} = (b, w)$  to create a signature  $\sigma^*$ . This requires one random oracle query over the message  $(x_0, x_1, \text{com}_0^*, \text{com}_1^*, m^*)$ . The randomness necessary to create the signature and the message  $m^*$  is computed by applying the inner random function to  $(\text{vk}, \sigma)$ .
5. The adversary outputs  $(m^*, \sigma^*)$  as its forgery.

Note that since the adversary includes the public key  $\text{vk}$  in the messages  $m_{\text{vk}}$ , our result would also hold if the signing process itself did not include  $\text{vk}$ ; according to our specification it currently does.

We observe that  $A$  obviously wins the UF-CMA experiment of  $\Gamma$  with overwhelming probability. We denote by  $C^{\mathcal{H}}(1^\lambda; \text{inst})$  the adversary against  $G$  in the non-programmable ROM obtained by letting  $R$  interact with  $A$  (see the left-hand side of Figure 9). By the properties of  $R$ , the advantage of  $C$  against  $G$  in the non-programmable ROM must be non-negligible.

*Zero-Knowledge.* Recall that we defined the zero-knowledge property for protocols w.r.t. relations  $R$  that have an efficient instance generator. Here, we need a stronger notion: Zero-knowledge must hold pointwise for every  $(x, w) \in R$ . The reason is that we will rely on the zero-knowledge property to argue that the reduction  $R$  does not learn any useful information from the signatures created by the all-powerful adversary  $A$ . The problem here is that the reduction may choose the instance  $\text{vk}_i = (x_0, x_1)$  in the execution of the  $i$ -th adversary adaptively and in dependence of the behavior of  $A$  in previous instances. The reduction may then also base its final output on this choice.

We therefore say that a protocol  $\Pi = (\text{P}, \text{V})$  w.r.t. a relation  $R$  is *pointwise HVCZK*, if there exist a uniform PPT algorithm  $S$  and a polynomial  $p$  with the following property: For every PPT distinguisher  $D$ , there exists a negligible function  $\mu: \mathbb{N} \rightarrow \mathbb{R}$  such that, for every  $\lambda \in \mathbb{N}$ , every  $(x, w) \in R$  with  $|x|, |w| \leq p(\lambda)$ , and every  $z \in \{0, 1\}^*$ ,  $D$  can distinguish verifier views  $\text{view}_{\text{V}}[\text{P}^{\mathcal{O}}(1^\lambda; x, w) \stackrel{\text{tr}}{\leftarrow} \text{V}^{\mathcal{O}}(1^\lambda; x)]$  in the honest interaction between  $\text{P}$  and  $\text{V}$  from the simulator's output  $S(1^\lambda; x)$  with advantage at most  $\mu(\lambda)$ , even if  $D$  receives  $z$  as auxiliary input.

Note that in the definition above, the relation and the language are still fixed, only the sampling process may vary. This seems to be a reasonable assumption which applies to known protocols, as the zero-knowledge simulator is usually independent of the generation process for the statement.

*Impossibility Result.* We now show that, if there exists a black-box reduction  $R$  as described above, our all-powerful adversary  $A$  induces an efficient algorithm  $B$  winning the game directly, such that the advantages of  $B$  and  $C$  are roughly the same. This is impossible by the assumed hardness of  $G$ , so that  $R$  cannot exist.

**Theorem 9.** *Let  $R_0$  and  $R_1$  be binary relations, and let  $\Pi_0$  and  $\Pi_1$  be two 3PC optimally sound pointwise HVCZK protocols w.r.t.  $R_0$  and  $R_1$ , such that the length functions satisfy  $\ell_0 = \ell_1 =: \ell$ . Denote by  $\Pi = \text{par-OR}[\Pi_0, \Pi_1, S_0, S_1]$  the corresponding parallel-OR protocol, and let  $\Gamma = \text{sFS}[\Pi, \mathcal{H}]$  be the parallel-OR signature scheme derived from  $\Pi$  in the ROM.*

*Assume that there exists a PPT black-box reduction  $R$  from the unforgeability of  $\Gamma$  to winning a game  $G = (\text{I}, \text{V}, \alpha)$ . Then there exists a PPT algorithm  $B$  which wins the game  $G$  with non-negligible advantage in the standard model.*

