Efficient Designated-Verifier Non-Interactive Zero-Knowledge Proofs of Knowledge

Pyrros Chaidos* and Geoffroy Couteau**

Abstract. We propose a framework for constructing efficient designated-verifier non-interactive zero-knowledge proofs (DVNIZK) for a wide class of algebraic languages over abelian groups, under standard assumptions. The proofs obtained via our framework are proofs of knowledge, enjoy statistical, and unbounded soundness (the soundness holds even when the prover receives arbitrary feedbacks on previous proofs). Previously, no efficient DVNIZK system satisfying *any* of those three properties was known. Our framework allows proving arbitrary relations between cryptographic primitives such as Pedersen commitments, ElGamal encryptions, or Paillier encryptions, in an efficient way. For the latter, we further exhibit the first non-interactive zero-knowledge proof system in the standard model that is more efficient than proofs obtained via the Fiat-Shamir transform, with still-meaningful security guarantees and under standard assumptions. Our framework has numerous applications, in particular for the design of efficient privacy-preserving non-interactive authentication.

Keywords. Zero-knowledge proofs, Non-interactive proofs.

1 Introduction

Zero-knowledge proof systems allow a prover to convince someone of the truth of a statement, without revealing anything beyond the fact that the statement is true. After their introduction in the seminal work of Goldwasser, Micali, and Rackoff [34], they have proven to be a fundamental primitive in cryptography. Among them, *non-interactive zero-knowledge proofs* (NIZK proofs), where the proof consists of a single flow from the prover to the verifier, are of particular interest, in part due to their tremendous number of applications in cryptographic primitives and protocols, and in part due to the theoretical and technical challenges that they represent.

For almost two decades after their introduction in [10], NIZKs coexisted in two types: inefficient NIZKs secure under standard assumptions (such as doubly enhanced trapdoor permutations [30]) in the common reference string model, and practically efficient NIZKs built from the Fiat-Shamir

* National & Kapodistrian University of Athens, Greece
** Karsruhe Institute of Technology, Germany. Part of this work was made while the second author was at École Normale Supérieure de Paris, France

heuristic [31, 47], which are secure in the random oracle model [6] (hence only heuristically secure in the standard model). This state of affairs changed with the arrival of pairing-based cryptography, from which a fruitful line of work (starting with the work of Groth, Ostrovsky, and Sahai [37, 38]) introduced increasingly more efficient NIZK proof systems in the standard model. That line of work culminated with the framework of Groth-Sahai proofs [39], which provided an efficient framework of pairing-based NIZKs for a large class of useful languages. Yet, one decade later, pairing-based NIZKs from the Groth-Sahai framework remain the only known efficient NIZK proof system in the standard model. Building efficient NIZKs in the standard model, without pairing-based assumptions, is a major open problem, and research in this direction has proven elusive.

1.1 Designated-Verifier Non-Interactive Zero-Knowledge

Parallel to the research on NIZKs, an alternative promising line of research has focused on *designated-verifier* non-interactive zero-knowledge proof systems (DVNIZKs). A DVNIZK retains most of the security properties of a NIZK, but is not publicly verifiable: only the owner of some secret information (the designated verifier) can check the proof. Nevertheless, DVNIZKs can replace publicly verifiable NIZKs in a variety of applications. In addition, unlike their publicly-verifiable counterpart, it is known that efficient DVNIZKs secure in the standard model for rich classes of languages can be constructed without pairing-based assumptions [17, 23, 43, 49]. However, to date, research in DVNIZKs has attracted less attention than NIZKs, the previously listed papers being (to our knowledge) the only existing works on this topic, and several important questions have been left open. We list the main open questions below.

Proofs Versus Arguments. A non-interactive zero-knowledge argument system is a NIZK in which the soundness property is only required to hold against computationally bounded adversaries. In a NIZK *proof* system, however, soundness is required to hold even against unbounded adversaries.

