

Non-interactive Commitment from Non-transitive Group Actions

Original

Non-interactive Commitment from Non-transitive Group Actions / D'Alconzo, Giuseppe; Flamini, Andrea; Gangemi, Andrea. - 14444:(2023), pp. 222-252. (29th Annual International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT) Guangzhou (China) December 4–8, 2023) [10.1007/978-981-99-8739-9_8].

Availability:

This version is available at: 11583/2984587 since: 2023-12-20T14:20:19Z

Publisher:

Springer

Published

DOI:10.1007/978-981-99-8739-9_8

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Springer postprint/Author's Accepted Manuscript (book chapters)

This is a post-peer-review, pre-copyedit version of a book chapter published in Advances in Cryptology – ASIACRYPT 2023. The final authenticated version is available online at: http://dx.doi.org/10.1007/978-981-99-8739-9_8

(Article begins on next page)

Non-Interactive Commitment from Non-Transitive Group Actions

Giuseppe D'Alconzo¹[0000-0001-7377-6617], Andrea
Flamini²[0000-0002-3872-7251], and Andrea Gangemi²[0000-0001-9689-8473]

¹ Department of Mathematical Sciences, Politecnico di Torino, Corso Duca degli
Abruzzi 24, 10129 Torino, Italy

² Department of Mathematics, University of Trento, Povo, 38123 Trento, Italy
giuseppe.dalconzo@polito.it, {andrea.flamini, andrea.gangemi}@unitn.it

Abstract. Group actions are becoming a viable option for post-quantum cryptography assumptions. Indeed, in recent years some works have shown how to construct primitives from assumptions based on isogenies of elliptic curves, such as CSIDH, on tensors or on code equivalence problems. This paper presents a bit commitment scheme, built on non-transitive group actions, which is shown to be secure in the standard model, under the decisional Group Action Inversion Problem. In particular, the commitment is computationally hiding and perfectly binding, and is obtained from a novel and general framework that exploits the properties of some orbit-invariant functions, together with group actions. Previous constructions depend on an interaction between the sender and the receiver in the commitment phase, which results in an interactive bit commitment. We instead propose the first non-interactive bit commitment based on group actions. Then we show that, when the sender is honest, the constructed commitment enjoys an additional feature, i.e., it is possible to tell whether two commitments were obtained from the same input, without revealing the input. We define the security properties that such a construction must satisfy, and we call this primitive *linkable commitment*. Finally, as an example, an instantiation of the scheme using tensors with coefficients in a finite field is provided. In this case, the invariant function is the computation of the rank of a tensor, and the cryptographic assumption is related to the Tensor Isomorphism problem.

Keywords: *Cryptographic group actions · Non-transitive group actions · Bit commitments · Linkable commitments · Tensors*

1 Introduction

Group Actions in Cryptography. Recent developments in quantum computing make the advent of a quantum machine suitable for cryptanalysis purposes a threat. Many cryptographic algorithms that are used nowadays can no longer be considered secure against a quantum attacker. Primitives relying on the hardness of the Discrete Logarithm or the Factorization problem are broken by the

well known Shor’s algorithm [33]. This leads to the birth of the Post-Quantum Cryptography, that aims to find and study protocols based on cryptographic assumptions that appear to be resistant to attacks performed by quantum computers. The most promising ones are based on lattices, multivariate polynomials, hash functions, error correcting codes and isogenies of elliptic curves. However, in order to increase the variety of probably secure assumptions, it is necessary to find new problems with useful features to build new cryptographic protocols. A recent line of study concerns equivalence problems and cryptographic group actions. The most known reference is given by Couveignes in 2006 [11] and was used in the setting of isogeny-based cryptography. Moreover, the explicit use of group actions can be found in the 1991 article of Brassard and Young [8]. More recently, the framework has been studied by Grigoriev and Shpilrain, [15], Alamati, De Feo, Montgomery and Patranabis [1] and Ji, Qiao, Feng and Yun [20], introducing some formal cryptographic assumptions. There are many group actions suitable for post-quantum cryptography, arising from different areas of mathematics and computer science. Some examples can be the class group action of CSIDH [9], the one induced by the general linear group on various objects [20, 30, 35], the action acting on polynomials [26] or the ones concerning linear codes [3, 30]. In the last years, cryptographic group actions have been employed to design many primitives such as sigma protocols and signature schemes (via the GMW scheme for Graph Isomorphism [14]), ring and group signatures [4, 5], key exchanges [9] and updatable encryption schemes [22].

Commitment Schemes. A commitment scheme is a cryptographic protocol between two parties, a sender and a receiver. The sender wants to commit to a value b without revealing it to the other party. To do this, he binds b to a commitment C that is sent to the receiver. In a second moment, the sender wants to reveal b and the receiver must be able to verify that it was the committed value behind C . A commitment must satisfy two security properties: it must not reveal any information about the committed value (hiding property), and the sender cannot reveal a different $b' \neq b$ that opens to the same commitment (binding property).

Commitment schemes are widely used, both as stand-alone protocols and as atomic parts of more involved mechanisms. For example, they are used in Zero-Knowledge proofs [23], digital auctions [25], signature schemes [21], multi-party computation [13], e-voting [12] and confidential transactions [29]. In this work, we will mainly focus on *bit* commitments, where the committed value b can be 1 or 0.

Related works. Bit commitments schemes are a component of many cryptographic algorithms. In 1991, Naor [24] showed how to obtain a bit commitment protocol starting from a pseudorandom generator. Bit commitments from group actions are known in literature. In 1991 Brassard and Young [8] present an interactive scheme from certified and uncertified group actions. In 2019 Ji, Qiao, Song and Yun [20] present, among other construction, two interactive bit commitment

schemes relying on cryptographic assumptions on non-abelian group actions. Finally, another famous commitment, which is however based on a pre-quantum assumption, is the *Pedersen commitment* [27]. This scheme has an interesting property: it can be shown that two commitments are created starting from the same value, without opening the commitments [29].

Our Contribution. We present a bit commitment scheme that is non-interactive, perfectly binding and computationally hiding in the standard model. This scheme is based on a group action framework that makes use of certain invariant functions. One of the innovative aspects of our proposal is that it concerns *non-transitive* group actions, while known cryptographic applications use transitive actions or they restrict to one orbit. The non-transitivity of the action used in this paper is crucial and necessary; in fact, we need to be able to exhibit two elements that are in two different orbits. Such elements are generated with the aid of the new group action framework, in which we endow the group action with a function that is constant inside the orbits. Given the group G acting on the set X via the action \star , an *invariant function* $f : X \rightarrow T$, with T be a set, has the following property

$$f(g \star x) = f(x), \quad \forall x \in X, g \in G.$$

The key point is that evaluating this function on a randomly chosen element is hard, while, for a particular subset of elements that we call *canonical elements*, it is easy to compute. Also, the fact that the function is constant inside the orbits guarantees that, if we consider two elements with distinct image, they must live in (and generate under the action of G) distinct orbits. This observation is crucial to prove our commitment scheme is perfectly binding. We call *Group Action with Canonical Elements* (GACE) a group action with the above properties. Moreover, the existence of decision problems about whether an element is randomly picked from a specific orbit or not enables us to prove that our commitment scheme is computationally hiding.

The structure of our construction enables an additional property that is shared with the Pedersen commitment. An honest sender generating two commitments of the same value b can prove to the receiver that they are in fact linked to the same message, without revealing it. We call this scheme a *linkable commitment* and we formally define the security properties that enable the adoption of such a primitive in cryptography. However, using some techniques from ring signature schemes [5], we show how to extend this property to the case of a possibly malicious sender in the Random Oracle Model.

This work is organized as follows: Section 2 recaps all the cryptographic tools that will be used in the rest of the paper, while Section 3 introduces the framework that we will use to design a non-interactive commitment scheme starting from cryptographic group actions. In particular, we introduce the concept of Group Action with Canonical Elements. Section 4 shows how to design a bit commitment starting from canonical elements, and its security is proved under the decisional Group Action Inversion Problem assumption, while in Section 5

we introduce the notion of linkable commitments and we show how our protocol is indeed a linkable one. Section 6 shows an instantiation of the framework with tensors, and finally, Section 7 concludes the work and gives some idea for further research.

2 Preliminaries

In the course of this paper, with $\Pr[A]$ we denote the probability of the event A . Let λ denote the security parameter, this means that the parameters of the cryptographic schemes instantiated with security parameter λ are chosen in such a way that the best known attack would break the scheme using at least 2^λ operations. A function $\mu(\lambda)$ is *negligible* in λ if for every positive integer c there exists a λ_0 such that for each $\lambda > \lambda_0$ we get $\mu(\lambda) < \frac{1}{\lambda^c}$. Finally, in the pseudocode “ \leftarrow ” denotes the random sampling, “ \leftarrow ” is a variable assignment and “ $=$ ” is the equality check.

2.1 Group Actions

This section introduces group actions, along with the complexity assumptions that must be made in order to use them in cryptographic protocols. Definitions reported here are mostly taken from [1]. We point out that through this work we do not need the action to be abelian, contrary to what is required in [11] or [1]. All the following definitions and constructions are meaningful also in the non-abelian case.

Definition 1. *A group G is said to act on a set X if there is a map $\star : G \times X \rightarrow X$ that satisfies the following properties:*

- Identity: *if e is the identity element of the group G , then $e \star x = x$ for every x in X .*
- Compatibility: *given g and h in G and x in X , we have that $(gh) \star x = g \star (h \star x)$.*

In this case, we say that the triple (G, X, \star) is a group action.

A group action (G, X, \star) may satisfy some algebraic properties that lead to the definition of classes of group actions, namely the action is *transitive* if for all x_1, x_2 in X there exists an element g in G such that $x_1 = g \star x_2$; moreover, the action is said *free* when the following holds: g is the identity element of G if and only if there is an x in X such that $g \star x = x$. Finally, we say that the action is *regular* if it is both free and transitive.

Note that, if the group action (G, X, \star) is regular and the group G is finite, then for every x in X the map $g \mapsto g \star x$ is a bijection and $|G| = |X|$. Furthermore, if the group action is regular, then we can define the element $\delta(x, y)$ of G as the unique element for which $x = \delta(x, y) \star y$. If the action is not transitive, instead, then there exist x and y in X such that $\delta(x, y)$ does not exist.

Alamati, De Feo, Montgomery and Patranabis also define the concept of *effective group action*: a formal definition can be found in [1], here we just report the key points.

Definition 2. A group action (G, X, \star) is effective if:

- the group G is finite and there exists a probabilistic polynomial time (PPT) algorithm for executing membership and equality testing, sampling, and for computing the group operation and the inverse of an element;
- the set X is finite and there exist PPT algorithms for computing membership testing and the unique representation of any element in X ;
- there exists an efficient algorithm to compute $g \star x$, for each g in G and x in X .