The idea is as follows. Algorithm  $B$  receives as input a challenge  $\text{inst}$  of the game  $G$ , and must compute a valid solution  $\text{sol}$  with non-negligible probability. The strategy of  $B$  is to run the reduction  $R$  on  $\text{inst}$  as a subroutine, and to efficiently simulate to  $R$  its interaction with  $A$ . To do so,  $B$  must be able to answer the two types of queries that  $R$  can make: Random oracle evaluations and forgery queries to  $A$ . The former are handled via lazy sampling, i.e.,  $B$  simulates a random oracle to  $R$ . If on the other hand  $R$  requests a forgery for a verification key  $\text{vk} = (x_0, x_1)$ ,  $B$  at first follows the definition of  $A$  and requests a signature for  $m_{\text{vk}}$ . This initial signature request ensures that the verification key  $\text{vk}$  must be such that  $x_0 \in L_0$  or  $x_1 \in L_1$  or both. Indeed, the reduction cannot program

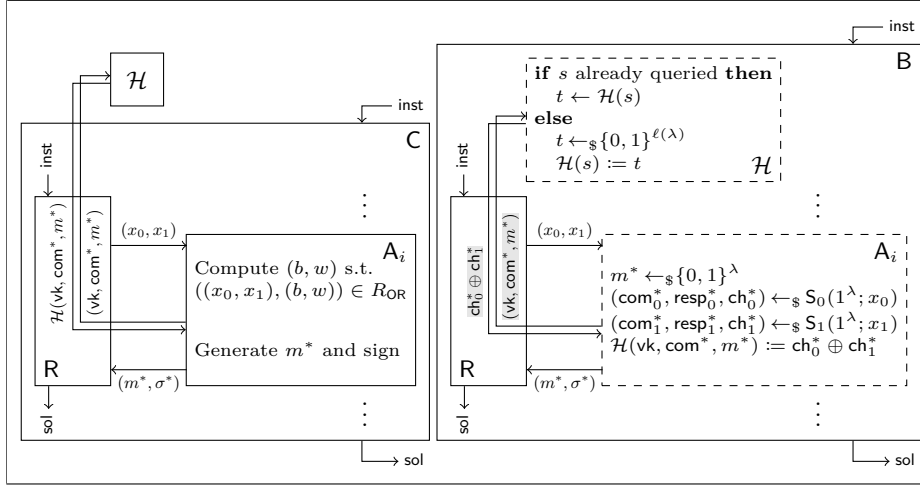


Fig. 9: Representation of the reduction R interacting with adversarial instances  $A_i$  in the ROM (left) and of the efficient solver B running R (right). The components simulated by B are dashed, and the queries of which R gets informed are highlighted in gray.

the random oracle (which is controlled by B) and, by special soundness of  $\Pi_0$  and  $\Pi_1$ , finding a valid signature when both  $x_0 \notin L_0$  and  $x_1 \notin L_1$  is infeasible for parallel-OR signatures. Hence, in the original experiment A will always be able to find a witness  $(b, w)$  for  $vk$  if it receives a valid signature.

Next, A will compute a forgery for the message  $m^*$ . Here B, instead of using  $w$  from the witness  $(b, w)$  to run  $P_b$  and compute  $com_b^*$  and  $resp_b^*$  in its forgery, uses the zero-knowledge simulator  $S_b$  for this part as well. Now both parts of the signature of  $m^*$  are independent of the actual witness. The algorithm B can now program the random oracle  $\mathcal{H}$  it is simulating to R, so that  $\mathcal{H}(vk, com^*, m^*)$  matches the XOR of the two challenges obtained from the two simulators.<sup>2</sup> By the strong zero-knowledge property of the base protocols, and since  $m^*$  contains sufficient randomness to make sure that we can still set the random oracle for R at this point, this is indistinguishable for the reduction. Finally, if at some point R returns a solution to the given instance, algorithm B terminates with the same output. In conclusion, we can now efficiently simulate A's behavior to R, so that the reduction together with this simulation technique yields our efficient algorithm B against game G (see the right-hand side of Figure 9).

<sup>2</sup> One could indeed argue why we are here allowed to program the random oracle in light of the discussion about non-programmability. One may think of this here as a restriction of the reduction, that it needs to be able to cope with such external oracles. Technically, it gives the required advantage over the reduction to make the meta-reduction argument work.



Let us stress that the impossibility result above does not hold for sequential-OR signatures. The difference lies in the observability of the reduction in both cases. In the parallel-OR case we still need to tell R which query  $\mathcal{H}(\text{vk}, \text{com}_0^*, \text{com}_1^*, m^*)$  the adversary has made to compute the forgery. But we have already argued that the simulated value  $\text{com}_b^*$  is indistinguishable from the prover's value  $\text{com}_b^*$  in the forgery, so that this query does not give any additional information to R. In the sequential-OR case, however, we would need to inform R which query A makes first, revealing which witness it has computed.

*Proof.* Consider an efficient reduction R interacting with instances of our all-powerful adversary A. Assume that the reduction calls at most  $q_A$  instances of A and makes at most  $q_{\mathcal{H}}$  calls to the random oracle. Since R is polynomial-time, both parameters are polynomially bounded. We can also assume that R never runs an instance for the same key vk and then the same signature  $\sigma$  twice, because it will just receive the same answers as before.

