

Fiat–Shamir for Highly Sound Protocols is Instantiable

Arno Mittelbach¹ Daniele Venturi²

¹ Cryptoplexity, Technische Universität Darmstadt, Germany

² Department of Computer Science, Sapienza University of Rome, Italy
mail@arno-mittelbach.de, venturi@di.uniroma1.it

Abstract. The Fiat–Shamir (FS) transformation (Fiat and Shamir, Crypto ’86) is a popular paradigm for constructing very efficient non-interactive zero-knowledge (NIZK) arguments and signature schemes using a hash function, starting from any three-move interactive protocol satisfying certain properties. Despite its wide-spread applicability both in theory and in practice, the known positive results for proving security of the FS paradigm are in the random oracle model, i.e., they assume that the hash function is modelled as an external random function accessible to all parties. On the other hand, a sequence of negative results shows that for certain classes of interactive protocols, the FS transform cannot be instantiated in the standard model.

We initiate the study of complementary positive results, namely, studying classes of interactive protocols where the FS transform *does* have standard-model instantiations. In particular, we show that for a class of “highly sound” protocols that we define, instantiating the FS transform via a q -wise independent hash function yields NIZK arguments and secure signature schemes. In the case of NIZK, we obtain a weaker “ q -bounded” zero-knowledge flavor where the simulator works for all adversaries asking an a-priori bounded number of queries q ; in the case of signatures, we obtain the weaker notion of random-message unforgeability against q -bounded random message attacks.

Our main idea is that when the protocol is highly sound, then instead of using random-oracle programming, one can use complexity leveraging. The question is whether such highly sound protocols exist and if so, which protocols lie in this class. We answer this question in the affirmative in the common reference string (CRS) model and under strong assumptions. Namely, assuming indistinguishability obfuscation and puncturable pseudorandom functions we construct a compiler that transforms any 3-move interactive protocol with instance-independent commitments and simulators (a property satisfied by the Lapidot-Shamir protocol, Crypto ’90) into a compiled protocol in the CRS model that is highly sound. We also present a second compiler, in order to be able to start from a larger class of protocols, which only requires instance-independent commitments (a property for example satisfied by the classical protocol for quadratic residuosity due to Blum, Crypto ’81). For the second compiler we require dual-mode commitments.

We hope that our work inspires more research on classes of (efficient) 3-move protocols where Fiat–Shamir is (efficiently) instantiable.

Keywords. Fiat–Shamir transform, non-interactive zero-knowledge, signature schemes, indistinguishability obfuscation, standard model

Contents

1	Introduction	3
1.1	Fiat–Shamir NIZK and Signatures	3
1.2	Positive and Negative Results	3
1.3	Our Contributions	5
1.4	Perspective	5
1.5	Related Work and Open Questions	6
1.6	Roadmap	7
2	Technical Overview	7
3	Preliminaries	10
3.1	Notation	10
3.2	q -Wise Independent Hashing	10
3.3	Interactive and Non-Interactive Arguments	11
3.4	Obfuscation	13
3.5	Puncturable Pseudorandom Functions	14
4	Fiat–Shamir NIZK	15
4.1	The Fiat–Shamir Transform	15
4.2	A Selective Variant of Fiat–Shamir	16
4.3	The FS-Collapse	18
4.4	Putting it Together	23
4.5	Obtaining the Required Properties	23
5	Obtaining Small Soundness-Error-to-Guessing Ratio	24
5.1	The Compiler	24
5.2	Security Analysis	25
6	Obtaining Instance-Independence	31
6.1	The Compiler	32
6.2	Security Analysis	33
7	Fiat–Shamir Signatures	36
7.1	Identification and Signature Schemes	36
7.2	Proof of Random-Message Unforgeability	38
7.3	Obtaining the Required Properties	40

1 Introduction

The Fiat–Shamir (FS) transformation [FS87] is a popular¹ technique to build efficient non-interactive zero-knowledge (NIZK) arguments and signature schemes, starting from three-round *public-coin* (3PC) protocols satisfying certain properties. In a 3PC protocol the prover starts by sending a commitment α , to which the verifier replies with a challenge β drawn at random from some space \mathcal{B} ; finally the prover sends a reply γ and the verifier’s verdict is computed as a predicate of the transcript (α, β, γ) .

1.1 Fiat–Shamir NIZK and Signatures

We briefly review both the main applications of the FS transform below.

NIZK. A NIZK is a non-interactive protocol in which the prover—holding a witness w for membership of a statement x in some NP -language L —can convince the verifier—holding just x —that $x \in L$, by sending a single message π . NIZK should satisfy three properties. First, *completeness* says that an honest prover holding a valid witness (almost) always convinces an honest verifier. Second, *soundness* says that a malicious prover should not be able to convince the honest verifier into accepting a *false* statement, i.e. a statement $x \notin L$; we speak of *arguments* (resp., *proofs*) when the soundness requirement holds for all computationally bounded (resp., computationally unbounded) provers. Third, *zero-knowledge* requires that a proof does not reveal anything about the witness beyond the validity of the statement being proven.

Apart from being a fascinating topic, NIZK have been demonstrated to be extremely useful for cryptographic applications (see, e.g., [GMW87, KZZ15, EL04, CL02, CHL05, DDN91]). NIZK require a setup assumption, typically in the form of a common reference string (CRS).

Starting with a 3PC protocol, the FS transform makes it a NIZK by having the prover compute the verifier’s challenge as a hash of the commitment α via some hash function H (with “hash key” hk); this results in a single message $\pi = (\alpha, \beta, \gamma)$, where $\beta = H(hk, \alpha)$, that is sent from the prover to the verifier.² (The description of the hash function, i.e. key hk , is included as part of the CRS.)

Signatures. Digital signatures are among the most important and well-studied cryptographic tools. Signature schemes allow a signer (holding a public/secret key pair (pk, sk)) to generate a signature σ on a message m , in such a way that anyone possessing the public key pk can verify the validity of (m, σ) . Signatures must be unforgeable, meaning that it should be hard to forge a signature on a “fresh” chosen message (even after seeing polynomially many signatures on possibly chosen messages).

Starting with a 3PC protocol, the FS transform makes it a signature by having the signer compute the verifier’s challenge as a hash of the commitment α , concatenated with the message m , via some hash function H (with “hash key” hk); this results in a signature $\sigma = (\alpha, \beta, \gamma)$, where $\beta = H(hk, \alpha || m)$.

1.2 Positive and Negative Results

We refer to the non-interactive system obtained by applying the FS transform to a 3PC protocol (i.e., a NIZK or a signature scheme) as the *FS collapse*. A fundamental question in cryptography is to understand what properties the initial 3PC protocol and the hash function should satisfy in

¹There are over 3.000 Google-Scholar-known citations to [FS87], as we type.

²The value β is typically omitted from the proof, as the verifier can compute it by itself.

order for the FS collapse to be a NIZK argument or a secure signature scheme. This question has been studied extensively in the literature; we briefly review the current state of affairs below.

Positive results. All security proofs for the FS transform follow the random oracle methodology (ROM) of Bellare and Rogaway [BR93], i.e., they assume that the function H behaves like an external random function accessible to all parties (including the adversary). In particular, a series of papers [FS87, Oka93, PS00, AABN02] establishes that the FS transform yields a secure NIZK/signature scheme in the ROM provided that the starting 3PC is a passively secure identification scheme.³ Later work proved that, for the case of NIZK, the FS transform yields even more sophisticated properties, namely, in the random oracle model, it achieves simulation-soundness [FKMV12] and simulation-extractability [BPW12].

Barak *et al.* [BLV03] put forward a new hash function property (called entropy preservation⁴) that allows to prove security of the FS transformation without random oracles; their result requires that the starting 3PC protocol is statistically sound, i.e. it is a *proof*. Dodis *et al.* [DRV12] show that such hash functions exist in case a conjecture on the existence of certain “condensers for leaky sources” turns out to be true. Canetti *et al.* [CCR15] recently study the correlation intractability of obfuscated pseudorandom functions and show a close connection between entropy preservation and correlation intractability, but it remains open whether their construction achieves entropy preservation or, in fact, whether entropy-preserving hash functions exist in the standard model. A negative indication to this question was recently presented by Bitansky *et al.* [BDG⁺13] who show that entropy-preservation security cannot be proven via a black-box reduction to a *cryptographic game*.

Negative results. It is often difficult to interpret what a proof in the ROM means in the standard model. This is not only because concrete hash functions seem far from behaving like random oracles, but stems from the fact that there exist cryptographic schemes that can be proven secure in the ROM, but are always insecure in the standard model no matter how we instantiate the hash function [CGH98].

The FS transformation is not an exception in this respect. In their study of “magic functions”, Dwork *et al.* [DNRS99] establish that whenever the initial 3PC protocol satisfies the zero-knowledge property, its FS collapse can never be (computationally) sound for any implementation of the hash function. Goldwasser and Kalai [GK03], building on previous work of Barak [Bar01], construct a specially-crafted 3PC *argument* for which the FS transform is insecure for any standard model implementation of the hash function; this in particular means that the random oracle in the FS transform cannot be universally instantiated on all 3PC arguments.

Recently, Bitansky *et al.* [BGW12] and Dachman-Soled *et al.* [DJKL12] (see also [BDG⁺13]) show an unprovability result that also covers 3PC *proofs*. More in detail, [BGW12] shows that the FS transform cannot always preserve soundness when starting with a 3PC proof, under a black-box reduction to any falsifiable assumption (even ones with an inefficient challenger). [DJKL12] shows a similar black-box separation (although only for assumptions with an efficient challenger) for any concrete proof that is honest-verifier zero-knowledge against sub-exponential size distinguishers.

³The focus of [FS87, Oka93, PS00, AABN02] was on obtaining signature schemes via the FS transform, but the same results hold for non-interactive zero-knowledge.

⁴Entropy preservation roughly says that for all efficient adversaries that get a uniformly random hash key hk and produce a correlated value α , the conditional Shannon entropy of $\beta = H(hk, \alpha)$ given α , but not hk , is sufficiently large.

1.3 Our Contributions

The negative results show that, for certain classes of interactive protocols, the FS transform cannot be instantiated in the standard model. We initiate the study of complementary positive results, namely, studying classes of interactive protocols where the FS transform *does* have a standard-model instantiation. We show that for a class of “highly sound” protocols that we define, instantiating the FS transform via a q -wise independent hash function yields both a NIZK argument in the CRS model and a secure signature scheme. In the case of NIZK, we obtain a weaker “ q -bounded” zero-knowledge flavor where the simulator works for all adversaries asking an a-priori bounded number of queries q ; in the case of signatures, we obtain the weaker notion of random-message unforgeability against q -bounded random message attacks, where the forger can only observe signatures on random messages and has to produce a forgery on a fresh random message.

We are not aware of natural protocols that are directly highly sound according to our definition. (But we will later discuss that, e.g., the Lapidot-Shamir protocol [LS91] partially satisfies our requirements.) Hence, the question is whether such highly sound protocols exist and, if so, which languages and protocols lie in this class. We answer this question in the affirmative in the CRS model and under strong assumptions. Namely, assuming indistinguishability obfuscation, puncturable pseudorandom functions and equivocal commitments, we build a sequence of two compilers that transform any three-move interactive protocol with instance-independent commitments,⁵ into a compiled protocol in the CRS model that satisfies the required properties. Noteworthy, our compilers are language-independent, and we know that assuming one-way permutations three-move interactive protocols with instance-independent commitments exist for all of NP . We refer the reader to Section 1.4 for a more in-depth interpretation of our results.

Our result avoids Dwork *et al.* [DNRS99], because we start from a protocol that is honest-verifier zero-knowledge rather than fully zero-knowledge. Note that our approach also circumvents the negative result of [BGW12] as our technique applies only to a certain class of 3PC arguments. Furthermore, we circumvent the black-box impossibility result [DJKL12] by using complexity leveraging and sub-exponential security assumptions.

1.4 Perspective

The main contribution from our perspective is to initiate the study of restricted positive standard-model results for the FS transform. Namely, we show that for the class of highly sound protocols, the FS transform can be instantiated via a q -wise independent hash function (both for the case of NIZK and signatures). This is particularly interesting given the negative results in [DNRS99, GK03, BDG⁺13].

An important complementary question is, of course, to study the class of highly sound protocols. Under strong assumptions, our compilers show that highly sound protocols exist for all languages in NP . However, the compilers yield protocols in the CRS model and, at least for the case of NIZK, as we discuss now, one has to take particular care in interpreting positive results about the FS transform applied to 3PC protocols in the CRS model.

It is well known that in the CRS model one can obtain a NIZK both for NP -complete languages [BFM88] and for specific languages [GS08]. Let L be a language. Given a standard 3PC protocol for proving membership of elements $x \in L$, and with transcripts (α, β, γ) , consider the following dummy “compiler” for obtaining a 3PC protocol for L in the CRS model. The first message α^* and the second message β^* of the compiled protocol are equal to the empty string ε ;

⁵This means that the first message is computed independently of the witness (or the secret key) and of the instance being proven (or the public key).

the third message is a NIZK proof γ^* that $x \in L$. Note that the FS transform is easily seen to be secure (without random oracles) on such a dummy protocol, the reason for this being that α^* and β^* play no role at all in the obtained 3PC! Further note that this artificial “compiler” actually ignores the original protocol, and hence it does not rely on any of the security features of the underlying protocol. Regrettably, the above example does not shed any light on the security of the FS transform and when it applies.

In turn, our result for FS NIZK has two interesting features. First, our instantiation of the FS transform works even if the starting 3PC is in the standard model (provided that it satisfies **P1-P3**). Second, our CRS-based compiler is very different from the above dummy compiler in that we do not simply “throw away” the initial 3PC but instead rely on all of its properties in order to obtain a 3PC satisfying **P1-P3**.

We remark that the above limitation does not apply to our positive result for FS signatures, since assuming the initial 3PC protocol works in the CRS model does not directly yield a dummy compiler as the one discussed above.

1.5 Related Work and Open Questions

Our approach of first compiling any “standard” 3PC protocol into one with additional properties that suffice for proving security of the FS transform is similar in spirit to the approach taken by Haitner [Hai09] who shows how to transform any interactive argument into one for which parallel repetition decreases the soundness error at an exponential rate.

Lindell recently used a similar idea to first transform a 3PC into a new protocol in the CRS model, and then show that the resulting 3PC when transformed with (a slightly modified version of) Fiat–Shamir satisfies zero-knowledge in the standard model [Lin15]. We note that the use of a CRS-enhanced interactive protocol is only implicit in Lindell’s work as he directly analyzes the collapsed non-interactive version. On the downside, to prove soundness Lindell still requires (non-programmable) random oracles. We note that one of our compilers is essentially equivalent to the compiler used by Lindell. Before Lindell’s work, interactive protocols in the CRS model have also been studied by Damgård who shows how to build 3-round concurrent zero-knowledge arguments for all *NP*-problems in the CRS model [Dam00].

It is also worth mentioning that using indistinguishability obfuscation and puncturable PRFs one can directly obtain a NIZK for all *NP* as shown by Sahai and Waters [SW14]. However, our main focus is not on constructions of NIZK, rather we aim at providing a better understanding of what can be proved for the FS transform without relying on random oracles. In this respect, our result shares similarities to the standard-model instantiation of Full-Domain Hash given in [HSW14].

In the case of NIZK, an alternative version of the FS transform is defined by having the prover hashing the statement x together with value α , in order to obtain the challenge β . The latter variant is sometimes called the *strong* FS transform (while the variant we analyze is known as the *weak* FS transform). Bernhard *et al.* [BPW12] show that the weak FS transform might lead to problems in certain applications where the statement to be proven can be chosen adversarially (this is the case, e.g., in the Helios voting protocol). Unfortunately, it seems hard to use our proof techniques to prove zero-knowledge of the strong FS collapse, because the simulator for zero-knowledge does not know the x values in advance.

Other FS-inspired transformations were considered in the literature. For instance Fischlin’s transformation [Fis05] (see also [DV14]) yields a simulation-sound NIZK argument with an online extractor; as mentioned above, Lindell [Lin15] defines a twist of the FS transform that allows to prove zero-knowledge in the CRS model, and soundness in the non-programmable random oracle model. It is an interesting direction for future research to apply our techniques to analyze the above

transformations without random oracles.

1.6 Roadmap

We provide a detailed informal overview of our main techniques in Section 2. In Section 3 we setup some notation and define the main cryptographic primitives on which we build. Section 4 contains our positive result for FS NIZK. We present our compilers for obtaining highly sound protocols (in the CRS model) in Section 5 and Section 6. Finally, we explain how to adapt our techniques to the case of FS signatures in Section 7.

2 Technical Overview

We first discuss the class of highly sound protocols for which the FS transform can be instantiated via a q -wise independent hash function. Then, we will explain how to obtain a compiler that transforms a large class of 3PC protocols into ones that are highly sound (in the CRS model). For the purpose of this overview we will only focus on the case of Fiat–Shamir NIZK, explaining only at the end how our techniques can be adapted to cover Fiat–Shamir signatures as well.

The security proof proceeds in two modular steps. In the first step, we prove completeness and soundness of a “selective” variant of the FS transform (which we define formally in Section 4.1); in the second step we analyze the standard FS transform using complexity leveraging. Details follow.

The selective FS transform. Consider a 3PC argument for a language L . For a hash family H , consider the following (interactive) selective adaptation of the FS transformation: The prover sends the commitment α as in the original protocol; the verifier, instead of sending the challenge $\beta \in \mathcal{B}$ directly, forwards a honestly generated hash key hk ; finally the prover uses (hk, α) to compute $\beta = H(hk, \alpha)$ and then obtains the response γ as in the original 3PC argument.

In Section 4 we prove that if the starting 3PC protocol has instance-independent commitments, is complete and computationally sound, so is the one obtained by applying the selective FS transform. The idea is to use a “programmable” q -wise independent hash function (e.g., a random polynomial of degree $q - 1$ over a finite field) to “program” the hash function up-front; note that commitment α is computed before the hash key is generated and hence, we can embed the challenge value β into the hash function such that it maps α to β and reduce to the soundness of the underlying 3PC argument.

Complexity leveraging. The second step in proving soundness of the FS collapse (we discuss zero-knowledge below) consists in applying complexity leveraging so that we can swap the order of α and β . Note however that if β is shorter than α , and if the soundness of the protocol is $2^{-|\beta|}$, then we lose too much through complexity leveraging. Moreover, if the commitment α depends on the witness, then it is not clear how to obtain an honestly generated α without already knowing a witness (which is needed in the reduction using complexity leveraging—see below). Hence, this step can only be applied to protocols satisfying certain properties as we discuss next. The first property is given below.

P1: The value α output by the prover is computed independently of the instance x being proven (and of the corresponding witness w).⁶

⁶The soundness proof actually only needs independence from the witness. However, for proving zero-knowledge, we will also need independence from the statement, so for convenience, we combine the two requirements into a single property.

Let Π be the initial 3PC argument, and denote by $\bar{\Pi}$ its corresponding FS collapse. Given a malicious prover P^* breaking soundness of $\bar{\Pi}$, we construct a prover P attacking soundness of the selective FS transform as follows. P computes α as the honest prover would do (note that this can be done thanks to **P1** above) and forwards α to the verifier; after receiving the challenge hash key hk , prover P runs P^* which outputs a proof (α^*, γ^*) . Prover P simply hopes that $\alpha^* = \alpha$, in which case it forwards γ^* to the verifier (otherwise it aborts). It follows that if the selective FS has soundness roughly $s(\lambda)$ (for security parameter λ), the soundness of $\bar{\Pi}$ is roughly $s(\lambda)/\delta(\lambda)$, where $\delta(\lambda)$ is the worst-case probability that the α computed by a honest prover is equal to a fixed value.

Note that for the above argument to give a meaningful bound, we need to assume that the ratio $\varrho(\lambda) := s(\lambda)/\delta(\lambda)$ is smaller than one. This leads to the second (non-standard) requirement that the initial 3PC argument should satisfy.

P2: $\varrho(\lambda) := s(\lambda)/\delta(\lambda) < 1$, where $s(\lambda)$ is the soundness error and $\delta(\lambda)$ is the worst-case collision probability associated to the commitment α .

Zero-knowledge. We assume that the initial 3PC is honest-verifier zero-knowledge (HVZK)—i.e., that it is zero-knowledge for honest verifiers. We need to show that $\bar{\Pi}$ satisfies zero-knowledge. Here, in addition to **P1**, we need property **P3** below.

P3: The value α output by the simulator is computed independently of the instance x being proven.

We now discuss the reduction for the zero-knowledge property and explain where **P1** and **P3** are used. We need to construct an efficient simulator that is able to simulate arguments for adaptively chosen (true) statements—without knowing a witness for such statements. The output of the simulator should result in a distribution that is computationally indistinguishable from the distribution generated by the real prover. The simulator gets extra power, as it can produce a “fake” CRS together with some trapdoor information tk (on which the simulator can rely) such that the “fake” CRS is indistinguishable from a real CRS.

In order to build some intuition, it is perhaps useful to recall the random-oracle-based proof for the zero-knowledge property of the FS transform. There, values α_i and β_i corresponding to the i -th adversarial query are computed by running the HVZK simulator and are later “matched” relying on the programmability of the random oracle. Roughly speaking, in our standard-model proof we take a similar approach, but we cannot use *adaptive* programming of the hash-function. Instead, we rely on **P1** and **P3** to program the hash function in advance. More specifically, the trapdoor information will consist of q random tapes r_i (one for simulating each proof queried by the adversary) and the corresponding q challenges β_i (that can be pre-computed as a function of r_i , relying on **P1**). Since the challenges have the correct distribution, we can use the underlying HVZK simulator to simulate the proofs; here is where we need **P3**, as the simulator has to pre-compute the values α_i in order to embed the β_i values on the correct points.

A caveat is that our simulator needs to know the value of q in advance; for this reason we only get a weaker *bounded* flavor of the zero-knowledge property where there exists a “universal” simulator that works for all adversaries asking q queries, for some a-priori fixed value of q . Note, however, that the CRS—as it contains the description of a q -wise independent hash function—needs to grow with q , and hence bound q should be seen as a parameter of the construction rather than a parameter of the simulator.

It is an interesting open problem whether this limitation can be removed, thus proving that actually our transformation achieves unbounded zero-knowledge.

Compilers. Wrapping up the above discussion, we can show that for 3PC protocols that satisfy completeness, computational soundness, HVZK and additionally **P1-P3**, the FS transform can be instantiated by a (programmable) q -wise independent hash function. We informally refer to protocols that satisfy all of the above properties as *highly sound* arguments.

