

Proposal: Commit-and-Prove Zero-Knowledge Proof Systems

Daniel Benarroch¹, Matteo Campanelli², and Dario Fiore²

¹ QED-it, Tel Aviv, Israel
daniel@qed-it.com

² IMDEA Software Institute, Madrid, Spain
matteo.campanelli@imdea.org, dario.fiore@imdea.org

1 Scope

Commit-and-Prove Zero-Knowledge Proof systems (CP-ZKPs) [Kil89, CLOS02] are a generalization of zero-knowledge proofs in which the prover proves statements about values that are committed.

The aim of this document³ is to stimulate a discussion on the formalization of CP-ZKPs. As we detail in the next section, the motivation for using CP-ZKPs is both theoretical and practical.

Following the ongoing standardization effort in the context of ZKPs [GKV⁺18], *our goal here is to address terminology and definitions for the notions of commitments and CP-ZKPs*. In particular, we believe that the starting point will be discussing a community standard for the definition of commitment schemes. In spite of being mentioned in [BCM⁺18], a more extended discussion on terminology, syntax and definitions of commitments is currently lacking. A second notion we believe is worth discussing is that of commit-and-prove ZKPs: there are several different definitions in literature; it would be important to agree on a single definition or outline a taxonomy of the existing variants.

Note. This is an updated version of the proposal submitted and discussed at the 2nd ZKP workshop. The updates reflect the feedback received during the discussion. The goal of resubmitting it is to propose a continuation of the discussion around the topic of commit-and-prove zero-knowledge proofs. We felt that there is interest around the topic and we believe that a second discussion may be useful, for example towards converging on specific standardization goals.

Status and next steps. We summarize here our subjective perception of the outcome from the discussion during the 2nd edition of the ZK proof workshop. Here is what we feel came out as a trend from the audience’s comments:

- yes, there exist relevant applications that would benefit from this approach;
- practitioners should have guidelines about how to align applications to the notion;
- hence, yes, the notion should be formalized and the definition proposed in this document may be the right approach for this formalization⁴.

What we find lacked from the last discussion was coming up with concrete next actions. We propose the following as next steps (overlapping with some of the comments in Section 5:

³This document is an edited and slightly extended version of the proposal accepted in the 2nd edition of the ZKProof Workshop. In this version we point out new applications and motivations as well as propose the standardizations of specific commitment schemes.

⁴This refers to a comment by Yael Kalai, Angelo de Caro and others (see minutes of the discussion). In this context we refer to the following feature of the commit-and-prove notion proposed here: that the commitment key should be as independent as possible of the properties we prove on inputs and of the proving scheme itself.)

- standardizing a versatile and widespread commitment scheme, Pedersen;
- coming up with a reference implementation for commit-and-prove. This should include a concrete scheme as a reference point (e.g. LegoGroth16, the commit-and-prove version of Groth16 [Gro16] in [CFQ19]) and should make explicit the interface for practitioners. The standardization team should discuss what this interface will look like.

2 Motivation

As observed in the ZKProof proceedings (Applications Track) [BCM⁺18], most applications of Zero Knowledge Proofs require the use of some commitment scheme in order to ensure the privacy and confidentiality of the users and their data, by proving knowledge of the opening. The use of the commit-and-prove paradigm has several interesting features:

- If the commitment is compressing one can distribute succinct and private representations of data that significantly reduce communication complexity and input size for verifiers (as well as their running time)⁵. An interesting observation is that a special case of this scenario are authenticated data structures (in their succinct variant) [Tam03], and in particular cryptographic primitives such as vector commitments [CF13, Mer88] and accumulators [Bd94] which are already in use (or considered to be put in use) in combination with ZKP systems in a variety of applications, such as for example [HBHW16, But18]. We can see these authenticated data structures as commit-and-prove systems where we first compress a large data structure (e.g. a vector or a set) and then prove something on its content (that a certain position has a certain value or that does not contain a certain string).
- One can publish commitments to data previous to generating proofs about them. For example, the Identity scheme in [BCM⁺18] requires that some issuer publishes a credential for the user, which will later be used to prove some attribute; in some cases even before the statement to be proven for a specific attribute is established. This means that, ideally, we would like to have some flexibility as to what statements are proven on the opening of the commitment, as well as to which ZKP schemes are used for the different statements.
- One can use commitments to make different proof systems *interoperable*. For example, one can prove two different statements about the same commitment using two distinct ZKP systems. This can be advantageous in order to exploit the different efficiency tradeoffs of existing systems (see e.g., [AGM18, CFQ19]), or simply because the public parameters of the ZKP systems are generated in different points in time or by different organizations. Interoperability *through time* is particularly important for legacy systems: commitments produced in the past and through a specific (legacy) commitment scheme can still be used as inputs to proof systems that will be developed in the future.