We start by making some simplifying assumptions about the reduction. First, we can assume that R only provides A with a valid signature to some verification key  $\text{vk} = (x_0, x_1)$  if  $x_0 \in L_0$  or  $x_1 \in L_1$  (or both). Indeed, since  $\Pi_0$  and  $\Pi_1$  are optimally sound, if both values are not in their language, then each commitment  $\text{com}_0$  for  $x_0$  and  $\text{com}_1$  for  $x_1$  only allows for at most one challenge,  $\text{ch}_0$  and  $\text{ch}_1$ , to have a valid response. But then, the probability that a random oracle query  $\mathcal{H}(\text{vk}, \text{com}_0, \text{com}_1, m_{\text{vk}})$  matches the unique value  $\text{ch}_0 \oplus \text{ch}_1$  is at most  $2^{-\ell(\lambda)}$ . The probability that such a random oracle query exists at all, either made by R or, if not, later made by any instance of the adversary A when verifying the signature, is therefore at most  $(q_{\mathcal{H}}(\lambda) + q_A(\lambda)) \cdot 2^{-\ell(\lambda)}$ . Given that A aborts if the signature it receives is not valid, we can from now on assume that each public key vk for which R requests a forgery (and must provide a signature) allows A to compute a witness  $(b, w)$ , and that R itself leaves the instance immediately if verification fails.

Second, we may assume that, whenever A creates a forgery for  $m^*$ , the random oracle has not been queried by any party yet about any value terminating in  $m^*$ . Indeed, since A applies the internal random function to compute  $m^*$  from vk and  $\sigma$ , and we assume that the reduction never runs the adversary twice on the same values, this can only happen if two random messages  $m^*$  of the adversary collide, or if the reduction has made such a query by chance. The probability for this is at most  $(q_{\mathcal{H}}(\lambda) + q_A(\lambda))^2 \cdot 2^{-\lambda}$ . Hence, we can from now on assume that this does not happen. In other words, if R stumbles upon such a value it immediately aborts.

We now define the algorithm B as explained in the overview above. On input  $(1^\lambda; \text{inst})$ , B runs the reduction on security parameter  $1^\lambda$  and instance inst as a subroutine, and simulates to R its interaction with A. The random oracle queries made by R are answered via lazy sampling. If on the other hand R calls an adversarial instance for a forgery under  $\text{vk} = (x_0, x_1)$ , B does the following:

1. It first requests a signature of  $m_{\text{vk}} = \text{vk}$  under vk to its signature oracle (provided by the reduction), and checks if the corresponding signature is valid. If not, it aborts the simulation of the current instance of A.

2. Assuming that  $R$  has provided a valid signature of  $m_{vk}$  under  $vk$ ,  $B$  does not compute a witness  $(b, w)$  for  $vk$  (as  $A$  would do). It still picks a random message  $m^* \in \{0, 1\}^\lambda$  and fresh coins for the signing process, though.
3. To compute the forgery for  $m^*$ , instead of invoking  $P_b(1^\lambda; x_b, w)$  to generate  $\text{com}_b$ ,  $B$  now runs the two simulators  $S_0(1^\lambda; x_0)$  and  $S_1(1^\lambda; x_1)$  to compute simulated views  $(\text{com}_0^*, \text{resp}_0^*, \text{ch}_0^*)$  and  $(\text{com}_1^*, \text{resp}_1^*, \text{ch}_1^*)$ .
4. Algorithm  $B$  saves  $\mathcal{H}(vk, \text{com}_0^*, \text{com}_1^*, m^*) := \text{ch}_0^* \oplus \text{ch}_1^*$  into the lookup table it keeps to simulate the random oracle to  $R$ , and informs  $R$  that the adversary  $A$  it is simulating has made a query  $(vk, \text{com}_0^*, \text{com}_1^*, m^*)$  to the random oracle, with answer  $\text{ch}_0^* \oplus \text{ch}_1^*$ .
5. Finally,  $B$  sends  $m^*$  and  $\sigma^* = (\text{com}_0^*, \text{com}_1^*, \text{resp}_0^*, \text{resp}_1^*)$  to  $R$  as the forgery computed by the simulated instance of  $A$ .