Currently, while several DVNIZK argument systems have been designed in the standard model without pairing-based assumptions, efficient DVNIZK proof systems without pairings remain an open question. In fact, to our knowledge, the only known constructions of (possibly inefficient) DVNIZK proofs rely on publicly-verifiable NIZK proofs.

Soundness Versus Knowledge Extraction. A non-interactive zero-knowledge proof (or argument) system is a NIZK *of knowledge* if it guar-

antees that, when the prover succeeds in convincing the verifier, he must *know* a witness for the truth of the statement. This is in constrast with the standard soundness notion, which only guarantees that the statement is true. Formally, this is ensured by requiring the existence of an efficient simulator that can extract a witness from the proof.

Non-interactive zero-knowledge proofs of knowledge are more powerful than standard NIZKs, and the knowledge-extractability property is crucial in many applications. In particular, they are necessary for the very common task of proving relations between values committed with a perfectly hiding commitment scheme, and they are a core component in privacy-preserving authentication mechanisms [4]. Currently, all known DVNIZK argument systems are not arguments of knowledge. Designing efficient DVNIZKs of knowledge without pairing-based assumptions remains an open question.

Bounded Soundness Versus Unbounded Soundness. The classical soundness security notion for non-interactive zero-knowledge proof systems states that if the statement is not true, no malicious prover can possibly convince the verifier of the truth of the statement with non-negligible probability. While this security notion is sufficient for publicly-verifiable NIZKs, it turns out to be insufficient when considering designated-verifier NIZKs, and corresponds only to a *passive* type of security notion. Indeed, the verification of a DVNIZK involves a secret value, known to the verifier. The fact that a DVNIZK satisfies the standard soundness notion does not preclude the possibility for a malicious prover to learn this secret value, e.g. by submitting a large number of proofs and receiving feedback on whether the proof was accepted or not. Intuitively, this is the same type of issue as for encryption schemes indistinguishable against chosen-plaintext attacks, which can be broken if the adversary is given access to a decryption oracle, or for signature schemes secure against key-only or known-message attacks, which can be broken if the adversary is given access to a signing oracle. Here, an adversary could possibly break the soundness of a DVNIZK if it is given access to a verification oracle.

In practice, this means that as soon as a proof system with bounded soundness is used for more than a logarithmic number of proofs, the soundness property is no longer guaranteed to hold. This calls for a stronger notion of soundness, *unbounded soundness*, which guarantees security even against adversaries that are given arbitrary access to a verification oracle.

Designing a DVNIZK with unbounded soundness has proven to be highly non-trivial. In fact, apart from publicly-verifiable NIZKs (which can be seen as particular types of DVNIZKs where the secret key of the verifier

is the empty string), the only known construction of DVNIZK claiming to satisfy unbounded soundness is the construction of [23], where the claim is supported by a proof of security in an idealized model. However, we found this claim to be flawed: there is an explicit attack against the unbounded soundness of any protocol obtained using the compiler of [23], which operates by using slightly malformed proofs to extract the verification key. In the full version of this work [16], we describe our attack, and identify the flaw in the proof of Theorem 5 in [23, Appendix A]. We have notified the authors of our finding and will update future versions of this work with their reply. To our knowledge, in all current constructions, the common reference string and the public key must be refreshed after a logarithmic number of proofs.