We find thus motivating to have a framework for properly building such applications and instantiating the commitment schemes with the corresponding ZKP schemes. We believe that the starting point of such a framework should be a formal definition of commitments and commit-and-prove zero-knowledge proof systems.

Finally, we believe that standardizing the definition for CP-ZKP (and, at a later time, its framework) offers yet one more advantage related to how we can program constraint systems. Today

⁵We point out that if this commitment is compressing we may benefit from its use even if it is not hiding. Indeed the verifier now just needs a “compact placeholder” to the input and this can be sufficient in some applications

there are several high-level languages that can be used to express constraints. In our experience, and analogously to other programming languages, it would seem beneficial to have pre-defined types for constraint system variables. These types would be common among most or all applications. The current proposal could lead to defining the variable type *commitment* and *opening*, which could lead to better security assurances and engineering practices.

3 Background

The commit-and-prove approach in zero-knowledge proofs dates back to the works of Kilian [Kil89] and Canetti et al. [CLOS02], and has been used extensively, implicitly or explicitly, in plenty of works.

As we detail in Section 4.4, the notion of CP-ZKPs could be seen as a specialization of ZKPs by considering languages that are parametrized by the commitment key, e.g., using informal notation,

$$\mathcal{L}_{ck} = \{c_x : c_x \text{ opens to } x \text{ and } x \in \mathcal{L}\}$$

When focusing on non-interactive zero-knowledge proof systems (NIZK), in which one generates a language-dependent CRS, there is a variety of CP-ZKP notions used in the literature, such as those where *the commitment key is generated together with the CRS*, e.g., [CFH⁺15], and those where *the commitment key is taken as an input in the NIZK CRS generation* [Lip16, EG14, CFQ19], which in turn include systems where the commitment key is the CRS itself (in which case the commitment must admit a trapdoor, e.g., [EG14, Lip16]).

Given the theoretical and practical relevance of the commit-and-prove approach, we believe it is important for the community to either agree on one or at least provide a taxonomy of the variants.

4 Definitions

In this section we give definitions for commitments, zero-knowledge proofs and commit-and-prove zero-knowledge proofs. We focus on the non-interactive setting although these definitions can be adapted to the interactive setting.

The definitions of CP-ZKPs proposed in this document are based on the ones recently used in [CFQ19]. We do not necessarily mean this to be *the* definition but rather to serve as a starting point for a related discussion.

Notation. We use $\lambda \in \mathbb{N}$ to denote the security parameter, and 1^λ to denote its unary representation. Throughout the paper we assume that all the algorithms of the cryptographic schemes take as input 1^λ , and thus we omit it from the list of inputs. For a distribution D , we denote by $x \leftarrow D$ the fact that x is being sampled according to D . We remind the reader that an ensemble $\mathcal{X} = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ is a family of probability distributions over a family of domains $\mathcal{D} = \{D_\lambda\}_{\lambda \in \mathbb{N}}$. We say two ensembles $\mathcal{D} = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}' = \{D'_\lambda\}_{\lambda \in \mathbb{N}}$ are statistically indistinguishable (denoted by $\mathcal{D} \approx_s \mathcal{D}'$) if $\frac{1}{2} \sum_x |D_\lambda(x) - D'_\lambda(x)| < \text{negl}(\lambda)$. If $\mathcal{A} = \{\mathcal{A}_\lambda\}$ is a (possibly non-uniform) family of circuits and $\mathcal{D} = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ is an ensemble, then we denote by $\mathcal{A}(\mathcal{D})$ the ensemble of the outputs of $\mathcal{A}_\lambda(x)$ when $x \leftarrow D_\lambda$. We say two ensembles $\mathcal{D} = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}' = \{D'_\lambda\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable (denoted by $\mathcal{D} \approx_c \mathcal{D}'$) if for every non-uniform polynomial time distinguisher \mathcal{A} we have $\mathcal{A}(\mathcal{D}) \approx_s \mathcal{A}(\mathcal{D}')$. We denote by $[n]$ the set of integers $\{1, \dots, n\}$ and by $[:n]$ the set $\{0, 1, \dots, n-1\}$. We use $(u_j)_{j \in [\ell]}$ to denote the tuple of elements (u_1, \dots, u_ℓ) .