Note that  $B$  is now efficient: The only potentially exponential step involving the witness search has been eliminated. We must now argue that  $B$ 's success probability in the standard model is close to the one of  $C$  in the ROM. This is done by carrying out a reduction to the pointwise zero-knowledge property of the protocols  $\Pi_0$  and  $\Pi_1$ , where zero-knowledge must hold for every  $(x, w) \in R$ , even in the presence of some auxiliary information  $z \in \{0, 1\}^*$  that may contain further information about  $(x, w)$ . The proof is done via a hybrid argument for hybrids  $\text{Hyb}_0, \dots, \text{Hyb}_{q_A}$ , where  $\text{Hyb}_i$  answers  $R$ 's forgery requests by running the (unbounded) algorithm  $A$  up to, and including, the  $i$ -th adversarial instance (as  $C$  would do), and then efficiently simulates  $A$  for the remaining instances (as  $B$  would do). Then the extreme hybrid  $\text{Hyb}_{q_A}$  corresponds to the original inefficient algorithm  $C$ , whereas the extreme hybrid  $\text{Hyb}_0$  coincides with  $B$ 's simulation.

The jump from hybrid  $\text{Hyb}_{i-1}$  to hybrid  $\text{Hyb}_i$  substitutes the honestly generated proof for  $x_b$  (where  $x_b$  is the instance that  $A$  finds a witness for) in the  $i$ -th adversarial instance with a simulated one, so that we can construct a reduction to the pointwise HVCZK property of  $\Pi_b$ . The main idea is to let the reduction interact with the inefficient forger  $A$  for the first  $i$  instances, up to the point where  $A_i$  has determined the witness  $(b, w)$  for  $x_b$ , and save all the state information into the auxiliary input  $z$ . This allows us to pick up the reduction later. We then leverage the pointwise HVCZK property of  $\Pi_b$ , with instance  $(x_b, w) \in R_b$ : The zero-knowledge distinguisher  $D_b$  receives a genuine or simulated view for  $x_b$  and the state information  $z$ , and continues to run the reduction, but now using  $B$ 's simulation for the remaining instances (so that  $D_b$  is efficient).

More formally, we use the pointwise HVCZK property of  $\Pi_b$  for the distinguisher  $D_b$ , the instance  $(x_b, w) \in R_b$ , and the auxiliary information  $z$  defined as follows. We let  $(\text{inst}, \text{sol}) \leftarrow_{\mathfrak{s}} \mathcal{I}(1^\lambda)$  generate an instance of  $\mathcal{G}$ , pick a random tape  $r$  for the reduction and a random index  $i$  between 1 and  $q_A$  for the jump in the hybrids, and then run the reduction (interacting with  $A$ ) on input  $\text{inst}$ , up to the point where  $A$  has computed a witness for one of the two instances in the  $i$ -th execution (on input  $vk = (x_0, x_1)$ ) and has generated the message  $m^*$ . All random oracle queries are answered via lazy sampling and a table  $H$  is maintained to record previously answered queries. Let  $S$  store all forgery attempts of  $A$ . Then we let  $(x_b, w) \in R_b$  be the instance and the corresponding witness

found by  $A$ , and we set  $z = (\text{inst}, \text{st}, r, i, x_{1-b}, b, w, m^*, H, S)$ . Note that if no witness can be found by  $A$ , or if  $A$  has stopped in this instance prematurely, then we simply set  $x_b$  and  $w$  to some fixed elements of the relation  $R_0$  and the output  $z$  as before. In any case,  $z$  is of polynomial size and can be processed by an efficient distinguisher, because  $q_{\mathcal{H}}$  and  $q_A$  are polynomially bounded.

The (efficient) distinguisher  $D_b$  against the pointwise HVCZK property of  $\Pi_b$  receives  $x_b$ , a real or simulated view  $(\text{com}_b^*, \text{resp}_b^*, \text{ch}_b^*)$  for  $x_b$ , and the auxiliary information  $z = (\text{inst}, \text{st}, r, i, x_{1-b}, b, w, m^*, H, S)$ . With these data  $D_b$  can re-run the reduction up to the interaction of  $R$  with the  $i$ -th adversarial instance and then inject the given transcript  $(\text{com}_b^*, \text{ch}_b^*, \text{resp}_b^*)$  into this instance (the transcript for  $x_{1-b}$  needed to complete the forgery is obtained via the simulator  $S_{1-b}(1^\lambda; x_{1-b})$ ). Algorithm  $D_b$  now completes the execution of the reduction, using lazy sampling and the table  $H$  to continue the consistent simulation of random oracle queries. In particular, in all subsequent signature forgeries it will use  $B$ 's efficient simulation technique, calling the simulators  $S_0$  and  $S_1$  to create the two transcripts and programming the random oracle accordingly. Note that the order of execution of these two simulators is irrelevant, because  $D_b$  only needs to inform the reduction about a single random oracle query. Finally,  $D_b$  takes the reduction's output  $\text{sol}$  and returns the decision bit  $V(1^\lambda; \text{inst}, \text{sol}, \text{st})$ .