Unfortunately we do not know of a natural highly sound 3PC argument. However, we do know of protocols that partially satisfy our requirements. Recall, for instance, the classical 3PC argument for quadratic residuosity due to Blum [Blu81] (all operations are modulo an integer N which is the product of two Blum integers): (i) The prover chooses a random r in \mathbb{Z}_N^* , and sends $\alpha = r^2$; (ii) The verifier selects a random bit $\beta \in \{0, 1\}$; (iii) The prover computes $\gamma = w^\beta \cdot r$, and finally the verifier checks that $\gamma^2 = x^\beta \cdot \alpha$. While the above clearly satisfies **P1** and moreover can be shown to achieve completeness, soundness, and HVZK, one can easily see that **P2** and **P3** are not directly met. **P2** is not met, because β consists only of a single bit and the soundness parameter is $\frac{1}{2}$, and to see that **P3** is not met, one needs to consider the simulator for this protocol which—for readers familiar with the protocol—computes its first message depending on the statement.

Another interesting example is given by the Lapidot-Shamir protocol for the NP -complete problem of graph Hamiltonicity [LS91] (see also [OV12, Appendix B]). Here, the prover’s commitment consists of a (statistically binding) commitment to the adjacency matrix of a random k -vertex cycle, where k is the size of the Hamiltonian cycle.⁷ Hence, the protocol clearly satisfies **P1**. Additionally the simulator fakes the prover’s commitment by either committing to a random k -vertex cycle, or by committing to the empty graph. Hence, the protocol also satisfies **P3**. As a corollary, we know that assuming non-interactive statistically binding commitment schemes (which follow from one-way permutations [Blu81]), for all languages in NP , there exist 3PC protocols that satisfy completeness, computational soundness, and HVZK, as well as **P1** and **P3**.

Motivated by the above examples, we turn to the question whether it is possible to compile a 3PC protocol (with completeness, soundness, and HVZK) satisfying either **P1** or **P1** and **P3**, into a highly sound argument. We refer the reader to Section 4.5 for a high-level overview how this can be achieved. We only mention here that our compilers rely on several cryptographic tools (including indistinguishability obfuscation, puncturable PRFs, complexity leveraging and equivocal commitment schemes), and yield a 3PC in the CRS model; note that this means that we obtain an interactive protocol with a CRS even if the original protocol was in the standard model. It is an intriguing open problem if a highly sound argument can be constructed in the standard model, or whether a CRS is, in fact, necessary.

The case of signatures. Finally, let us explain how our techniques can be adapted to the case of FS signatures. To this end, we introduce a notion of *highly-sound* canonical identification schemes that need to satisfy similar requirements to the properties **P1**, **P2**, and **P3** discussed above (for the case of 3PC arguments).

Recall that in order to apply our main technique, we need to program the q -wise independent hash function up-front. For this reason we are only able to show that our standard-model instantiation of FS signatures achieves the weaker notion of random-message unforgeability against random-message attacks (RUF-RMA)—in which the adversary can only observe signatures on randomly chosen messages, and also has to forge on an additional fresh random message. While strictly weaker than standard existential unforgeability against chosen-message attacks (EUF-CMA), RUF-RMA is still useful for some applications (e.g., to secure authentication [FHN⁺12, NVZ14]). We refer the reader directly to Section 7 for the details.

⁷Note that the value k can be included in the language, and thus considered as public.

3 Preliminaries

3.1 Notation

By $\lambda \in \mathbb{N}$, we denote the security parameter that we give to all algorithms implicitly in unary representation 1^λ . By $\{0, 1\}^\ell$ we denote the set of all bit-strings of length ℓ , and by $\{0, 1\}^*$ the set of all bit-strings of finite length. If $x, y \in \{0, 1\}^*$ are two bit strings, then $x\|y$ denotes concatenation. The length of x is denoted by $|x|$. We denote vectors of strings in bold face, for example, \mathbf{x} and denote the i -th component by $\mathbf{x}[i]$. For a finite set X , we denote the action of sampling x uniformly at random from X by $x \leftarrow_{\$} X$, and denote the cardinality of X by $|X|$. We denote by $[i]$ the set $\{1, \dots, i\}$. Algorithms are assumed to be randomized, unless otherwise stated. In particular polynomial-time refers to *deterministic* polynomial-time computable algorithms, while PPT refers to probabilistic polynomial-time. We write $x \leftarrow_{\$} \mathcal{A}(\cdot)$ to denote that probabilistic algorithm \mathcal{A} is run on freshly sampled random coins and produces output x . We write $x \leftarrow \mathcal{A}(\cdot; r)$ to denote that \mathcal{A} runs on coins r . Similarly, we write $x \leftarrow \mathcal{A}(\cdot)$ to denote that deterministic algorithm \mathcal{A} outputs x . We say a function $\text{negl}(\lambda)$ is negligible if $\text{negl}(\lambda) \in \lambda^{-\omega(1)}$. We say a function $\text{poly}(\lambda)$ is polynomial if $\text{poly}(\lambda) \in \lambda^{\mathcal{O}(1)}$.

Function families. We formalize families of functions F by considering a tuple of algorithms $F.\text{KGen}$, $F.\text{kl}$, $F.\text{Eval}$, $F.\text{il}$ and $F.\text{ol}$. Algorithm $F.\text{KGen}$ is a PPT algorithm taking the security parameter 1^λ and outputting a key $\mathbf{k} \in \{0, 1\}^{F.\text{kl}(\lambda)}$ where $F.\text{kl} : \mathbb{N} \rightarrow \mathbb{N}$ denotes the key length. Functions $F.\text{il} : \mathbb{N} \rightarrow \mathbb{N}$ and $F.\text{ol} : \mathbb{N} \rightarrow \mathbb{N}$ denote the input and output length functions associated to F and for any $x \in \{0, 1\}^{F.\text{il}(\lambda)}$ and $\mathbf{k} \leftarrow_{\$} F.\text{KGen}(1^\lambda)$ we have that $F.\text{Eval}(\mathbf{k}, x) \in \{0, 1\}^{F.\text{ol}(\lambda)}$, where the PPT algorithm $F.\text{Eval}$ denotes the “evaluation” function associated to F . Depending on the function, Eval may be renamed to a more speaking name and additional algorithms might be added. If the functionality is randomized then we let $F.\text{rl}(\lambda)$ denote the randomness length.

Asymptotic security. In this paper we allow adversaries to be probabilistic polynomial-time (PPT) and ask that the success probability be smaller than some function $\epsilon(\lambda)$ in the security parameter λ . However, when we fix a function $\epsilon(\lambda)$, then for finitely many λ , a specific PPT adversary might be more successful than $\epsilon(\lambda)$. Hence, when defining ϵ -security for a scheme, we say that for all PPT adversaries \mathcal{A} , the advantage function $\text{Adv}_{\mathcal{A}}(\lambda)$ is asymptotically smaller than ϵ , denoted $\text{Adv}_{\mathcal{A}}(\lambda) \stackrel{\text{asym}}{\leq} \epsilon(\lambda)$, which means that there is some value λ_0 such that for all $\lambda \geq \lambda_0$, it holds that $\text{Adv}(\lambda) \leq \epsilon(\lambda)$.

For two random variables X and Y , we say that they are ϵ -indistinguishable, denoted $X \approx_{\epsilon} Y$, if for all PPT distinguishers the distinguishing advantage is asymptotically smaller than ϵ .

3.2 q -Wise Independent Hashing

We recall the standard notion of a q -wise independent hash function.

Definition 3.1 (q -Wise Independent Hashing) *A family of functions $H := (H.\text{KGen}, H.\text{kl}, H.\text{Eval}, H.\text{il}, H.\text{ol})$ is called q -wise independent if for all $\lambda \in \mathbb{N}$ and $x_1, \dots, x_q \in \{0, 1\}^{H.\text{il}(\lambda)}$ and all $y_1, \dots, y_q \in \{0, 1\}^{H.\text{ol}(\lambda)}$ we have that*

$$\Pr \left[H.\text{Eval}(\mathbf{hk}, x_1) = y_1 \wedge \dots \wedge H.\text{Eval}(\mathbf{hk}, x_q) = y_q : \mathbf{hk} \leftarrow_{\$} H.\text{KGen}(1^\lambda) \right] \leq 2^{-q \cdot H.\text{ol}(\lambda)}.$$

We call a q -wise independent hash function programmable if there exists an additional procedure $H.\text{KGen}(1^\lambda, \mathbf{x}, \mathbf{y})$ such that for any two q -size vectors \mathbf{x} and \mathbf{y} , with $\mathbf{x}[i] \in \{0, 1\}^{H.\text{il}(\lambda)}$ and $\mathbf{y}[i] \in$

$\{0, 1\}^{\text{H.ol}(\lambda)}$ for all $i \in [q]$, we have that

$$\Pr \left[\text{H.Eval}(\text{hk}, \mathbf{x}[i]) = \mathbf{y}[i] : \text{hk} \leftarrow_{\$} \text{H.KGen}(1^\lambda, \mathbf{x}, \mathbf{y}) \right] = 1$$

and furthermore the distributions $\text{H.KGen}(1^\lambda)$ and $\text{H.KGen}(1^\lambda, \mathbf{x}, \mathbf{y})$ are identical where the second distribution is also over the uniformly random choice of vectors \mathbf{x} and \mathbf{y} .

We note that the constant function $\text{H}(\text{hk}, x) := \text{hk}$ is 1-wise independent and programmeable. Furthermore, note that we can construct a programmable q -wise independent hash function by considering polynomials of degree $q - 1$ over finite fields. Programmability is obtained by using polynomial interpolation.

3.3 Interactive and Non-Interactive Arguments

Let $R : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}$ be a polynomial-time computable relation together with a polynomial $p(\cdot)$, defining the NP-language

$$L_R := \{x : \exists w \text{ s.t. } |w| < p(|x|) \text{ and } R(x, w) = 1\}.$$

In the rest of the paper, we will drop the bound $p(\cdot)$ for ease of presentation. An interactive argument system for R , consists of three PPT algorithms $(\text{K}, \text{P}, \text{V})$. Algorithm K takes as input 1^λ and outputs a common reference string (CRS) $\text{crs} \in \{0, 1\}^*$. Later the prover P interacts with the verifier V to convince him into accepting a common input $x \in L_R$ (where both P and V are also given crs); the honest prover additionally holds a witness w for x , i.e. $R(x, w) = 1$. At the end of the protocol execution, the verifier outputs a bit (representing his decision); we write $\langle \text{P}(w), \text{V} \rangle(\text{crs}, x)$ for the random variable corresponding to the verifier's verdict. Similarly, we write $\text{P}(\text{crs}, x, w) \stackrel{\text{def}}{=} \text{V}(\text{crs}, x)$ for the random variable corresponding to transcripts of honest protocol executions.⁸

An interactive argument should satisfy at least two properties, completeness and soundness. Completeness says that an honest prover (holding a valid witness) is able to convince the verifier.

Definition 3.2 (Completeness) *Let $\Pi = (\text{K}, \text{P}, \text{V})$ be an interactive argument system for a polynomial-time computable relation R . We say that Π satisfies c -completeness if for all (x, w) such that $R(x, w) = 1$ we have*

$$\Pr \left[\langle \text{P}(w), \text{V} \rangle(\text{crs}, x) = 1 : \text{crs} \leftarrow_{\$} \text{K}(1^\lambda) \right] \stackrel{\text{asym}}{\geq} 1 - c(\lambda),$$

where the probability is taken over the randomness of algorithms P , V and K .

Soundness informally says that, whenever $x \notin L_R$, no computationally bounded prover can convince the verifier into accepting x .

Definition 3.3 (Soundness) *Let $\Pi = (\text{K}, \text{P}, \text{V})$ be an interactive argument system for a polynomial-time computable relation R . We say that Π satisfies s -soundness if for all PPT algorithms P^* , and for any $x \notin L_R$, we have that*

$$\Pr \left[\langle \text{P}^*, \text{V} \rangle(\text{crs}, x) = 1 : \text{crs} \leftarrow_{\$} \text{K}(1^\lambda) \right] \stackrel{\text{asym}}{\leq} s(\lambda),$$

where the probability is taken over the randomness of algorithms P^* , V and K .

⁸We stress that interactive arguments typically do not require a CRS. Looking ahead, the reason for defining interactive arguments in the CRS model is that our compilers in Section 5 and 6 will produce an interactive argument in the CRS model (to be used in our instantiation of the FS transform).

Completeness and soundness do not quantify how much information an interactive argument reveals about the witness, which in turn can be covered by notions such as witness indistinguishability and zero-knowledge. In this paper we will use different flavors of the zero-knowledge property. We will postpone the actual definitions to the place in the paper where they are actually used.

Standard model interactive arguments. We can cast the case where the interactive argument is in the standard-model, i.e., it does not rely on a CRS (which is typically the case), by saying that the algorithm K returns the empty string; similarly P and V do not take the CRS as input (or take an empty string as additional input). When we write $\Pi = (P, V)$, we denote an interactive argument in the standard model. Adapting the definitions of completeness and soundness to the standard model works by replacing the CRS generation algorithm by an algorithm that outputs the empty string.

Non-interactive arguments. We speak of non-interactive arguments in case the protocol consists of a single message π sent from the prover to the verifier. Non-interactive arguments that satisfy a zero-knowledge property typically require a setup assumption, such as a CRS.⁹ Syntactically a non-interactive argument system for a polynomial-time computable relation R consists of three PPT algorithms $\bar{\Pi} := (K, P, V)$ specified as follows: (i) Algorithm K takes as input 1^λ and outputs a CRS $\text{crs} \in \{0, 1\}^*$; (ii) Algorithm P takes as input (crs, x, w) such that $R(x, w) = 1$ and outputs a proof π ; (iii) Algorithm V takes as input (crs, x, π) and outputs a bit indicating whether π is a valid proof for x (under crs) or not.

A non-interactive argument should satisfy three main properties, which are analogous to the definitions of completeness, soundness and (honest-verifier) zero-knowledge for interactive arguments. We define these properties below.

Definition 3.4 (Completeness of non-interactive arguments) *Let $\bar{\Pi} = (K, P, V)$ be a non-interactive argument system for a polynomial-time computable relation R . We say that $\bar{\Pi}$ satisfies c -completeness if for all (x, w) such that $R(x, w) = 1$ we have*

$$\Pr \left[V(\text{crs}, x, \pi) = 1 : \text{crs} \leftarrow_{\$} K(1^\lambda); \pi \leftarrow_{\$} P(\text{crs}, x, w) \right] \stackrel{\text{asym}}{\geq} 1 - c(\lambda),$$

where the probability is taken over the randomness of algorithms P, V and K .

Definition 3.5 (Soundness of non-interactive arguments) *Let $\bar{\Pi} = (K, P, V)$ be a non-interactive argument system for a polynomial-time computable relation R . We say that $\bar{\Pi}$ satisfies s -soundness if for all for all PPT algorithms P^* , and for any $x \notin L_R$, we have that*

$$\Pr \left[V(\text{crs}, x, \pi) = 1 : \text{crs} \leftarrow_{\$} K(1^\lambda); \pi \leftarrow_{\$} P^*(\text{crs}, x) \right] \leq s(\lambda),$$

where the probability is taken over the randomness of algorithms P^*, V and K .

Definition 3.6 (q -Bounded Computational Zero-Knowledge) *Let $\bar{\Pi} = (K, P, V)$ be a non-interactive argument system for a polynomial-time computable relation R . We say that $\bar{\Pi}$ satisfies q -bounded ϵ -computational zero-knowledge if for all binary PPT adversaries \mathcal{A} there exists a PPT simulator $\mathcal{S} := (\mathcal{S}', \mathcal{S}'')$ such that $\text{rNIZK}_{\mathcal{A}}^{\bar{\Pi}}(\lambda) \approx_{\epsilon} \text{sNIZK}_{\mathcal{S}', \mathcal{S}''}^{\bar{\Pi}}(\lambda)$, where experiments rNIZK and sNIZK are defined below*

⁹In particular, assuming a CRS is necessary for obtaining non-interactive zero knowledge [Gol01, Chapter 4].

$\text{rNIZK}_{\mathcal{A}}^{\overline{\Pi}}(\lambda)$	$\text{sNIZK}_{\mathcal{S}', \mathcal{S}''}^{\overline{\Pi}}(\lambda)$
$\text{crs} \leftarrow_{\S} \mathcal{K}(1^\lambda)$	$(\text{crs}, \text{tk}) \leftarrow_{\S} \mathcal{S}'(1^\lambda)$
return $\mathcal{A}^{\text{PROV}(\text{crs}, \cdot, \cdot)}(1^\lambda, \text{crs})$	return $\mathcal{A}^{\text{SIMU}(\text{crs}, \text{tk}, \cdot, \cdot)}(1^\lambda, \text{crs})$
$\text{PROV}(\text{crs}, x, w)$	$\text{SIMU}(\text{crs}, \text{tk}, x, w)$
if $R(x, w) = 1$ then	if $R(x, w) = 1$ then
return $\pi \leftarrow_{\S} \mathcal{P}(\text{crs}, x, w)$	return $\pi \leftarrow_{\S} \mathcal{S}''(\text{crs}, \text{tk}, x)$
else return \perp	else return \perp

and \mathcal{A} can ask $q(\lambda)$ queries to its oracle. Additionally we say that $\overline{\Pi}$ satisfies unbounded computational non-interactive zero-knowledge if indistinguishability of the above experiments holds for an arbitrary polynomial $q(\lambda)$.

Note that we quantify over all *binary* adversaries, that is, we consider only adversaries that output a bit. This is without loss of generality, but makes notation in later game-hop proofs easier. For brevity, we sometimes write that $\overline{\Pi}$ is a (c, s, q, ϵ) -NIZK to denote that $\overline{\Pi}$ satisfies c -completeness, s -soundness, and q -bounded ϵ -computational zero-knowledge.

3.4 Obfuscation

Indistinguishability obfuscation [BGI⁺01, BGI⁺12] intuitively captures that the obfuscation of two functionally equivalent circuits cannot be distinguished. We here give a game-based definition, following the definitional framework of [BST14] which captures indistinguishability obfuscation based notions via the IO security game and a class of samplers Sam .

Definition 3.7 (Obfuscation Scheme) A PPT algorithm \mathcal{O} is called an obfuscation scheme if, on input the security parameter 1^λ and a description of a circuit C , it returns (a description of) a circuit \overline{C} such that $\forall x : C(x) = \overline{C}(x)$. We call a PPT algorithm Sam a circuit sampler if on input the security parameter 1^λ sampler Sam outputs (C_0, C_1, aux) where C_0 and C_1 are descriptions of circuits and aux is a string. If Sam is a circuit sampler and \mathcal{O} is an obfuscation scheme we define the advantage $\text{Adv}_{\mathcal{O}, \text{Sam}, \mathcal{D}}^{\text{io}}(\cdot)$ for a distinguisher \mathcal{D} relative to game IO:

$$\text{Adv}_{\mathcal{O}, \text{Sam}, \mathcal{D}}^{\text{io}}(\lambda) := 2 \cdot \Pr \left[\text{IO}_{\mathcal{D}, \text{Sam}}^{\mathcal{O}}(\lambda) \right] - 1$$

$\text{IO}_{\mathcal{D}, \text{Sam}}^{\mathcal{O}}(\lambda)$

$b \leftarrow_{\S} \{0, 1\}$
 $(C_0, C_1, \text{aux}) \leftarrow_{\S} \text{Sam}(1^\lambda)$
 $\overline{C} \leftarrow_{\S} \mathcal{O}(1^\lambda, C_b)$
 $b' \leftarrow_{\S} \mathcal{D}(1^\lambda, \overline{C}, \text{aux})$
return $(b = b')$

If \mathcal{S} is a class of circuit samplers, we call an obfuscation scheme \mathcal{O} $\epsilon_{\mathcal{O}}$ -secure for \mathcal{S} , if for all $\text{Sam} \in \mathcal{S}$ and all PPT distinguishers \mathcal{D} advantage $\text{Adv}_{\mathcal{O}, \text{Sam}, \mathcal{D}}^{\text{io}}(\lambda) \stackrel{\text{asym}}{\leq} \epsilon_{\mathcal{O}}(\lambda)$.

We can now capture indistinguishability obfuscation via restricting the class of samplers to so-called *equality samplers*. As we only use efficient samplers we can further restrict the class of samplers.

Definition 3.8 (Equality circuit sampler) We call a PPT algorithm Sam an equality circuit sampler if for all security parameters $\lambda \in \mathbb{N}$ it outputs a triple (C_0, C_1, aux) consisting of two circuit descriptions and a string such that with overwhelming probability over the coins of Sam we have that the circuits C_0 and C_1 have the same size, number of inputs and number of outputs and are functionally equivalent, that is

$$\Pr_{(C_0, C_1, \text{aux}) \leftarrow \text{Sam}(1^\lambda)}[|C_0| = |C_1| \wedge \forall x : C_0(x) = C_1(x)] \geq 1 - \text{negl}(\lambda).$$

A beautiful result that we will use is the relationship between differing-inputs obfuscation and indistinguishability obfuscation proved by Boyle, Chung and Pass [BCP14], who show that any general purpose indistinguishability obfuscator is also a differing-inputs obfuscator for circuits that differ only on a few (at most polynomially many) inputs. We first define differing-inputs obfuscation by restricting samplers to be differing-inputs samplers.

Definition 3.9 (Differing-inputs circuit sampler) Let Sam be a circuit sampler. We call Sam a differing-inputs circuit sampler if advantage $\text{Adv}_{\text{Sam}, \text{Ext}}^{\text{diff}}(\cdot)$ is negligible for all PPT algorithms Ext and where the advantage is defined as (relative to game Diff on the right):

$$\text{Adv}_{\text{Sam}, \text{Ext}}^{\text{diff}}(\lambda) := \Pr \left[\text{Diff}_{\text{Sam}}^{\text{Ext}}(\lambda) \right]$$

$$\frac{\text{Diff}_{\text{Sam}}^{\text{Ext}}(\lambda)}{(C_0, C_1, \text{aux}) \leftarrow \text{Sam}(1^\lambda)$$

$$x \leftarrow \text{Ext}(1^\lambda, C_0, C_1, \text{aux})$$

$$\text{return } (C_0(x) \neq C_1(x))$$

With that we are ready to formulate the result due to Boyle, Chung and Pass [BCP14].