4.1 Relations

Let $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of polynomial-time decidable relations R on pairs (x, w) where $x \in \mathcal{D}_x$ is called the *statement* or *input*, and $w \in \mathcal{D}_w$ the *witness*. We write $R(x, w) = 1$ to denote that R holds on (x, w) , else we write $R(x, w) = 0$. When discussing schemes that prove statements on committed values we assume that \mathcal{D}_w can be split in two subdomains $\mathcal{D}_u \times \mathcal{D}_\omega$. Finally we sometimes use an even finer grained specification of \mathcal{D}_u assuming we can split it over ℓ arbitrary domains $(\mathcal{D}_1 \times \dots \times \mathcal{D}_\ell)$ for some arity ℓ . In our security definitions we assume relations to be generated by a *relation generator* $\mathcal{RG}(1^\lambda)$ that, on input the security parameter 1^λ , outputs R together with some side information, an auxiliary input aux_R , that is given to the adversary. We define \mathcal{RG}_λ as the set of all relations that can be returned by $\mathcal{RG}(1^\lambda)$.

4.2 NIZKs

We recall the definition of non-interactive zero-knowledge arguments of knowledge (NIZK, for short).

Definition 4.1 (NIZK). A NIZK for $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ is a tuple of three algorithms $\Pi = (\text{KeyGen}, \text{Prove}, \text{VerProof})$ that work as follows and satisfy the notions of completeness, knowledge soundness and zero-knowledge defined below

- $\text{KeyGen}(R) \rightarrow (\text{ek}, \text{vk})$ takes the security parameter λ and a relation $R \in \mathcal{R}_\lambda$, and outputs a common reference string consisting of an evaluation and a verification key.
- $\text{Prove}(\text{ek}, x, w) \rightarrow \pi$ takes an evaluation key for a relation R , a statement x , and a witness w such that $R(x, w)$ holds, and returns a proof π .
- $\text{VerProof}(\text{vk}, x, \pi) \rightarrow b$ takes a verification key, a statement x , and either accepts ($b = 1$) or rejects ($b = 0$) the proof π .

Completeness. For any $\lambda \in \mathbb{N}$, $R \in \mathcal{R}_\lambda$ and (x, w) such that $R(x, w)$, it holds $\Pr[(\text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R), \pi \leftarrow \text{Prove}(\text{ek}, x, w) : \text{VerProof}(\text{vk}, x, \pi) = 1] = 1$.

Knowledge Soundness. Π has knowledge soundness for \mathcal{RG} and auxiliary input distribution \mathcal{Z} , denoted $\text{KSND}(\mathcal{RG}, \mathcal{Z})$ for brevity, if for every (non-uniform) efficient adversary \mathcal{A} there exists a (non-uniform) efficient extractor \mathcal{E} such that $\Pr[\text{Game}_{\mathcal{RG}, \mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{KSND}} = 1] = \text{negl!}$. We say that Π is knowledge sound if there exists benign \mathcal{RG} and \mathcal{Z} such that Π is $\text{KSND}(\mathcal{RG}, \mathcal{Z})$.

$\text{Game}_{\mathcal{RG}, \mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{KSND}} \rightarrow b$

$(R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda)$
 $\text{crs} := (\text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R)$
 $\text{aux}_Z \leftarrow \mathcal{Z}(R, \text{aux}_R, \text{crs})$
 $(x, \pi) \leftarrow \mathcal{A}(R, \text{crs}, \text{aux}_R, \text{aux}_Z)$
 $w \leftarrow \mathcal{E}(R, \text{crs}, \text{aux}_R, \text{aux}_Z)$
 $b \leftarrow \text{VerProof}(\text{vk}, x, \pi) = 1 \wedge R(x, w) = 0$

Composable Zero-Knowledge. A scheme Π satisfies composable zero-knowledge for a relation generator \mathcal{RG} if there exists a simulator $\mathcal{S} = (\mathcal{S}_{\text{kg}}, \mathcal{S}_{\text{prv}})$ such that both following conditions hold for all adversaries \mathcal{A} :