Observe that  $D_b$  runs in polynomial time, because it does not need to invoke any super-polynomial subroutines like  $A$ . If  $D_b$  receives a real view  $(\text{com}_b^*, \text{resp}_b^*, \text{ch}_b^*)$  in the  $i$ -th instance, then  $\text{ch}_b^*$  is truly random and independent, and therefore programming the (simulated) random oracle to  $\mathcal{H}(\text{vk}, \text{com}_0^*, \text{com}_1^*, m^*) := \text{ch}_0^* \oplus \text{ch}_1^*$  is perfectly sound. Hence, for real transcripts  $D_b$  simulates the hybrid  $\text{Hyb}_i$  with the first  $i$  instances according to  $C$ 's strategy, and the following instances with the simulated mode of  $B$ .

If on the other hand the transcript is simulated by  $S_b$ , then both parts of the signature are simulated. This means that both  $\text{ch}_0^*$  and  $\text{ch}_1^*$  are indistinguishable from random strings to an efficient adversary, which again implies that programming  $\mathcal{H}(\text{vk}, \text{com}_0^*, \text{com}_1^*, m^*) := \text{ch}_0^* \oplus \text{ch}_1^*$  is sound for  $R$ . In this case, only the first  $i - 1$  instances follow  $C$ 's method; starting from the  $i$ -th adversarial instance we have two simulated proofs, each simulated individually. Hence, this corresponds to the  $(i - 1)$ -th hybrid  $\text{Hyb}_{i-1}$ .

Let  $\mu_b: \mathbb{N} \rightarrow \mathbb{R}$  be the negligible function bounding the distinguishing advantage of  $D_b$  in the pointwise HVCZK experiment of  $\Pi_b$ . It follows via a standard hybrid argument that any change in the reduction's behavior translates into a distinguisher against the pointwise HVCZK property of  $\Pi_0$  and  $\Pi_1$  (times the number of queries  $q_A$ ). The advantage of our algorithm  $B$  in breaking the game is thus at least

$$\epsilon(\lambda) - (q_{\mathcal{H}}(\lambda) + q_A(\lambda))^2 \cdot 2^{-\lambda} - (q_{\mathcal{H}}(\lambda) + q_A(\lambda)) \cdot 2^{-\ell(\lambda)} - q_A(\lambda)(\mu_0(\lambda) + \mu_1(\lambda)),$$

where  $\epsilon$  is the advantage of  $C$ . Since  $\epsilon$  is non-negligible by assumption, so must be  $B$ 's advantage. But this contradicts the presumed hardness of  $G$ .  $\square$

## 6 Security in the Quantum Random Oracle Model

In this section we give an outline of the security proof for signatures derived from the sequential-OR construction in the QROM. More details can be found in the full version [34].

While treating quantum random oracles is a clear qualitative extension in terms of the security guarantees (especially if we work with quantum-resistant primitives), we have to sacrifice two important features of our proof in the classical case. One is that the bound we obtain is rather loose. The other point is that we need to program the random oracle in the security reduction. Both properties are currently shared by all proofs in the quantum random oracle model, e.g., programmability appears in form of using pairwise independent hash functions or semi-constant distributions (see [60]). Hopefully, progress in this direction will also carry over to the case of sequential-OR signatures.

Our starting point is the “measure-and-reprogram” technique of Don et al. [27] for Fiat-Shamir protocols in the QROM. They show that it is possible to turn a quantum adversary  $A$  into an algorithm  $R^A$  such that  $R^A$  measures one of the  $q_{\mathcal{H}}$  quantum queries of  $A$  to the random oracle, yielding some classical query  $\text{com}'$ . The choice of this query is made at random. Algorithm  $R^A$  returns either correctly  $\mathcal{H}(\text{com}')$  or an independent and random value  $\Theta$  to this now classical query, the choice being made at random. Algorithm  $R^A$  continues the execution of  $A$  but always returns  $\Theta$  for  $\text{com}'$  from then on. Algorithm  $R^A$  eventually returns the output  $(\text{com}, \text{resp})$  of  $A$ .