Informally, a group action is said effective if it can be manipulated easily and it can be computed in practical time. An example of non-effective group actions is the set of polynomials in m variables of bounded degree n over a finite field, with the symmetric group \mathcal{S}_m , permuting the variables. It can be seen that the unique representation is given by the algebraic normal form, but it cannot be computed in polynomial time in n and m .

In the rest of this work, even when not explicitly written, we will consider effective group actions.

2.2 Cryptographic assumptions on group actions

The presented definition leads to efficient group actions, which can be used to build cryptographic protocols. However, in order to use them in cryptography we need to define some suitable computational assumptions. In [1], the authors report some computational assumptions on group actions, for instance the following embraces the fact that, given two random elements $x, y \in X$ in the same orbit, then it must be intractable to compute $\delta(x, y)$.

Definition 3. Let λ be a parameter indexing G and X . Being \mathcal{D}_G and \mathcal{D}_X two distributions over G and X respectively, then the group action (G, X, \star) is $(\mathcal{D}_G, \mathcal{D}_X)$ -one-way if for all PPT adversaries \mathcal{A} there is a negligible function $\mu(\lambda)$ such that

$$\Pr[\mathcal{A}(x, g \star x) \star x = g \star x] \leq \mu(\lambda),$$

where x is sampled according to \mathcal{D}_X and g according to \mathcal{D}_G .

In this paper, we assume that \mathcal{D}_G and \mathcal{D}_X are the uniform distributions over G and X , and we refer to this assumption as *One-way group action assumption*.

Another assumption that can be used when working with group actions is the *Group Action Pseudo Randomness* (GA-PR) problem, defined in [20]. It can be seen as a generalisation of the Decisional Diffie-Hellman assumption. An equivalent assumption can be found in [1], and a group action with this property is called *weakly pseudorandom*. For example, in [20], the authors state that it can be applied to the general linear group action on tensors. Let us now define more formally the problem on which the GA-PR assumption is based.

whether t lives in the same orbit of s or not, cannot win the game with probability greater than $\frac{3}{4}$.

The observation above motivates the introduction of an assumption which we refer to as *decisional Group Action Inversion Problem* (dGA-IP). The dGA-IP problem, also known as Isomorphism Problem [20], is the decisional variant of the group action inversion problem presented in [34], applied to the case in which the set X is given by only two orbits. If the restriction on the two orbits is removed, a large number of similar problems can be found in literature [16, 17, 28].

Definition 5. *The dGA-IP game is presented in Figure 2, where pp is given by the tuple (G, X, \star, t_0, t_1) , with t_0 and t_1 elements that lie in distinct orbits under the action of G . We define the advantage of an adversary \mathcal{A} of dGA-IP as*

$$\text{Adv}(\mathcal{A}, \text{dGA-IP}) = \left| \Pr[\mathcal{A} \text{ wins dGA-IP}(\text{pp})] - \frac{1}{2} \right|.$$

The dGA-IP assumption states that for all PPT adversaries \mathcal{A} there is a negligible function $\mu(\lambda)$, with λ being the security parameter, such that

$$\text{Adv}(\mathcal{A}, \text{dGA-IP}) \leq \mu(\lambda),$$

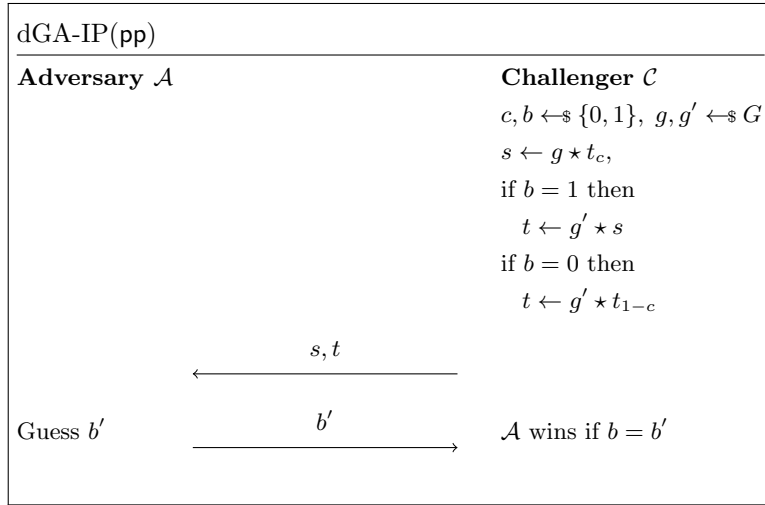


Fig. 2. decisional Group Action Inversion Problem game.

This game, compared to 2GA-PR, reflects more clearly the fact that it is hard to distinguish whether two elements in X live in the same orbit or not, and an adversary with unbounded computational power would win this game with probability 1.

2.3 Commitment schemes

A commitment scheme is a cryptographic scheme that allows one party to commit to a value m by sending a commitment com , and then to reveal m by opening the commitment at a later point in time.

Definition 6. A commitment scheme on a message space \mathcal{M} is a triple of PPT algorithms $(\text{PGen}, \text{Commit}, \text{Open})$ such that:

1. $\text{PGen}(1^\lambda)$ takes as input a security parameter λ in unary and returns public parameters pp ;
2. $\text{Commit}(\text{pp}, m)$ takes as input the public parameters pp , a message m in \mathcal{M} and returns the commitment com and the opening material r ;
3. $\text{Open}(\text{pp}, m, \text{com}, r)$ takes as input the public parameters pp , the message m , the commitment com and the opening material r and returns **accept** if com is the commitment of m or **reject** otherwise.

In the rest of this work we omit the public parameters pp in the inputs of Commit and Open .

To be suitable in cryptography, commitment schemes must satisfy the *hiding* and *binding* properties. Hiding means that com reveals nothing about m and binding means that it is not possible to create a commitment com that can be opened in two different ways. These properties are formally defined.

Definition 7. Let $\Pi_{\text{Com}} = (\text{PGen}, \text{Commit}, \text{Open})$ be a commitment scheme and let $\text{Hiding}(\Pi_{\text{Com}})$ be the hiding game represented in Figure 3. We define the advantage of an adversary \mathcal{A} of $\text{Hiding}(\Pi_{\text{Com}})$ as

$$\text{Adv}(\mathcal{A}, \text{Hiding}(\Pi_{\text{Com}})) = \left| \Pr[\mathcal{A} \text{ wins } \text{Hiding}(\Pi_{\text{Com}})] - \frac{1}{2} \right|.$$

A commitment scheme Π_{Com} is computationally hiding if for all PPT adversaries \mathcal{A} there is a negligible function $\mu(\lambda)$, with λ being the security parameter, such that

$$\text{Adv}(\mathcal{A}, \text{Hiding}(\Pi_{\text{Com}})) \leq \mu(\lambda),$$

If, for every pair m_0, m_1 , the commitments com_0 and com_1 have the same distribution, where $(\text{com}_i, r_i) = \text{Commit}(m_i)$ for $i = 0, 1$, we say that the commitment is perfectly hiding.

Note that, in the case of a bit commitment, the adversary does not send m_0 and m_1 , and the bit chosen by the challenger is the committed bit in com .

Definition 8. A commitment scheme $\Pi_{\text{Com}} = (\text{PGen}, \text{Commit}, \text{Open})$ is computationally binding if for all PPT adversaries \mathcal{A} there is a negligible function $\mu(\lambda)$, with λ being the security parameter, such that

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{PGen}(1^\lambda), \\ (\text{com}, m_0, r_0, m_1, r_1) \leftarrow \mathcal{A}(\text{pp}) \end{array} \middle| \begin{array}{l} m_0 \neq m_1, \\ \text{Open}(m_0, \text{com}, r_0) = \text{accept}, \\ \text{Open}(m_1, \text{com}, r_1) = \text{accept} \end{array} \right] \leq \mu(\lambda).$$

If for every adversary \mathcal{A} it holds that $\mu(\lambda) = 0$, we say that the commitment scheme is perfectly binding.

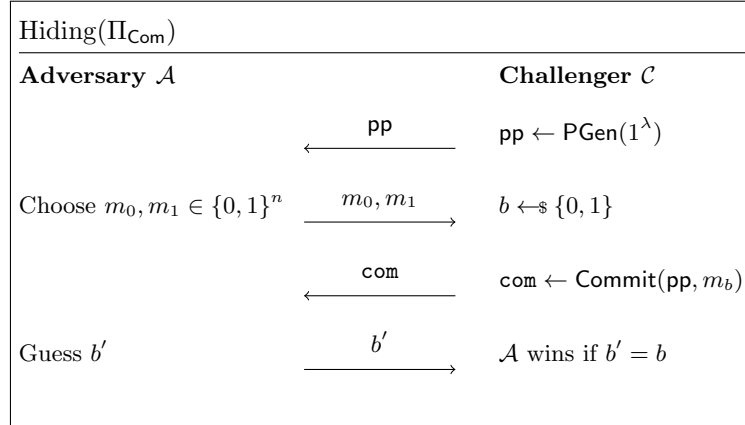


Fig. 3. Hiding game for commitment schemes.

Commitment schemes from Group Actions Previous commitments were known from cryptographic group actions. Brassard and Young [8] propose two kind of bit commitments from what they call *certified* and *uncertified* group actions. A certified group action is an action from the group G over the set X such that checking that two elements are in the same orbit is an easy task. On the contrary, the same verification could not be polynomial-time for an uncertified group action. Since the problem of deciding whether two elements of X are in the same orbit is assumed to be hard in this work, we will focus on the latter case. Given a group action from G on X , the computationally binding and perfectly hiding bit commitment presented in [8] is as follows.

- The receiver randomly generates x_0 from X and g from G . Then sets x_1 as $g \star x_0$. He sends to the sender the pair (x_0, x_1) and a proof π that they are in the same orbit.
- The sender wants to commit to the bit b . First, he checks that the proof π is valid, then he picks h from G and sends com = $h \star x_b$ to the receiver, keeping secret h .
- To open the committed bit b , the sender reveals b and h to the receiver, which checks that com is equal to $h \star x_b$.

The first thing to notice is that this is an interactive bit commitment, since the sender needs the receiver's cooperation for the creation of the commitment. Secondly, the communication cost is at least as big as the proof of the statement that x_0 and x_1 are in the same orbit. This is an NP-statement (the witness is given by g) and hence admits an interactive proof (even a non-interactive one, using the Fiat-Shamir heuristic and the Random Oracle Model), but it can be very large in communication.