1.2 Our Contribution

In this work, we first introduce a framework for designated-verifier NIZKs on group-dependent languages, in the spirit of the Groth-Sahai framework for NIZKs on languages related to pairing-friendly elliptic curves. Our framework only requires that the underlying abelian group on which it is instantiated has order M, where $\mathbb{Z}_M$ is the plaintext-space of an homomorphic cryptosystem with specific properties, and allows to prove a wide variety statements formulated in terms of the operation associated to this abelian group. In particular, we do not need to rely on pairings. The DVNIZKs obtained with our framework are efficient, as they only require a few group elements and ciphertexts. The zero-knowledge property of our schemes reduces to the IND-CPA security of the underlying encryption scheme. Additionally, our DVNIZKs enjoy the following properties: they are (adaptively) *knowledge-extractable*; their knowledge-extractability holds *statistically*; their knowledge-extractability is *unbounded*. We stress that previously, no efficient construction of DVNIZK in the standard model satisfying *any* of the above properties was known. The third property, unbounded soundness, was only claimed to hold for the construction of [23], and this claim was formalized with a proof in an idealized model, but as previously mentioned, we found this claim to be flawed. We also point out that in the Groth-Sahai framework, witness extraction is limited either to statements about group elements, or to statements about exponents committed in a bit-by-bit fashion (making the proof highly inefficient). In contrast, our proof system allows to efficiently extract large exponents, without harming the efficiency of the proof. In addition to the above properties, our DVNIZKs satisfy some other useful properties: they are multi-theorem [30], randomizable [3], and same-string zero-knowledge [27]

(*i.e.*, the common reference string used by the prover and the simulator are the same).

Second, our framework comes with a dual variant, where the role of the encryption scheme and the abelian group are reversed, to prove statements, not about elements of the abelian group, but about the underlying homomorphic encryption scheme. This dual variant leads to DVNIZKs satisfying adaptive statistical unbounded soundness, but not knowledge-extractability (i.e. the dual variant does not give proofs of knowledge).

Third, we show that if one is willing to give up unbounded soundness for efficiency, our techniques can be used to construct extremely efficient DVNIZKs with bounded-soundness. The DVNIZKs that we obtain this way are more efficient than *any* previously known construction of non-interactive zero-knowledge proofs, even when considering NIZKs in the random oracle model using the Fiat-Shamir transform: the proofs we obtain are shorter than the proofs obtained via the Fiat-Shamir transform by almost a factor two. To our knowledge, this is the first example of a NIZK construction in the standard model which (conditionally) improves on the Fiat-Shamir paradigm.

Instantiating the Encryption Scheme. Informally, the security properties we require from the underlying scheme are the following: it must be additively homomorphic, with plaintext space $\mathbb{Z}_M$, random source $\mathbb{Z}_R$, and $\gcd(M, R) = 1$, and it must be *decodable*, which means that a plaintext m can be efficiently recovered from an encryption of m with random coin 0. A natural candidate for the above scheme is the Paillier encryption scheme [45] (and its variants, such as Damgård-Jurik [26]). This gives rise to efficient DVNIZK proofs of knowledge over abelian groups of composite order (e.g. subgroups of $\mathbb{F}_p^*$, with order a prime $p = k \cdot n + 1$ for a small k and an RSA modulus n, or composite-order elliptic curves), as well as efficient DVNIZKs for proving relations between Paillier ciphertexts (using the dual variant of our framework). Alternatively, the scheme can also be instantiated with the more recent Castagnos-Laguillaumie encryption scheme [15] to get DVNIZKs over prime-order abelian groups.

Our framework captures many useful zero-knowledge proofs of knowledge that are commonly used in cryptography. This includes DVNIZK proofs of knowledge of a discrete logarithm, of correctness of a Diffie-Hellman tuple, of multiplicative relationships between Pedersen commitments or ElGamal ciphertexts (or variants thereof), among many others. Our results show that, in the settings where a designated-verifier is sufficient, one can build efficient non-interactive zero-knowledge proofs of knowledge for most

statements of interest, under well-known assumptions and with strong security properties, without having to rely on pairing-friendly groups.

1.3 Our Method

It is known that linear relations (*i.e.*, membership in linear subspaces) can be non-interactively verified, using the homomorphic properties of cryptographic primitives over abelian groups. Indeed, DVNIZK proofs for linear languages can be constructed, e.g., from hash proof systems [33, 41]. In [39], pairings provide exactly the additional structure needed to evaluate degree-two relations, which can be easily generalized to arbitrary relations.