KEYS INDISTINGUISHABILITY

$$\Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda); \\ \text{crs} \leftarrow \text{KeyGen}(R) \\ \mathcal{A}(\text{crs}, \text{aux}_R) = 1 \end{array} \right] \approx \Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda); \\ (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R) \\ \mathcal{A}(\text{crs}, \text{aux}_R) = 1 \end{array} \right]$$

PROOF INDISTINGUISHABILITY For all (x, w) such that $R(x, w) = 1$,

$$\Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda); \\ (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R); \\ \pi \leftarrow \text{Prove}(\text{ek}, x, w) \\ \mathcal{A}(\text{crs}, \text{aux}_R, \pi) = 1 \end{array} \right] \approx \Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda); \\ (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R); \\ \pi \leftarrow \mathcal{S}_{\text{prv}}(\text{crs}, \text{td}_k, x) \\ \mathcal{A}(\text{crs}, \text{aux}_R, \pi) = 1 \end{array} \right]$$

Remark 4.1 (On auxiliary inputs). In the notion of knowledge soundness defined above we consider two kinds of auxiliary inputs, aux_R generated together with the relation by \mathcal{RG} , and aux_Z that is generated from some distribution \mathcal{Z} that may depend on the common reference string that in turns depends on R . Notice that although this notion is implied by a notion where auxiliary inputs can be arbitrary, our aim is a precise formalization of auxiliary inputs; this is useful to justify why certain auxiliary inputs should be considered benign, as required to avoid known impossibility results [BCPR14, BP15]. Finally, we also note that our notion is also implied by SNARKs that admit black-box extractors (as may be the case for those relying on random oracles [Mic00]).

Remark 4.2 (On definition of zkSNARKs). One can define zero-knowledge succinct non-interactive arguments (zkSNARKs) as NIZKs enjoying an additional property, *succinctness*, i.e. if the running time of VerProof is $\text{poly}(\lambda + |x| + \log |w|)$ and the proof size is $\text{poly}(\lambda + \log |w|)$.

Remark 4.3 (On notions of knowledge-soundness). Above we use a non black-box definition of extractability. Although this is virtually necessary in the case of zkSNARKs, NIZKs can also satisfy stronger notions of (knowledge) soundness.

4.3 Commitment Schemes

We recall the notion of non-interactive commitment schemes.

Definition 4.2. A commitment scheme is a tuple of algorithms $\text{Com} = (\text{Setup}, \text{Commit}, \text{VerCommit})$ that work as follows and satisfy the notions of correctness, binding and hiding defined below.

- $\text{Setup}(1^\lambda) \rightarrow \text{ck}$ takes the security parameter and outputs a commitment key ck . This key includes descriptions of the input space \mathcal{D} , commitment space \mathcal{C} and opening space \mathcal{O} .
- $\text{Commit}(\text{ck}, u) \rightarrow (c, o)$ takes the commitment key ck and a value $u \in \mathcal{D}$, and outputs a commitment c and an opening o .
- $\text{VerCommit}(\text{ck}, c, u, o) \rightarrow b$ takes as input a commitment c , a value u and an opening o , and accepts ($b = 1$) or rejects ($b = 0$).

A commitment scheme $\text{Com} = (\text{Setup}, \text{Commit}, \text{VerCommit})$ must satisfy the following properties:

Correctness. For all $\lambda \in \mathbb{N}$ and any input $u \in \mathcal{D}$ we have:

$$\Pr \left[\begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda) \\ (c, o) \leftarrow \text{Commit}(\text{ck}, u) \end{array} : \text{VerCommit}(\text{ck}, c, u, o) = 1 \right] = 1$$

Binding. For every polynomial-time adversary \mathcal{A}

$$\Pr \left[\begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda) \\ (c, u, o, u', o') \leftarrow \mathcal{A}(\text{ck}) \end{array} : \begin{array}{l} \text{VerCommit}(\text{ck}, c, u', o') \\ \wedge \text{VerCommit}(\text{ck}, c, u, o) \\ \wedge u \neq u' \end{array} \right] = \text{negl}$$

Hiding. For $\text{ck} \leftarrow \text{Setup}(1^\lambda)$ and every values $u, u' \in \mathcal{D}$, the following two distributions are statistically close: $\text{Commit}(\text{ck}, u) \approx \text{Commit}(\text{ck}, u')$.