In [20], Ji, Qiao, Song and Yun propose two bit commitment protocols. The first is a slight generalization of the protocol from [8], using non-abelian group

actions. The obtained protocol has the same drawbacks noticed above: it is interactive and has a large communication cost. The second proposal concerns the use of the following pseudorandom function

$$f : X \times G \rightarrow X \times X, \quad (x, g) \mapsto (x, g \star x)$$

and, after applying the Blum-Micali amplification [7], the authors build an interactive bit commitment scheme using the construction from [24]. In this construction it is needed that $|X| \geq |G|$, and the obtained bit commitment is statistically binding and computationally hiding.

3 Our Framework

The goal of this section is to design a non-interactive commitment scheme using assumptions from cryptographic group actions. We will focus on non-abelian and non-transitive actions. To develop such a commitment scheme, we first analyze the issues arising from an initial construction, then we define a framework that we use to circumvent these problems.

3.1 A first attempt

Based on the non-transitivity of the group action (G, X, \star) , we can do a first attempt in building a *non-interactive* bit commitment scheme. We give its description using a trusted third party (TTP), and then we analyze how to remove it.

Given the action (G, X, \star) , the TTP chooses and publishes two elements x_0 and x_1 of X lying in different orbits. The sender, to commit a bit b , generates a random g in G and sets as the commitment of b the value $\text{com} = g \star x_b$. The opening material is g . In other words, the sender picks a random element in the orbit of x_b . In the opening phase, given b , com and g , the receiver accepts if com is equal to $g \star x_b$ and rejects otherwise. Informally, the hiding property is given by the fact that checking whether com is in the orbit of x_0 or x_1 is hard, while the binding property follows from the impossibility of going from an orbit to another via the action of G .

In the following we try to remove the TTP and analyze some possible scenarios.

1. **The sender generates and publishes x_0 and x_1 .** In this case we can see that a malicious sender can generate x_0 and x_1 in the same orbit via $x_1 = h \star x_0$. He commits to $g \star x_0$ and, during the opening phase, he could open to both 0 and 1 using g or gh^{-1} . In this case, the binding property does not hold.
2. **The sender generates and publishes x_0, x_1 together with a proof π that they are in different orbits.** Given a proof π that x_0 and x_1 are not in the same orbit, we obtain that the protocol is hiding and binding, under the assumption that deciding whenever two elements share the orbit is hard.

In this scenario, the hard task is the generation of the proof π . In fact, the language

$$L = \{(y_0, y_1) \in X \times X \mid y_0 \text{ and } y_1 \text{ are in different orbits}\}$$

is in coNP . Unless we have a computationally unbounded prover [14] (and this is not the case), it means that known techniques fails to generate a short non-interactive proof for L which would enable the design of a non-interactive commitment scheme. Since interactive bit commitments based on group actions are known [8, 20], we do not further study this case.

3. **The receiver generates and publishes x_0 and x_1 .** We are again in the case of interactive bit commitments, and we remand to the known schemes based on group actions.

With such techniques, we have seen that there are some tricky aspects that are hard to deal with. For example, we need to build a proof for a language in coNP , and the absence of a witness (as we are used to, when we work in NP) is the first obstacle. To overcome such difficulties, we introduce a general framework on group actions that ease the design of the non-interactive bit commitment sketched above. The trick is the definition of an invariant function that is constant inside the orbits and hard to compute for a randomly chosen element. However, we assume that there is a set of representative elements for which the computation of such a function is easy. This avoids the need of a proof for the above language L . These concepts will be formalized in the next subsection.

3.2 Group Actions with Canonical Elements

In this section, we introduce the concepts of invariant functions and canonical elements, and we present the cryptographic assumptions linked to them.

Definition 9. *Given a group action (G, X, \star) and a function $f : X \rightarrow T$, we say that f is invariant under the action of G if $f(g \star x) = f(x)$ for every g in G and every x in X . We say that f is fully invariant if $f(x) = f(y)$ if and only if there exists g in G such that $y = g \star x$.*

In the following, we can assume that f is surjective, restricting the set T to the image $f(X)$. To exploit the properties of invariant functions while keeping the dGA-IP hard, we want the function f to be hard to compute on a large class of elements of X . At the same time, we want to define particular elements of X on which the computation of f is feasible.

Definition 10. *Let $f : X \rightarrow T$ be a surjective invariant function for the action (G, X, \star) and let $T' \subset T$. Suppose that there exists a polynomial-computable map*

$$\langle \cdot \rangle : T' \rightarrow X, \quad t \mapsto \langle t \rangle$$

such that the function $f \circ \langle \cdot \rangle$ is the identity on the subset T' of T . We call $\langle \cdot \rangle$ the canonical representation of T' in X and $\langle t \rangle$ the canonical t -element (with

respect to f and $\langle \cdot \rangle$. If $T' = T$, we say that $\langle \cdot \rangle$ is complete. Moreover, we say that $(G, X, \star, f, \langle \cdot \rangle)$ is a Group Action with Canonical Element (GACE) if the following hold:

1. if $O(z)$ is the orbit of z in X , then for any PPT adversary \mathcal{A} there is a negligible function μ such that

$$\Pr[\mathcal{A}(x) = f(x)] \leq \frac{1}{|T'|} + \mu(|x|),$$

where x is sampled uniformly random from $\bigsqcup_{t \in T'} O(\langle t \rangle)$;

2. there is a PPT algorithm that for any t in T' computes $f(\langle t \rangle)$.

In other words, the definition above says that, for every t in T' , we have $f(\langle t \rangle) = t$ and the function f is hard to compute in general, but is instead easy to calculate on canonical elements. Moreover, the construction of such $\langle t \rangle$ is a polynomial-time task.

In the following constructions, whenever a random element of X is needed, we pick a random canonical element $\langle t \rangle$, a random g from G and compute $g \star \langle t \rangle$. In this way, instead of using the whole X , we always work with the union of the orbits of the canonical elements. In other words, the set on which the group G acts becomes

$$X' = \bigsqcup_{t \in T'} O(\langle t \rangle).$$

This implies that the GACE $(G, X', \star, f, \langle \cdot \rangle)$ has a fully invariant function f and the canonical representation $\langle \cdot \rangle$ is complete. Given a fully invariant function f , the problem of determining whether two elements have the same image under f is equivalent to deciding whether they lie in the same orbit (dGA-IP).

4 The Commitment Scheme

4.1 Bit commitment scheme from a GACE

The first application of our framework is a bit commitment scheme. Given a Group Action with Canonical Elements, we design the commitment scheme described in Figure 4, following the attempts shown in Subsection 3.1. The bit commitment is proven secure under both the dGA-IP assumption that we have introduced in this paper and the 2GA-PR assumption; the security proof under the latter assumption can be found in Appendix A.

Theorem 1. *The bit commitment scheme in Figure 4 is perfectly binding.*

Proof. Without loss of generality, we can assume $m_0 = 0$ and $m_1 = 1$. Suppose there exists an adversary \mathcal{A} that on input $\text{pp} = (G, X, \star, f, \langle \cdot \rangle, t_0, t_1)$ returns the tuple com, r_0, r_1 such that

$$\text{Open}(0, \text{com}, r_0) = \text{Open}(1, \text{com}, r_1) = \text{accept}$$

PGen (1^λ)	
1 : choose $(G, X, \star, f, \langle \cdot \rangle)$	
2 : $t_0 \leftarrow \$ T'$	
3 : $t_1 \leftarrow \$ T' \setminus \{t_0\}$	
4 : return $(G, X, \star, f, \langle \cdot \rangle, t_0, t_1)$	
Commit (b)	Open (b, c, g)
1 : $g \leftarrow \$ G$	1 : if $g^{-1} \star c = \langle t_b \rangle$
2 : $c \leftarrow g \star \langle t_b \rangle$	2 : return accept
3 : return (c, g)	3 : else return reject

Fig. 4. Bit commitment scheme from a GACE.

with positive probability. This means that $r_0 \star \langle t_0 \rangle = \text{com} = r_1 \star \langle t_1 \rangle$, and then $r_1^{-1} r_0 \star \langle t_0 \rangle = \langle t_1 \rangle$. Therefore, $\langle t_0 \rangle$ and $\langle t_1 \rangle$ are in the same orbit, but this is a contradiction and such an adversary \mathcal{A} cannot exist.

Theorem 2. *The bit commitment scheme in Figure 4 is computationally hiding under the decisional Group Action Inversion Problem assumption.*

Proof. The dGA-IP assumption states that every adversary of the dGA-IP game has at most negligible advantage. We prove that the existence of an adversary of the game Hiding(Π_{Com}) with advantage at least $\epsilon(\lambda)$, where $\epsilon(\lambda)$ is a non-negligible function, implies the existence of an adversary \mathcal{A} of the dGA-IP game with advantage $2\epsilon^2(\lambda)$, which is non-negligible.

The proof is divided in 3 parts: firstly, we describe our adversary \mathcal{A} of the dGA-IP game. It will exploit two instances of an adversary of the Hiding(Π_{Com}) game, therefore we must show that it correctly simulates the challenger of such a game. Finally, we quantify a lower bound to the advantage of the adversary \mathcal{A} .

1. *Reduction description.*

The adversary \mathcal{A} of the dGA-IP game (see Figure 5) receives from the challenger two set elements s and t , generated according to the dGA-IP game. \mathcal{A} creates two instances of the adversary of Hiding(Π_{Com}) game having non-negligible advantage, namely \mathcal{A}_1 and \mathcal{A}_2 . Then, the adversary \mathcal{A} provides \mathcal{A}_1 with s and \mathcal{A}_2 with t separately. The two hiding commitment adversaries \mathcal{A}_1 and \mathcal{A}_2 return respectively the bits b_0 and b_1 as outputs of their internal routine. Finally, the dGA-IP adversary \mathcal{A} returns to the challenger the bit b' which is set to 1 if $b_0 = b_1$, otherwise it is set to 0.

2. *\mathcal{A} correctly simulates the Hiding(Π_{Com}) challenger.*

We show that \mathcal{A} correctly simulates the challenger of the Hiding(Π_{Com}) game, so that it is possible to quantify the probability of success of the adversaries

\mathcal{A}_1 and \mathcal{A}_2 . The elements s and t which \mathcal{A} uses as input to \mathcal{A}_1 and \mathcal{A}_2 are generated as follows:

- s is a random element in the orbit generated by $\langle t_c \rangle$, with c chosen uniformly at random in $\{0, 1\}$;
- when $b = 1$, t is chosen uniformly at random in the same orbit of s (note that $g' \star s = g'g \star \langle t_c \rangle$ is random as long as $g' \leftarrow G$), otherwise, if $b = 0$, t is chosen at random in the orbit of $\langle t_{1-c} \rangle$.