Don et al. [27] now show that, for any quantum adversary  $A$  making at most  $q_{\mathcal{H}}$  quantum random oracle queries, there exists a (quantum) algorithm  $R^A$  such that, for every fixed  $\text{com}_0$  and every predicate  $A$ , there exists a negligible function  $\mu_{\text{com}_0} : \mathbb{N} \rightarrow \mathbb{R}$  such that

$$\begin{aligned} & \Pr[\text{com} = \text{com}_0 \wedge A(1^\lambda; \text{com}, \Theta, \text{resp}) : (\text{com}, \text{resp}) \leftarrow_{\mathfrak{s}} R^{A, \mathcal{H}}(1^\lambda; \Theta)] \\ & \geq \frac{1}{O(q_{\mathcal{H}}(\lambda)^2)} \cdot \Pr\left[A(1^\lambda; \text{com}, \mathcal{H}(\text{com}), \text{resp}) : (\text{com}, \text{resp}) \leftarrow_{\mathfrak{s}} A^{\mathcal{H}}(1^\lambda)\right] - \mu_{\text{com}_0}(\lambda), \end{aligned}$$

where  $\sum_{\text{com}_0} \mu_{\text{com}_0}(\lambda) = \frac{1}{q_{\mathcal{H}}(\lambda) \cdot 2^{\ell(\lambda)+1}}$  for the output size  $\ell$  of the random oracle.

We will apply the above measure-and-reprogram technique twice in order to capture the two (classical) queries in which the adversary asks for the two commitments  $\text{com}_0^*$  and  $\text{com}_1^*$  for the forgery. However, we do not know if the strategy can be safely applied multiple times in general. Fortunately, we can apply the technique in our setting once without actually reprogramming the random oracle, only turning one of the queries into a classical one, and then view this as a special adversary  $B$  which still works with the given random oracle model. In doing so we lose a factor of approximately  $1/q^2$  in the success probability, where  $q(\lambda) = q_{\mathcal{H}}(\lambda) + 2 + 2q_s(\lambda)$  counts the number of hash queries made by both the adversary and the signature scheme. Then we can apply the technique once more to  $B$ , losing another factor  $1/q^2$ . Finally, we need to take into account that we actually obtain the matching commitments in the two measured

queries, costing us another factor  $1/q$ . Eventually, we get an algorithm  $R$  which makes two classical queries about the two commitments in the forgery with high probability, but with a loose factor of  $1/q^5$  compared to the original success probability of the forger.

Note that we now have a forger making two classical queries about the commitments  $\text{com}_{a^*}^*$  and  $\text{com}_{1-a^*}^*$  in the forgery in this order, but where we reprogram the random oracle reply in the second query about  $\text{com}_{1-a^*}^*$  to  $\Theta$ . In our sequential-OR construction this value  $\Theta$  describes the (now reprogrammed) challenge for the first commitment. In particular, the forgery then satisfies  $V_{a^*}(1^\lambda; x_{a^*}, \text{com}_{a^*}^*, \Theta, \text{resp}_{a^*}^*) = 1$  for the commitment  $\text{com}_{a^*}^*$  chosen *before*  $\Theta$  is determined. If  $x_{a^*}$  was a no-instance, this should be infeasible by the optimal soundness property. The last step in the argument is then similar to the classical setting, showing that if  $R$  is forced to use the “wrong order” and queries about a no-instance first with sufficiently high probability, its success probability will be small by the witness indistinguishability of the protocol and the decisional hardness of the problems (but this time against quantum algorithms).

Overall, we get:

**Theorem 10.** *Let  $R_0$  and  $R_1$  be decisional hard relations against quantum algorithms, and let  $\Pi_0$  and  $\Pi_1$  be two 3PC optimally sound SCZK protocols w.r.t.  $R_0$  and  $R_1$ , where zero-knowledge holds with respect to quantum distinguishers, such that the length functions satisfy  $\ell_0 = \ell_1 =: \ell$ . Consider the signature scheme  $\Gamma$  obtained from the protocol  $\Pi = \text{seq-OR}[\Pi_0, \Pi_1, S_0, S_1, \mathcal{H}]$  as depicted in Figure 8. Then  $\Gamma$  is an UF-CMA-secure signature scheme in the quantum random oracle model. More precisely, for any polynomial-time quantum adversary  $A$  against the UF-CMA-security of  $\Gamma$  making at most  $q_{\mathcal{H}}$  quantum queries to the random oracle  $\mathcal{H}$  and at most  $q_s$  signature queries, there exist a negligible function  $\mu: \mathbb{N} \rightarrow \mathbb{R}$  and polynomial-time quantum algorithms  $C, V^*, D_0$  and  $D_1$  such that*

$$\begin{aligned} \text{Adv}_{A,\Gamma}^{\text{UF-CMA}}(\lambda) \leq & O((q_{\mathcal{H}}(\lambda) + q_s(\lambda) + 2)^5) \cdot \left( \text{Adv}_{V^*,C,\Pi}^{\text{mqCWI}}(\lambda) + \text{Adv}_{D_0,R_0}^{\text{DHR}}(\lambda) \right. \\ & \left. + \text{Adv}_{D_1,R_1}^{\text{DHR}}(\lambda) + 2^{-\ell(\lambda)+1} \right) + \mu(\lambda). \end{aligned}$$

## Acknowledgments

We thank the anonymous reviewers for valuable comments. We thank Serge Fehr and Tommaso Gagliardini for helpful discussions. This work was funded by the Deutsche Forschungsgemeinschaft (DFG) – SFB 1119 – 236615297.