Remark 4.4. In literature there exists more than one syntax for commitment schemes. For example, some definitions replace the predicate **VerCommit** above with a procedure that, given a commitment c and some opening information o , outputs the committed value u (or \perp when appropriate); or other definitions define $\text{Commit}(\text{ck}, u; r)$ as a deterministic function where the input r is the randomness, which also acts as opening information, and checking that c is a commitment for message u and opening r consists simply in recomputing $\text{Commit}(\text{ck}, u; r)$.

4.4 Definition of Commit-and-Prove NIZKs

In a nutshell, a *commit-and-prove* NIZK (CP-NIZK) is a NIZK that can prove knowledge of (x, w) such that $R(x, w)$ holds with respect to a witness $w = (u, \omega)$ such that u opens a commitment c_u . Our formal definitions below essentially add some syntactic sugar to this idea in order to explicitly handle relations in which the input domain \mathcal{D}_u is more fine grained and splits over ℓ subdomains. We call these subdomains *commitment slots* following [CFQ19]. Intuitively, each item in the commitment slot represents an input (or a vector of inputs) to the relation which the prover has previously committed to. We assume that the description of the splitting is part of R 's description.

Definition 4.3 (CP-NIZK). Let $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of relations R over $\mathcal{D}_x \times \mathcal{D}_u \times \mathcal{D}_\omega$ such that \mathcal{D}_u splits over ℓ arbitrary domains $(\mathcal{D}_1 \times \dots \times \mathcal{D}_\ell)$ for some arity parameter $\ell \geq 1$. Let $\text{Com} = (\text{Setup}, \text{Commit}, \text{VerCommit})$ be a commitment scheme (as per Definition 4.2) whose input space \mathcal{D} is such that $\mathcal{D}_i \subset \mathcal{D}$ for all $i \in [\ell]$.

A *commit and prove* NIZK for Com and $\{\mathcal{R}_\lambda\}_{\lambda \in \mathbb{N}}$ is a NIZK for a family of relations $\{\mathcal{R}_\lambda^{\text{Com}}\}_{\lambda \in \mathbb{N}}$ such that:

- every $\mathbf{R} \in \mathcal{R}^{\text{Com}}$ is represented by a pair (ck, R) where $\text{ck} \in \text{Setup}(1^\lambda)$ and $R \in \mathcal{R}_\lambda$;
- \mathbf{R} is over pairs (\mathbf{x}, \mathbf{w}) where the statement is $\mathbf{x} := (x, (c_j)_{j \in [\ell]}) \in \mathcal{D}_x \times \mathcal{C}^\ell$, the witness is $\mathbf{w} := ((u_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, \omega) \in \mathcal{D}_1 \times \dots \times \mathcal{D}_\ell \times \mathcal{O}^\ell \times \mathcal{D}_\omega$, and the relation \mathbf{R} holds iff

$$\bigwedge_{j \in [\ell]} \text{VerCommit}(\text{ck}, c_j, u_j, o_j) = 1 \wedge R(x, (u_j)_{j \in [\ell]}, \omega) = 1$$

Furthermore, when we say that CP is knowledge-sound for a relation generator \mathcal{RG} and auxiliary input generator \mathcal{Z} (denoted $\text{KSND}(\mathcal{RG}, \mathcal{Z})$, for short) we mean it is a knowledge-sound NIZK for the relation generator $\mathcal{RG}_{\text{Com}}(1^\lambda)$ that runs $\text{ck} \leftarrow \text{Setup}(1^\lambda)$ and $(R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda)$, and returns $((\text{ck}, R), \text{aux}_R)$.

We denote a CP-NIZK as a tuple of algorithms $\text{CP} = (\text{KeyGen}, \text{Prove}, \text{VerProof})$. For ease of exposition, an explicit syntax for CP’s algorithms is as follows.

- $\text{KeyGen}(\text{ck}, R) \rightarrow \text{crs} := (\text{ek}, \text{vk})$
- $\text{Prove}(\text{ek}, x, (c_j)_{j \in [\ell]}, (u_j)_{j \in [\ell]}, (o_j)_{j \in [\ell]}, \omega) \rightarrow \pi$
- $\text{VerProof}(\text{vk}, x, (c_j)_{j \in [\ell]}, \pi) \rightarrow b \in \{0, 1\}$

Comparing with existing definitions In the context of non-interactive proof systems, notion of commit-and-prove schemes have appeared explicitly in a few works in the literature [Fuc11, EG14, CFH⁺15, Lip16, CFQ19]. The CP-NIZK notion given above is the one that aims to make the proof system and the commitment scheme as decoupled as possible (and in particular it makes less requirements on the commitment). It is in fact a specialization of the NIZK notion when considering specific families of relations that include verifying openings of commitments.