In particular, the orbit of s is chosen uniformly at random via the selection of c ; then, given c , the orbit of t is chosen uniformly at random via b . This guarantees that \mathcal{A} correctly simulates the challenger of the Hiding(II_{Com}) game, who must choose, in the first step, whether to create a commitment to 0 or to 1. Therefore, the adversaries $\mathcal{A}_1, \mathcal{A}_2$ win their games with probability greater than $\frac{1}{2} + \epsilon(\lambda)$.

3. Measurement of \mathcal{A} 's advantage.

Finally, we compute a lower bound to the probability of success of \mathcal{A} that we have described in the dGA-IP game.

We observe that the adversaries \mathcal{A}_1 and \mathcal{A}_2 do not interact, so the events that they win their games can be considered independent as long as their inputs are also independent.

It is possible to show that the selection of the inputs is independent, since the selection process of s and t is performed picking at random the orbit $O(s)$ of s by sampling the bit c , and the orbit $O(t)$ of t by sampling the bit b (actually the bit that determines the orbit of t is interpreted according to the value of s , but this is not relevant as long as the bit b is chosen at random).

Then, the canonical elements of the sampled orbits are randomized by sampling two random group elements $g, g' \in G$ and computing the action of such elements (or of the element $g'g$ instead of g' , if $b = 1$, which is a random element as long as g' is random) on the canonical elements.

Given that the inputs to \mathcal{A}_1 and \mathcal{A}_2 are independent and that the two adversaries perform their operations regardless of the existence of each other, the events that \mathcal{A}_1 wins its game and \mathcal{A}_2 wins its game are independent.

For the sake of brevity, we refer to the event that \mathcal{A}_1 wins or loses its game as (\mathcal{A}_1 wins) or (\mathcal{A}_1 loses) and we do the same for \mathcal{A}_2 and \mathcal{A} : the game they are playing will be clear from the context.

Finally, we compute the lower bound of the probability of advantage of \mathcal{A} . To do that, we observe that \mathcal{A} wins the game when $b' = b$ and this happens either when both \mathcal{A}_1 and \mathcal{A}_2 win, or when they both lose.

In fact, when $b = 0$ then $O(t) \neq O(s)$; therefore, $b_0 \neq b_1$ happens if and only if both \mathcal{A}_1 and \mathcal{A}_2 win or when they both lose. The same holds when $b = 1$.

Therefore,

$$\begin{aligned}
 \Pr[\mathcal{A} \text{ wins}] &= \\
 &\Pr[(\mathcal{A}_1 \text{ wins} \wedge \mathcal{A}_2 \text{ wins}) \vee (\mathcal{A}_1 \text{ loses} \wedge \mathcal{A}_2 \text{ loses})] = \\
 &\Pr[(\mathcal{A}_1 \text{ wins} \wedge \mathcal{A}_2 \text{ wins})] + \Pr[(\mathcal{A}_1 \text{ loses} \wedge \mathcal{A}_2 \text{ loses})] = \\
 &\Pr[(\mathcal{A}_1 \text{ wins})] \Pr[(\mathcal{A}_2 \text{ wins})] + \Pr[(\mathcal{A}_1 \text{ loses})] \Pr[(\mathcal{A}_2 \text{ loses})] \geq \\
 &\left(\frac{1}{2} + \epsilon(\lambda)\right)^2 + \left(\frac{1}{2} - \epsilon(\lambda)\right)^2 = \frac{1}{2} + 2\epsilon(\lambda)^2.
 \end{aligned}$$

Since $\epsilon(\lambda)$ is a non-negligible function, we have defined an adversary \mathcal{A} of the dGA-IP game that has a non-negligible advantage. This contradicts the dGA-IP assumption, therefore the adversary of $\text{Hiding}(\Pi_{Com})$ with non-negligible advantage does not exist and the commitment scheme Π_{Com} satisfies the hiding property.

The two previous results can be summarized in the following corollary.

Corollary 1. *The bit commitment scheme in Figure 4 is secure under the decisional Group Action Inversion Problem assumption.*

We also have expanded the security analysis of the hiding property of the commitment scheme under to the 2GA-PR assumption requiring that the two orbits O_0 and O_1 used to instantiate the bit commitment have *similar size*, i.e.

$$|\Pr[x \in O_0] - \Pr[x \in O_1]| = \nu(\lambda)$$

for a randomly chosen x in $O_0 \cup O_1$ and a negligible function $\nu(\lambda)$.

We have proved the following theorem.

Theorem 3. *If the bit commitment scheme in Figure 4 is instantiated using two orbits of similar size, it is secure under the 2GA-PR assumption.*

Proof. The commitment scheme satisfies the property of perfect binding, as shown in Theorem 1. The proof of the computationally hiding property can be found in Appendix A.

Finally, Appendix B shows that $\text{Hiding}(\Pi_{Com})$ reduces to dGA-IP, also. This allows us describe the relation between the dGA-IP and 2GA-PR assumptions.

Corollary 2. *The 2GA-PR problem reduces to dGA-IP when it is instantiated with two orbits of similar size.*

We summarize the reductions between the hiding game of the commitment scheme and the two assumptions in Figure 6.

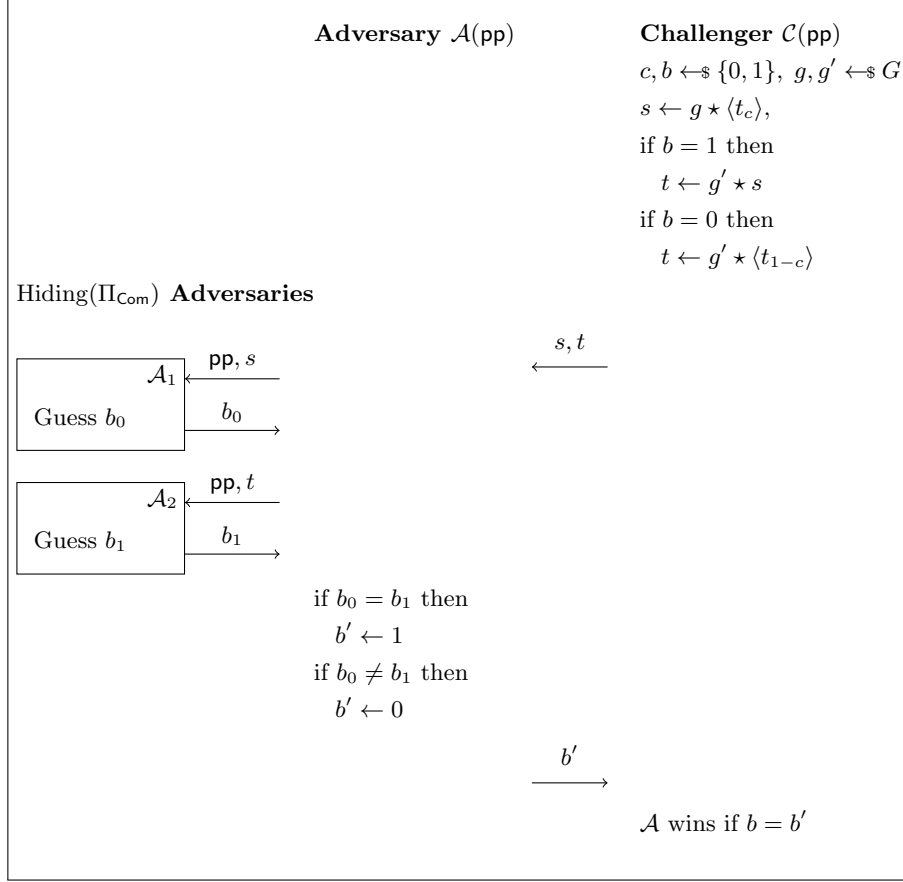


Fig. 5. Reduction from dGA-IP(pp) to the hiding game for the bit commitment scheme.

5 Linkable Commitments

The proposed bit commitment has the following additional feature. Given two commitments com_0 and com_1 , if we suppose that the sender is honest, there is a way to prove that their committed value is the same. Based on this notion, we define the concept of *linkable commitment*. We require that the sender is honest to be assured that the commitments lie either in the orbit of $\langle t_0 \rangle$ or $\langle t_1 \rangle$. To the best of our knowledge, this property has not been formally defined before. However it is well known that, for example, Pedersen commitments enjoy this property which is used, among other things, in the Monero's RingCT protocol [29].

Definition 11. Let $\Pi_{\text{Com}} = (\text{PGen}, \text{Commit}, \text{Open})$ be a commitment scheme. Let m_0 and m_1 be two messages and let $(\text{com}_0, r_0) = \text{Commit}(m_0)$ and $(\text{com}_1, r_1) =$




Fig. 6. Reductions between games and problems. “ $A \rightarrow B$ ” means that solving B implies solving A . The reductions represented by a dashed line require the extra hypothesis about the similarity of the orbits.

$\text{Commit}(m_1)$. We say that Π_{Com} is linkable if there exist the two following PPT algorithms:

1. $\text{LinkMaterial}(r_0, r_1)$, whose output is a value r_L ;
2. $\text{Link}(\text{com}_0, \text{com}_1, r_L)$, that returns 1 if $m_0 = m_1$ and 0 otherwise.

In order to be secure, a linkable bit commitment must satisfy some security properties for these two additional algorithms Link and LinkMaterial as well. First, we want that the linking material r_L does not reveal any information about the committed value. This means that an adversary that has access to two commitments of m and the linking material r_L does not learn anything about m . We call this property *linkable-hiding*. Then, it must not be possible to link two commitments that are obtained starting from two distinct values. A linkable commitment with this property is said *linkable-binding*. Finally, we focus on how the value r_L can be generated. We want that, if a user (somehow) knows that two commitments are linked without knowing their opening material, he can not generate a proof of that (via the linking material). In other words, being m a message, and being $(\text{com}_0, r_0) = \text{Commit}(m)$ and $(\text{com}_1, r_1) = \text{Commit}(m)$, no one can generate a value r_L such that $\text{Link}(\text{com}_0, \text{com}_1, r_L) = 1$ without knowledge of any information regarding the opening materials r_0 and r_1 . This additional property is called *link secrecy*.

We formalize these new properties in the following definition.