Theorem 3.10 ([BCP14]) Let iO be an indistinguishability obfuscator for all circuits in P/poly . Let Sam be a differing-inputs circuit sampler for which there exists a polynomial $d : \mathbb{N} \rightarrow \mathbb{N}$, such that

$$\Pr \left[|\{x : C_0(x) \neq C_1(x)\}| \leq d(\lambda) \mid (C_0, C_1, \text{aux}) \leftarrow \text{Sam}(1^\lambda) \right] \geq 1 - \text{negl}(\lambda).$$

Then iO is a differing-inputs obfuscator for Sam , i.e., obfuscator iO is $\{\text{Sam}\}$ -secure.

3.5 Puncturable Pseudorandom Functions

A key ingredient in the compilers are so-called puncturable pseudorandom functions (PRFs) [SW14]. A family of puncturable PRFs is a function family that additionally comes with a PPT *puncturing* algorithm Pntr which on input a polynomial-size set $S \subseteq \{0, 1\}^{\text{il}(\lambda)}$, outputs a special key k_S .

Definition 3.11 (Puncturable PRF) A family of functions $F := (F.\text{KGen}, F.\text{Pntr}, F.\text{kl}, F.\text{Eval}, F.\text{il}, F.\text{ol})$ is called a ϵ_{prf} -secure, puncturable PRF if the following holds.

FUNCTIONALITY PRESERVED UNDER PUNCTURING. For every PPT adversary \mathcal{A} such that $\mathcal{A}(1^\lambda)$ outputs a polynomial-size set $S \subseteq \{0, 1\}^{\text{F.il}(\lambda)}$, it holds for all $x \in \{0, 1\}^{\text{F.il}(\lambda)} \setminus S$ that:

$$\Pr \left[F.\text{Eval}(k, x) = F.\text{Eval}(k_S, x) : k \leftarrow F.\text{KGen}(1^\lambda), k_S \leftarrow F.\text{Pntr}(k, S) \right] = 1.$$

PSEUDORANDOM AT PUNCTURED POINTS. For every PPT adversary $(\mathcal{A}_1, \mathcal{A}_2)$, the advantage $\text{Adv}_{F, \mathcal{A}_1, \mathcal{A}_2}^{\text{pprf}}(\cdot)$ is asymptotically smaller than ϵ_{prf} , i.e.,

$$\text{Adv}_{F, \mathcal{A}_1, \mathcal{A}_2}^{\text{pprf}}(\lambda) = 2 \cdot \Pr \left[\text{pPRF}_{\mathcal{A}_1, \mathcal{A}_2}^F(\lambda) \right] - 1 \stackrel{\text{asym}}{\leq} \epsilon_{\text{prf}}$$

where game pPRF is defined as

$\text{pPRF}_{\mathcal{A}_1, \mathcal{A}_2}^F(\lambda)$	$\text{CHALLENGE}(x)$
$S \leftarrow \emptyset; b \leftarrow_{\$} \{0, 1\}$	if $x \in S$ then return \perp
$k \leftarrow_{\$} \text{F.KGen}(1^\lambda)$	$S \leftarrow S \cup \{x\}$
$\text{state} \leftarrow_{\$} \mathcal{A}_1^{\text{CHALLENGE}}(1^\lambda)$	if $b = 0$ then
$k^* \leftarrow_{\$} \text{F.Pntr}(k, S)$	$y \leftarrow \text{F.Eval}(k, x)$
$b' \leftarrow_{\$} \mathcal{A}_2(1^\lambda, \text{state}, k^*)$	else $y \leftarrow_{\$} \{0, 1\}^{\text{F.ol}(\lambda)}$
return $(b = b')$	return y

4 Fiat–Shamir NIZK

We show that under specific assumptions on the underlying protocol, a q -wise independent hash function is enough to instantiate the random oracle in the Fiat–Shamir collapse yielding a secure NIZK with q -bounded computational zero-knowledge.

After recalling the standard FS transform in Section 4.1, we present a “selective” interactive variant of the transformation, and establish its completeness and soundness in Section 4.2. Later, in Section 4.3, we put forward three properties of the initial 3PC argument that allow to prove completeness, soundness, and q -bounded computational zero-knowledge of the FS-collapse in the standard model; the proof of completeness and soundness reduce directly to the completeness and soundness of the above selective FS transform. Our main theorem is summarized in Section 4.4. Finally, in Section 4.5, we take a closer look at the required properties and discuss how to achieve them.

4.1 The Fiat–Shamir Transform

The Fiat–Shamir (FS) transform [FS87] is a generic way to remove interaction from certain argument systems, using a hash-function. For the rest of the paper, we consider only interactive arguments consisting of three messages—which we denote by (α, β, γ) —where the first message is sent by the prover. We also focus on so-called *public-coin* protocols where the verifier’s message β is uniformly random over some space \mathcal{B} (e.g., $\beta \in \{0, 1\}^k$ for some $k \in \mathbb{N}$). We call this a 3PC argument system for short.

For 3PC arguments it is convenient to think of the prover algorithm as being split into two sub-algorithms $\text{P} := (\text{P}_0, \text{P}_1)$, where P_0 takes as input a pair (x, w) and outputs the prover’s first message α (the so-called commitment) and P_1 takes as input (x, w) as well as the verifier’s challenge β to produce the prover’s second message γ (the so-called response). In general P_0 and P_1 are allowed to share the same random tape, which we denote by $r \in \{0, 1\}^*$. In a similar fashion we can think of the verifier’s algorithm as split into two sub-algorithms $\text{V} = (\text{V}_0, \text{V}_1)$, where V_0 outputs a uniformly random value $\beta \in \mathcal{B}$ and V_1 is deterministic and corresponds to the verifier’s verdict (i.e., V_1 takes as input x and a transcript (α, β, γ) and returns a decision bit $d \in \{0, 1\}$).

Non-interactive version. The FS transform allows to remove interaction from any 3PC argument system for a polynomial-time computable relation R as specified below (see also Fig. 1). Let $\Pi = (\text{K}, \text{P}, \text{V})$ be the initial 3PC argument system. Additionally, consider a family of hash functions H consisting of algorithms H.KGen , H.kl , H.Eval , H.il and H.ol (see Section 3.1); here H.il and H.ol correspond, respectively, to the bit lengths of messages α and β (as a function of the security parameter λ).

The FS collapse of Π using H is a triple of algorithms $\bar{\Pi}_{\text{FS}, \text{H}} := (\text{K}_{\text{FS}}, \text{P}_{\text{FS}}, \text{V}_{\text{FS}})$ defined as follows.

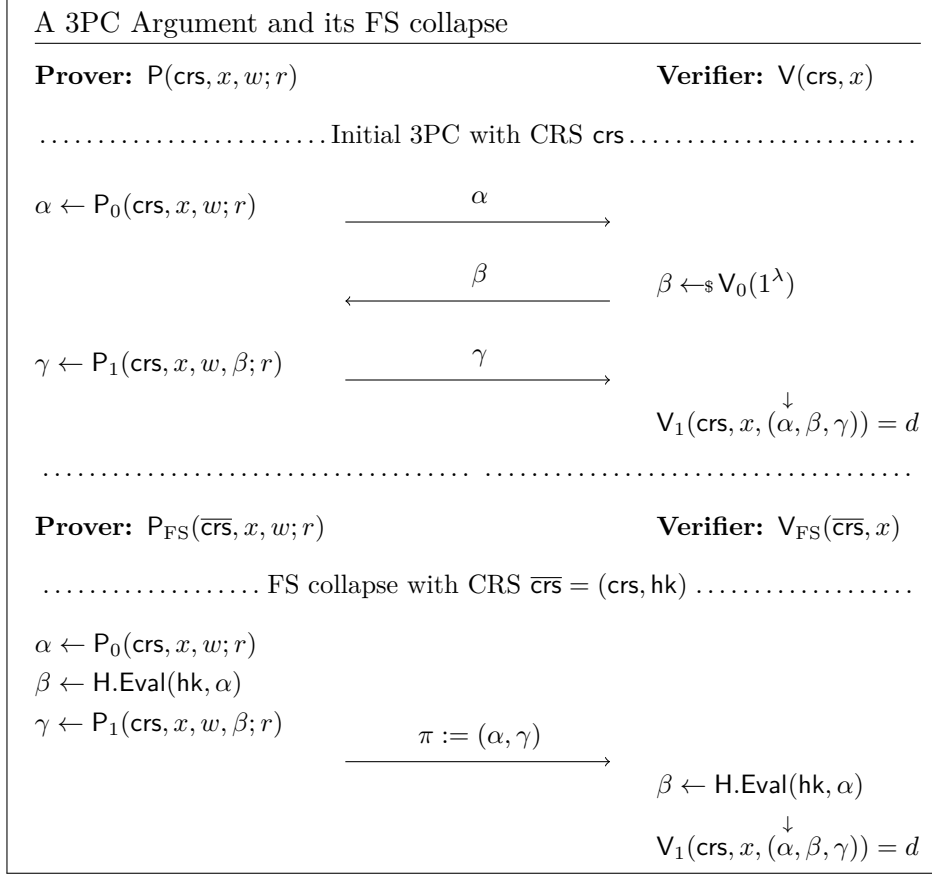


Figure 1: Message flow of a typical 3PC argument system and its corresponding FS collapse. In case the initial 3PC is in the standard model we simply have $\Pi = (P, V)$ and $\overline{\text{crs}}$ contains only the hash key. Note also that we consider public-coin protocols and thus do not specify the randomness of the verifier (the randomness of V_0 is β and V_1 is deterministic given β).


- Algorithm K_{FS} takes as input the security parameter, samples $\text{hk} \leftarrow_{\$} \text{H.KGen}(1^\lambda)$, $\text{crs} \leftarrow_{\$} \text{K}(1^\lambda)$, and publishes $\overline{\text{crs}} := (\text{crs}, \text{hk})$.
- Algorithm P_{FS} takes as input $(\overline{\text{crs}}, x, w)$ and runs $P_0(\text{crs}, x, w)$ in order to obtain the commitment $\alpha \in \{0, 1\}^{\text{H.il}(\lambda)}$; next P_{FS} defines the challenge as $\beta := \text{H.Eval}(\text{hk}, \alpha)$ and runs $P_1(\text{crs}, x, w, \beta)$ in order to obtain the response γ . Finally P_{FS} outputs $\pi := (\alpha, \gamma)$.
- Algorithm V_{FS} takes as input $(\overline{\text{crs}}, x, \pi)$ and returns 1 if and only if verifier $V_1(\text{crs}, x, (\alpha, \beta, \gamma)) = 1$ where $\beta = \text{H.Eval}(\text{hk}, \alpha)$.

In a nutshell the result of Fiat and Shamir says that whenever $\Pi = (P, V)$ is a (standard model) 3PC argument satisfying completeness, computational soundness, and computational honest-verifier zero-knowledge (in addition to a basic requirement on the min-entropy of the prover's commitment), its FS collapse $\overline{\Pi}_{\text{FS}, \text{H}}$ is a NIZK argument system if H is modeled as a random oracle.

4.2 A Selective Variant of Fiat–Shamir

As an intermediate step in the proof of soundness of our standard model instantiation of the FS transform, we will consider a selective variant of the FS transform of a 3PC argument system which

basically translates into allowing the hash function to depend on the commitment α . Note that this selective variant is still *interactive* since we consider the prover to be split into two algorithms, where the first algorithm is identical to P_0 and the second algorithm first computes β using the received hash key and later runs P_1 in order to obtain γ ; similarly the verifier is split into two algorithms, where the first algorithm now generates the hash key (instead of sampling β directly) and the second algorithm is identical to V_1 :



Note that the verifier in the above protocol accepts if and only if (α, β, γ) is an accepting proof for x and moreover $\beta = \text{H.Eval}(\text{hk}, \alpha)$. We write $\Pi_{\text{sel-FS,H}}$ for the above selective (interactive) version of the FS transform, and define its completeness and soundness properties below.

Definition 4.1 (Completeness of the Selective FS Transform) *Let $\Pi = (\text{K}, (P_0, P_1), (V_0, V_1))$ be a 3PC argument system for a polynomial-time computable relation R , and let $\Pi_{\text{sel-FS,H}}$ be the corresponding selective FS transform using hash function family H . We say that $\Pi_{\text{sel-FS,H}}$ satisfies c -completeness if for all (x, w) such that $R(x, w) = 1$ we have*

$$\Pr \left[V_1(\text{crs}, x, (\alpha, \beta, \gamma)) = 1 : \begin{array}{l} \text{crs} \leftarrow_{\$} \text{K}(1^\lambda); \alpha \leftarrow_{\$} P_0(\text{crs}, x, w); \\ \text{hk} \leftarrow_{\$} \text{H.KGen}(1^\lambda); \\ \beta = \text{H.Eval}(\text{hk}, \alpha); \gamma \leftarrow_{\$} P_1(\text{crs}, x, w, \beta) \end{array} \right] \stackrel{\text{asym}}{\geq} 1 - c(\lambda),$$

where the probability is taken over the randomness of algorithms P_0, P_1, V, K and over the choice of the hash key.

Definition 4.2 (Soundness of the Selective FS Transform) *Let $\Pi = (\text{K}, (P_0, P_1), (V_0, V_1))$ be a 3PC argument system for a polynomial-time computable relation R , and let $\Pi_{\text{sel-FS,H}}$ be the corresponding selective FS transform using hash function family H . We say that $\Pi_{\text{sel-FS,H}}$ satisfies s -soundness if for all PPT algorithms $P^* = (P_0^*, P_1^*)$, and for all $x \notin L_R$, we have that*

$$\Pr \left[V_1(\text{crs}, x, (\alpha, \beta, \gamma)) = 1 : \begin{array}{l} \text{crs} \leftarrow_{\$} \text{K}(1^\lambda); \alpha \leftarrow_{\$} P_0^*(\text{crs}, x); \\ \text{hk} \leftarrow_{\$} \text{H.KGen}(1^\lambda); \\ \beta = \text{H.Eval}(\text{hk}, \alpha); \gamma \leftarrow_{\$} P_1^*(\text{crs}, x, \text{hk}, \alpha, \beta) \end{array} \right] \stackrel{\text{asym}}{\leq} s(\lambda),$$

where the probability is taken over the randomness of algorithms P^*, V, K and over the choice of the hash key.

Completeness and soundness for selective FS. We can now move on to state our first result: If H is a 1-wise independent hash function, then the selective FS transform instantiated with H

maintains completeness and computational soundness of the starting 3PC argument.¹⁰ Note that already the constant function $H(hk, x) = hk$ is 1-wise independent and thus fulfills the requirements of the following theorem.

Theorem 4.3 *Let $\Pi = (K, P, V)$ be a 3PC argument system for a polynomial-time computable relation R , that is c -complete and s -sound, and let H be a 1-wise independent hash function. Then, the selective FS transform $\Pi_{\text{sel-FS}, H}$ of Π using H is c -complete and s -sound for relation R .*

Proof. The proof for completeness and soundness follows directly from noting that β is distributed uniformly at random in $\{0, 1\}^{H.o(\lambda)}$ over the choice of the hash key, and as the hash key is chosen independently of α the proof reduces directly to the completeness and soundness of the interactive version of the underlying 3PC. \square

4.3 The FS-Collapse

We now consider the standard FS collapse and discuss each property (completeness, soundness, and zero-knowledge) in turn. We reduce soundness and completeness to the soundness and completeness of the *selective FS* transform, and reduce zero-knowledge directly to the (instance-independent honest-verifier) zero-knowledge property of the underlying 3PC argument. Instance-independence is a new property for protocols that we define in this section.

Note that, for our final theorem, we require the starting 3PC protocol to satisfy three “non-standard” requirements (that we introduce along the way), including for example, the previously mentioned instance-independence property.

4.3.1 Completeness and Soundness

We start by showing that if the underlying 3PC argument satisfies completeness, so does the resulting FS non-interactive argument.

Lemma 4.4 *Let $\Pi = (K, P, V)$ be a 3PC argument system for a polynomial-time computable relation R satisfying c -completeness. Then, assuming H is a 1-wise independent hash function, the FS collapse $\bar{\Pi}_{\text{FS}, H}$ of Π using H satisfies c -completeness.*

Proof. The proof follows by noting that when both the prover and the verifier of the non-interactive protocol are honest, the probability that the verifier accepts is the same as in the (interactive) selective variant of the FS transform applied to Π . The statement then follows from Theorem 4.3. \square

To capture soundness of the FS-collapse we need an additional property of the underlying 3PC protocol, namely, a gap between the probability δ with which one can guess the first message of the protocol and the soundness error s . We can interpret that $s/\delta < 1$ as follows: even if we allow a loss of δ to guess the first message, then still there remains a level of security that can be leveraged to obtain soundness. We next capture the guessing probability δ and the soundness-error-to-guessing-gap ϱ . We begin with defining the guessing probability, that we call *collision probability of commitments* and note that it is a worst-case notion, i.e., it is the probability of the least likely value that the prover outputs.

¹⁰Note that we do not prove zero-knowledge of the selective FS transform; this is because we will later prove directly non-interactive zero-knowledge of the FS collapse.

Definition 4.5 (Worst-Case Collision Probability of Commitments) Let λ be a security parameter, and consider a 3PC argument system $\Pi = (\mathsf{K}, (\mathsf{P}_0, \mathsf{P}_1), \mathsf{V})$ for a polynomial-time computable relation R . We say that Π has worst-case collision probability $\delta(\lambda)$ when

$$\min_{(x,w) \in R, \alpha \in \text{Im}(\mathsf{P}_0(\text{crs}, x, w))} \Pr \left[\mathsf{P}_0(\text{crs}, x, w) = \alpha : \text{crs} \leftarrow_{\S} \mathsf{K}(1^\lambda) \right] = \delta(\lambda).$$

The soundness of our standard model instantiation depends on the gap between the soundness error and the worst-case collision probability.

Definition 4.6 (Soundness-Error-to-Guessing Ratio) Let λ be a security parameter, and consider a 3PC argument system $\Pi = (\mathsf{K}, \mathsf{P}, \mathsf{V})$ for a polynomial-time computable relation R with $\delta(\lambda)$ worst-case collision probability and $s(\lambda)$ -soundness. The soundness-error-to-guessing ratio (SEGR) associated to Π is defined as $\varrho(\lambda) := \frac{s(\lambda)}{\delta(\lambda)}$.

To prove soundness we need one more property of the underlying 3PC. While the above SEGR property is rather non-standard and we do not know if it can be realized in the standard model (recall that our compilers yield protocols in the CRS model) the following property is shared by many 3PC protocols. We require that the prover—we will later in Section 4.3.2 see an analogous definition for the HVZK simulator, which however seems to be less standard—chooses its commitment α independently of the instance x and the witness w . We call this property *instance-independent commitment*.

Definition 4.7 (Instance-Independent Commitments) Let $\Pi = (\mathsf{K}, \mathsf{P} = (\mathsf{P}_0, \mathsf{P}_1), \mathsf{V})$ be a 3PC argument system for a polynomial-time computable relation R . We say that Π has instance-independent commitments if $\mathsf{P}_0(\text{crs}, x, w; r) := \mathsf{P}_0(\text{crs}; r)$ for any choice of randomness r , instance x and witness w .

One example of a 3PC protocol that has instance-independent commitments is the 3PC argument due to Blum for the quadratic residuosity [Blu81]. Another example is given by the Lapidot-Shamir protocol for graph Hamiltonicity [LS91]. Armed with a “sub-one” soundness-error-to-guessing ratio and instance-independent commitments we can now quantify the soundness of our instantiation of the FS-collapse.

Lemma 4.8 Let $\Pi = (\mathsf{K}, \mathsf{P}, \mathsf{V})$ be a 3PC argument system for a polynomial-time computable relation R , with worst-case collision probability δ and SEGR ϱ , s -soundness and instance-independent commitments. Then, assuming H is a 1-wise independent hash function, the FS collapse $\bar{\Pi}_{\mathsf{FS}, \mathsf{H}} = (\mathsf{K}_{\mathsf{FS}}, \mathsf{P}_{\mathsf{FS}}, \mathsf{V}_{\mathsf{FS}})$ of Π using H satisfies ϱ -soundness.

Proof. Let algorithm $\mathsf{P}_{\mathsf{FS}}^*$ be an adversary for the non-interactive FS collapse. Let $x \notin L_R$ and let

$$\mu(\lambda) := \Pr \left[\mathsf{V}_{\mathsf{FS}}((\text{crs}, \text{hk}), x, (\alpha^*, \gamma^*)) = 1 : \begin{array}{l} \text{crs} \leftarrow_{\S} \mathsf{K}(1^\lambda); \text{hk} \leftarrow_{\S} \mathsf{H}.\mathsf{KGen}(1^\lambda); \\ (\alpha^*, \gamma^*) \leftarrow_{\S} \mathsf{P}_{\mathsf{FS}}^*((\text{crs}, \text{hk}), x) \end{array} \right]$$

the advantage of $\mathsf{P}_{\mathsf{FS}}^*$ in breaking soundness of the FS collapse. We show how to use $\mathsf{P}_{\mathsf{FS}}^*$ to construct a malicious prover $\mathsf{P}^* := (\mathsf{P}_0^*, \mathsf{P}_1^*)$ breaking soundness of $\Pi_{\text{sel-FS}, \mathsf{H}}$ as follows. The prover P_0^* runs the honest (instance-independent¹¹) prover P_0 on $(1^\lambda, \text{crs})$ to get a value α and sends α to the verifier. It gets back a hash-function key hk that is independent from the value α that P_0^* sent to the verifier in the first message. Now, prover P_1^* runs prover $\mathsf{P}_{\mathsf{FS}}^*$ on $((\text{crs}, \text{hk}), x)$ to obtain a pair (α^*, γ^*) . If $\alpha^* = \alpha$, then P_1^* passes γ^* to the verifier. Else, P_1^* aborts.

¹¹We note that for this witness-independence would actually suffice.

Observe that the success probability of P^* is lower bounded by the success probability of P_{FS}^* times the probability that α^* is equal to α . If the selective FS transform of Π has soundness $s'(\lambda)$ we obtain

$$\mu(\lambda)\delta(\lambda) \stackrel{\text{asym}}{\leq} s'(\lambda) = s(\lambda)$$

where the last equality is due to Theorem 4.3. The statement now follows by a division of the inequality by δ and the definition of soundness-error-to-guessing ratio. \square

4.3.2 Zero-Knowledge

To quantify the zero-knowledge property of our instantiation of the FS collapse we need the analogue to instance-independent commitments for honest-verifier zero-knowledge (HVZK) simulators, that is, we require that also simulators can choose α and β independently of the instance. Note that while instance-independent commitments has nothing to do with the challenge β , it follows from the definition of 3PC protocols that β is chosen independently of the instance by verifier V_0 . Additionally, we note that we can only prove bounded HVZK, that is, we require that the adversary can only make q -many oracle queries where q is an arbitrary polynomial. We define the property in the CRS model and, again, note that the standard-model version of this definition is obtained by replacing the crs with an empty string.