One difference with all the other definitions [Fuc11, EG14, CFH⁺15, Lip16] is that ours does not require the commitment to be a trapdoor commitment. In addition to this, when comparing with the definitions of Fuchsbauer [Fuc11] and Escala and Groth [EG14], we leave the possibility that the proof system has its own common reference string and that the relation to be proven takes other inputs in addition to those that are explicitly committed. Our notion is closer to the one given by Lipmaa [Lip16] (with the exception that again we do not need commitments to be trapdoor).

Finally, in comparison to the notion of commit-and-prove SNARKs defined in [CFH⁺15] for the Geppetto scheme, the main differences are the following. First, our commitment key can be generated without fixing a priori a relation (or a set of relations, e.g., a multi-QAP). Second, in the model of [CFH⁺15] one needs to commit to data using a commitment key corresponding to a specific portion of the input (in their lingo a “bank”), whereas in our model one can just commit to a vector of data, and only at proving time one assigns that data to a specific input slot. Third, in our notion the commitment scheme does not need to be a trapdoor commitment.

5 Miscellaneous Comments for the Standardization Process

- We propose that the standardization process should start from the syntax and security semantics of a commitment scheme according to the syntax used in this document (where opening is explicitly modeled as a relation).
- Hiding may not be necessary for all applications that do not require full privacy as discussed in Section 2 and Footnote 5.
- The standardization process may take into consideration to standardize specific commitment schemes used in commit-and-prove constructions, especially those that are relatively simple, more mature and already widely deployed (e.g., Pedersen commitments). This process may also include commit-and-prove variants that are similarly widespread, e.g., Merkle trees and accumulators.

References

- AGM18. Shashank Agrawal, Chaya Ganesh, and Payman Mohassel. Non-interactive zero-knowledge proofs for composite statements. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 643–673. Springer, Heidelberg, August 2018.

- BCM⁺18. Daniel Benarroch, Ran Canetti, Andrew Miller, Shashank Agrawal, Tony Arcieri, Vipin Bharathan, Josh Cincinnati, Joshua Daniel, Anuj Das Gupta, Angelo De Caro, Michael Dixon, Maria Dubovitskaya, Nathan George, Brett Hemenway Falk, Hugo Krawczyk, Jason Law, Anna Lysyanskaya, Zaki Manian, Eduardo Morais, Neha Narula, Gavin Pacini, Jonathan Rouach, Kartheek Solipuram, Mayank Varia, Douglas Wikstrom, and Aviv Zohar. Applications track proceeding. Technical report, ZKProof Standards, Berkeley, CA, May 2018. <https://zkproof.org/documents.html>.
- BCPR14. Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In David B. Shmoys, editor, *46th ACM STOC*, pages 505–514. ACM Press, May / June 2014.
- Bd94. Josh Cohen Benaloh and Michael de Mare. One-way accumulators: A decentralized alternative to digital signatures (extended abstract). In Tor Hellesest, editor, *EUROCRYPT'93*, volume 765 of *LNCS*, pages 274–285. Springer, Heidelberg, May 1994.
- BP15. Elette Boyle and Rafael Pass. Limits of extractability assumptions with distributional auxiliary input. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 236–261. Springer, Heidelberg, November / December 2015.
- But18. Vitalik Buterin. On-chain scaling to potentially 500 tx/sec through mass tx validation, 2018. <https://ethresear.ch/t/on-chain-scaling-to-potentially-500-tx-sec-through-mass-tx-validation/3477>.
- CF13. Dario Catalano and Dario Fiore. Vector commitments and their applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 55–72. Springer, Heidelberg, February / March 2013.
- CFH⁺15. Craig Costello, Cédric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. Geppetto: Versatile verifiable computation. In *2015 IEEE Symposium on Security and Privacy*, pages 253–270. IEEE Computer Society Press, May 2015.
- CFQ19. Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: Modular Design and Composition of Succinct Zero-Knowledge Proofs. Cryptology ePrint Archive, Report 2019/142, 2019. <http://eprint.iacr.org/>.
- CLOS02. Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *34th ACM STOC*, pages 494–503. ACM Press, May 2002.
- EG14. Alex Escala and Jens Groth. Fine-tuning Groth-Sahai proofs. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 630–649. Springer, Heidelberg, March 2014.
- Fuc11. Georg Fuchsbauer. Commuting signatures and verifiable encryption. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 224–245. Springer, Heidelberg, May 2011.
- GKV⁺18. Jens Groth, Yael Kalai, Muthu Venkitasubramaniam, Nir Bitansky, Ran Canetti, Henry Corrigan-Gibbs, Shafi Goldwasser, Charanjit Jutla, Yuval Ishai, Rafail Ostrovsky, Omer Paneth, Tal Rabin, Mariana Raykova, Ron Rothblum, Alessandra Scafuro, Eran Tromer, and Douglas Wikström. Security track proceeding. Technical report, ZKProof Standards, Berkeley, CA, May 2018. <https://zkproof.org/documents.html>.
- Gro16. Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.
- HBHW16. Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash protocol specification. *Tech. rep. 2016–1.10. Zerocoin Electric Coin Company, Tech. Rep.*, 2016.
- Kil89. J. Kilian. Uses of randomness in algorithms and protocols. PhD Thesis. Massachusetts Institute of Technology, 1989.
- Lip16. Helger Lipmaa. Prover-efficient commit-and-prove zero-knowledge SNARKs. In David Pointcheval, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 16*, volume 9646 of *LNCS*, pages 185–206. Springer, Heidelberg, April 2016.
- Mer88. Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, pages 369–378. Springer, Heidelberg, August 1988.
- Mic00. Silvio Micali. Computationally sound proofs. *SIAM Journal on Computing*, 30(4):1253–1298, 2000.
- Tam03. Roberto Tamassia. Authenticated data structures. In Giuseppe Di Battista and Uri Zwick, editors, *Algorithms - ESA 2003*, pages 2–5, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.