Definition 12. Let $\text{HidingLink}(\Pi_{\text{Com}})$ be the game described in Figure 7. We define the advantage of an adversary \mathcal{A} of the game $\text{HidingLink}(\Pi_{\text{Com}})$ as

$$\text{Adv}(\mathcal{A}, \text{HidingLink}(\Pi_{\text{Com}})) = \left| \Pr[\mathcal{A} \text{ wins } \text{HidingLink}(\Pi_{\text{Com}})] - \frac{1}{2} \right|.$$

Let λ be the security parameter. A linkable bit commitment $\Pi_{\text{Com}} = (\text{PGen}, \text{Commit}, \text{Open}, \text{LinkMaterial}, \text{Link})$ is said

- computationally linkable-hiding if for all PPT adversaries \mathcal{A} there is a negligible function $\mu(\lambda)$ such that

$$\text{Adv}(\mathcal{A}, \text{HidingLink}(\Pi_{\text{Com}})) \leq \mu(\lambda);$$

- *computationally linkable-binding* if for all PPT adversaries \mathcal{A} there is a negligible function $\mu(\lambda)$ such that

$$\Pr \left[\begin{array}{c} \text{pp} \leftarrow \text{PGen}(1^\lambda), \\ (m_0, \text{com}_0, m_1, \text{com}_1, r_L) \leftarrow \mathcal{A}(\text{pp}) \end{array} \middle| \begin{array}{c} m_0 \neq m_1, \\ \text{Link}(\text{com}_0, \text{com}_1, r_L) = 1 \end{array} \right] \leq \mu(\lambda);$$

- *computationally link secret* if for all PPT adversaries \mathcal{A} there is a negligible function $\mu(\lambda)$ such that

$$\Pr[\mathcal{A} \text{ wins LinkSecrecy}(\Pi_{\text{Com}})] \leq \mu(\lambda),$$

where $\text{LinkSecrecy}(\Pi_{\text{Com}})$ is the linking secrecy game in Figure 8.

In the above definitions, whenever $\mu(\lambda) = 0$, we say that the property is perfect.

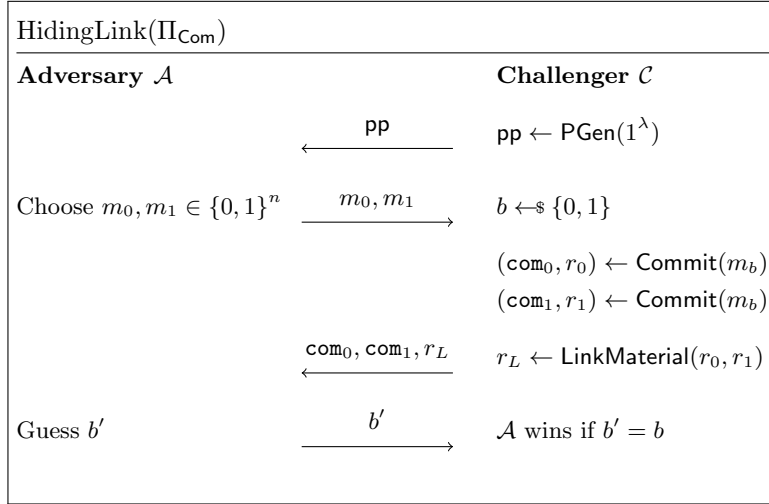


Fig. 7. Linkable-hiding game.

5.1 Linkable bit commitment from GACE

Using the bit commitment shown in Subsection 4.1, we can endow the scheme to obtain a linkable bit commitment. This extension is natural, since the commitments of a chosen message are in the orbit of that message, and showing that they are linked reduces to exhibit a group element which sends one into the other.

Theorem 4. *The bit commitment scheme in Figure 4 endowed with the algorithms in Figure 9 is a secure linkable bit commitment scheme under the One-Way Group Action and dGA-IP assumptions.*

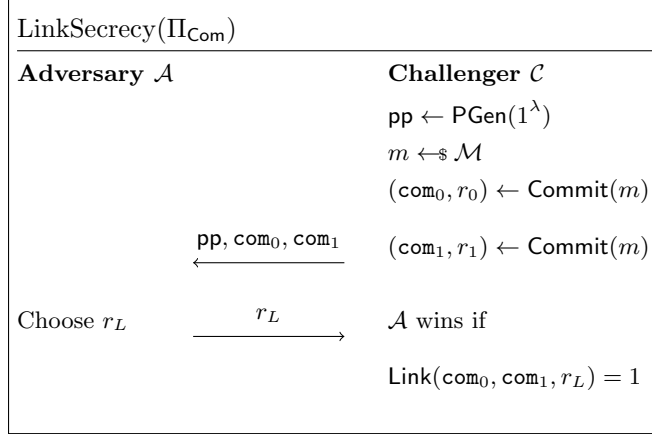


Fig. 8. Link secrecy game.




Fig. 9. Algorithm for linking commitment from a GACE.

Proof. We have already proven in Theorem 1 that the bit commitment in Figure 4 is secure under the dGA-IP assumption. Now, we prove that the linkable commitment scheme is secure, namely it is computationally linkable-hiding, perfectly linkable-binding and computationally link secret.

- **Linkable-hiding.** We show that the Hiding game reduces to the HidingLink game. The idea is to let the adversary of the Hiding(Π_{Com}) game to simulate the HidingLink game challenger by creating a new random commitment (and the linking material) to the same message of the commitment it has received from its challenger. Now we explain it in greater detail.

Let \mathcal{A}' be an adversary that wins the HidingLink game with non-negligible advantage $\epsilon(\lambda)$. We can define an adversary \mathcal{A} for the Hiding game that wins with a non-negligible advantage. Since we are in the binary case, the challenger \mathcal{C} picks a message b and sends to \mathcal{A} the commitment com of b . Now \mathcal{A} picks a random element g in G and computes $\text{com}' = g \star \text{com}$, that is a valid and randomly generated commitment to b . \mathcal{A} queries to \mathcal{A}' , the adversary of the HidingLink game, the commitments com , com' and the linking material g . Note that \mathcal{A} correctly simulates the challenger of the HidingLink game since the bit b and com are chosen at random from \mathcal{C} , com' is chosen at random from \mathcal{A} and the linking material is valid.

\mathcal{A}' returns a bit b' which \mathcal{A} sends to \mathcal{C} as its guess. If \mathcal{A}' correctly guesses the bit committed to in com and com' then clearly also \mathcal{A} wins its game. Therefore the advantage of \mathcal{A} is the same of the one of \mathcal{A}' and is non-negligible.

We can conclude that, since the commitment Π_{Com} is computationally hiding under the dGA-IP assumption, it is also computationally linkable-hiding.

- **Perfectly linkable-binding.** Suppose that an adversary returns with positive probability a tuple $(m_0, m_1, \text{com}_0, \text{com}_1, r_L)$ such that $m_0 \neq m_1$ and $\text{Link}(\text{com}_0, \text{com}_1, r_L) = 1$. By construction, there exist two elements g_0 and g_1 in G such that

$$\text{com}_0 = g_0 \star \langle m_0 \rangle \text{ and } \text{com}_1 = g_1 \star \langle m_1 \rangle$$

From $\text{Link}(\text{com}_0, \text{com}_1, r_L) = 1$ we have that $r_L \star \text{com}_1 = \text{com}_0$, and hence com_0 and com_1 are in the same orbit. Since $m_0 = f(\text{com}_0) = f(\text{com}_1) = m_1$, where f is the invariant function in the GACE, we have a contradiction. Hence, there are no adversaries that can output such a tuple with positive probability.

- **Computationally link secret.** We show that, if a PPT adversary \mathcal{A} , on input com_0 and com_1 , can find r_L such that $\text{Link}(\text{com}_0, \text{com}_1, r_L) = 1$, then it contradicts the One-way group action assumption. Essentially, if com_0 and com_1 are commitments to m_0 , then they are in the same orbit of $\langle m_0 \rangle$. Finding an r_L in G such that $\text{Link}(\text{com}_0, \text{com}_1, r_L) = 1$ means finding an element of G sending com_1 to com_0 , and this is intractable by hypothesis.

Remark 1. Observe that, if an inadmissible value is committed, for instance an element x that is not in the orbit of $\langle t_0 \rangle$ nor $\langle t_1 \rangle$, then the linkability continues to work. In fact, two commitments of the above x can be linked. Therefore we refer to the above scheme as a *honest sender* linkable commitment. To cover even the case where the sender may commit to an inadmissible value, some techniques from ring signature schemes can be used. Using the framework of Beullens, Katsumata and Pintore [5], a proof of the legitimacy of the commitment can be generated in the random oracle model. In the commit phase, the sender generates (com, r) from $\text{Commit}(b)$, then attaches to com a non-interactive proof of the OR-relation

$$\{(\text{com}, g) \mid \text{com} = g \star \langle t_0 \rangle \text{ or } \text{com} = g \star \langle t_1 \rangle\}.$$

We refer to [5] for the details. However, this proof needs many repetitions to achieve a reasonable security level, leading to a huge cost in communication.

6 An Instantiation with Tensors

6.1 3-tensors and group actions

Let n be a positive integer and let \mathbf{V} be the tensor space given by $\mathbb{F}_q^n \otimes \mathbb{F}_q^n \otimes \mathbb{F}_q^n$. Let $\{e_1, \dots, e_n\}$ be a base of \mathbb{F}_q^n , hence an element M of \mathbf{V} can be written as

$$M = \sum_{i,j,k} M(i, j, k) e_i \otimes e_j \otimes e_k, \quad (1)$$

where $M(i, j, k)$ are elements in \mathbb{F}_q . A *rank one* (or *decomposable*) tensor is an element of the form $a \otimes b \otimes c$, where a, b, c are in \mathbb{F}_q^n . Given a tensor M , its *rank* is the minimal non-negative integer r such that there exist M_1, \dots, M_r rank one tensors for which $M = \sum_{i=1}^r M_i$, and we write $\text{rk}(M) = r$. In general, computing the rank of a tensor is an hard task [18, 31, 32].

A group action can be defined on the vector space \mathbf{V} of tensors from the group $G = \text{GL}(n) \times \text{GL}(n) \times \text{GL}(n)$ as follows:

$$\begin{aligned} \star : G \times \mathbf{V} &\rightarrow \mathbf{V}, \\ \left((A, B, C), \sum_{i,j,k} M(i, j, k) e_i \otimes e_j \otimes e_k \right) &\mapsto \sum_{i,j,k} M(i, j, k) A e_i \otimes B e_j \otimes C e_k. \end{aligned}$$

It can be shown that this action does not change the rank of a tensor. However, if it is extended to non-invertible matrices, this property does not hold: for example the zero matrix sends every tensor into the zero tensor.