An alternative road was taken in [23] and subsequent works, to obtain non-interactive zero-knowledge proofs for a wide variety of relations, in the designated-verifier setting. To illustrate, let us consider a prover interacting with a verifier, with a common input $(g_1, g_2, h_1, h_2) \in \mathbb{G}^4$ in some group $\mathbb{G}$ of order p, where p is a λ-bit prime. The prover wants to show that (h_1, h_2) have the same discrete logarithm in the basis (g_1, g_2), *i.e.*, there exists x such that $(h_1, h_2) = (g_1^x, g_2^x)$. The standard interactive zero-knowledge proof for this statement proceeds as follows:[1]

1. The prover picks $r \xleftarrow{\$} \{0,1\}^{3\lambda}$, and sends $(a_1, a_2) \leftarrow (g_1^r, g_2^r)$.
2. The verifier picks and sends a uniformly random challenge $e \xleftarrow{\$} \mathbb{Z}_p$.
3. The prover computes and sends $d \leftarrow e \cdot x + r$. The verifier accepts the proof if and only if $(g_1^d, g_2^d) = (h_1^e a_1, h_2^e a_2)$.

The idea of [23] is to squash this interactive protocol into a (designated-verifier) non-interactive proof, by giving the challenge to the prover in advance. As knowing the challenge before sending the first flow gives the prover the ability to cheat, the challenge is encrypted with an additively homomorphic encryption scheme. That way, the prover cannot see the challenge; yet, he can still compute an encryption of the value d homomorphically, using the encryption of e. The verifier, who is given the secret verification key, can decrypt the last flow and perform the above check. Thus, the proof is a tuple (a_1, a_2, c_d), where c_d is an encryption of d computed from (x, r) and an encryption c_e of the challenge e.

Although natural, this intuitive approach has proven quite tough to analyze. In [23], the authors had to rely on a new complexity-leveraging-type assumption tailored to their scheme, which (informally) states that the simulator cannot break the security of the encryption scheme, even if he is

[1] More formally, this proof only satisfies zero-knowledge against honest verifiers, but this property is sufficient for the construction of [23].

powerful enough to break the problem underlying the protocol (in the above example, the discrete logarithm problem over $\mathbb{G}$). Even in the bounded setting, analyzing the soundness guarantees of the protocols obtained by this compilation technique (and its variants) is non-trivial, and it has been the subject of several subsequent works [17, 43, 49]. Additionally, in the unbounded setting, where we must give an efficient simulator that can successfully answer to the proofs submitted by any malicious prover, this compilation technique breaks down. Furthermore, for DVNIZKs constructed with this method, soundness holds only computationally, and security does not guarantee that the simulator can extract a witness for the statement.

Our core idea to overcome all of the above issues is to implement the same strategy in a slightly different way: rather than encrypting the challenge e as the *plaintext* of an homomorphic encryption scheme, we encrypt it as the *random coin* of an encryption scheme which is also homomorphic over the coins. To understand how this allows us to improve over all previous constructions, suppose that we have an encryption scheme Enc which is homomorphic over both the plaintext and the random coins, with plaintext space $\mathbb{Z}_M$ and random source $\mathbb{Z}_R$, and that M is coprime to R. Consider the previously described protocol for proving equality of two discrete logarithms. Given an encryption $\mathsf{Enc}(0; e)$ of 0, where the challenge is the random coin, a prover holding (x, r) can compute and send $\mathsf{Enc}(x; \rho)$ and $\mathsf{Enc}(r; -e\rho)$, for some random ρ. This allows the verifier, who knows e, to compute $\mathsf{Enc}(x \cdot e + r; 0)$, from which she can extract $d = x \cdot e + r \bmod M$ (note that the verifier only needs to know e; unlike in previous work, she does not need to know the decryption key of Enc). Observe that the extracted value depends only on e *modulo* M. At the same time, however, the ciphertext $E(0; e)$ only leaks e *modulo* R, even to an unbounded adversary. By picking e to be sufficiently large ($e > MR$), as M is coprime to R, the verifier can ensure that this leaks *no information* (statistically) about $e \bmod M$. Therefore, we can use a statistical argument to show that the prover cannot cheat when the verification using d succeeds. To allow for efficient simulation of the verifier, we simply give to the simulator the secret key of the scheme, which will allow him to extract all encrypted values, and to check the validity of the equations, without knowing $e \bmod M$. As the simulator is able to extract the values encrypted with Enc, the scheme can be proven to be (statistically) knowledge-extractable. Contrary to previous constructions, the verification key is a random coin rather than the secret key of an encryption scheme. The secret key is only used to extract information in the simulated game.