Definition 4.9 (q -Bounded Instance-Independent HVZK) *Let $\Pi = (K, P, V)$ be a 3PC argument system in the CRS model for a polynomial-time computable relation R , with instance-independent commitments. We say that Π satisfies q -bounded instance-independent ϵ -computational honest-verifier zero-knowledge (instance-independent (q, ϵ) -HVZK for short) if there exists a PPT simulator $\mathcal{S} := (\mathcal{S}', (\mathcal{S}_0'', \mathcal{S}_1''))$ such that for all binary PPT adversaries \mathcal{A} we have that $\text{rIPS}_{\mathcal{A}}^{\Pi}(\lambda) \approx_{\epsilon} \text{sIPS}_{\mathcal{S}', (\mathcal{S}_0'', \mathcal{S}_1''), \mathcal{A}}^{\Pi}(\lambda)$, where experiments rIPS and sIPS are defined below:*

$\text{rIPS}_{\mathcal{A}}^{\Pi}(\lambda)$	$\text{PROV}(\mathbf{r}_P, \mathbf{r}_V, \text{crs}, x, w, b, i)$	$\text{SIMU}(\mathbf{r}_S, \text{crs}, \text{tk}, x, w, b, i)$
for $i = 1, \dots, q(\lambda)$ do $\mathbf{r}_P[i] \leftarrow_{\$} \{0, 1\}^{P.r(\lambda)}$ $\mathbf{r}_V[i] \leftarrow_{\$} \{0, 1\}^{V.r(\lambda)}$ $\text{crs} \leftarrow_{\$} K(1^\lambda)$ return $\mathcal{A}^{\text{PROV}(\mathbf{r}_P, \mathbf{r}_V, \text{crs}, \cdot, \cdot, \cdot)}(1^\lambda, \text{crs})$	if $R(x, w) = 0 \vee i \notin [q(\lambda)]$ then return \perp if $T[i, b] \neq \perp$ then return $T[i, b]$ if $b = 0$ then $\alpha \leftarrow P_0(\text{crs}; \mathbf{r}_P[i])$ $\beta \leftarrow V_0(\text{crs}; \mathbf{r}_V[i])$ // i.e. $\beta = \mathbf{r}_V[i]$ $\gamma \leftarrow \perp$ else $\alpha \leftarrow P_0(\text{crs}; \mathbf{r}_P[i])$ $\beta \leftarrow V_0(\text{crs}; \mathbf{r}_V[i])$ // i.e. $\beta = \mathbf{r}_V[i]$ $\gamma \leftarrow P_1(\text{crs}, x, w; \mathbf{r}_P[i])$ $T[i, b] \leftarrow (\alpha, \beta, \gamma)$ return (α, β, γ)	if $R(x, w) = 0 \vee i \notin [q(\lambda)]$ then return \perp if $T[i, b] \neq \perp$ then return $T[i, b]$ if $b = 0$ then $(\alpha, \beta) \leftarrow \mathcal{S}_0''(\text{crs}, \text{tk}; \mathbf{r}_S[i])$ $\gamma \leftarrow \perp$ else $(\alpha, \beta, \gamma) \leftarrow \mathcal{S}_1''(\text{crs}, \text{tk}, x; \mathbf{r}_S[i])$ $T[i, b] \leftarrow (\alpha, \beta, \gamma)$ return (α, β, γ)
$\text{sIPS}_{\mathcal{S}', (\mathcal{S}_0'', \mathcal{S}_1''), \mathcal{A}}^{\Pi}(\lambda)$ for $i = 1, \dots, q(\lambda)$ do $\mathbf{r}_S[i] \leftarrow_{\$} \{0, 1\}^{\mathcal{S}.r(\lambda)}$ $(\text{crs}, \text{tk}) \leftarrow_{\$} \mathcal{S}'(1^\lambda)$ return $\mathcal{A}^{\text{SIMU}(\mathbf{r}_S, \text{crs}, \text{tk}, \cdot, \cdot, \cdot)}(1^\lambda, \text{crs})$		

and \mathcal{A} can ask $q(\lambda)$ queries to its oracle and outputs a single bit. Additionally we say that Π satisfies instance-independent ϵ -HVZK if indistinguishability of the above experiments holds for an arbitrary polynomial $q(\lambda)$.

Note that a standard q -bound variant (i.e., without instance-independence) is obtained by fixing bit b in oracles PROV and SIMU to 1 and provide P_0 with x and w as additional input, as then the oracles either return a complete honest transcript or a complete simulated transcript.

One example of a protocol readily satisfying (unbounded) instance-independent HVZK is given by the Lapidot-Shamir protocol for graph Hamiltonicity [LS91].

We are now in a position to quantify the zero-knowledge property of our instantiation of the FS collapse.

Lemma 4.10 *Let $\Pi = (\mathsf{K}, \mathsf{P}, \mathsf{V})$ be a 3PC argument system for a polynomial-time computable relation R , such that Π has instance-independent commitments and satisfies q -bounded instance-independent ϵ -HVZK. Then, assuming H is a programmable q -wise independent hash function, the FS collapse $\overline{\Pi}_{\mathsf{FS}, \mathsf{H}} = (\mathsf{K}_{\mathsf{FS}}, \mathsf{P}_{\mathsf{FS}}, \mathsf{V}_{\mathsf{FS}})$ of Π using H satisfies q -bounded ϵ -computational zero-knowledge.*

Proof. For ease of notation let us write $\overline{\Pi} := \overline{\Pi}_{\mathsf{FS}, \mathsf{H}}$. We start with the real distribution $\mathsf{rNIZK}_{\overline{\Pi}}^{\mathcal{A}}(\lambda)$, where the CRS is defined as $\overline{\mathsf{crs}} = (\mathsf{crs}, \mathsf{hk})$ for $\mathsf{crs} \leftarrow_{\mathcal{S}} \mathsf{K}(1^\lambda)$ and $\mathsf{hk} \leftarrow_{\mathcal{S}} \mathsf{H.KGen}(1^\lambda)$. Here the adversary \mathcal{A} , given the CRS, can ask q adaptive queries (x_i, w_i) to oracle PROV which replies with $\pi \leftarrow_{\mathcal{S}} \mathsf{P}_{\mathsf{FS}}(\overline{\mathsf{crs}}, x_i, w_i)$ (provided that $R(x_i, w_i) = 1$).

We describe a series of computationally close hybrids, starting from $\mathsf{rNIZK}_{\overline{\Pi}}^{\mathcal{A}}(\lambda)$. The hybrids are depicted in Fig. 2 and are described below.

$\mathsf{rNIZK}_1(\lambda)$: This is identical to the real experiment, but now the randomness r_i used to generate the q proofs π_i corresponding to \mathcal{A} 's queries is pre-sampled. Additionally, values $\alpha_i = \mathsf{P}_0(1^\lambda; r_i)$ are pre-computed—note that this is possible because of instance-independent commitments—which allows us to also pre-compute values β_i as $\beta_i = \mathsf{H.Eval}(\mathsf{hk}, \alpha_i)$ where hk is the hash key. All of these values are stored in a trapdoor $\overline{\mathsf{tk}}$ which is given to the hybrid oracle (which is now a mixture between PROV and SIMU). We write $\mathsf{P.rl}$ for the length of the random tape required for $(\mathsf{P}_0, \mathsf{P}_1)$. Note that the PROV oracle now additionally takes as input $\overline{\mathsf{tk}} = (\beta, r_1, \dots, r_q)$ and uses r_i as random tape of both P_0 and P_1 (and thus of P_{FS}).

Observe that all of the above steps (pre-computing values) are clearly just syntactical changes, and thus $\mathsf{rNIZK}_{\overline{\Pi}}^{\mathcal{A}}(\lambda) \equiv \mathsf{rNIZK}_1(\lambda)$.

q-wise

$\mathsf{rNIZK}_2(\lambda)$: In the second and last step we replace all values β_i with uniformly random values sampled from the range of hash function H and the key hk is chosen via programming the hash function. Down to the programmable q -wise independence property of H the distribution corresponding to $\mathsf{rNIZK}_1(\lambda)$ and $\mathsf{rNIZK}_2(\lambda)$ are identical, i.e., $\mathsf{rNIZK}_2(\lambda) \equiv \mathsf{rNIZK}_1(\lambda)$.

Simulator. Let $\mathcal{S} = (\mathcal{S}', \mathcal{S}'' = (\mathcal{S}'_0, \mathcal{S}'_1))$ be the instance-independent HVZK simulator for the underlying 3PC protocol (cf. Definition 4.9). Let q denote a bound on the number of oracle queries of adversary \mathcal{A} . We write $\mathcal{S}''.\mathsf{rl}$ to denote the length of the random tape required for algorithm \mathcal{S}'' .

We construct simulator $\mathsf{Sim} = (\mathsf{Sim}', \mathsf{Sim}'')$ for q -bounded computational zero-knowledge (cf. Definition 3.6). Simulator Sim' first runs \mathcal{S}' in order to obtain a pair $(\mathsf{crs}, \mathsf{tk})$. Afterwards Sim' chooses q random strings r_1, \dots, r_q of length $|r_i| = \mathcal{S}''.\mathsf{rl}(\lambda)$. It then runs \mathcal{S}''_0 on randomness r_i to obtain q pairs (α_i, β_i) , and chooses a “programmed” hash key $\mathsf{hk} \leftarrow_{\mathcal{S}} \mathsf{H.KGen}(1^\lambda, \alpha, \beta)$. It outputs key hk together with crs as common reference string $\overline{\mathsf{crs}}$ for the FS collapse, and tk together with the list of all r_i 's as trapdoor $\overline{\mathsf{tk}}$. It thus perfectly simulates the setup in rNIZK_2 .

Upon input $\overline{\mathsf{crs}}$ and $\overline{\mathsf{tk}}$, together with the i -th query x_i , simulator Sim'' extracts randomness r_i from tk and calls \mathcal{S}'' on input $(\mathsf{crs}, \mathsf{tk}, x_i)$ and with random coins r_i to obtain a proof $\pi_i = (\alpha, \beta, \gamma)$ which it returns. We give the pseudocode of $\mathsf{Sim} = (\mathsf{Sim}', \mathsf{Sim}'')$ in Figure 3. Note that the adversary can only make q queries to its oracle and we provide index i as explicit input to simulator Sim'' .

pre-compute values

<p style="text-align: center;">$r\text{NIZK}_1(\lambda)$ $\xrightarrow{q\text{-wise independence}}$ $r\text{NIZK}_2(\lambda)$</p> <hr style="border: 0.5px solid black;"/> <pre> 1 : crs \leftarrow $\mathcal{K}(1^\lambda)$ 2 : hk \leftarrow $\mathcal{H}.\text{KGen}(1^\lambda)$ 3 : for $i = 1, \dots, q$ do 4 : $\mathbf{r}[i] \leftarrow \mathcal{S}\{0, 1\}^{\text{P.r}(\lambda)}$ 5 : $\alpha[i] \leftarrow \text{P}_0(1^\lambda; \mathbf{r}[i])$ 6 : $\beta[i] \leftarrow \mathcal{H}.\text{Eval}(\text{hk}, \alpha[i])$ 7 : 8 : $\overline{\text{tk}} \leftarrow (\beta, \mathbf{r})$ 9 : $\overline{\text{crs}} \leftarrow (\text{crs}, \text{hk})$ 10 : return $\mathcal{A}^{\text{PROV}(\overline{\text{crs}}, \overline{\text{tk}}, \cdot, \cdot)}(1^\lambda, \overline{\text{crs}})$ </pre>	<pre> crs \leftarrow $\mathcal{K}(1^\lambda)$ for $i = 1, \dots, q$ do $\mathbf{r}[i] \leftarrow \mathcal{S}\{0, 1\}^{\text{P.r}(\lambda)}$ $\alpha[i] \leftarrow \text{P}_0(1^\lambda; \mathbf{r}[i])$ $\beta[i] \leftarrow \mathcal{S}\{0, 1\}^{\mathcal{H}.\text{ol}(\lambda)}$ $\text{hk} \leftarrow \mathcal{H}.\text{KGen}(1^\lambda, \alpha, \beta)$ $\overline{\text{tk}} \leftarrow (\beta, \mathbf{r})$ $\overline{\text{crs}} \leftarrow (\text{crs}, \text{hk})$ return $\mathcal{A}^{\text{PROV}(\overline{\text{crs}}, \overline{\text{tk}}, \cdot, \cdot)}(1^\lambda, \text{crs})$ </pre>	<p style="text-align: center;">$\text{PROV}(\overline{\text{crs}}, \overline{\text{tk}}, x_i, w_i)$</p> <hr style="border: 0.5px solid black;"/> <pre> if $R(x_i, w_i) = 1$ then $(\text{crs}, \text{hk}) \leftarrow \overline{\text{crs}}$ $\overline{\text{tk}} \leftarrow (\beta, \mathbf{r})$ $\alpha[i] \leftarrow \text{P}_0(1^\lambda; \mathbf{r}[i])$ $\gamma_i \leftarrow \text{P}_1(\text{crs}, x_i, w_i, \beta[i]; \mathbf{r}[i])$ return $\pi_i \leftarrow (\alpha[i], \gamma[i])$ else return \perp </pre>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 2: The game hops needed for the proof of Lemma 4.10. The highlighted lines mark those lines that change from step to step. By $\mathcal{H}.\text{KGen}(1^\lambda, \alpha, \beta)$ in line 7 of $r\text{NIZK}_2$ we denote the “programming” of the q -wise independent hash function such that $\mathcal{H}.\text{Eval}(\text{hk}, \alpha[i]) = \beta[i]$ for all $i \in [q]$. Note that by choosing \mathcal{H} to be a $(q - 1)$ degree polynomial over an appropriate finite field we can do the programming via polynomial interpolation.

<p style="text-align: center;">$\text{Sim}'(1^\lambda)$</p> <hr style="border: 0.5px solid black;"/> <pre> $(\text{crs}, \text{tk}) \leftarrow \mathcal{S}'(1^\lambda)$ for $i = 1, \dots, q(\lambda)$ do $\mathbf{r}[i] \leftarrow \mathcal{S}\{0, 1\}^{\mathcal{S}''.\text{r}(\lambda)}$ $(\alpha[i], \beta[i]) \leftarrow \mathcal{S}''_0(1^\lambda; \mathbf{r}[i])$ $\text{hk} \leftarrow \mathcal{H}.\text{KGen}(1^\lambda, \alpha, \beta)$ $\overline{\text{crs}} := (\text{crs}, \text{hk})$ $\overline{\text{tk}} := (\text{tk}, \mathbf{r})$ return $(\overline{\text{crs}}, \overline{\text{tk}})$ </pre>	<p style="text-align: center;">$\text{Sim}''(\overline{\text{crs}}, \overline{\text{tk}}, x, i)$</p> <hr style="border: 0.5px solid black;"/> <pre> if $i \notin [q(\lambda)]$ then return \perp if $T[i] \neq \perp$ then return $T[i]$ $(\text{crs}, \text{hk}) \leftarrow \overline{\text{crs}}$ $\overline{\text{tk}} \leftarrow (\text{tk}, \mathbf{r})$ $(\alpha, \beta, \gamma) \leftarrow \mathcal{S}''(\text{crs}, \text{tk}, x; \mathbf{r}[i])$ $T[i] \leftarrow (\alpha, \beta, \gamma)$ return (α, β, γ) </pre>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 3: The HVZK simulator for the proof of Lemma 4.10.

Analysis. Note that simulator Sim' can pre-compute the values α_i using the HVZK simulator \mathcal{S}''_0 of the underlying 3PC protocol. This is because \mathcal{S}'' is instance-independent. It remains to show that for all q -query adversaries \mathcal{A} , the distributions $r\text{NIZK}_2(\lambda)$ and $s\text{NIZK}_{\text{Sim}'\text{Sim}''}(\lambda)$ are computationally close which follows by q -bounded HVZK of the initial 3PC protocol. For this note that games $r\text{NIZK}_2(\lambda)$ and $s\text{NIZK}_{\text{Sim}'\text{Sim}''}(\lambda)$ differ only in how values α, β and γ are computed. In game $r\text{NIZK}_2(\lambda)$ they are computed with the honest prover while in game $s\text{NIZK}_{\text{Sim}'\text{Sim}''}(\lambda)$ they are computed using the instance-independent HVZK simulator \mathcal{S} . Hence, an adversary against the instance-independent HVZK property of the underlying protocol can perfectly simulate games $r\text{NIZK}_2(\lambda)$ and $s\text{NIZK}_{\text{Sim}'\text{Sim}''}(\lambda)$, by running the steps of $r\text{NIZK}_2(\lambda)$ and using its oracle to obtain α_i, β_i and later using its oracle to complete proofs and obtain γ_i . If the adversary is connected to PROV it perfectly simulates game $r\text{NIZK}_2(\lambda)$ and otherwise it perfectly simulates $s\text{NIZK}_{\text{Sim}'\text{Sim}''}(\lambda)$. Thus, if the underlying protocol is instance-independent (q, ϵ) -HVZK then

$$|\Pr[r\text{NIZK}_2(\lambda) = 1] - \Pr[s\text{NIZK}_{\text{Sim}'\text{Sim}''}(\lambda) = 1]| \stackrel{\text{asym}}{\leq} \epsilon(\lambda).$$

4.4 Putting it Together

Combining the results in the previous section we obtain the following theorem stating that the FS transform is instantiable with a q -wise independent hash function given that the underlying 3PC satisfies three properties **P1-P3**.

Theorem 4.11 *Let $\Pi = (\mathsf{K}, \mathsf{P}, \mathsf{V})$ be a 3PC argument system for a polynomial-time computable relation R . Let H be a programmable q -wise independent hash function. Assume that Π is c -complete and s -sound and additionally satisfies the following three properties:*

- (**P1**) *instance-independent commitment;*
- (**P2**) *SEGR $\varrho < 1$;*
- (**P3**) *instance-independent (q, ϵ) -HVZK.*

Then, the FS collapse $\bar{\Pi}_{\mathsf{FS}, \mathsf{H}}$ of Π using H is a (c, s', q, ϵ) -NIZK for the relation R , with $s' = \varrho$.

4.5 Obtaining the Required Properties

It remains the question which 3PC arguments (if any) satisfy properties **P1-P3**. While we know of protocols directly satisfying **P1** (e.g., the Blum protocol [Blu81]), and of at least one candidate satisfying both **P1** and **P3** (i.e., the Lapidot-Shamir protocol [LS91]), we do not know any 3PC argument already satisfying all properties.

Hence, we turn to the question how to compile a 3PC argument into one satisfying all the properties we need. We do so using two compilers, as outlined below:

- Given a 3PC argument satisfying **P1** and **P3**, we show a compiler yielding a 3PC argument that additionally satisfies **P2**, that is, it has a small soundness-error-to-guessing ratio while retaining properties **P1** and **P3**. This compiler requires a CRS and relies heavily on indistinguishability obfuscation (which we formally introduce in Section 3.4), and is presented in details in Section 5.
- Given a 3PC argument satisfying **P1** and special HVZK (properties, for example, present in the classical 3PC argument for quadratic residuosity due to Blum [Blu81]), we show a compiler yielding a 3PC argument that additionally satisfies **P3**, that is, it has instance-independent HVZK while retaining property **P1**. This compiler—which is presented in details in Section 6—also requires a CRS, and is inspired by the recent work of Lindell [Lin15]. It relies on so-called *dual-mode commitments* [CV07] which can be set up either to be perfectly binding or to be equivocal.

Intuitively, the protocol is changed such that the prover, instead of sending α , sends a commitment c to α which it opens in the last message. As in an honest setup the commitment is perfectly binding, soundness and completeness follow easily; for zero-knowledge the simulator can setup the commitment scheme such that it is equivocal, which allows it to choose its first message as a simulated commitment and later open this to the message α as obtained by the underlying simulator. Note that, in particular, this allows the simulator to choose its first message independently of the instance x .

Open questions are whether we can similarly find compilers that do not require a CRS and whether there exist compilers also for protocols which do not already satisfy **P1** and special HVZK.

5 Obtaining Small Soundness-Error-to-Guessing Ratio

In this section we present a compiler that turns a 3PC argument (possibly in the CRS model) with instance-independent commitments and HVZK (Definition 4.7) into a 3PC argument which has the soundness-error-to-guessing ratio (Definition 4.5) needed for the complexity leveraging in Lemma 4.8. We note that the resulting protocol will be in the CRS model regardless whether the starting protocol is in the CRS model or in the standard model. The idea for the compiler is to provide a mechanism that allows to produce many challenges β given only a single commitment α . To this effect the CRS will contain two obfuscated circuits to help the prover and the verifier run the protocol. For obfuscation we use an indistinguishability obfuscator (which we formally introduce in Section 3.4). The first circuit C_0 is used by the prover to generate a *pre-commitment* α^* which it sends over to the verifier. The verifier will then use the second circuit C_1 and run it on α^* to obtain multiple commitments. For this $C_1[k, \text{crs}]$ has a PRF key and the crs for algorithm P_0 of the underlying protocol hardcoded, and computes ℓ commitments as follows:

```


$$\frac{C_1[k, \text{crs}](\alpha^*)}{\text{for } i = 1, \dots, \ell \text{ do}$$


$$r^* \leftarrow \text{F.Eval}(k, \alpha^* + i)$$


$$\alpha[i] \leftarrow P_0(\text{crs}; r^*)$$


$$\text{return } \alpha$$


```

Using C_1 the compiled verifier V^* can generate ℓ real commitments $\alpha[1]$ to $\alpha[\ell]$ given the single (short) pre-commitment α^* . The verifier will then run the underlying verifier V on all these commitments to receive $\beta_1, \dots, \beta_\ell$ which it sends back to the prover.

In order to correctly continue the prover's computation (which was started on the verifier's side) the compiled prover P^* needs to somehow obtain the randomnesses r^* used within C_1 . For this, we will build a backdoor into C_1 which allows to obtain the randomness r^* if one knows the randomness that was used to generate α^* . Once the prover has recovered randomnesses r_1^*, \dots, r_ℓ^* it can run the underlying prover P on this randomness and the corresponding challenges β_i to get correct values γ_i which it sends back to the verifier. In a final step verifier V^* runs the original verifier on the implicit transcripts $(\alpha_i, \beta_i, \gamma_i)_{i=1, \dots, \ell}$ and returns 1 if and only if the original verifier returns 1 on all the transcripts.