6.2 GACE and bit commitment from tensors

Given the group action defined above, we want to build a Group Action with Canonical Element. Since the computation of the rank is supposed to be hard, we set $T = \mathbb{N}$ and

$$f : \mathbf{V} \rightarrow \mathbb{N}, M \mapsto \text{rk}(M).$$

In order to define the function $\langle \cdot \rangle$, we need to do some observations. From Eq. (1), we see that the rank of a tensor is at most n^3 and with a simple trick it can be shown that it is at most n^2 . Actually, the maximal rank is strictly less than this value. As showed in [19], the maximal rank attainable by a tensor in \mathbf{V} is between $\frac{1}{3}n^2$ and $\frac{3}{4}n^2$. Moreover, an open problem in this field is to exhibit the explicit construction of a high-rank tensor. Even if there are some results [2, 6, 36], we are not able to construct a tensor of any given rank. Luckily, there is a set of integers for which we can easily exhibit tensors of a given rank. Let $T' = \{1, \dots, n\}$ and we can define the function

$$\begin{aligned} \langle \cdot \rangle : T' &\rightarrow \mathbf{V}, \\ r &\mapsto \sum_{i=1}^r e_i \otimes e_i \otimes e_i. \end{aligned}$$

We can see that $f(\langle r \rangle) = r$ for any r in $T' = \{1, \dots, n\}$, hence the tuple $(G, \mathbf{V}, \star, f, \langle \cdot \rangle)$ is a GACE. In fact, computing the rank of a random tensor of promised rank between 1 and n is hard, while recognize the rank of $\langle r \rangle$ is easy.

The non-interactive bit commitment scheme we present is based on the general one in Figure 4. During the parameter generation phase, we choose $n - 1$ and n as elements of T' encoding the bits 0 and 1, respectively.

Concretely, given a security parameter λ , a prime power q , an integer n and the tensor space $\mathbf{V} = \mathbb{F}_q^n \otimes \mathbb{F}_q^n \otimes \mathbb{F}_q^n$, the public parameters are

$$(G, \mathbf{V}, \star, f, \langle \cdot \rangle, n-1, n).$$

Let us analyze the assumptions on this particular group action. The dGA-IP assumption for tensors is related to the Tensor Isomorphism problem [16, 17], which is complete for a large class of problems and it is conjectured hard even for a quantum computer. The One-Way assumption on tensors is linked to the computational version of the dGA-IP problem: given two tensors in the same orbit, find the group element that links them. This problem is believed to be hard and it is directly used in various cryptosystem [10, 20], while other constructions use polynomially equivalent problems [35]. When we consider just the orbits of rank n and $n-1$, these assumptions seem to remain intractable.

Summarizing, to commit to a bit b , the sender picks a random g in G and obtains the commitment com equal to $g \star \langle n-1 \rangle$ if $b = 0$ or $g \star \langle n \rangle$ if $b = 1$. The opening material is given by g . To open the commitment com , the sender communicates to the receiver b and g and the latter checks that $g^{-1} \star \text{com}$ is equal to $\langle n-1 \rangle$ or $\langle n \rangle$. There is one additional check to take care during the opening phase: the receiver must verify the membership of g to G . In fact, if $g = (A, B, C)$ and A, B or C are non-invertible, then g can send a tensor of rank n to a tensor of rank $n-1$, breaking the binding property.

Analogously, a linkable bit commitment can be designed on tensors with the constructions given in Subsection 5.1.

7 Conclusions

In this work, we have presented a framework based on group actions that makes use of invariant functions and canonical elements, namely a Group Action with Canonical Element (GACE). The considered invariant function must be hard to compute on a large class of elements, but at the same time its computation on the canonical elements must be feasible. Then, we showed how to design a bit commitment based on this framework that is proven secure in the standard model. More in detail, breaking the hiding assumption of our commitment scheme means breaking independently both 2GA-PR and dGA-IP. This leads to the first non-interactive bit commitment relying on group actions.

One of the most interesting aspects of our construction is that it requires the action to be non-transitive. This is somehow novel in the cryptographic group action literature, where previous schemes rely on transitive action or they restrict to a single orbit. Concretely, in our framework we need to exhibit two elements that belong to two different orbits.

Moreover, we introduce the notion of linkable commitment and we prove that our bit commitment can be easily extended to a linkable one. Finally, we show an instantiation of our framework and commitment using tensors on finite fields. In this case, the invariant function is the tensor rank, and the cryptographic assumption is linked to the computational version of the dGA-IP problem.

As a future work, a commitment based on more orbits or new cryptographic schemes starting from a GACE could be investigated. At the same time, it would be interesting to find other GACEs to concretely instantiate the framework.

Acknowledgements The authors are members of GNSAGA of INdAM. The first and the third authors are members of CryptTO, the group of Cryptography and Number Theory of Politecnico di Torino. The first author acknowledges support from TIM S.p.A. through the PhD scholarship. The second author acknowledges support from Eustema S.p.A. through the PhD scholarship. The third author acknowledges support from Ripple’s University Blockchain Research Initiative.

A 2GA-PR reduces to Hiding(Π_{Com})

The reduction used to prove the hiding property under the 2GA-PR assumption is exactly the same given in the proof of Theorem 2, and the main difference between the proof of the hiding property under the dGA-IP assumption and the following is that the outcome of the adversaries of the Hiding(Π_{Com}) game \mathcal{A}_1 and \mathcal{A}_2 are not independent anymore, but are only conditionally independent once the input values (s and t) are fixed.

In fact, in the 2GA-PR game $\Pr[O(s) = O(t)] = \frac{3}{4}$ which means that the selection of the value of t , input to \mathcal{A}_2 depends on the selection of s , given in input to \mathcal{A}_1 .

Theorem 5. *The bit commitment scheme in Figure 4 instantiated with two orbits of similar size is computationally hiding under the 2GA-PR assumption.*

For simplicity, in the following proof we assume that the cardinality of the two orbits is the same, that is, the probability of picking an element at random inside any orbit is $\frac{1}{2}$. The proof can be easily generalized to the case where the probability of falling into one orbit is negligibly greater than the probability of falling into the other. In other words, the proof holds whenever there exists a negligible function $\nu(\lambda)$ such that, given the two orbits O_0 and O_1 ,

$$|\Pr[x \in O_0] - \Pr[x \in O_1]| = \nu(\lambda)$$

for a randomly chosen x in $O_0 \cup O_1$. This assumption seems admissible and not too strict for cryptographic purposes.

Proof. We must prove that the hiding property holds for Π_{Com} . We show that, given an adversary of the Hiding(Π_{Com}) game with non-negligible advantage, we can build an adversary of the 2GA-PR game with non-negligible advantage (recall that the advantage of \mathcal{A} is defined as $\mathbf{Adv}(\mathcal{A}, 2\text{GA-PR}(\text{pp})) = \Pr[\mathcal{A} \text{ wins } 2\text{GA-PR}(\text{pp})] - \frac{1}{2}$).

1. *Reduction description.*

To define \mathcal{A} , we use two independent instances of the same adversary $\mathcal{A}_1, \mathcal{A}_2$ of the hiding game as we did in the proof of Theorem 2; then, we perform the same reduction, as it is presented in Figure 10.

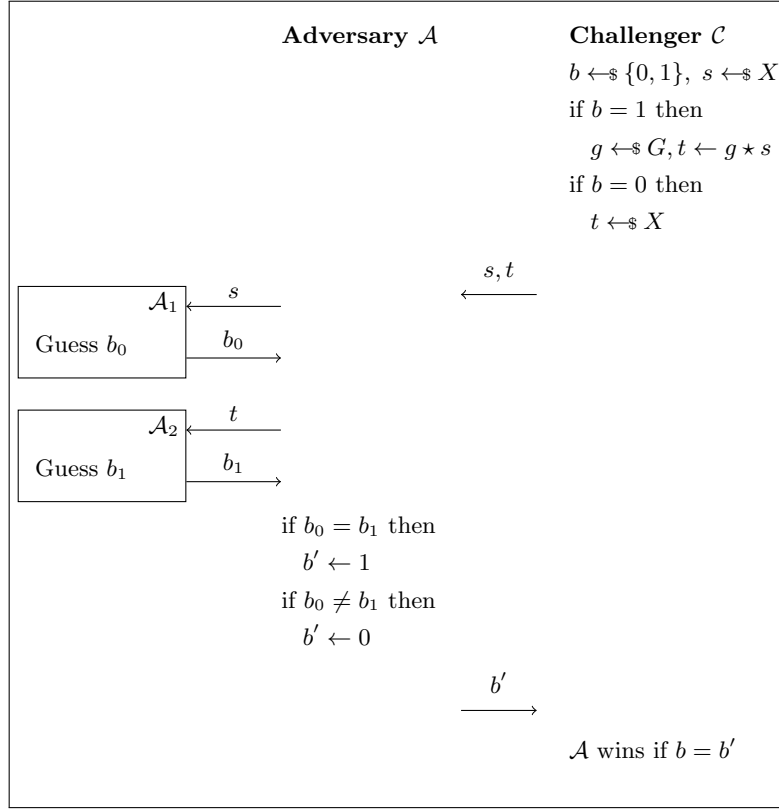


Fig. 10. Reduction from 2GA-PR to the hiding game for the bit commitment scheme.

2. \mathcal{A} correctly simulates the $\text{Hiding}(\Pi_{\text{Com}})$ challenger.

The adversary \mathcal{A} correctly simulates the challenger of $\text{Hiding}(\Pi_{\text{Com}})$ with respect to the adversaries \mathcal{A}_1 and \mathcal{A}_2 separately, in fact both s and t are uniformly sampled from the set of commitment to 0 and 1. Therefore, \mathcal{A}_1 and \mathcal{A}_2 will output the right bit with advantage $\epsilon(\lambda)$.

3. *Measurement of \mathcal{A} 's advantage.*

From now on, when we consider the orbits $O(s)$ and $O(t)$ of s and t respectively, they will assume binary values according to the relation used in the bit commitment scheme Π_{Com} : $O(s) = 1$ if s lives in the orbit of commitments to 1, and $O(s) = 0$ if s lives in the orbit of commitments to 0. The same holds for $O(t)$.

Before computing the lower bound of the advantage of the adversary \mathcal{A} , we state the following remark.

Remark 2. The outcomes of the games performed by \mathcal{A}_1 and \mathcal{A}_2 in the reduction of Figure 10 are not independent since the values given as inputs to them are dependent values (note that t is in the same orbit of s with

probability $\frac{3}{4}$). However, it is still true that the outcomes of the adversaries \mathcal{A}_1 and \mathcal{A}_2 are independent if conditioned to fixed input values.