A Other Flavors of Commit-and-Prove NIZKs

In this section we describe a variant of NIZKs (proposed in [CFQ19]) that lies in between standard NIZKs and CP-NIZKs. We believe it is worth mentioning it since it captures a class of existing

schemes that are not explicitly commit-and-prove but they implicitly have a weak form of the commit flavor.

This class of schemes are called *NIZKs with commit-carrying proofs* (or commit-carrying NIZKs, cc-NIZKs for short). In a nutshell, a cc-NIZK is like a NIZK in which the proof contains a commitment to the portion u of the witness. Formalizing this idea requires to make explicit the commitment scheme associated to the NIZK, as well as the commitment key that is part of the common reference string. In [CFQ19] Campanelli, Fiore and Querol discuss how many of the existing NIZK constructions satisfy this property. In particular, this holds for popular zkSNARKs like [Gro16] They also show how cc-NIZKs can be lifted to become full fledged, composable, CP-NIZKs.

Definition A.1. *cc-NIZK*

A *cc-NIZK* is a tuple $\text{cc}\Pi$ of algorithms working as follows:

- $\text{KeyGen}(R) \rightarrow (\text{ck}, \text{ek}, \text{vk})$: the key generation takes as input the security parameter λ and a relation $R \in \mathcal{R}_\lambda$, and outputs a common reference string that includes a commitment key, an evaluation key and verification key.
- $\text{Prove}(\text{ek}, x, w) \rightarrow (c, \pi; o)$: the proving algorithm takes as input an evaluation key, a statement x and a witness $w := (u, \omega)$ such that the relation $R(x, u, \omega)$ holds, and it outputs a proof π , a commitment c and opening o such that $\text{VerCommit}(\text{ck}, c, u, o) = 1$.
- $\text{VerProof}(\text{vk}, x, c, \pi) \rightarrow b$: the verification algorithm takes a verification key, a statement x , a commitment c , and either accepts ($b = 1$) or rejects ($b = 0$) the proof π .
- $\text{VerCommit}(\text{ck}, c, u, o) \rightarrow b$: the commitment verification algorithm takes as input a commitment key, a commitment c , a message u and an opening o and accepts ($b = 1$) or rejects ($b = 0$).