We will next present a formal description of the compiler and then show that it achieves the claimed properties and retains soundness, completeness and zero-knowledge.

5.1 The Compiler

Let $\Pi = (\mathbf{K}, \mathbf{P}, \mathbf{V})$ be a 3PC argument system where the prover generates instance-independent commitments and that satisfies instance-independent HVZK. Let rl denote an upper bound on the randomness used by the prover (i.e., $\mathbf{P.rl}$) and HVZK simulator (i.e., $\mathcal{S.rl}$). Let F_1 be a puncturable pseudorandom function which is length doubling. Let F_2 be a puncturable pseudorandom function with $F_2.\text{il} = F_1.\text{ol}$ and with $F_2.\text{ol} = \text{rl}$. Let ℓ be a polynomial. We construct an argument system $\Pi^* = (\mathbf{K}^*, \mathbf{P}^*, \mathbf{V}^*)$ in the CRS model as follows. On input the security parameter \mathbf{K}^* will construct an obfuscation of the following two circuits:





Figure 4: The compiled protocol from Section 5.1 to turn a 3PC protocol into one that has a small soundness-error-to-guessing ratio (in the CRS model).



Note that we assume that the underlying protocol is in the CRS model and has a setup algorithm \mathbf{K} . If this is not the case one recovers the transformation for a 3PC in the standard model by assuming that \mathbf{K} outputs the empty string ε . The compiled 3PC $\Pi^* = (\mathbf{K}^*, \mathbf{P}^*, \mathbf{V}^*)$ is then constructed as in Figure 4.

5.2 Security Analysis

It remains to show that the compiled protocol is computationally sound, achieves (bounded) instance-independent HVZK, is complete, and that it has instance-independent commitments and a sufficient soundness-error-to-guessing ratio:

Theorem 5.1 *Let $\Pi = (K, P, V)$ be a 3PC argument system for a polynomial-time computable relation R such that Π is c -complete and s -sound and has instance-independent commitments and satisfies q -bounded instance-independent HVZK. Let iO be an indistinguishability obfuscator and F_1 and F_2 puncturable pseudorandom functions. Let ℓ be a polynomial. Then, in the CRS model, the compiled protocol $\Pi^* = (K^*, P^*, V^*)$ is $(\ell \cdot c)$ -complete, $(2 \cdot s^{-\ell} + 2^{F_1 \cdot \text{il}(\lambda)} s^{-\ell})$ -sound, has a worst-case collision probability of $2^{-F_1 \cdot \text{il}(\lambda)}$, and satisfies q/ℓ -bounded instance-independent HVZK. Furthermore the compiled protocol has instance-independent commitments.*

We discuss each of these properties in turn.

Completeness. Consider an honest protocol execution of the compiled protocol that does not end with an accepting vote of the verifier. As the final verification uses the underlying verifier and also α, β , and γ are constructed using the underlying algorithms P and V this thus yields also an honest protocol execution of the underlying protocol where the verifier does not accept. There are two things to note: the compiled protocol “internally” runs the underlying protocol multiple times (ℓ times) and algorithm P_0 is run on pseudorandom coins rather than on truly random coins. Thus, if the underlying protocol is c -complete the compiled protocol is at most $(\ell \cdot c)$ -complete plus the distinguishing probability for the pseudorandom function.

Soundness. We analyze the soundness in two steps: Firstly, if the original protocol has soundness s , then its ℓ -parallel repetition version has soundness s^ℓ , as the protocol is public-coin [PV07, HPWP10, CL09, CP15]. In a second step, we consider that $\alpha[1], \dots, \alpha[\ell]$ are generated via applying the first stage of the prover P_0 to the output of an obfuscated pPRF rather than using truly uniformly random coins. We loose a factor of $2^{-|\alpha^*|}$ for the number of choices that the prover can make for the input to the obfuscated pPRF and else reduce to iO and the pPRF via a standard puncturing argument. Hence, when using iO that is $2^{-|\alpha^*| \ell \log s}$ -secure, and using a pPRF that is $2^{-|\alpha^*| \ell \log s}$ -secure, adding up soundness error s^ℓ , iO -security $2^{-|\alpha^*| \ell \log s}$ and pPRF-security $2^{-|\alpha^*| \ell \log s}$ and multiplying them all by $2^{|\alpha^*|}$, we obtain that the soundness of the compiled protocol is $2 \cdot s^{-\ell} + 2^{|\alpha^*|} s^{-\ell}$. Note that the length of α^* is independent from the number of repetition ℓ and thus, we can make soundness as small as we want.

Soundness-error-to-guessing ratio. For the collision probability of commitments (see Definition 4.5) note that the commitment of the compiled protocol α^* is generated by evaluating pseudorandom function F_1 on a random value τ . As α^* is independent of witness and instance we have that in the worst case every input $\tau \in F_1 \cdot \text{il}(\lambda)$ is mapped to a different output which gives us a worst-case collision probability of $2^{-F_1 \cdot \text{il}(\lambda)}$. Note that in particular, the worst-case collision-probability is independent of the soundness amplification parameter ℓ (yielding a SEGR $\varrho(\lambda)$ which is bounded away from 1).

Instance-independence. The compiled protocol has instance independent commitments (see Definition 4.7). Furthermore, the protocol retains instance-independent HVZK simulators as we discuss below.

Zero-knowledge. Finally, we show that the compiled protocol satisfies instance-independent q/ℓ -bounded HVZK if the underlying protocol satisfies instance-independent q -bounded HVZK (see Definition 4.9). We show the proof for the simplified setting where the amplification parameter ℓ is set to 1. The proof for the general case is analogous. For the argument we assume the existence of

an injective one-way function po with domain $\text{po.il} = F_1.\text{il}$. The one-way function will be used as a simple point obfuscation scheme. To obfuscate a point x we store $\bar{p}_x \leftarrow \text{po}(x)$ which allows us to later compare a point x' by simply checking if $\bar{p}_x = \text{po}(x')$.

Let $\Pi = (\mathbf{K}, \mathbf{P}, \mathbf{V})$ be a 3PC protocol which satisfies the conditions for the compiler and which is instance-independent HVZK. We need to show that the compiled protocol $\Pi^* = (\mathbf{K}^*, \mathbf{P}^*, \mathbf{V}^*)$ achieves bounded instance-independent HVZK in the CRS model. For this we will use several game hops where the first game is the real world setting rIPS and the last game is identical to the simulated setting sIPS (with a simulator to be defined). We first describe the individual games and present the accompanying pseudocode in Figures 7, 8, and 9 (starting on page 47). We present a formal analysis of the individual game hops after the description of all games and denote the reduction target for each step next to the game description.

Game₁(λ): The game is identical to the real world setting rIPS where the crs is honestly generated by \mathbf{K}^* and the adversary on querying oracle PROV on an instance x in the language with witness w , obtains a transcript of an honest execution between prover $\mathbf{P}^*(1^\lambda, \text{crs}, x, w)$ and verifier $\mathbf{V}^*(1^\lambda, \text{crs}, x)$. Without loss of generality we assume that prover \mathbf{P}^* on the i -th query uses random coins τ_i sampled uniformly at random from $\{0, 1\}^{\text{P.r}(\lambda)} = \{0, 1\}^{F_1.\text{il}(\lambda)}$.

iO

Game₂(λ): The game is identical to the previous game except that setup \mathbf{K}^* now punctures k_1 on all values τ . For this it chooses q random values τ_1, \dots, τ_q in $\{0, 1\}^{F_1.\text{il}(\lambda)}$ (if the values are not distinct we abort and define that the adversary wins). Additionally simple forms of “point obfuscations” are generated for all q points τ_i by running each τ_i through one-way function po to obtain

$$\bar{p}_{\tau_i} \leftarrow \text{po}(\tau_i).$$

It then punctures key k_1 on values τ_1, \dots, τ_q to receive punctured key k_1^* and hardcodes this into both circuits. To not change the functionality on these circuits the original PRF values are hardcoded, that is $\alpha_i^* \leftarrow F_1.\text{Eval}(k_1, \tau_i)$ is hardcoded into both circuits C_0 and C_1 . The branching and test operation whether an input is for a punctured value is done with the point obfuscation.

In addition, on the i -th call to oracle PROV prover \mathbf{P}^* is run on randomness τ_i . Note that τ_i is the only randomness of the prover and prover \mathbf{P}^* then choses its commitment as α_i .

pPRF
bad α^*

Game₃(λ): The game is identical to the previous game except that now values α_i^* are chosen uniformly at random in $\{0, 1\}^{F_1.\text{ol}(\lambda)}$.

Game₄(λ): The game is identical to the game before except that we define that the adversary wins in case any of the values $\alpha^*[i]$ is chosen to be in the image of the PRF for key k_1 . Note that since F_1 is length doubling this happens only with probability $q \cdot 2^{-\lambda}$.

iO

Game₅(λ): As before but we again use the unpunctured key k_1 . The hardcoded values for τ_1, \dots, τ_q (i.e. their point obfuscations) as well as the α_i^* remain hardcoded.

diO+BCP14+OWF

Game₆(λ): We now change the if branch in circuit C_1 such that it does not depend on hardcoded values \bar{p}_τ anymore. Note that the if-branch was changed in the second step in order not to change the functionality of the circuits. This step changes the functionality by reverting the if-branch to check only whether input α^* is different from $F_1.\text{Eval}(k_1, \tau)$ (i.e., the original check).

\downarrow iO
 \downarrow pPRF
 \downarrow iO
 \downarrow hvzk
Game₇(λ): The game is identical to the previous game but now key k_2^* is punctured on values $\alpha_1^*, \dots, \alpha_q^*$. To not change the functionality on these circuits the original PRF values are hardcoded, that is $\mathbf{r}^*[i] \leftarrow F_2.\text{Eval}(k_2, \alpha_i^*)$ is hardcoded into circuit C_1 .
Game₈(λ): The game is identical to the previous game but now hardcoded values r_i^* are sampled uniformly at random.
Game₉(λ): In this step we precompute the actual values α instead of having values \mathbf{r}^* hardcoded. That is, we change circuit K^* to compute additionally $\alpha[i] \leftarrow P_0(\text{crs}; \mathbf{r}^*[i])$ and hardcode these into C_1 .
Game₁₀(λ): Let $\mathcal{S} = (\mathcal{S}', \mathcal{S}'' = (\mathcal{S}''_0, \mathcal{S}''_1))$ denote the honest verifier zero knowledge simulator of the underlying 3PC protocol. The setting is identical to the previous game, but now we “switch to oracle SIMU”. The crs is generated as before with three exceptions: (1) the simulator of the underlying protocol \mathcal{S}' is run to generate a CRS and trapdoor crs and tk for the underlying protocol; (2) Values α are computed with simulator \mathcal{S}''_0 of the underlying protocol (with coins $\mathbf{r}^*[i]$); (3) As we are now in the simulated setting we generate the trapdoor $\overline{\text{tk}}$ to contain the randomness values r_1^*, \dots, r_q^* as well as pre-commitments $\alpha_1^*, \dots, \alpha_q^*$.

On the i -th query $(\text{crs}, \text{tk}, x, w, b, i)$ oracle SIMU answers as follows: if $b = 0$ then it runs simulator \mathcal{S}''_0 of the underlying 3PC protocol on input its CRS and trapdoor and on randomness r_i to obtain α and β (which are identical as the precomputed ones) which are returned to the adversary. If, on the other hand $b = 1$ it runs simulator \mathcal{S}' of the underlying 3PC protocol on input its CRS and trapdoor as well as instance x and with randomness r_i^* to obtain a proof $\pi = (\alpha, \beta, \gamma)$. Note that by definition, as \mathcal{S} has instance-independent commitments \mathcal{S} will generate a commitment and challenge as $(\alpha, \beta) \leftarrow \mathcal{S}''_0(\text{crs}, \text{tk}, r_i^*)$ (where crs and tk are the CRS and trapdoor by the underlying protocol). It replaces α for $\alpha^*[i]$ and returns proof $\pi^* := (\alpha^*[i], \beta, \gamma)$.

We give a complete description of simulator $\text{Sim} = (\text{Sim}', \text{Sim}'' = (\text{Sim}''_0, \text{Sim}''_1))$ for the compiled protocol as pseudocode in Figure 5.

Analysis. First note that simulator Sim is indeed instance-independent. What remains to show is that the view of the adversary between Game₁ and the final game Game₁₀ does only negligibly change. We discuss each step in turn below.

Game₁ to Game₂. By construction the generated circuits C_0 and C_1 are equivalent in both games and hence a distinguisher for the two games yields a distinguisher against the indistinguishability obfuscator. Note that we are not dealing with a single circuit but two circuits and thus we have that

$$|\Pr[\text{Game}_1(\lambda) = 1] - \Pr[\text{Game}_2(\lambda) = 1]| \leq 2 \cdot \text{Adv}_{\text{iO}, \text{Sam}, \text{D}}^{\text{iO}}(\lambda)$$

where sampler Sam and distinguisher D are the adversary induced by games Game₁ and Game₂.

Game₂ to Game₃. By construction the only change is that the answers on “punctured points” are now chosen as uniformly random values. Thus, a distinguisher between the games induces a distinguisher against the puncturable pseudorandom function F_1 .

$$|\Pr[\text{Game}_2(\lambda) = 1] - \Pr[\text{Game}_2(\lambda) = 1]| \leq \text{Adv}_{F_1, \mathcal{A}_1, \mathcal{A}_2}^{\text{pprf}}(\lambda).$$

$\text{Sim}'(1^\lambda)$ $(\text{crs}, \text{tk}) \leftarrow \mathcal{S}'(1^\lambda)$ $k_1 \leftarrow \mathcal{F}_1.\text{KGen}(1^\lambda)$ $k_2 \leftarrow \mathcal{F}_2.\text{KGen}(1^\lambda)$ for $i = 1, \dots, q$ do $\tau[i] \leftarrow \mathcal{S}\{0, 1\}^{\text{F.il}(\lambda)}$ $\bar{\mathbf{p}}_\tau[i] \leftarrow \text{po}(\tau[i])$ $\alpha^*[i] \leftarrow \mathcal{S}\{0, 1\}^{\text{F.ol}(\lambda)}$ $\mathbf{r}^*[i] \leftarrow \mathcal{S}\{0, 1\}^{\text{F.ol}(\lambda)}$ $(\alpha[i], \beta_{\text{tmp}}) \leftarrow \mathcal{S}_0''(\text{crs}, \text{tk}; \mathbf{r}^*[i])$ $k_2^* \leftarrow \mathcal{F}_2.\text{Pntr}(k_2, \{\alpha^*[1], \dots, \alpha^*[q]\})$ $\bar{C}_0 \leftarrow \mathcal{S}\text{iO}(C_0[k_1, \bar{\mathbf{p}}, \alpha^*])$ $\bar{C}_1 \leftarrow \mathcal{S}\text{iO}(C_1[k_1, \alpha^*, k_2^*, \alpha])$ $\bar{\text{crs}} \leftarrow (\text{crs}, \bar{C}_0, \bar{C}_1)$ $\bar{\text{tk}} \leftarrow (\text{tk}, \alpha^*, \mathbf{r}^*)$ return $(\bar{\text{crs}}, \bar{\text{tk}})$	$\text{Sim}_0''(1^\lambda, \bar{\text{crs}}, \bar{\text{tk}})$ $(\text{crs}, \bar{C}_0, \bar{C}_1) \leftarrow \bar{\text{crs}}$ $(\text{tk}, \alpha^*, \mathbf{r}^*) \leftarrow \bar{\text{tk}}$on i -th query..... $(\alpha, \beta) \leftarrow \mathcal{S}_0''(\text{crs}, \text{tk}; \mathbf{r}^*[i])$ return $(\alpha[i], \beta)$ $\text{Sim}_1''(1^\lambda, \bar{\text{crs}}, \bar{\text{tk}}, x)$ $(\text{crs}, \bar{C}_0, \bar{C}_1) \leftarrow \bar{\text{crs}}$ $(\text{tk}, \alpha^*, \mathbf{r}^*) \leftarrow \bar{\text{tk}}$on i -th query..... $\gamma \leftarrow \mathcal{S}_1''(\text{crs}, \text{tk}, x; \mathbf{r}^*[i])$ return γ	$C_0[k_1, \bar{\mathbf{p}}, \alpha^*](\tau)$ if $\exists i \in [q] : \bar{\mathbf{p}}[i] = \text{po}(\tau)$ then return $\alpha^*[i]$ $\alpha^* \leftarrow \mathcal{F}_1.\text{Eval}(k_1, \tau)$ return α^* $C_1[k_1, \bar{\mathbf{p}}, \alpha^*, k_2^*, \alpha](\alpha^*, \tau)$ if $\exists i \in [q] : \alpha^* = \alpha^*[i]$ $\alpha \leftarrow \alpha[i]$ else $r^* \leftarrow \mathcal{F}_2.\text{Eval}(k_2^*, \alpha^*)$ $\alpha \leftarrow \mathcal{P}_0(r^*)$ if $\alpha^* \neq \mathcal{F}_1.\text{Eval}(k_1, \tau)$ then $r^* \leftarrow \perp$ return (α, r^*)
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 5: The complete pseudocode for q -bounded HVZK simulator Sim for the compiled protocol. Note that simulator Sim_1'' runs the underlying simulator on the same random coins as simulator Sim_0'' thus “generating” the same values (α, β) .

Game₃ to Game₄. By the fundamental lemma of the game playing technique [BR06] games **Game₃** and **Game₄** are identical unless event bad_{α^*} occurs which we can upper bound by $q \cdot 2^{-\text{F.il}(\lambda)}$ as the PRF by definition has stretch 2.

$$|\Pr[\text{Game}_2(\lambda) = 1] - \Pr[\text{Game}_2(\lambda) = 1]| \leq q \cdot 2^{-\text{F.il}(\lambda)}.$$

Game₄ to Game₅. Noting that from **Game₄** to **Game₅** the circuits only change syntactically but not functionally (the unpunctured key k_1 is used instead of k_1^*) allows us to perform an analysis analogously to the first game hop.

$$|\Pr[\text{Game}_4(\lambda) = 1] - \Pr[\text{Game}_5(\lambda) = 1]| \leq 2 \cdot \text{Adv}_{\text{IO}, \text{Sam}, \text{D}}^{\text{io}}(\lambda).$$

Game₅ to Game₆. We will reduce the distinguishing advantage of any distinguisher between the two games **Game₅** and **Game₆** to the security of the indistinguishability obfuscator iO and the security of the injective one-way function po . For this we rely on the result of Boyle *et al.* [BCP14] who relate indistinguishability obfuscation and restricted differing-inputs obfuscation. We recall their result as Theorem 3.10 on page 14. For this we consider a circuit sampler Sam that runs the steps of **Game₅** up to and including line 12. It outputs as auxiliary information circuit \bar{C}_0 and string crs , that is, it sets $\text{aux} \leftarrow (\bar{C}_0, \text{crs})$. As circuits it constructs the circuit C_1 from line 13 once as in **Game₅** and once as in **Game₆**.

Additionally we construct a diO distinguisher that gets as input an obfuscation \bar{C}_1 of either of the two circuits and the auxiliary information $\text{aux} \leftarrow (\bar{C}_0, \text{crs})$. It sets $\bar{\text{crs}} \leftarrow (\text{crs}, \bar{C}_0, \bar{C}_1)$ and then runs the game distinguisher on input $\bar{\text{crs}}$ and outputs whatever it outputs.

If \bar{C}_1 is as in **Game₅** then together sampler and distinguisher perfectly simulate game **Game₅** and otherwise they perfectly simulate **Game₆**. We want to argue that

$$|\Pr[\text{Game}_6(\lambda) = 1] - \Pr[\text{Game}_7(\lambda) = 1]| \leq \text{Adv}_{\text{IO}, \text{Sam}, \text{D}}^{\text{io}}(\lambda) + [\text{BCP14}] + q \cdot \text{Adv}_{\text{po}, \mathcal{A}}^{\text{owf}}(\lambda)$$

where [BCP14] denotes the loss from using the indistinguishability obfuscator in a diO setting and $\text{Adv}_{\text{po}, \mathcal{A}}^{\text{owf}}(\lambda)$ is a factor due to the use of one-way function po . For this we need to show that sampler Sam is differing-inputs and the circuits differ only on polynomially many points. First note that the two circuits, as prepared by sampler Sam , differ only on inputs $\tau[i]$ for $i \in [q]$. We can further reduce the advantage of any extractor in the Diff game to the inversion advantage $\text{Adv}_{\text{po}, \mathcal{A}}^{\text{owf}}(\lambda)$ of an adversary \mathcal{A} against one-way function po . For this consider that all values in τ are sampled uniformly at random and no extra information about these values is available in aux . Thus an adversary \mathcal{A} against the one-wayness of po that gets as input a random image y can simply choose $q - 1$ additional values $\text{po}(\tau_i)$ and construct auxiliary information and circuits as does sampler Sam but using y as one of the values in τ . An extractor that is successful in finding a differing input will with probability $1/q$ have inverted y which concludes the argument.

Game₆ to Game₇. Again, the circuits are only changed syntactically allowing for an analogous analysis as in the first game hop.

$$|\Pr[\text{Game}_6(\lambda) = 1] - \Pr[\text{Game}_7(\lambda) = 1]| \leq 2 \cdot \text{Adv}_{\text{IO}, \text{Sam}, \text{D}}^{\text{io}}(\lambda).$$

Game₇ to Game₈. Similarly to the second game hop the hardcoded values on punctured inputs are now chosen uniformly at random allowing for an analysis analogous to the second game hop.

$$|\Pr[\text{Game}_7(\lambda) = 1] - \Pr[\text{Game}_8(\lambda) = 1]| \leq \text{Adv}_{\text{F}_2, \mathcal{A}_1, \mathcal{A}_2}^{\text{pprf}}(\lambda).$$

Game₈ to Game₉. Precomputing values α for the punctured target points does not change the functionality of circuit C_1 and an analysis analogously to the first game hop allows us to reduce the distinguishing probability to the security of the indistinguishability obfuscator. Note that here only C_1 is changed and thus:

$$|\Pr[\text{Game}_8(\lambda) = 1] - \Pr[\text{Game}_9(\lambda) = 1]| \leq \text{Adv}_{\text{IO}, \text{Sam}, \text{D}}^{\text{io}}(\lambda).$$

Game₉ to Game₁₀. On the i -th query simulator Sim runs the simulator \mathcal{S} of the underlying 3PC protocol on random coins $\mathbf{r}^*[i]$ to obtain (α, β, γ) . As by assumption \mathcal{S} is instance-independent it will generate (α, β, γ) such that $(\alpha, \beta) = \mathcal{S}_0''(\text{crs}, \text{tk}; \mathbf{r}^*[i])$ where crs and tk are the CRS and trapdoor of the underlying simulator. Simulator Sim returns as protocol transcript $(\alpha^*[i], \beta, \gamma)$. By construction C_1 will map $\alpha^*[i]$ to α thus presenting the adversary with an identical simulation as in the previous step except that (α, β, γ) are now generated by the HVZK simulator \mathcal{S} . Thus a distinguisher between the two games can be used to construct a distinguisher against the HVZK simulator.