For the sake of generality, we need to consider the case in which the advantage of the adversaries \mathcal{A}_1 and \mathcal{A}_2 in playing $\text{Hiding}(\Pi_{\text{Com}})$ game is not uniformly distributed on the possible outputs. That is, it is possible that

$$\Pr[\mathcal{A}_1 \text{ wins} \mid O(s) = 1] = \frac{1}{2} + \epsilon(\lambda) + \Delta,$$

$$\Pr[\mathcal{A}_1 \text{ wins} \mid O(s) = 0] = \frac{1}{2} + \epsilon(\lambda) - \Delta,$$

with Δ possibly a negative value. The same holds for $\Pr[\mathcal{A}_2 \text{ wins} \mid O(t) = b]$, with $b \in \{0, 1\}$.

Now, we can start with the computation of the lower bound of the advantage of \mathcal{A} in winning the 2GA-PR game.

The probability that \mathcal{A} wins the 2GA-PR game can be computed as follows, partitioning the event in three disjoint events:

$$\begin{aligned} \Pr[\mathcal{A} \text{ wins}] &= \Pr[b' = b] = \\ & \Pr \left[\underbrace{(b = 0 \wedge O(s) \neq O(t)) \wedge b' = b}_{\text{Event A}} \right] + \\ & \Pr \left[\underbrace{(b = 0 \wedge O(s) = O(t)) \wedge b' = b}_{\text{Event B}} \right] + \\ & \Pr \left[\underbrace{b = 1 \wedge b' = b}_{\text{Event C}} \right]. \end{aligned}$$

We now separately quantify the three probabilities as follows. We recall that according to the event we are considering, the event $b = b'$ can be translated in terms of success of the adversaries \mathcal{A}_1 and \mathcal{A}_2

- **Event A:** when $b = 0$ and $O(s) \neq O(t)$, then $b = b'$ when both \mathcal{A}_1 and \mathcal{A}_2 win or when both of them lose. Therefore, it holds that

$$\begin{aligned} \Pr[b = 0 \wedge O(s) \neq O(t) \wedge b' = b] &= \\ & \Pr[b = 0 \wedge O(s) \neq O(t) \wedge \mathcal{A}_1 \text{ wins} \wedge \mathcal{A}_2 \text{ wins}] + \quad (2) \\ & \Pr[b = 0 \wedge O(s) \neq O(t) \wedge \mathcal{A}_1 \text{ loses} \wedge \mathcal{A}_2 \text{ loses}]. \end{aligned}$$

We can compute this probability by considering the general case $\Pr[b = 0 \wedge O(s) \neq O(t) \wedge \mathcal{A}_1 \text{ outcome} \wedge \mathcal{A}_2 \text{ outcome}]$ and then substituting **outcome** with **wins** or **loses** accordingly with the formula above.

It holds that

$$\begin{aligned} & \Pr[b = 0 \wedge O(s) \neq O(t) \wedge \mathcal{A}_1 \text{ outcome} \wedge \mathcal{A}_2 \text{ outcome}] = \\ & \sum_{c=0}^1 \Pr[b = 0 \wedge O(s) = c \wedge O(t) = 1 - c \wedge \mathcal{A}_1 \text{ outcome} \wedge \mathcal{A}_2 \text{ outcome}] = \\ & \sum_{c=0}^1 \left(\Pr[\mathcal{A}_1 \text{ outcome} \wedge \mathcal{A}_2 \text{ outcome} \mid b = 0 \wedge O(s) = c \wedge O(t) = 1 - c] \cdot \right. \\ & \quad \left. \Pr[b = 0 \wedge O(s) = c \wedge O(t) = 1 - c] \right). \end{aligned}$$

Since the outcomes of \mathcal{A}_1 and \mathcal{A}_2 are independent once their input values are fixed, we have that

$$\begin{aligned} & \Pr[\mathcal{A}_1 \text{ outcome} \wedge \mathcal{A}_2 \text{ outcome} \mid b = 0 \wedge O(s) = c \wedge O(t) = 1 - c] = \\ & \prod_{i=1}^2 \Pr[\mathcal{A}_i \text{ outcome} \mid b = 0 \wedge O(s) = c \wedge O(t) = 1 - c], \end{aligned}$$

with $c \in \{0, 1\}$.

Since the outcome of \mathcal{A}_1 only depends on the value of $O(s)$ and the outcome of \mathcal{A}_2 depends only on $O(t)$, then

$$\begin{aligned} & \Pr[\mathcal{A}_1 \text{ outcome} \wedge \mathcal{A}_2 \text{ outcome} \mid b = 0 \wedge O(s) = c \wedge O(t) = 1 - c] = \\ & \Pr[\mathcal{A}_1 \text{ outcome} \mid O(s) = c] \Pr[\mathcal{A}_2 \text{ outcome} \mid O(t) = 1 - c] \end{aligned}$$

Therefore, since $\Pr[b = 0 \wedge O(s) = \bar{b} \wedge O(t) = 1 - \bar{b}] = \frac{1}{8}$ with $\bar{b} \in \{0, 1\}$ then

$$\begin{aligned} & \Pr[b = 0 \wedge O(s) \neq O(t) \wedge \mathcal{A}_1 \text{ outcome} \wedge \mathcal{A}_2 \text{ outcome}] = \\ & \frac{1}{8} \left(\Pr[\mathcal{A}_1 \text{ outcome} \mid O(s) = 1] \cdot \Pr[\mathcal{A}_2 \text{ outcome} \mid O(t) = 0] + \right. \\ & \quad \left. \Pr[\mathcal{A}_1 \text{ outcome} \mid O(s) = 0] \cdot \Pr[\mathcal{A}_2 \text{ outcome} \mid O(t) = 1] \right). \end{aligned}$$

We can finally compute the initial probability given in Eq. (2), by substituting **outcome** with **wins** and **loses** and obtaining

$$\Pr[b = 0 \wedge O(s) \neq O(t) \wedge b' = b] = \frac{1}{8} + \frac{1}{2}\epsilon^2(\lambda) - \frac{1}{2}\Delta^2. \quad (3)$$

- **Event B**: when $b = 0$ and $O(s) = O(t)$, then $b = b'$ when either \mathcal{A}_1 wins and \mathcal{A}_2 loses or when \mathcal{A}_1 loses and \mathcal{A}_2 wins. Therefore, it holds that

$$\begin{aligned} & \Pr[b = 0 \wedge O(s) = O(t) \wedge b' = b] = \\ & \Pr[b = 0 \wedge O(s) = O(t) \wedge \mathcal{A}_1 \text{ wins} \wedge \mathcal{A}_2 \text{ loses}] + \\ & \Pr[b = 0 \wedge O(s) = O(t) \wedge \mathcal{A}_1 \text{ loses} \wedge \mathcal{A}_2 \text{ wins}]. \end{aligned} \quad (4)$$

Since in this case the input of \mathcal{A}_1 and \mathcal{A}_2 are in the same orbit, then we can state

$$\begin{aligned} \Pr[b = 0 \wedge O(s) = O(t) \wedge b' = b] &= \\ 2 \Pr[b = 0 \wedge O(s) = O(t) \wedge \mathcal{A}_1 \text{ wins} \wedge \mathcal{A}_2 \text{ loses}] &= \\ 2 \sum_{c=0}^1 \Pr[b = 0 \wedge O(s) = c \wedge O(t) = c \wedge \mathcal{A}_1 \text{ wins} \wedge \mathcal{A}_2 \text{ loses}]. \end{aligned}$$

Using arguments similar to the ones used for **Event A**, that is the conditional independence of the outcomes of the adversaries once the inputs are fixed, the fact that the output of \mathcal{A}_1 (resp. \mathcal{A}_2) depends only on $O(s)$ (resp. on $O(t)$) and finally that $\Pr[b = 0 \wedge O(s) = c \wedge O(t) = c] = \frac{1}{8}$, for $c \in \{0, 1\}$, we can write the Eq. (4) as follows

$$\Pr[b = 0 \wedge O(s) = O(t) \wedge b' = b] = \frac{1}{8} - \frac{1}{2}\epsilon^2(\lambda) - \frac{1}{2}\Delta^2. \quad (5)$$

– **Event C**: when $b = 1$, $O(s) = O(t)$, then $b = b'$ when both \mathcal{A}_1 and \mathcal{A}_2 win or when both of them lose. Therefore, it holds that

$$\begin{aligned} \Pr[b = 1 \wedge b' = b] &= \\ \Pr[b = 1 \wedge \mathcal{A}_1 \text{ wins} \wedge \mathcal{A}_2 \text{ wins}] + & \quad (6) \\ \Pr[b = 1 \wedge \mathcal{A}_1 \text{ loses} \wedge \mathcal{A}_2 \text{ loses}]. \end{aligned}$$

As in the computation of the probability of **Event A**, we must compute $\Pr[b = 1 \wedge \mathcal{A}_1 \text{ outcome} \wedge \mathcal{A}_2 \text{ outcome}]$. Using similar arguments as before, and noticing that $\Pr[b = 1 \wedge O(s) = c \wedge O(t) = c] = \frac{1}{4}$ with $c \in \{0, 1\}$, it can be shown that

$$\begin{aligned} \Pr[b = 1 \wedge \mathcal{A}_1 \text{ outcome} \wedge \mathcal{A}_2 \text{ outcome}] &= \\ \frac{1}{4} \sum_{c=0}^1 \Pr[\mathcal{A}_1 \text{ outcome} \mid O(s) = c] \Pr[\mathcal{A}_2 \text{ outcome} \mid O(t) = c] \end{aligned}$$

Therefore, substituting **outcome** with **loses** and **wins**, and using the probabilities of success of adversaries \mathcal{A}_1 and \mathcal{A}_2 , from Eq. (6) we obtain

$$\Pr[b = 1 \wedge b' = b] = \frac{1}{4} + \epsilon^2(\lambda) + \Delta^2. \quad (7)$$

Combining the partial results derived analysing **Event A**, **Event B** and **Event C** from Equations (3),(5) and (7) respectively, we obtain the final result

$$\Pr[\mathcal{A} \text{ wins}] = \frac{1}{2} + \epsilon^2(\lambda),$$

which proves that we have built an adversary for the 2GA-PR game which wins with non-negligible advantage. Therefore, an adversary who wins the hiding game with non-negligible advantage does not exist due to the 2GA-PR assumption. This means that the binary commitment scheme we have described results to be perfectly binding and computationally hiding.

B Hiding(Π_{Com}) reduces to dGA-IP

Theorem 6. *The Hiding(Π_{Com}) game reduces to dGA-IP game.*

Proof. We show how the existence of an adversary of dGA-IP problem with non-negligible advantage allows the creation of an adversary of the Hiding(Π_{Com}) game with non-negligible advantage.