Example: DVNIZK Proof of Knowledge of a Discrete Logarithm. We illustrate our method with the classical example of proving knowledge of a discrete logarithm. For concreteness, we describe an explicit protocol using the Paillier encryption scheme; therefore, this section assumes some basic knowledge of the Paillier encryption scheme. All necessary preliminaries can be found in Section 2. Let $\mathbb{G}$ be a group of order n, where $n = p \cdot q$ is an RSA modulus (*i.e.*, a product of two strong primes). Let g be a generator of $\mathbb{G}$, and let T be a group element. A prover P wishes to prove to a verifier V that he knows a value $t \in \mathbb{Z}_n$ such that $g^t = T$.

Let $h \leftarrow u^n \bmod n^2$, where u denotes an arbitrary generator of $\mathbb{J}_n$, the subgroup of elements of $\mathbb{Z}_n^*$ with Jacobi symbol 1. The Paillier encryption of a message $m \in \mathbb{Z}_n$ with randomness $r \in \mathbb{Z}_{\varphi(n)/2}$ is $\mathsf{Enc}(m; r) = (1 + n)^m h^r \bmod n^2$. The public key of the DVNIZK is $E = h^e \in \mathbb{Z}_{n^2}^*$, for a random $e \gg n \cdot \varphi(n)/2$; observe that this is exactly $\mathsf{Enc}(0; e)$. The secret key is e. The DVNIZK proceeds as follows:

The prover P picks $x \xleftarrow{\$} \mathbb{Z}_n$ and a Paillier random coin r, and computes $X \leftarrow g^x$, $T' \leftarrow (1 + n)^t h^r \bmod n^2$, and $X' \leftarrow (1 + n)^x E^{-r} \bmod n^2$. The verifier V computes $D \leftarrow T^e X \bmod n^2$ and $D' \leftarrow (T')^e X' \bmod n^2$. Then, she checks that D' is of the form $(1 + n)^d \bmod n^2$. If so, V computes $d \bmod n$ from D', and checks that $D = g^d$. V accepts iff both checks succeeded.

Let us provide an intuition of the security of this scheme. Correctness follows easily by inspection. Zero-knowledge comes from the fact that T' hides t, under the IND-CPA security of Paillier. For statistical knowledge extractability, note E only reveals $e \bmod \varphi(n)$ to an unbounded adversary, which leaks (statistically) no information on $e \bmod n$ as $\varphi(n)$ is coprime to n. This ensures the value t' encrypted in T' must be equal to t, otherwise the verification equations would uniquely define $e \bmod n$, which is statistically unknown to the prover. The simulator knows $\varphi(n)$ (but not $e \bmod n$) and gets t by decrypting T'.

1.4 Applications

A natural application of non-interactive zero-knowledge proofs of knowledge is the design of privacy-preserving non-interactive authentication schemes. This includes classical authentication protocols, but also P-signatures [4] and their many applications, such as anonymous credentials [4], group signatures [20], electronic cash [19], or anonymous authentication [48]. Our framework can lead to a variety of efficient new constructions of designated-verifier variants for the above applications without pairings, whereas all previous constructions either had to rely on the random oracle model,

or use pairing-based cryptography.[2] In many scenarios of non-interactive authentication, the designated-verifier property is not an issue.