## References

1. M. Abdalla, J. H. An, M. Bellare, and C. Namprempre. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In *EUROCRYPT 2002*, pages 418–433, 2002.

2. M. Abdalla, P.-A. Fouque, V. Lyubashevsky, and M. Tibouchi. Tightly secure signatures from lossy identification schemes. *Journal of Cryptology*, 29(3):597–631, 2016.
3. M. Abe, M. Ohkubo, and K. Suzuki. 1-out-of-n signatures from a variety of keys. In *ASIACRYPT 2002*, pages 415–432, 2002.
4. E. Alkim, N. Bindel, J. A. Buchmann, Ö. Dagdelen, E. Eaton, G. Gutoski, J. Krämer, and F. Pawlega. Revisiting TESLA in the quantum random oracle model. In *PQCrypto 2017*, pages 143–162, 2017.
5. A. Ambainis, A. Rosmanis, and D. Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *55th FOCS*, pages 474–483, 2014.
6. C. Bader, D. Hofheinz, T. Jager, E. Kiltz, and Y. Li. Tightly-secure authenticated key exchange. In *TCC 2015, Part I*, pages 629–658, 2015.
7. S. Bai and S. D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA 2014*, pages 28–47, 2014.
8. C. Baum, H. Lin, and S. Oechsner. Towards practical lattice-based one-time linkable ring signatures. In *ICICS 18*, pages 303–322, 2018.
9. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS 93*, pages 62–73, 1993.
10. D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry. Random oracles in a quantum world. In *ASIACRYPT 2011*, pages 41–69, 2011.
11. D. Boneh and R. Venkatesan. Breaking RSA may not be equivalent to factoring. In *EUROCRYPT’98*, pages 59–71, 1998.
12. B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334, 2018.
13. J. Camenisch. Efficient and generalized group signatures. In *EUROCRYPT’97*, pages 465–479, 1997.
14. J. Camenisch, M. Drijvers, T. Gagliardoni, A. Lehmann, and G. Neven. The wonderful world of global random oracles. In *EUROCRYPT 2018, Part I*, pages 280–312, 2018.
15. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145, 2001.
16. R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally composable security with global setup. In *TCC 2007*, pages 61–85, 2007.
17. R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisited (preliminary version). In *30th ACM STOC*, pages 209–218, 1998.
18. R. Canetti, A. Jain, and A. Scafuro. Practical UC security with a global random oracle. In *ACM CCS 2014*, pages 597–608, 2014.
19. D. Chaum and E. van Heyst. Group signatures. In *EUROCRYPT’91*, pages 257–265, 1991.
20. M. Ciampi, G. Persiano, A. Scafuro, L. Siniscalchi, and I. Visconti. Improved OR-composition of sigma-protocols. In *TCC 2016-A, Part II*, pages 112–141, 2016.
21. M. Ciampi, G. Persiano, A. Scafuro, L. Siniscalchi, and I. Visconti. Online/offline OR composition of sigma protocols. In *EUROCRYPT 2016, Part II*, pages 63–92, 2016.
22. R. Cramer and I. Damgård. Fast and secure immunization against adaptive man-in-the-middle impersonation. In *EUROCRYPT’97*, pages 75–87, 1997.
23. R. Cramer, I. Damgård, and B. Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *CRYPTO’94*, pages 174–187, 1994.