Completeness. For any $\lambda \in \mathbb{N}$, $R \in \mathcal{R}_\lambda$ and (x, w) such that $R(x, w)$, it holds

$$\Pr \left[\begin{array}{l} (\text{ck}, \text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R); \\ (c, \pi; o) \leftarrow \text{Prove}(\text{ek}, x, w) \end{array} : \text{VerProof}(\text{vk}, x, c, \pi) \right] = 1$$

Knowledge Soundness. Let \mathcal{RG} be a relation generator such that $\mathcal{RG}_\lambda \subseteq \mathcal{R}_\lambda$. $\text{cc}\Pi$ satisfies knowledge soundness for \mathcal{RG} and auxiliary input distribution \mathcal{Z} , or $\text{ccKSND}(\mathcal{RG}, \mathcal{Z})$, if there exists a (non-uniform) efficient extractor \mathcal{E} that for every (non-uniform) efficient adversary \mathcal{A} is such that $\Pr[\text{Game}_{\mathcal{RG}, \mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{ccKSND}} = 1] = \text{negl}$. We say that $\text{cc}\Pi$ is knowledge sound if there exist benign \mathcal{RG} and \mathcal{Z} such that $\text{cc}\Pi$ is $\text{ccKSND}(\mathcal{RG}, \mathcal{Z})$.

$$\begin{array}{l} \text{Game}_{\mathcal{RG}, \mathcal{Z}, \mathcal{A}, \mathcal{E}}^{\text{ccKSND}} \rightarrow b \in \{0, 1\} \\ \hline (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda) \\ \text{crs} := (\text{ck}, \text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R) \\ \text{aux}_Z \leftarrow \mathcal{Z}(R, \text{aux}_R, \text{crs}) \\ (x, c, \pi) \leftarrow \mathcal{A}(R, \text{crs}, \text{aux}_R, \text{aux}_Z) \\ (u, o, \omega) \leftarrow \mathcal{E}^{\mathcal{A}}(R, \text{crs}, \text{aux}_R, \text{aux}_Z) \\ b \leftarrow \text{VerProof}(\text{vk}, x, c, \pi) = 1 \wedge (\text{VerCommit}(\text{ck}, c, u, o) = 0 \vee R(x, u, \omega) = 0) \end{array}$$

Composable Zero-Knowledge. A scheme $\text{cc}\Pi$ has composable zero-knowledge for a relation generator \mathcal{RG} if for every adversary \mathcal{A} there exists a simulator $\mathcal{S} = (\mathcal{S}_{\text{kg}}, \mathcal{S}_{\text{prv}})$ such that both following conditions hold for all adversaries \mathcal{A} :

KEYS INDISTINGUISHABILITY.

$$\begin{aligned} & \Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda); \\ \text{crs} \leftarrow \text{KeyGen}(R) \end{array} : \mathcal{A}(\text{crs}, \text{aux}_R) = 1 \right] \\ & \approx \Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda); \\ (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R) \end{array} : \mathcal{A}(\text{crs}, \text{aux}_R) = 1 \right] \end{aligned}$$

PROOF INDISTINGUISHABILITY. For all (x, w) ,

$$\begin{aligned} & \Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda); \\ (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R); \\ (c, \pi; o) \leftarrow \text{Prove}(\text{ek}, x, w) \end{array} : \begin{array}{l} \mathcal{A}(\text{crs}, \text{aux}_R, c, \pi) = 1 \wedge \\ R(x, w) = 1 \end{array} \right] \\ & \approx \Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda); \\ (\text{crs}, \text{td}_k) \leftarrow \mathcal{S}_{\text{kg}}(R); \\ (c, \pi) \leftarrow \mathcal{S}_{\text{prv}}(\text{crs}, \text{td}_k, x) \end{array} : \begin{array}{l} \mathcal{A}(\text{crs}, \text{aux}_R, c, \pi) = 1 \wedge \\ R(x, w) = 1 \end{array} \right] \end{aligned}$$

Binding. For every polynomial-time adversary \mathcal{A} the following probability is $\text{negl}(\lambda)$:

$$\Pr \left[\begin{array}{l} (R, \text{aux}_R) \leftarrow \mathcal{RG}(1^\lambda) \\ \text{crs} := (\text{ck}, \text{ek}, \text{vk}) \leftarrow \text{KeyGen}(R) \\ (c, u, o, u', o') \leftarrow \mathcal{A}(R, \text{crs}, \text{aux}_R) \end{array} : \begin{array}{l} \text{VerCommit}(\text{ck}, c, u', o') \\ \wedge \text{VerCommit}(\text{ck}, c, u, o) \\ \wedge u \neq u' \end{array} \right]$$

Remark A.1. While our definitions consider the case where the proof contains a commitment to a portion u of the witness $w = (u, \omega)$, notice that this partition of the witness is arbitrary and thus this notion also captures those constructions where the commitment is to the entire witness if one thinks of a void ω .