In addition, the aforementioned applications build upon the Groth-Sahai framework for NIZKs. However, Groth-Sahai NIZKs only satisfy a restricted notion of extractability, called f-extractability in [4]. As a result, constructions of privacy-preserving authentication mechanisms from Groth-Sahai NIZKs require a careful security analysis. Our framework leads to fully extractable zero-knowledge proofs, which could potentially simplify this. We note that our DVNIZKs are additionally randomizable, which has applications for delegatable anonymous credential schemes [3].

Other potential applications of our framework include round-efficient two-party computation protocols secure against malicious adversaries, electronic voting (see e.g. [17]), as well as designated-verifier variants of standard cryptographic primitives, such as verifiable encryption [13], or verifiable pseudorandom-functions [5]. Potential applications to the construction of adaptive oblivious transfers can also be envisioned: in [35], the authors mention that an adaptive oblivious transfer protocol can be designed by replacing the interactive zero-knowledge proofs of the protocol of [14] by non-interactive one. They raise two issues to this approach, namely, that Groth-Sahai proofs are only witness-indistinguishable for the required class of statements, and that they only satisfy a weak form of extractability. None of these restrictions apply to our DVNIZK constructions.

1.5 Related Work

Non-interactive zero-knowledge proofs were first introduced in [10]. Efficient publicly-verifiable non-interactive zero-knowledge proofs can be constructed in the random oracle model [31, 32, 47], or in the non-programmable random oracle model [42] (using a common reference string in addition). The latter construction was improved in [21]. In the standard model, the main construction of efficient publicly-verifiable NIZKs is the Groth-Sahai framework [39].

Designated-verifier non-interactive zero-knowledge arguments where first introduced in [46], where it was shown that the existence of semantically secure encryption implies the existence of DVNIZK arguments with bounded soundness; however, the construction is highly inefficient and

[2] These applications typically require a proof-friendly signature scheme, but designated-verifier variants of such scheme can easily be constructed (without pairings) from algebraic MACs [18, 40], by committing to the secret key of the MAC and proving knowledge of the committed value with a DVNIZK; such statements are naturally handled by our framework.

therefore only of theoretical interest. Furthermore, even putting aside efficiency consideration, the construction is inherently limited to arguments (as opposed to proofs) with bounded soundness (as opposed to unbounded soundness).

Designated-verifier NIZKs for linear languages can be constructed from hash proof systems [22, 33, 41]. Such NIZKs are perfectly zero-knowledge and statistically adaptively sound, but are not proofs of knowledge and are restricted to very specific statements, captured by linear equations.

Efficient designated-verifier NIZKs for more general statements were first described in [23]. The authors describe a general compiler that converts any three-round (honest-verifier) zero-knowledge protocol satisfying some (mild) requirements into a DVNIZK. However, the construction has several drawbacks: the soundness only holds under a very specific complexity-leveraging assumption, and only against adversaries making at most $O(\log \lambda)$ proofs (as already mentioned, the paper claims that the construction enjoy unbounded soundness as well, but this claim is flawed, see the full version [16]). In addition, the proofs obtained with this compiler are not proofs of knowledge.

In subsequent works [17, 49], variations of the compilation technique of [23] are described, where the complexity-leveraging assumption was replaced by more standard assumptions (although achieving a more restricted type of soundness) by relying on encryption schemes with additional properties. Eventually, [43] removes some of the constraints of the constructions of [17], and provides new protocols that can be compiled using the transformation. However, all the constructions obtained in these papers are only computationally sound, do not enjoy unbounded soundness, and are not proofs of knowledge; this strongly limits their scope, and in particular, prevents them from being used in the previously discussed applications.

1.6 Organization

In Section 2, we introduce our notation, and necessary primitives. We refer the reader to the full version of this work [16] for classical preliminaries on commitments and cryptosystems. Section 2 also describes the notion of a DVNIZK-friendly encryption scheme, which is central to our framework. In Section 3, we introduce our framework for building DVNIZKs of knowledge over an abelian group, illustrate it with practical examples, and prove its security. In Section 4, we describe the dual variant of our framework for proving statements over plaintexts of a DVNIZK-friendly encryption scheme. In the full version of this work [16], we additionally describe optimizations

on the efficiency of DVNIZKs for relations between plaintexts of a DVNIZK-friendly scheme, by eschewing unbounded soundness, as well as our attack on the unbounded soundness of [23].