24. Ö. Dagdelen, M. Fischlin, and T. Gagliardoni. The Fiat-Shamir transformation in a quantum world. In *ASIACRYPT 2013, Part II*, pages 62–81, 2013.
25. I. Damgård. On  $\Sigma$ -protocols. *Lecture Notes, University of Aarhus, Department for Computer Science*, 2002.
26. Y. Dodis, V. Shoup, and S. Walfish. Efficient constructions of composable commitments and zero-knowledge proofs. In *CRYPTO 2008*, pages 515–535, 2008.
27. J. Don, S. Fehr, C. Majenz, and C. Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In *CRYPTO 2019, Part II*, pages 356–383, 2019.
28. L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS-Dilithium: A lattice-based digital signature scheme. *IACR TCHES*, 2018(1):238–268, 2018. <https://tches.iacr.org/index.php/TCHES/article/view/839>.
29. L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehle. Crystals-dilithium: Algorithm specifications and supporting documentation, 2019. <https://pq-crystals.org/dilithium/index.shtml>.
30. M. F. Esgin, R. Steinfeld, J. K. Liu, and D. Liu. Lattice-based zero-knowledge proofs: New techniques for shorter and faster constructions and applications. In *CRYPTO 2019, Part I*, pages 115–146, 2019.
31. U. Feige and A. Shamir. Witness indistinguishable and witness hiding protocols. In *22nd ACM STOC*, pages 416–426, 1990.
32. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO’86*, pages 186–194, 1987.
33. M. Fischlin and N. Fleischhacker. Limitations of the meta-reduction technique: The case of Schnorr signatures. In *EUROCRYPT 2013*, pages 444–460, 2013.
34. M. Fischlin, P. Harasser, and C. Janson. Signatures from sequential-OR proofs. *IACR Cryptology ePrint Archive*, 2020.
35. M. Fischlin, A. Lehmann, T. Ristenpart, T. Shrimpton, M. Stam, and S. Tessaro. Random oracles with(out) programmability. In *ASIACRYPT 2010*, pages 303–320, 2010.
36. M. Fukumitsu and S. Hasegawa. Impossibility on the provable security of the Fiat-Shamir-type signatures in the non-programmable random oracle model. In *ISC 2016*, pages 389–407, 2016.
37. M. Fukumitsu and S. Hasegawa. Black-box separations on fiat-shamir-type signatures in the non-programmable random oracle model. *IEICE Transactions*, 101-A(1):77–87, 2018.
38. J. A. Garay, P. D. MacKenzie, and K. Yang. Strengthening zero-knowledge protocols using signatures. In *EUROCRYPT 2003*, pages 177–194, 2003.
39. K. Gjøsteen and T. Jager. Practical and tightly-secure digital signatures and authenticated key exchange. In *CRYPTO 2018, Part II*, pages 95–125, 2018.
40. S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.*, 17(2):281–308, 1988.
41. L. C. Guillou and J.-J. Quisquater. A practical zero-knowledge protocol fitted to security microprocessor minimizing both transmission and memory. In *EUROCRYPT’88*, pages 123–128, 1988.
42. D. Hofheinz and T. Jager. Tightly secure signatures and public-key encryption. In *CRYPTO 2012*, pages 590–607, 2012.
43. M. Jakobsson, K. Sako, and R. Impagliazzo. Designated verifier proofs and their applications. In *EUROCRYPT’96*, pages 143–154, 1996.
44. T. E. Jedisur. Mumblewimble, 2016. <https://download.wpsoftware.net/bitcoin/wizardry/mumblewimble.txt>.

45. E. Kiltz, V. Lyubashevsky, and C. Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In *EUROCRYPT 2018, Part III*, pages 552–586, 2018.
46. Y. Lindell. An efficient transform from sigma protocols to NIZK with a CRS and non-programmable random oracle. In *TCC 2015, Part I*, pages 93–109, 2015.
47. J. K. Liu, V. K. Wei, and D. S. Wong. Linkable spontaneous anonymous group signature for ad hoc groups (extended abstract). In *ACISP 04*, pages 325–335, 2004.
48. Q. Liu and M. Zhandry. Revisiting post-quantum Fiat-Shamir. In *CRYPTO 2019, Part II*, pages 326–355, 2019.
49. V. Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT 2012*, pages 738–755, 2012.
50. G. Maxwell and A. Poelstra. Borromean ring signatures, 2015. <https://pdfs.semanticscholar.org/4160/470c7f6cf05ffc81a98e8fd67fb0c84836ea.pdf>.
51. M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011.
52. T. Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In *CRYPTO'92*, pages 31–53, 1993.
53. P. Paillier and D. Vergnaud. Discrete-log-based signatures may not be equivalent to discrete log. In *ASIACRYPT 2005*, pages 1–20, 2005.
54. A. Poelstra. Mumblewimble, 2016. <https://download.wpsoftware.net/bitcoin/wizardry/mumblewimble.pdf>.
55. D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000.
56. R. L. Rivest, A. Shamir, and Y. Tauman. How to leak a secret. In *ASIACRYPT 2001*, pages 552–565, 2001.
57. C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.
58. N. van Saberhagen. Cryptonote v 2.0, 2013. <https://cryptonote.org/whitepaper.pdf>.
59. D. Venturi. Zero-knowledge proofs and applications, 2015. <http://wwwusers.di.uniroma1.it/~venturi/misc/zero-knowledge.pdf>.
60. M. Zhandry. Secure identity-based encryption in the quantum random oracle model. In *CRYPTO 2012*, pages 758–775, 2012.
61. Z. Zhang, Y. Chen, S. S. M. Chow, G. Hanaoka, Z. Cao, and Y. Zhao. Black-box separations of hash-and-sign signatures in the non-programmable random oracle model. In *ProvSec 2015*, pages 435–454, 2015.