To construct an attacker against HVZK note that in line 8 of games Game_9 and Game_{10} the randomness \mathbf{r}^* is not encoded in any of the circuits any longer and only used for running P_0 and \mathcal{S}_0'' , respectively. Thus we construct an adversary \mathcal{A} that gets as input the crs and then runs game Game_9 up to but not including line 8. (The for loop is executed for the lines before line 8.) It then calls its oracle on all indexes $i = 1, \dots, q$ and with bit $b = 0$ (and x set to, for example, \perp , see also Definition 4.9). The oracle returns $(\alpha_i, \beta_i)_{i=1, \dots, q}$ which allows adversary \mathcal{A} to construct vector α as is done in line 9. It then runs the remaining steps of the setup to obtain $(\text{crs}, \overline{C}_0, \overline{C}_1)$ which it passes on to the distinguisher. Queries of the distinguisher are passed on to its own oracle and it outputs whatever the distinguisher outputs.

If the oracle implements PROV then the adversary perfectly simulates game Game_9 and otherwise it perfectly simulates Game_{10} . Thus, if the underlying protocol achieves (q, ϵ) -bounded instance-independent HVZK then

$$|\Pr[\text{Game}_8(\lambda) = 1] - \Pr[\text{Game}_9(\lambda) = 1]| \stackrel{\text{asym}}{\leq} \epsilon.$$

Note that for the general case with $\ell > 1$ one needs to simulate $q \cdot \ell$ oracle queries and thus loses a factor of $\ell\epsilon$.

6 Obtaining Instance-Independence

In this section, we present a compiler that turns a 3PC protocol with HVZK and instance-independent commitments into a 3PC protocol in the CRS model that has instance-independent commitments *and* instance-independent simulators, that is, the HVZK simulator produces α and β independently of the instance.

The idea is inspired by Lindell's compiler [Lin15]. Namely, we replace α by a commitment α^* to α where the deployed commitment scheme can come in one of two modes: if honestly generated the commitment will be *perfectly binding* thus allowing us to directly argue that the resulting compiled protocol retains soundness and completeness. On the other hand, the commitment scheme can be initialized to be *equivocal* (looking indistinguishably from the honest commitment setup) such that a simulator can open a commitment to arbitrary values. This way, the simulator can first commit to an arbitrary α^* and then, using the trapdoor in the CRS, it can open α^* to some arbitrary value α . In particular, in the reduction to the HVZK property, the verifier can choose α^* before knowing the statement that the simulator of the underlying protocol needs in order to produce α .

Such *dual-mode* commitment schemes were studied and constructed by Catalano and Visconti [CV07] who called them *hybrid commitments*. We here give the definition due to Lindell [Lin15]. We write $\text{C.Com}(\text{pp}, m)$ taking public parameters pp and message m to obtain a commitment c and opening δ , and write $\text{C.Vf}(\text{pp}, m, c, \delta)$ to denote verification of a commitment. We also let $\text{C.il} : \mathbb{N} \rightarrow \mathbb{N}$ (resp., $\text{C.ol} : \mathbb{N} \rightarrow \mathbb{N}$) denote the input (resp., output) length functions corresponding to C .

Definition 6.1 (Dual-Mode Commitment) *A dual-mode commitment scheme is a tuple of PPT algorithms $(\text{C.GenPP}, \text{C.Com}, \text{C.Vf}, \text{C.il}, \text{C.ol})$ with PPT commitment simulator $(\mathcal{S}_{\text{com}}.\text{GenPP}, \mathcal{S}_{\text{com}}.\text{Com}, \mathcal{S}_{\text{com}}.\text{Open})$ such that the following holds.*

PUBLIC PARAMETERS. *On input the security parameter 1^λ algorithm GenPP outputs public parameters pp .*

PERFECT COMPLETENESS. *For all security parameters $\lambda \in \mathbb{N}$ and all messages $m \in \{0, 1\}^{\text{C.il}(\lambda)}$ it holds that*

$$\Pr \left[\text{C.Vf}(\text{pp}, c, m, \delta) = 1 : \text{pp} \leftarrow_s \text{GenPP}(1^\lambda), (c, \delta) \leftarrow_s \text{C.Com}(\text{pp}, m) \right] = 1.$$

PERFECTLY BINDING. *For all security parameters $\lambda \in \mathbb{N}$ and all public parameters $\text{pp} \leftarrow_s \text{GenPP}(1^\lambda)$ algorithm C.Com is a perfectly-binding non-interactive commitment scheme, that is, for all security parameters $\lambda \in \mathbb{N}$, all messages $m, m' \in \{0, 1\}^{\text{C.il}(\lambda)}$ such that $m \neq m'$, and all openings $\delta' \in \{0, 1\}^*$ we have that*

$$\Pr \left[\text{C.Vf}(\text{pp}, c, m', \delta') = 0 : \text{pp} \leftarrow_s \text{GenPP}(1^\lambda), (c, \delta) \leftarrow_s \text{C.Com}(\text{pp}, m) \right] = 1$$

EQUIVOCAL. For every PPT adversary \mathcal{A} , advantage $\text{Adv}_{\mathcal{C}, \mathcal{S}_{\text{com}}, \mathcal{A}}^{\text{com}}(\lambda)$ defined as

$$\text{Adv}_{\mathcal{C}, \mathcal{S}_{\text{com}}, \mathcal{A}}^{\text{com}}(\lambda) := 2 \cdot \Pr \left[\text{COM}_{\mathcal{C}, \mathcal{S}_{\text{com}}}^{\mathcal{A}}(\lambda) \right] - 1$$

should be negligible where game COM is defined as

$\text{COM}_{\mathcal{C}, \mathcal{S}_{\text{com}}}^{\mathcal{A}}(\lambda)$	$\text{COM}(m)$
$b \leftarrow_{\$} \{0, 1\}$	if $b = 0$ then
$\text{pp}_0 \leftarrow_{\$} \text{C.GenPP}(1^\lambda)$	$(c, \delta) \leftarrow_{\$} \text{C.Com}(\text{pp}_0, m)$
$(\text{pp}_1, \text{tk}_{\text{com}}) \leftarrow_{\$} \mathcal{S}_{\text{com}}.\text{GenPP}(1^\lambda)$	else
$b' \leftarrow_{\$} \mathcal{A}^{\text{COM}}(\text{pp}_b)$	$c \leftarrow_{\$} \mathcal{S}_{\text{com}}.\text{Com}(\text{pp}_1, \text{tk}_{\text{com}})$
return $b = b'$	$\delta \leftarrow_{\$} \mathcal{S}_{\text{com}}.\text{Open}(\text{pp}_1, \text{tk}_{\text{com}}, c, m)$
	return (c, δ)

Additionally to dual-mode commitments we require that the underlying 3PC protocol satisfies so-called *special HVZK* that essentially allows the simulator to simulate arguments for a given challenge $\beta \in \mathcal{B}$.

Definition 6.2 (Special HVZK) Let $\Pi = ((P_0, P_1), V)$ be a 3PC argument system for a polynomial-time computable relation R . We say that Π satisfies *special ϵ -HVZK* if there exists a PPT machine \mathcal{S} such that for all $x \in L_R$ and for all $\beta \in \mathcal{B}$ the following two distributions are ϵ -computationally indistinguishable

$$\{(\alpha, \beta, \gamma) : \alpha \leftarrow_{\$} P_0(x, w); \gamma \leftarrow_{\$} P_1(x, w, \beta)\} \quad \text{and} \quad \{(\alpha, \beta, \gamma) : (\alpha, \gamma) \leftarrow_{\$} \mathcal{S}(x, \beta)\}.$$

In other words, for any $\beta \in \mathcal{B}$, the simulator outputs an argument (α, β, γ) with this challenge, which is indistinguishable from a real argument with challenge β .

As observed by Fischlin [Fis05] common protocols obey this special zero-knowledge notion and if, furthermore, the challenge size is logarithmic in the security parameter then assuming special HVZK is without loss of generality.

6.1 The Compiler

Given a 3PC protocol $\Pi = (P, V)$ with HVZK simulator \mathcal{S} which satisfies the conditions above, we construct a compiled protocol $\Pi^* = (K^*, P^*, V^*)$ as follows. Let $C = (C.\text{GenPP}, C.\text{Com}, C.\text{Open}, C.\text{il}, C.\text{ol})$ be a dual-mode commitment scheme, where $C.\text{GenPP}$ is the honest public parameter generator for the perfectly binding variant of the commitment scheme. On input the security parameter setup algorithm K^* simply runs the public parameter generation of the commitment scheme.

```

K*(1λ)
-----
pp ←$ C.GenPP(1λ)
crs ← pp
return crs

```

We present the compiled protocol in Figure 6. What remains to show is that the compiled protocol retains instance-independent commitments, completeness, soundness and that it achieves instance-independent HVZK simulators.

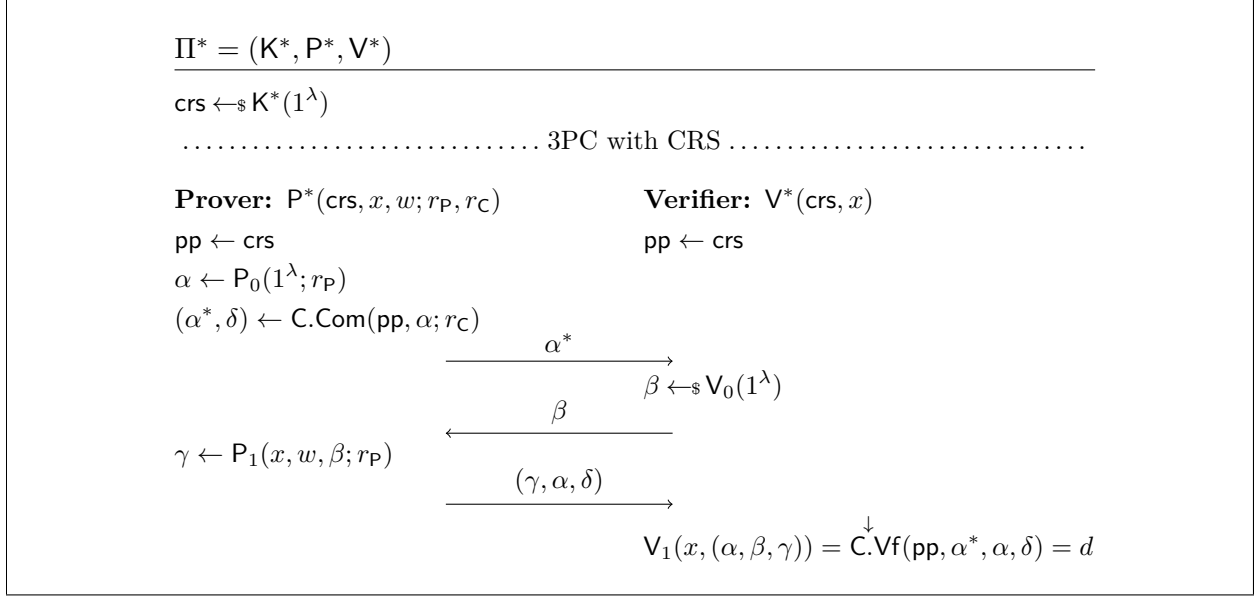


Figure 6: The compiled protocol achieving instance-independent HVZK simulators.

6.2 Security Analysis

Theorem 6.3 *Let $\Pi = (\mathsf{K}, \mathsf{P}, \mathsf{V})$ be a 3PC argument system for a polynomial-time computable relation R such that Π is c -complete and s -sound and has instance-independent commitments and satisfies special HVZK. Let $\mathsf{C} = (\mathsf{C.GenPP}, \mathsf{C.Com}, \mathsf{C.Open}, \mathsf{C.il}, \mathsf{C.ol})$ be a dual-mode commitment scheme. Then, in the CRS model, the compiled protocol $\Pi^* = (\mathsf{K}^*, \mathsf{P}^*, \mathsf{V}^*)$ is c -complete, s -sound and satisfies unbounded instance-independent HVZK. Furthermore the compiled protocol has instance-independent commitments.*

We discuss the various properties in turn.

Instance-independent commitments. As instead of the original commitment α now we use α^* which is a commitment to α , the compiled protocol retains instance-independent prover commitments.

Completeness. If the public parameters are generated honestly for the commitment scheme, the scheme is perfectly binding. Hence, if on an honest execution the verifier does not accept it would not accept for the same execution on the underlying protocol. Hence, if the underlying protocol is c -complete then so is the compiled protocol.

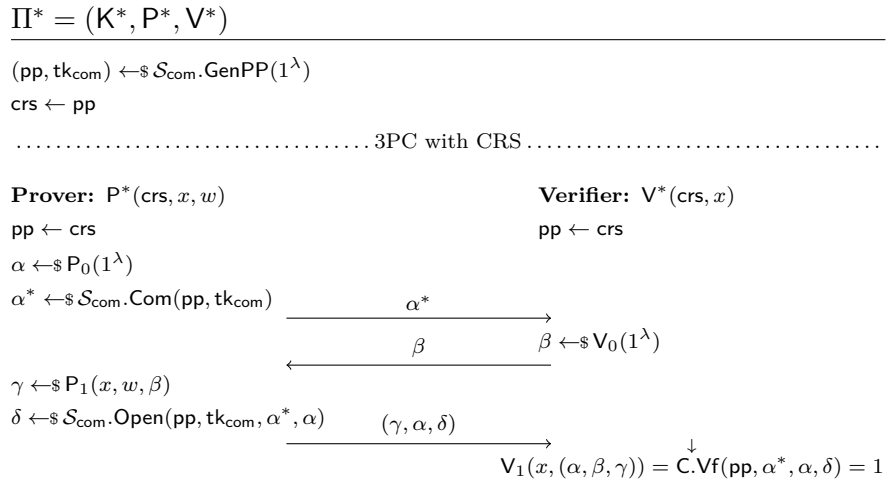
Soundness. Again soundness directly follows from the underlying protocol and the fact that the commitment scheme is perfectly binding if the public parameters are honestly created. Thus, if the underlying 3PC protocol has s -soundness then so does the compiled protocol.

Zero-knowledge. We show that if the underlying protocol satisfies special-HVZK then the compiled protocol satisfies q -bounded instance-independent HVZK (see Definition 4.9). Note that, in fact we can prove an unbounded variant, but for our purpose the bounded variant suffices as our FS-collapse only supports bounded zero-knowledge. Further note that we require the underlying protocol to satisfy special-HVZK but claim that the compiled protocol does not achieve *special* any

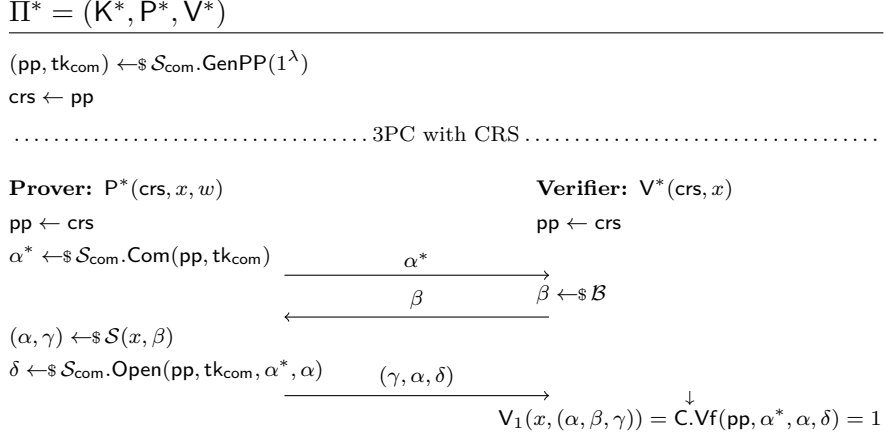
longer. We note that the compiled protocol, in fact, achieves a *special*-instance-independent HVZK variant but again, as we do not need this for our result we chose not to additionally formalize it.

The proof will be down to the security of the commitment scheme and the special-HVZK of the underlying protocol and consists of two game hops.


- Game₁(λ): The first game is equivalent to real world setting $\text{rIPS}_{\mathcal{A}}^{\Pi^*}$ where the adversary has access to the PROV oracle which runs honest executions of the protocol to obtain the transcripts for the adversary.
- Game₂(λ): The setting is similar to the previous game with the exception that now the commitment is returning fake commitments. In order for this to work we switch to the SIMU oracle and let the trapdoor contain the commitment trapdoor. That is, the setup algorithm is changed to run the simulated commitment setup and then the first message of the prover is replaced by a simulated commitment. To make the view consistent on calling SIMU with $b = 1$ the simulated commitment will be opened to the correct α as obtained by the underlying protocol. For this note that at this point the adversary already supplied a valid instance and witness. That is the protocol transcript is still computed with the underlying prover and simulator. The current setting is best visualized with an adapted protocol:



- Game₃(λ): The setting now switches to use the HVZK simulator \mathcal{S} of the underlying protocol. Calls to SIMU with bit $b = 0$ are answered as before with a simulated commitment. Note that while in the last setting β was chosen by the verifier \mathbf{V}_0 we now choose a random value $\beta \in \mathcal{B}$. (Technically this is not a change, since we consider 3PC protocols.) If later for the same index a call to SIMU with bit $b = 1$ is made, then the special HVZK simulator is run on challenge β (that was chosen before) and instance x to obtain $(\alpha, \gamma) \leftarrow_{\mathcal{S}} \mathcal{S}(x, \beta)$. Then, as before the commitment is opened to α and $(\alpha^*, \beta, (\gamma, \alpha, \delta))$ is returned to the adversary. Before we present the pseudocode of the simulator, we again present a view of the *adapted protocol* to visualize the setting:



We next present the pseudocode of the simulator $\text{Sim} = (\text{Sim}', (\text{Sim}''_0, \text{Sim}''_1))$. Note that it is assumed that Sim''_0 and Sim''_1 are always run on matching random coins, that is, Sim''_1 can recover (α^*, β) as output by Sim''_0 (see Definition 4.9).



What is left to show is that the two game hops are negligibly close. We prove this below.

Game₁ to Game₂. The difference between the settings in **Game₁** and **Game₂** matches exactly the commitment game. That is, a distinguisher can be turned into an adversary against the commitment scheme. The commitment scheme adversary takes as input the public parameters \mathbf{pp} which it forwards to the distinguisher between games **Game₁** and **Game₂**. On an oracle query for index i with bit b the adversary distinguishes between b being set to 0 or to 1 and answers in these cases as follows:

$b = 0$ It runs the underlying prover \mathbf{P}_0 on input the security parameter to obtain message α . It then sends α to its commitment oracle to receive commitment α^* and opening δ . Finally, it runs verifier \mathbf{V}_0 on input the security parameter to receive β and returns (α^*, β) to the distinguisher.

$b = 1$ It recovers α, α^* and β (or if the corresponding call with $b = 0$ was not yet made, it performs the above steps). It then calls prover \mathbf{P}_1 on input x, w and β to receive γ . (Note that \mathbf{P}_1 is run on the same random coins as prover \mathbf{P}_0 before, and is thus able to recover α .) Finally, it returns to the distinguisher $(\alpha^*, \beta, (\gamma, \alpha, \delta))$.

It is easy to see that if the commitment oracle returns simulated commitments then the adversary perfectly simulates **Game₂** and otherwise **Game₁**. Thus, we have:

$$|\Pr[\text{Game}_1(\lambda) = 1] - \Pr[\text{Game}_2(\lambda) = 1]| \leq \text{Adv}_{\mathcal{C}, \mathcal{S}_{\text{com}}, \mathcal{A}}^{\text{com}}(\lambda).$$

Game₂ to Game₃. The difference between the two games matches exactly the setup of the special-HVZK setting. That is, a distinguisher between the two games can be turned into an adversary \mathcal{A} against the special-HVZK property of the underlying protocol. For this, adversary \mathcal{A} runs the simulated setup of the commitment scheme to obtain $(\text{pp}, \text{tk}_{\text{com}}) \leftarrow \mathcal{S}_{\text{com}}.\text{GenPP}(1^\lambda)$ and forwards pp as CRS to the distinguisher. On an oracle query for index i with bit b the adversary \mathcal{A} distinguishes between b being set to 0 or to 1 and answers these cases as follows:

$b = 0$ It runs commitment simulator $\mathcal{S}_{\text{com}}.\text{Com}$ to obtain a simulated commitment α^* . It then picks a random value $\beta \in \mathcal{B}$ and returns (α^*, β) .

$b = 1$ It recovers (α^*, β) (or freshly generates them if no corresponding oracle call with bit $b = 0$ has been made). It then calls its own oracle on x, w and β to receive a transcript (α, β, γ) . (Note that Definition 6.2 is not stated in an oracle setting but it can be restated as an adversary that gets oracle access to either the left distribution or the right distribution that it can call on inputs x, w and β for values x in the language and where w is a valid witness to this effect.) Adversary \mathcal{A} then uses the commitment simulator to open the commitment α^* to α and it returns $(\alpha^*, \beta, (\gamma, \alpha, \delta))$.

If the adversary receives honest protocol transcripts then it perfectly simulates **Game₂** and otherwise it perfectly simulates **Game₃**. Thus, the distinguishing probability between the two settings can be upper bounded by:

$$|\Pr[\text{Game}_2(\lambda) = 1] - \Pr[\text{Game}_3(\lambda) = 1]| \stackrel{\text{asym}}{\leq} \epsilon$$

assuming that the underlying 3PC protocol is special ϵ -HVZK. This concludes the proof.

7 Fiat–Shamir Signatures

We explain how to extend our techniques in order to obtain a standard model instantiation of FS signatures, under similar complexity assumptions as in the case of FS NIZK. In particular we will identify a certain class of so-called highly sound identification (ID) schemes, such that we can instantiate the hash function in the corresponding FS collapse via a q -wise independent hash function. On the positive side, we remark that the obtained signature scheme is in the standard model (i.e., without a CRS) even if the starting identification scheme is in the CRS model; the reason is that in the case of signatures we can always include the CRS as part of the public key of the verifier. On the negative side, the obtained signature scheme satisfies only a weak unforgeability flavour.