1. *Reduction description.*

The adversary \mathcal{A} of the Hiding(Π_{Com}) game (see Figure 11) receives from the challenger a commitment c to a randomly generated bit b . \mathcal{A} generates a commitment c' to a random bit b' and sends c, c' to \mathcal{A}' , the adversary to the dGA-IP game with non-negligible advantage. \mathcal{A} receives a response b_0 from \mathcal{A}' and returns to the Hiding(Π_{Com}) challenger the bit b' if $b_0 = 1$ (i.e. \mathcal{A}' has guessed that c and c' are in the same orbit), otherwise \mathcal{A} returns $1 - b'$.

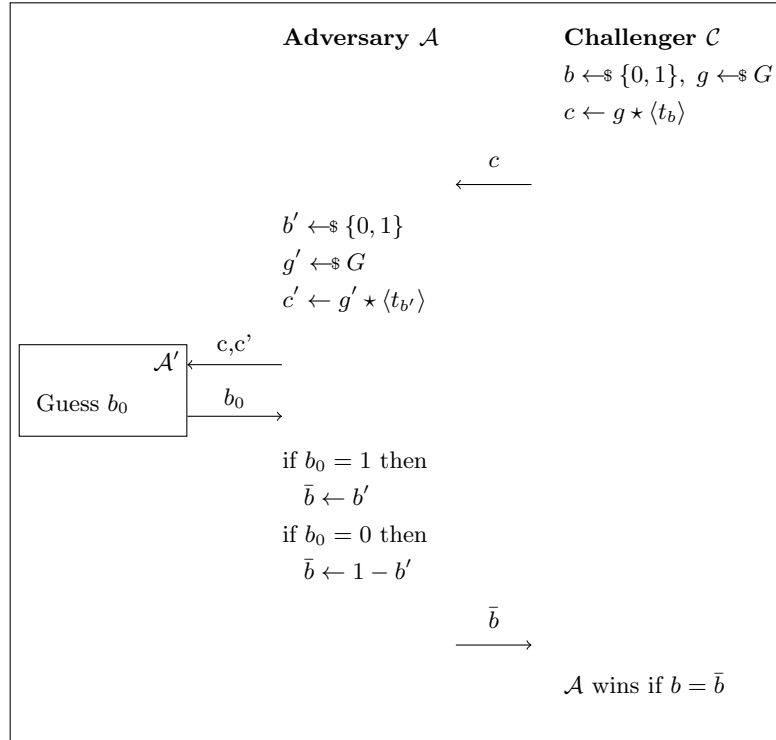


Fig. 11. Reduction from the hiding game for the bit commitment scheme to dGA-IP.

2. *\mathcal{A} correctly simulates the dGA-IP challenger.*

The adversary \mathcal{A} receives a commitment to a random unknown bit b . Therefore, in order to simulate the dGA-IP challenger, it generates a random bit b'

and a commitment to such bit. In this way, \mathcal{A} generates couples of elements in X that live in the same orbit with probability $\frac{1}{2}$ as it does the dGA-IP challenger.

3. *Measurement of \mathcal{A} 's advantage.*

The adversary \mathcal{A} wins exactly with the same probability of \mathcal{A}' , since every time \mathcal{A}' guesses the right answer to the dGA-IP game, \mathcal{A} learns the orbit in which the element c lies since it knows the orbit of c' . Therefore, if \mathcal{A}' wins the dGA-IP game with non-negligible advantage, also \mathcal{A} wins the Hiding(Π_{Com}) game with non negligible advantage.

References

1. Alapati, N., De Feo, L., Montgomery, H., Patranabis, S.: Cryptographic group actions and applications. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 411–439. Springer (2020)
2. Alexeev, B., Forbes, M.A., Tsimerman, J.: Tensor rank: Some lower and upper bounds. In: 2011 IEEE 26th Annual Conference on Computational Complexity. pp. 283–291. IEEE (2011)
3. Barenghi, A., Biasse, J.F., Persichetti, E., Santini, P.: On the computational hardness of the code equivalence problem in cryptography. *Advances in Mathematics of Communications* **17**(1), 23–55 (2023)
4. Beullens, W., Dobson, S., Katsumata, S., Lai, Y.F., Pintore, F.: Group signatures and more from isogenies and lattices: generic, simple, and efficient. *Designs, Codes and Cryptography* pp. 1–60 (2023)
5. Beullens, W., Katsumata, S., Pintore, F.: Calamari and Falafel: logarithmic (linkable) ring signatures from isogenies and lattices. In: *Advances in Cryptology—ASIACRYPT 2020: 26th International Conference on the Theory and Application of Cryptology and Information Security*, Daejeon, South Korea, December 7–11, 2020, Proceedings, Part II. pp. 464–492. Springer (2020)
6. Bläser, M.: Explicit tensors. *Perspectives in Computational Complexity: The Somnath Biswas Anniversary Volume* pp. 117–130 (2014)
7. Blum, M., Micali, S.: How to generate cryptographically strong sequences of pseudo random bits. In: *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*, pp. 227–240 (2019)
8. Brassard, G., Yung, M.: One-way group actions. In: *Advances in Cryptology—CRYPTO'90: Proceedings* 10. pp. 94–107. Springer (1991)
9. Castryck, W., Lange, T., Martindale, C., Panny, L., Renes, J.: CSIDH: an efficient post-quantum commutative group action. In: *Advances in Cryptology—ASIACRYPT 2018: 24th International Conference on the Theory and Application of Cryptology and Information Security*, Brisbane, QLD, Australia, December 2–6, 2018, Proceedings, Part III 24. pp. 395–427. Springer (2018)
10. Chou, T., Niederhagen, R., Persichetti, E., Randrianarisoa, T.H., Reijnders, K., Samardjiska, S., Trimoska, M.: Take your MEDS: Digital Signatures from Matrix Code Equivalence. *Cryptology ePrint Archive* (2022)
11. Couveignes, J.M.: Hard homogeneous spaces. *Cryptology ePrint Archive* (2006)
12. Darwish, A., El-Gendy, M.M.: A new cryptographic voting verifiable scheme for e-voting system based on bit commitment and blind signature. *Int J Swarm Intel Evol Comput* **6**(158), 2 (2017)

13. Frederiksen, T.K., Pinkas, B., Yanai, A.: Committed mpc: Maliciously secure multiparty computation from homomorphic commitments. In: *Public-Key Cryptography–PKC 2018: 21st IACR International Conference on Practice and Theory of Public-Key Cryptography*, Rio de Janeiro, Brazil, March 25–29, 2018, Proceedings, Part I 21. pp. 587–619. Springer (2018)
14. Goldreich, O., Micali, S., Wigderson, A.: Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM (JACM)* **38**(3), 690–728 (1991)
15. Grigoriev, D., Shpilrain, V.: Authentication schemes from actions on graphs, groups, or rings. *Annals of Pure and Applied Logic* **162**(3), 194–200 (2010)
16. Grochow, J.A., Qiao, Y.: Isomorphism problems for tensors, groups, and cubic forms: completeness and reductions. arXiv preprint arXiv:1907.00309 (2019)
17. Grochow, J.A., Qiao, Y.: On the complexity of isomorphism problems for tensors, groups, and polynomials I: Tensor Isomorphism-completeness. In: *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2021)
18. Hästad, J.: Tensor rank is NP-complete. In: *International Colloquium on Automata, Languages, and Programming*. pp. 451–460. Springer (1989)
19. Howell, T.D.: Global properties of tensor rank. *Linear Algebra and its Applications* **22**, 9–23 (1978)
20. Ji, Z., Qiao, Y., Song, F., Yun, A.: General linear group action on tensors: a candidate for post-quantum cryptography. In: *Theory of Cryptography Conference*. pp. 251–281. Springer (2019)
21. Juels, A., Luby, M., Ostrovsky, R.: Security of blind digital signatures. In: *Advances in Cryptology—CRYPTO’97: 17th Annual International Cryptology Conference Santa Barbara, California, USA August 17–21, 1997 Proceedings*. pp. 150–164. Springer (2006)
22. Leroux, A., Roméas, M.: Updatable encryption from group actions. *Cryptology ePrint Archive* (2022)
23. Lyubashevsky, V., Nguyen, N.K., Seiler, G.: Shorter lattice-based zero-knowledge proofs via one-time commitments. In: *Public-Key Cryptography–PKC 2021: 24th IACR International Conference on Practice and Theory of Public Key Cryptography, Virtual Event, May 10–13, 2021, Proceedings, Part I*. pp. 215–241. Springer (2021)
24. Naor, M.: Bit commitment using pseudorandomness. *Journal of cryptology* **4**, 151–158 (1991)
25. Ostrovsky, R., Persiano, G., Visconti, I.: Simulation-based concurrent non-malleable commitments and decommitments. In: *Theory of Cryptography: 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15–17, 2009. Proceedings 6*. pp. 91–108. Springer (2009)
26. Patarin, J.: Hidden fields equations (HFE) and isomorphisms of polynomials (IP): Two new families of asymmetric algorithms. In: *International Conference on the Theory and Applications of Cryptographic Techniques*. pp. 33–48. Springer (1996)
27. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: *Advances in Cryptology—CRYPTO’91: Proceedings*. pp. 129–140. Springer (2001)
28. Petrank, E., Roth, R.M.: Is code equivalence easy to decide? *IEEE Transactions on Information Theory* **43**(5), 1602–1604 (1997)
29. Poelstra, A., Back, A., Friedenbach, M., Maxwell, G., Wuille, P.: Confidential assets. In: *Financial Cryptography and Data Security: FC 2018 International Work-*

- shops, BITCOIN, VOTING, and WTSC, Nieuwpoort, Curaçao, March 2, 2018, Revised Selected Papers 22. pp. 43–63. Springer (2019)
30. Reijnders, K., Samardjiska, S., Trimoska, M.: Hardness estimates of the Code Equivalence Problem in the Rank Metric. *Cryptology ePrint Archive* (2022)
 31. Schaefer, M., Štefankovič, D.: The complexity of tensor rank. *Theory of Computing Systems* **62**, 1161–1174 (2018)
 32. Shitov, Y.: How hard is the tensor rank? *arXiv preprint arXiv:1611.01559* (2016)
 33. Shor, P.W.: Algorithms for quantum computation: discrete logarithms and factoring. In: *Proceedings 35th annual symposium on foundations of computer science*. pp. 124–134. IEEE (1994)
 34. Stolbunov, A.: *Cryptographic schemes based on isogenies* (2012)
 35. Tang, G., Duong, D.H., Joux, A., Plantard, T., Qiao, Y., Susilo, W.: Practical Post-Quantum Signature Schemes from Isomorphism Problems of Trilinear Forms. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. pp. 582–612. Springer (2022)
 36. Weitz, B.: An improvement on ranks of explicit tensors. *arXiv preprint arXiv:1102.0580* (2011)