We recall the definition of ID and signature schemes in Section 7.1. Section 7.2 contains our positive result for FS signatures in the standard model. Finally, in Section 7.3, we explain how to obtain the desired properties by relying on similar tools as the ones used in the compilers from Section 5 and Section 6.

7.1 Identification and Signature Schemes

The FS transform can be used in order to generically obtain signature schemes starting from ID schemes satisfying certain properties.

Canonical ID schemes. An ID scheme consists of three PPT algorithms (K, P, V) . Algorithm K takes as input 1^λ and outputs a CRS $\text{crs} \in \{0, 1\}^*$, together with a pair of keys (pk, sk) , where pk is called the public key and sk is called the secret key. (In case the ID scheme is in the standard model, then we simply set $\text{crs} = \varepsilon$.) Later the prover P interacts with the verifier V to convince him he knows the secret key sk corresponding to pk (where both P and V are also given crs). At the end of the protocol execution, the verifier outputs a bit (representing his decision); we write $\langle P(\text{sk}), V \rangle(\text{crs}, \text{pk})$ for the random variable corresponding to the verifier’s verdict. Similarly, we write $P(\text{crs}, \text{pk}, \text{sk}) \stackrel{r}{\rightrightarrows} V(\text{crs}, \text{pk})$ for the random variable corresponding to transcripts of honest protocol executions.

For applying the FS transform one is typically interested in so-called *canonical* ID schemes. Intuitively, canonical ID schemes are the counterpart of 3PC arguments (cf. Section 4.1). In particular, we can think of the prover algorithm as being split into two sub-algorithms $P := (P_0, P_1)$, where P_0 takes as input a pair (pk, sk) and outputs the prover’s first message α (the so-called commitment) and P_1 takes as input (pk, sk) as well as the verifier’s challenge β to produce the prover’s second message γ (the so-called response). In general P_0 and P_1 are allowed to share the same random tape, which we denote by $r \in \{0, 1\}^*$. In a similar fashion we can think of the verifier’s algorithm as split into two sub-algorithms $V = (V_0, V_1)$, where V_0 outputs a uniformly random value $\beta \in \mathcal{B}$ and V_1 is deterministic and corresponds to the verifier’s verdict (i.e., V_1 takes as input pk and a transcript (α, β, γ) and returns a decision bit $d \in \{0, 1\}$).

A canonical ID scheme typically satisfies three properties known as completeness, soundness, and HVZK, which are analogues to the corresponding properties of 3PC argument systems. We provide an informal definition of completeness and soundness below (we discuss HVZK later in this section).

- **Completeness.** The honest prover P (holding sk) convinces the verifier V (holding pk) with overwhelming probability (over the randomness of K, P, V).
- **Soundness.** For all PPT provers P^* , the probability that P^* convinces V on common input (crs, pk) is bounded by $s(\lambda) \in \text{negl}(\lambda)$.¹²

Signature schemes. A signature scheme is a triple of PPT algorithms (K, S, V) . Algorithm K takes as input the security parameter 1^λ and outputs a pair of keys (pk, sk) . Algorithm S takes as input a pair (sk, m) , and outputs a signature σ on message $m \in \{0, 1\}^*$. Algorithm V takes as input pk and a pair (m, σ) and returns a decision bit $d \in \{0, 1\}$.

We say that a signature scheme satisfies completeness if for all messages $m \in \{0, 1\}^*$ we have that $V(\text{pk}, (m, S(\text{sk}, m))) = 1$ with overwhelming probability over the generation of (pk, sk) .

The standard security notion for signature schemes is called existential unforgeability against chosen-message attacks (EUF-CMA). Roughly speaking this notion demands that it should be hard to forge a signature without knowing the secret key, even when given access to an oracle, signing polynomially many messages of the adversary’s choice. Here, we will instead stick to a strictly weaker notion called random-message unforgeability against random message attack (RUF-RMA). The difference is that now both signature queries and the final forgery are for messages that are not under the adversary’s control (in particular they are uniformly random messages).

¹²The definition of soundness we consider is very weak, in that for instance, it is achieved by the naive protocol where P forwards sk to V . Such a protocol is not passively secure, since an adversary observing a single honest execution can later impersonate the prover. Note, however, that we additionally require a canonical ID scheme to satisfy HVZK, which implies passive security.

Definition 7.1 (RUF-RMA) Let $\Pi = (\mathsf{K}, \mathsf{S}, \mathsf{V})$ be a signature scheme. We say that Π satisfies q -bounded random-message unforgeability against random-message attacks (q -bounded RUF-RMA) if for all PPT adversaries \mathcal{A} asking at most q oracle queries the following holds:

$$\Pr \left[\mathsf{V}(\mathsf{pk}, (m^*, \sigma^*)) = 1 : (\mathsf{pk}, \mathsf{sk}) \leftarrow_{\$} \mathsf{K}(1^\lambda); m^* \leftarrow_{\$} \{0, 1\}^*; \sigma^* \leftarrow \mathcal{A}^{\mathsf{Sign}(\mathsf{sk})}(1^\lambda, \mathsf{pk}, m^*) \right] \leq \text{negl}(\lambda),$$

where oracle $\mathsf{Sign}(\mathsf{sk})$, upon each query, samples a fresh random message $m_i \leftarrow_{\$} \{0, 1\}^*$ and returns (m_i, σ_i) with $\sigma_i \leftarrow_{\$} \mathsf{S}(\mathsf{sk}, m_i)$.

FS signatures. The FS transform allows to turn canonical ID schemes into signature schemes, as specified below. Let $\Pi = (\mathsf{K}, \mathsf{P}, \mathsf{V})$ be the initial canonical ID scheme. Additionally, consider a family of hash functions H consisting of algorithms $\mathsf{H.KGen}$, $\mathsf{H.kl}$, $\mathsf{H.Eval}$, $\mathsf{H.il}$ and $\mathsf{H.ol}$ (see Section 3.1); here $\mathsf{H.il}$ and $\mathsf{H.ol}$ correspond, respectively, to the bit lengths of messages $\alpha||m$ and β (as a function of the security parameter λ).

The FS collapse of Π using H is a triple of algorithms $\overline{\Pi}_{\mathsf{FS}, \mathsf{H}} := (\mathsf{K}_{\mathsf{FS}}, \mathsf{S}_{\mathsf{FS}}, \mathsf{V}_{\mathsf{FS}})$ defined as follows.

- Algorithm K_{FS} takes as input the security parameter, samples $\mathsf{hk} \leftarrow_{\$} \mathsf{H.KGen}(1^\lambda)$, $(\mathsf{crs}, \mathsf{pk}, \mathsf{sk}) \leftarrow_{\$} \mathsf{K}(1^\lambda)$, and outputs $\overline{\mathsf{pk}} := (\mathsf{crs}, \mathsf{hk}, \mathsf{pk})$ and $\overline{\mathsf{sk}} := (\mathsf{crs}, \mathsf{hk}, \mathsf{pk}, \mathsf{sk})$.
- Algorithm S_{FS} takes as input $\overline{\mathsf{sk}} := (\mathsf{crs}, \mathsf{hk}, \mathsf{pk}, \mathsf{sk})$ and a message $m \in \{0, 1\}^*$, and runs $\mathsf{P}_0(\mathsf{crs}, \mathsf{pk}, \mathsf{sk})$ in order to obtain the commitment $\alpha \in \{0, 1\}^{\mathsf{H.il}(\lambda) - |m|}$; next P_{FS} defines the challenge as $\beta := \mathsf{H.Eval}(\mathsf{hk}, \alpha||m)$ and runs $\mathsf{P}_1(\mathsf{crs}, \mathsf{pk}, \mathsf{sk}, \beta)$ in order to obtain the response γ . Finally P_{FS} outputs $\sigma := (\alpha, \gamma)$.
- Algorithm V_{FS} takes as input $\overline{\mathsf{pk}} := (\mathsf{crs}, \mathsf{hk}, \mathsf{pk})$ and a pair (m, σ) , and returns 1 if and only if verifier $\mathsf{V}_1(\mathsf{crs}, \mathsf{pk}, (\alpha, \beta, \gamma)) = 1$ where $\beta = \mathsf{H.Eval}(\mathsf{hk}, \alpha||m)$.

7.2 Proof of Random-Message Unforgeability

In a nutshell the result of Fiat and Shamir (for the case of signatures) says that whenever $\Pi = (\mathsf{P}, \mathsf{V})$ is a (standard-model) canonical ID scheme satisfying completeness, computational soundness, and computational HVZK (in addition to a basic requirement on the min-entropy of the prover's commitment), its FS collapse $\overline{\Pi}_{\mathsf{FS}, \mathsf{H}}$ is an EUF-CMA signature scheme if H is modeled as a random oracle.

Here we show that the FS transform admits a very simple standard-model instantiation when starting from a special class of canonical ID schemes $\Pi = (\mathsf{K}, \mathsf{P}, \mathsf{V})$ satisfying three additional properties. The obtained signature scheme will satisfy q -bounded RUF-RMA. Since most of the technical details are identical to our standard-model instantiation for FS NIZK, we only highlight the main differences here. The properties we need are defined below:

P1': The first property requires that the commitment α can be computed independently of $(\mathsf{pk}, \mathsf{sk})$, i.e. $\alpha \leftarrow_{\$} \mathsf{P}_0(\mathsf{crs})$. This property is analogous to “instance-independent commitments” for argument systems (cf. Definition 4.7).

P2': We can define the worst-case collision probability of commitments $\delta(\lambda)$ in the same way as we did for argument systems (cf. Definition 4.5). The second property requires that the SEGR $\varrho(\lambda) := s(\lambda)/\delta(\lambda)$ is negligible in the security parameter, where $s(\lambda)$ is the soundness parameter of the ID scheme. The definition of SEGR is analogous to the definition of SEGR for argument systems (cf. Definition 4.6).

P3': There exists a PPT simulator $\mathcal{S} := (\mathcal{S}', \mathcal{S}'')$ such that the following two distributions are computationally indistinguishable.

- Adversary $\mathcal{A}(\text{crs}, \text{pk})$ asking up-to q queries to an oracle returning transcripts of honest executions between $\text{P}(\text{crs}, \text{pk}, \text{sk})$ and $\text{V}(\text{crs}, \text{pk})$, where $(\text{crs}, \text{pk}, \text{sk})$ are generated by running algorithm K .
- Adversary $\mathcal{A}(\text{crs}, \text{pk})$ asking up-to q queries to \mathcal{S}'' , where crs is generated by running \mathcal{S}' (also holding a trapdoor tk) and (pk, sk) are generated using K . Here, \mathcal{S}'' computes a simulated transcript (α, β, γ) knowing only pk (and some trapdoor information); furthermore the simulator can compute (α, β) independently of pk .

This property is analogous to “ q -bounded instance-independent HVZK” for argument systems (cf. Definition 4.9).

We will call canonical ID schemes satisfying properties **P1'**-**P3'** above (besides completeness and soundness) *highly sound* canonical ID schemes.

Theorem 7.2 *Let $\Pi = (\text{K}, \text{P}, \text{V})$ be a highly sound canonical ID scheme and H be a programmable q -wise independent hash function. Then, the FS collapse $\overline{\Pi}_{\text{FS}, \text{H}}$ of Π using H yields a signature scheme satisfying q -RUF-RMA.*

Proof (sketch). Assume there exists a PPT adversary \mathcal{A} and some polynomial $p(\cdot)$ such that for infinitely many values of $\lambda \in \mathbb{N}$ the following holds:

$$\Pr \left[\text{V}_{\text{FS}}(\overline{\text{pk}}, (m^*, \sigma^*)) = 1 : (\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow_{\$} \text{K}_{\text{FS}}(1^\lambda); m^* \leftarrow_{\$} \{0, 1\}^*; \sigma^* \leftarrow \mathcal{A}^{\text{Sign}_{\text{FS}}(\overline{\text{sk}})}(1^\lambda, \overline{\text{pk}}, m^*) \right] \geq 1/p(\lambda),$$

where oracle $\text{Sign}_{\text{FS}}(\overline{\text{sk}})$ internally runs the signing algorithm $\text{S}_{\text{FS}}(\overline{\text{sk}}, \cdot)$.

We consider a series of computationally indistinguishable games, which are outlined below.

Game₁(λ): This is identical to the RUF-RMA experiment, but now the randomness r_i used to generate the q signatures σ_i corresponding to \mathcal{A} 's signature queries, and the q messages $m_i \leftarrow_{\$} \{0, 1\}^*$, are pre-sampled. Additionally, values $\alpha_i = \text{P}_0(1^\lambda; r_i)$ are pre-computed—note that this is possible because of property **P1'**—which allows us to also pre-compute values β_i as $\beta_i = \text{H.Eval}(\text{hk}, \alpha_i \| m_i)$ where hk is the hash key. All of these values are stored in a trapdoor $\overline{\text{tk}}$ which is given to the hybrid oracle answering signature queries.

We say that **Game₁(λ)** = 1 if and only if the forgery returned by \mathcal{A} is accepting. Observe that all of the above steps (pre-computing values) are clearly just syntactical changes, and thus $\Pr[\text{Game}_1(\lambda) = 1] \geq 1/p(\lambda)$.

q-wise

Game₂(λ): In the second step we replace all values β_i with uniformly random values sampled from the range of hash function H and the key hk is chosen via programming the hash function. Down to the programmable q -wise independence property of H the distribution corresponding to **Game₁(λ)** and **Game₂(λ)** are identical, i.e., **Game₂(λ)** \equiv **Game₁(λ)**.

HVZK

Game₃(λ): In the third step we change the way signature queries are answered. In particular, instead of running $\text{S}_{\text{FS}}(\overline{\text{sk}}, \cdot)$, we define a simulator Sim that answers signature queries only given as input $\overline{\text{pk}}$ (and some trapdoor information). The simulator pre-computes all values and then relies on the HVZK simulator \mathcal{S} (here is where we use property **P3'**). The description of Sim is essentially identical (with minor modifications) to the simulator Sim defined in the proof of Lemma 4.10, and is therefore omitted.

Down to the q -bounded instance-independent HVZK property of Π , we can write:

$$|\Pr[\text{Game}_3(\lambda) = 1] - \Pr[\text{Game}_2(\lambda) = 1]| \in \text{negl}(\lambda).$$

Combining the above equations, we have that there exists a polynomial $\bar{p}(\cdot)$ such that, for infinitely many values of $\lambda \in \mathbb{N}$, $\Pr[\text{Game}_3(\lambda) = 1] \geq 1/\bar{p}(\lambda)$.

We now use this fact to contradict the soundness of the underlying ID scheme. This last step of the proof is analogous to the proof of soundness for FS NIZK, and consists of two sub-steps that are outlined below:

- In the first step, we consider a selective variant of the FS transform, in a similar way as we did in Section 4.3 for the case of argument systems. In the selective variant, which we denote by $\Pi_{\text{sel-FS,H}}$, the verifier V sends a random hash key hk together with a random message $m^* \leftarrow_{\$} \{0,1\}^*$ that is hashed by the prover together with the commitment α in order to generate the challenge β (and the corresponding answer γ).

Note that the above described selective variant of the FS transform still yields a public-coin ID scheme. Furthermore, it is easy to prove that $\Pi_{\text{sel-FS,H}}$ satisfies both completeness and soundness, provided that the original ID scheme does. This proof is almost identical to the proof of Theorem 4.3, and is therefore omitted.

- In the second step, we use the adversary \mathcal{A} (with non-negligible advantage in game $\text{Game}_3(\lambda)$) to construct a prover P^* that breaks the soundness of $\Pi_{\text{sel-FS,H}}$ with non-negligible advantage. Prover P^* receives as input a pair (crs, pk) , and initially runs the honest (instance-independent) prover $P_0(\text{crs})$ in order to obtain a value α that it forwards to the verifier. The verifier replies with a random hash key hk and random message $m^* \in \{0,1\}^*$. At this point, P^* initializes the adversary \mathcal{A} with a simulated public key $\overline{\text{pk}} = (\text{crs}', \text{hk}, \text{pk})$ (where crs' is the simulated CRS coming from Sim) and replies to \mathcal{A} 's signature queries as specified in $\text{Game}_3(\lambda)$ (i.e., by running the simulator Sim). Finally, it forwards m^* to \mathcal{A} obtaining a forgery $\sigma^* = (\alpha^*, \gamma^*)$. In case $\alpha^* = \alpha$, then P^* passes γ^* to the verifier. Else, P^* aborts.

Clearly, the probability of P^* breaking soundness is lower bounded by the success probability of \mathcal{A} times the probability that α^* is equal to α . It follows that, if $\Pi_{\text{sel-FS,H}}$ has soundness parameter $s(\lambda)$, the SEGR $\varrho(\lambda) := s(\lambda)/\delta(\lambda) \geq 1/\bar{p}(\lambda)$, contradicting property **P2'**.

The above two facts imply the statement of the theorem. □

7.3 Obtaining the Required Properties

It remains to construct a highly sound canonical ID scheme. As we briefly explain now, the latter can be done using similar techniques as the ones we used to construct highly sound 3PC arguments.

Let us start with any (standard-model) canonical ID scheme Π that satisfies completeness, soundness, HVZK, and moreover has instance-independent commitments (i.e., property **P1'**). Many canonical ID schemes satisfy this requirement, including the ones by Schnorr [Sch91] and Guillou-Quisquater [GQ90]. Hence, we can apply the two compilers described in Section 5 and Section 6 in order to obtain a highly sound canonical ID scheme Π'' as follows:

- First, we transform Π into an ID scheme Π' that additionally satisfies **P3'** (i.e., the HVZK simulator is instance-independent). This is achieved by having the prover commit to the value α that would have been sent in Π using an equivocal commitment, in a similar fashion as we did in our compiler from Fig. 6.

- Second, we transform Π' into an ID scheme Π'' that additionally satisfies $\mathbf{P2}'$. This is achieved by providing a mechanism that allows to produce many challenges β given only a single commitment α , in a similar way as we did in the compiler from Fig. 4.

We remark that, for the case of FS signatures, the obtained signature scheme will be in the standard model even though the starting highly sound canonical ID scheme is in the CRS model.

Acknowledgments

We are grateful to Christina Brzuska for her active participation in this research. Her feedback and suggestions played an essential part in the development of this work.

We thank Nils Fleischhacker and Markulf Kohlweiss for helpful comments on the presentation. We are grateful to an anonymous reviewer of TCC 2016 for pointing out that the constant hash function already suffices for obtaining a 1-bounded NIZK assuming properties $\mathbf{P1-P3}$ and thereby inspiring using a q -wise independent hash-function as instantiation. Before, we used a more complicated construction based on indistinguishability obfuscation and puncturable PRFs. We also thank the reviewer for pointing out the Blum-Lapidot-Shamir protocol, and we thank Ivan Visconti for helpful discussions and clarifications on the protocol itself.

References

- [AABN02] Michel Abdalla, Jee Hea An, Mihir Bellare, and Chanathip Namprempre. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 418–433, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer, Berlin, Germany. (Cited on page 4.)
- [Bar01] Boaz Barak. How to go beyond the black-box simulation barrier. In *42nd Annual Symposium on Foundations of Computer Science*, pages 106–115, Las Vegas, Nevada, USA, October 14–17, 2001. IEEE Computer Society Press. (Cited on page 4.)
- [BCP14] Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In Yehuda Lindell, editor, *TCC 2014: 11th Theory of Cryptography Conference*, volume 8349 of *Lecture Notes in Computer Science*, pages 52–73, San Diego, CA, USA, February 24–26, 2014. Springer, Berlin, Germany. (Cited on pages 14, 29, and 30.)
- [BDG⁺13] Nir Bitansky, Dana Dachman-Soled, Sanjam Garg, Abhishek Jain, Yael Tauman Kalai, Adriana López-Alt, and Daniel Wichs. Why “Fiat-Shamir for proofs” lacks a proof. In Amit Sahai, editor, *TCC 2013: 10th Theory of Cryptography Conference*, volume 7785 of *Lecture Notes in Computer Science*, pages 182–201, Tokyo, Japan, March 3–6, 2013. Springer, Berlin, Germany. (Cited on pages 4 and 5.)
- [BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th Annual ACM Symposium on Theory of Computing*, pages 103–112, Chicago, Illinois, USA, May 2–4, 1988. ACM Press. (Cited on page 5.)

- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Berlin, Germany. (Cited on page 13.)
- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6:1–6:48, May 2012. (Cited on page 13.)
- [BGW12] Nir Bitansky, Sanjam Garg, and Daniel Wichs. Why Fiat-Shamir for proofs lacks a proof. Cryptology ePrint Archive, Report 2012/705, 2012. <http://eprint.iacr.org/2012/705>. (Cited on pages 4 and 5.)
- [Blu81] Manuel Blum. Coin flipping by telephone. In Allen Gersho, editor, *Advances in Cryptology – CRYPTO’81*, volume ECE Report 82-04, pages 11–15, Santa Barbara, CA, USA, 1981. U.C. Santa Barbara, Dept. of Elec. and Computer Eng. (Cited on pages 9, 19, and 23.)
- [BLV03] Boaz Barak, Yehuda Lindell, and Salil P. Vadhan. Lower bounds for non-black-box zero knowledge. In *44th Annual Symposium on Foundations of Computer Science*, pages 384–393, Cambridge, Massachusetts, USA, October 11–14, 2003. IEEE Computer Society Press. (Cited on page 4.)
- [BPW12] David Bernhard, Olivier Pereira, and Bogdan Warinschi. How not to prove yourself: Pitfalls of the Fiat-Shamir heuristic and applications to Helios. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology – ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 626–643, Beijing, China, December 2–6, 2012. Springer, Berlin, Germany. (Cited on pages 4 and 6.)
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73, Fairfax, Virginia, USA, November 3–5, 1993. ACM Press. (Cited on page 4.)
- [BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426, St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Berlin, Germany. (Cited on page 29.)
- [BST14] Mihir Bellare, Igors Stepanovs, and Stefano Tessaro. Poly-many hardcore bits for any one-way function and a framework for differing-inputs obfuscation. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014, Part II*, volume 8874 of *Lecture Notes in Computer Science*, pages 102–121, Kaoshiung, Taiwan, R.O.C., December 7–11, 2014. Springer, Berlin, Germany. (Cited on page 13.)
- [CCR15] Ran Canetti, Yilei Chen, and Leonid Reyzin. On the correlation intractability of obfuscated pseudorandom functions. Cryptology ePrint Archive, Report 2015/334, 2015. <http://eprint.iacr.org/>. (Cited on page 4.)

- [CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th Annual ACM Symposium on Theory of Computing*, pages 209–218, Dallas, Texas, USA, May 23–26, 1998. ACM Press. (Cited on page 4.)
- [CHL05] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. Compact e-cash. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 302–321, Aarhus, Denmark, May 22–26, 2005. Springer, Berlin, Germany. (Cited on page 3.)
- [CL02] Jan Camenisch and Anna Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 61–76, Santa Barbara, CA, USA, August 18–22, 2002. Springer, Berlin, Germany. (Cited on page 3.)
- [CL09] Kai-Min Chung and Feng-Hao Liu. Tight parallel repetition theorems for public-coin arguments. *Electronic Colloquium on Computational Complexity (ECCC)*, 16:109, 2009. (Cited on page 26.)
- [CP15] Kai-Min Chung and Rafael Pass. Tight parallel repetition theorems for public-coin arguments using KL-divergence. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015: 12th Theory of Cryptography Conference, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 229–246, Warsaw, Poland, March 23–25, 2015. Springer, Berlin, Germany. (Cited on page 26.)
- [CV07] Dario Catalano and Ivan Visconti. Hybrid commitments and their applications to zero-knowledge proof systems. *Theoretical computer science*, 374(1):229–260, 2007. (Cited on pages 23 and 31.)
- [Dam00] Ivan Damgård. Efficient concurrent zero-knowledge in the auxiliary string model. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 418–430, Bruges, Belgium, May 14–18, 2000. Springer, Berlin, Germany. (Cited on page 6.)
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *23rd Annual ACM Symposium on Theory of Computing*, pages 542–552, New Orleans, Louisiana, USA, May 6–8, 1991. ACM Press. (Cited on page 3.)
- [DJKL12] Dana Dachman-Soled, Abhishek Jain, Yael Tauman Kalai, and Adriana López-Alt. On the (in)security of the Fiat-Shamir paradigm, revisited. *IACR Cryptology ePrint Archive*, 2012:706, 2012. (Cited on pages 4 and 5.)
- [DNRS99] Cynthia Dwork, Moni Naor, Omer Reingold, and Larry J. Stockmeyer. Magic functions. In *40th Annual Symposium on Foundations of Computer Science*, pages 523–534, New York, New York, USA, October 17–19, 1999. IEEE Computer Society Press. (Cited on pages 4 and 5.)
- [DRV12] Yevgeniy Dodis, Thomas Ristenpart, and Salil P. Vadhan. Randomness condensers for efficiently samplable, seed-dependent sources. In Ronald Cramer, editor, *TCC 2012: 9th Theory of Cryptography Conference*, volume 7194 of *Lecture Notes in Computer Science*, pages 618–635, Taormina, Sicily, Italy, March 19–21, 2012. Springer, Berlin, Germany. (Cited on page 4.)

- [DV14] Özgür Dagdelen and Daniele Venturi. A second look at Fischlin’s transformation. In David Pointcheval and Damien Vergnaud, editors, *AFRICACRYPT 14: 7th International Conference on Cryptology in Africa*, volume 8469 of *Lecture Notes in Computer Science*, pages 356–376, Marrakesh, Morocco, May 28–30, 2014. Springer, Berlin, Germany. (Cited on page 6.)
- [EL04] Edith Elkind and Helger Lipmaa. Interleaving cryptography and mechanism design: The case of online auctions. In Ari Juels, editor, *FC 2004: 8th International Conference on Financial Cryptography*, volume 3110 of *Lecture Notes in Computer Science*, pages 117–131, Key West, USA, February 9–12, 2004. Springer, Berlin, Germany. (Cited on page 3.)
- [FHN⁺12] Sebastian Faust, Carmit Hazay, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Angela Zottarel. Signature schemes secure against hard-to-invert leakage. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology – ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 98–115, Beijing, China, December 2–6, 2012. Springer, Berlin, Germany. (Cited on page 9.)
- [Fis05] Marc Fischlin. Communication-efficient non-interactive proofs of knowledge with online extractors. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 152–168, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Berlin, Germany. (Cited on pages 6 and 32.)
- [FKMV12] Sebastian Faust, Markulf Kohlweiss, Giorgia Azzurra Marson, and Daniele Venturi. On the non-malleability of the Fiat-Shamir transform. In Steven D. Galbraith and Mridul Nandi, editors, *Progress in Cryptology - INDOCRYPT 2012: 13th International Conference in Cryptology in India*, volume 7668 of *Lecture Notes in Computer Science*, pages 60–79, Kolkata, India, December 9–12, 2012. Springer, Berlin, Germany. (Cited on page 4.)
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology – CRYPTO’86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194, Santa Barbara, CA, USA, August 1987. Springer, Berlin, Germany. (Cited on pages 3, 4, and 15.)
- [GK03] Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the Fiat-Shamir paradigm. In *44th Annual Symposium on Foundations of Computer Science*, pages 102–115, Cambridge, Massachusetts, USA, October 11–14, 2003. IEEE Computer Society Press. (Cited on pages 4 and 5.)
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th Annual ACM Symposium on Theory of Computing*, pages 218–229, New York City, New York, USA, May 25–27, 1987. ACM Press. (Cited on page 3.)
- [Gol01] Oded Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, Cambridge, UK, 2001. (Cited on page 12.)
- [GQ90] Louis C. Guillou and Jean-Jacques Quisquater. A “paradoxical” indentity-based signature scheme resulting from zero-knowledge. In Shafi Goldwasser, editor, *Advances in Cryptology – CRYPTO’88*, volume 403 of *Lecture Notes in Computer Science*, pages

216–231, Santa Barbara, CA, USA, August 21–25, 1990. Springer, Berlin, Germany. (Cited on page 40.)

- [GS08] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 415–432, Istanbul, Turkey, April 13–17, 2008. Springer, Berlin, Germany. (Cited on page 5.)
- [Hai09] Iftach Haitner. A parallel repetition theorem for any interactive argument. In *50th Annual Symposium on Foundations of Computer Science*, pages 241–250, Atlanta, Georgia, USA, October 25–27, 2009. IEEE Computer Society Press. (Cited on page 6.)
- [HPWP10] Johan Håstad, Rafael Pass, Douglas Wikström, and Krzysztof Pietrzak. An efficient parallel repetition theorem. In Daniele Micciancio, editor, *TCC 2010: 7th Theory of Cryptography Conference*, volume 5978 of *Lecture Notes in Computer Science*, pages 1–18, Zurich, Switzerland, February 9–11, 2010. Springer, Berlin, Germany. (Cited on page 26.)
- [HSW14] Susan Hohenberger, Amit Sahai, and Brent Waters. Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 201–220, Copenhagen, Denmark, May 11–15, 2014. Springer, Berlin, Germany. (Cited on page 6.)
- [KZZ15] Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. End-to-end verifiable elections in the standard model. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part II*, volume 9057 of *Lecture Notes in Computer Science*, pages 468–498, Sofia, Bulgaria, April 26–30, 2015. Springer, Berlin, Germany. (Cited on page 3.)
- [Lin15] Yehuda Lindell. An efficient transform from sigma protocols to NIZK with a CRS and non-programmable random oracle. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015: 12th Theory of Cryptography Conference, Part I*, volume 9014 of *Lecture Notes in Computer Science*, pages 93–109, Warsaw, Poland, March 23–25, 2015. Springer, Berlin, Germany. (Cited on pages 6, 23, and 31.)
- [LS91] Dror Lapidot and Adi Shamir. Publicly verifiable non-interactive zero-knowledge proofs. In Alfred J. Menezes and Scott A. Vanstone, editors, *Advances in Cryptology – CRYPTO’90*, volume 537 of *Lecture Notes in Computer Science*, pages 353–365, Santa Barbara, CA, USA, August 11–15, 1991. Springer, Berlin, Germany. (Cited on pages 5, 9, 19, 21, and 23.)
- [NVZ14] Jesper Buus Nielsen, Daniele Venturi, and Angela Zottarel. Leakage-resilient signatures with graceful degradation. In Hugo Krawczyk, editor, *PKC 2014: 17th International Conference on Theory and Practice of Public Key Cryptography*, volume 8383 of *Lecture Notes in Computer Science*, pages 362–379, Buenos Aires, Argentina, March 26–28, 2014. Springer, Berlin, Germany. (Cited on page 9.)
- [Oka93] Tatsuaki Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In Ernest F. Brickell, editor, *Advances in Cryptology – CRYPTO’92*, volume 740 of *Lecture Notes in Computer Science*, pages 31–53, Santa Barbara, CA, USA, August 16–20, 1993. Springer, Berlin, Germany. (Cited on page 4.)

- [OV12] Rafail Ostrovsky and Ivan Visconti. Simultaneous resettability from collision resistance. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:164, 2012. (Cited on page 9.)
- [PS00] David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000. (Cited on page 4.)
- [PV07] Rafael Pass and Muthuramakrishnan Venkatasubramanian. An efficient parallel repetition theorem for Arthur-Merlin games. In David S. Johnson and Uriel Feige, editors, *39th Annual ACM Symposium on Theory of Computing*, pages 420–429, San Diego, California, USA, June 11–13, 2007. ACM Press. (Cited on page 26.)
- [Sch91] Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991. (Cited on page 40.)
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 475–484, New York, NY, USA, May 31 – June 3, 2014. ACM Press. (Cited on pages 6 and 14.)

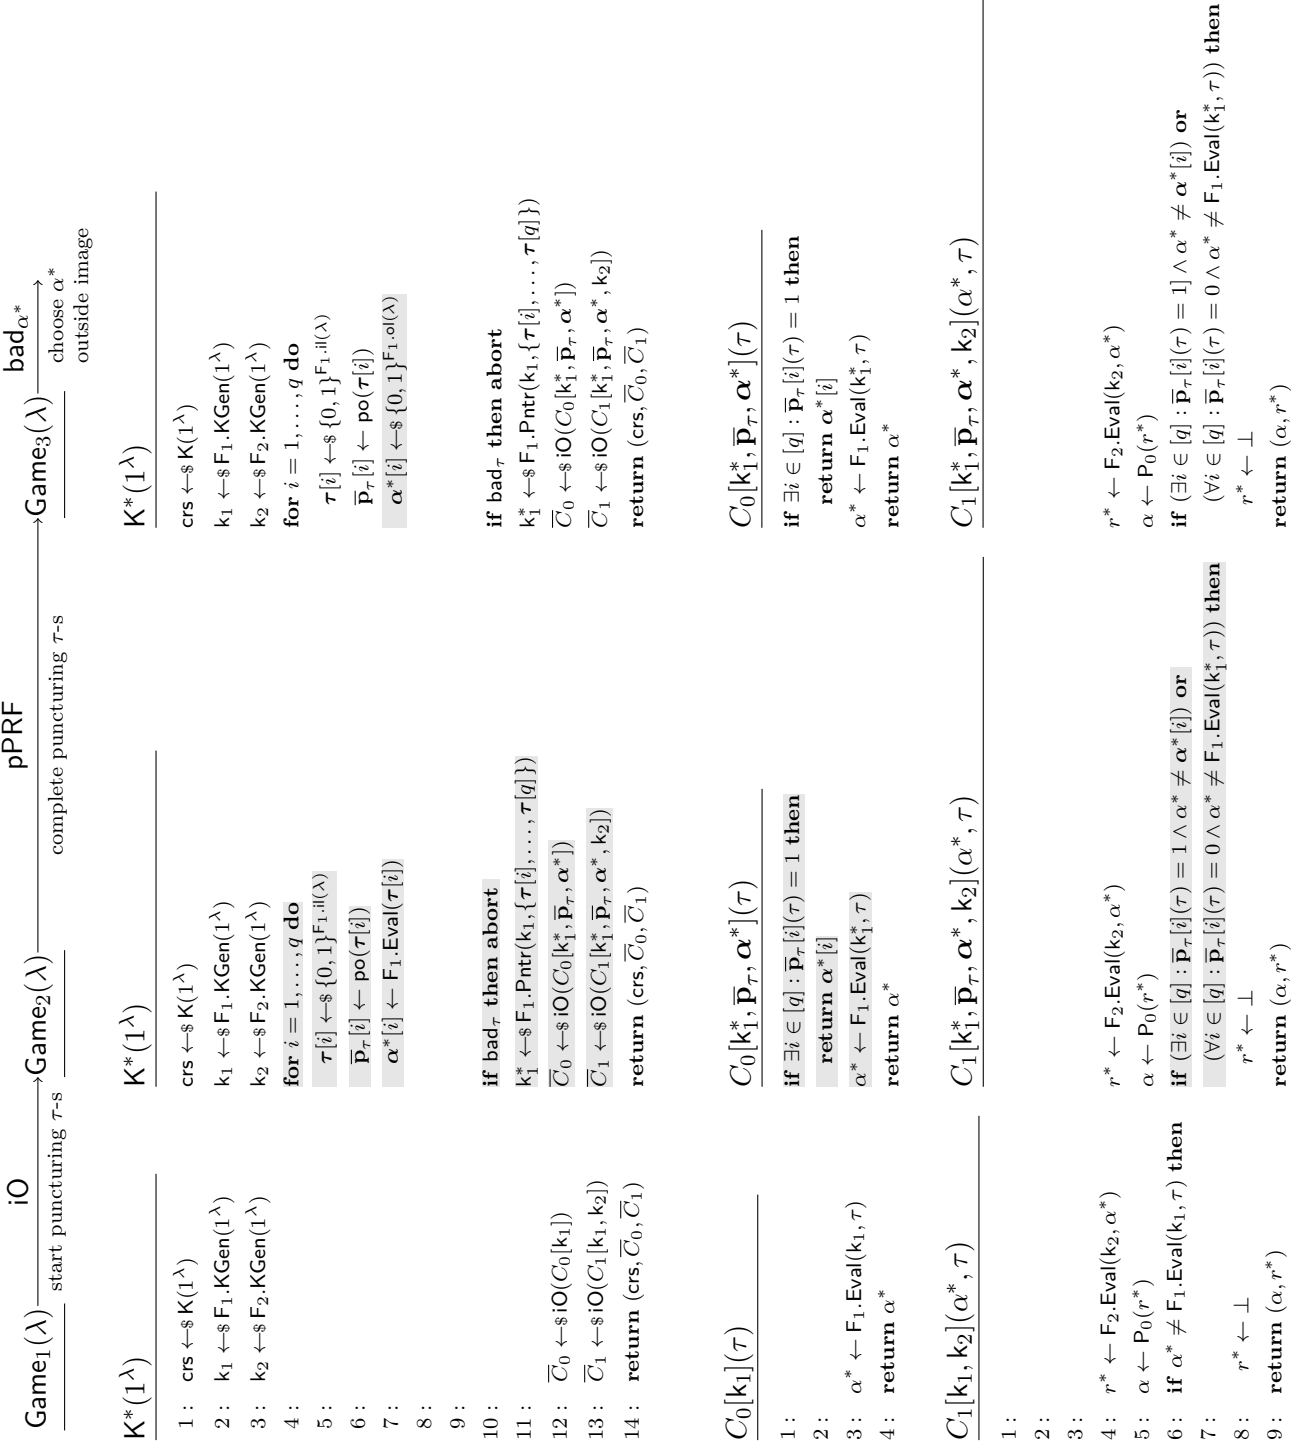


Figure 7: Pseudocode for proof of zero-knowledge for the compiled protocol from Section 5.

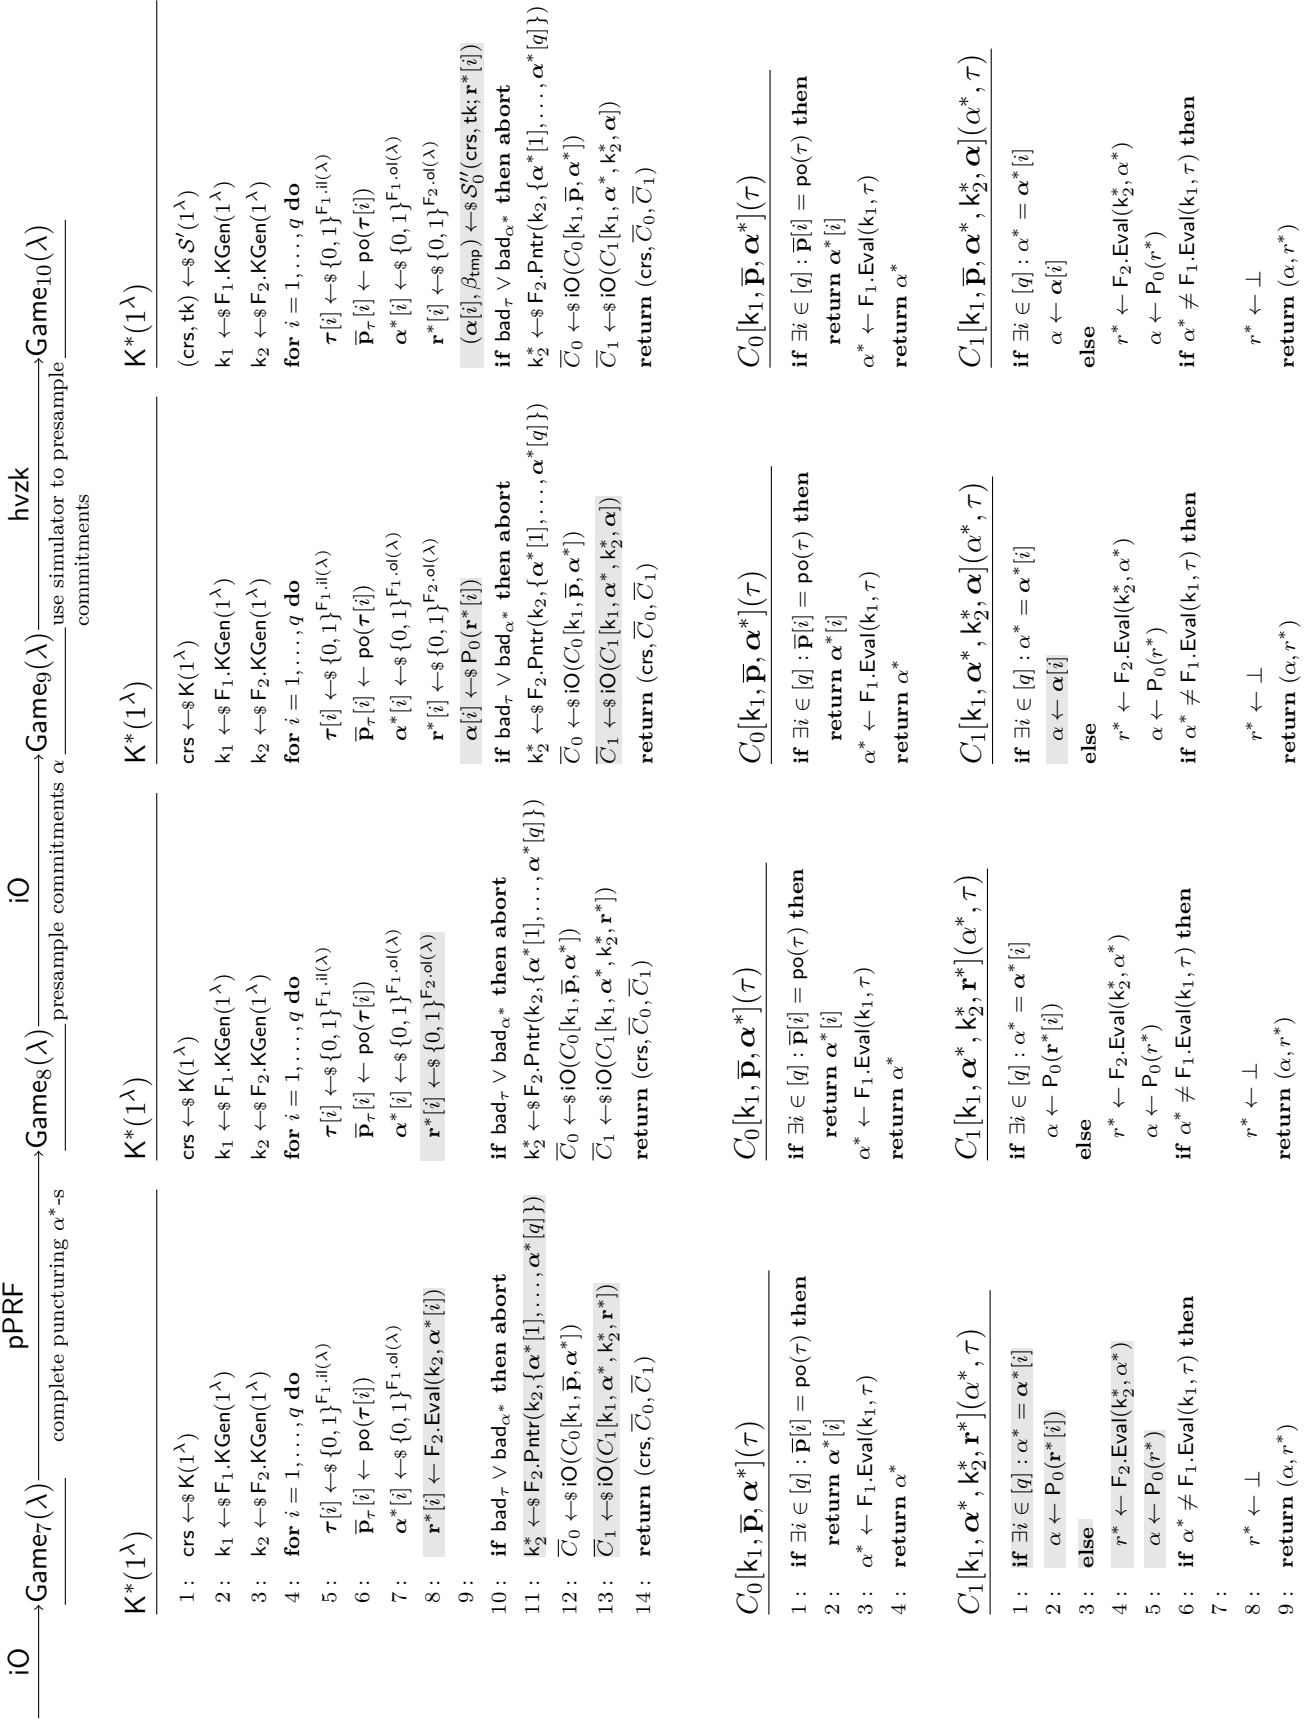


Figure 9: Pseudocode for proof of zero-knowledge for the compiled protocol from Section 5.