2 Preliminaries

Throughout this paper, λ denotes the security parameter. A probabilistic polynomial time algorithm (PPT, also denoted *efficient* algorithm) runs in time polynomial in the (implicit) security parameter λ. A positive function f is *negligible* if for any polynomial p there exists a bound $B > 0$ such that, for any integer $k \geq B$, $f(k) \leq 1/|p(k)|$. An event depending on λ occurs with *overwhelming probability* when its probability is at least $1 - \mathrm{negl}(\lambda)$ for a negligible function negl. Given a finite set S, the notation $x \xleftarrow{\$} S$ means a uniformly random assignment of an element of S to the variable x. We represent adversaries as interactive probabilistic Turing machines; the notation $\mathscr{A}^{\mathcal{O}}$ indicates that the machine $\mathscr{A}$ is given oracle access to $\mathcal{O}$. Adversaries will sometime output an arbitrary state st to capture stateful interactions.

Abelian Groups and Modules. We use additive notation for groups for convenience, and write $(\mathbb{G}, \boldsymbol{+})$ for an abelian group of order k. When it is clear from the context, we denote 0 its neutral element (otherwise, we denote it $0_{\mathbb{G}}$). We denote by $\bullet$ the scalar-multiplication algorithm (*i.e.* for any $(x, G) \in \mathbb{Z}_k \times \mathbb{G}$, $x \bullet G = G\boldsymbol{+}G\boldsymbol{+}\ldots\boldsymbol{+}G$, where the sum contains x terms). Observe that we can naturally view $\mathbb{G}$ as a $\mathbb{Z}_k$-module $(\mathbb{G}, \boldsymbol{+}, \bullet)$, for the ring $(\mathbb{Z}_k, +, \cdot)$. For simplicity, we write $\boldsymbol{-}G$ for $(-1) \bullet G$. We use lower case to denote elements of $\mathbb{Z}_k$, upper case to denote elements of $\mathbb{G}$, and bold notations to denote vectors. We extend the notations $(\boldsymbol{+}, \boldsymbol{-})$ to vectors and matrices in the natural way, and write $\boldsymbol{x} \bullet \boldsymbol{G}$ to denote the scalar product $x_1 \bullet G_1 \boldsymbol{+} \ldots \boldsymbol{+} x_t \bullet G_t$ (where $\boldsymbol{x}, \boldsymbol{G}$ are vectors of the same length t). For a vector $\boldsymbol{v}$, we denote by $\boldsymbol{v}^{\intercal}$ its transpose. By $\mathsf{GGen}(1^\lambda)$, we denote a probabilistic efficient algorithm that, given the security parameter λ, generates an abelian group $\mathbb{G}$ such that the best known algorithm for solving discrete logs in G takes time 2^λ. In the following, we write $(\mathbb{G}, k) \xleftarrow{\$} \mathsf{GGen}(1^\lambda)$. Additionally, we denote by $\mathsf{GGen}(1^\lambda, k)$ a group generation algorithm that allows us to select the order k beforehand.

RSA Groups. A *strong prime* is a prime $p = 2p' + 1$ such that p' is also a prime. We call *RSA modulus* a product $n = pq$ of two strong primes. We

denote by φ Euler's totient function; it holds that $\varphi(n) = (p-1)(q-1)$. We denote by $\mathbb{J}_n$ the cyclic subgroup of $\mathbb{Z}_n^*$ of elements with Jacobi symbol 1 (the order of this group is $\varphi(n)/2$), and by QR_n the cyclic subroup of squares of $\mathbb{Z}_n^*$ (which is also a subgroup of $\mathbb{J}_n$ and has order $\varphi(n)/4$). By $\mathsf{Gen}(1^\lambda)$, we denote a probabilistic efficient algorithm that, given the security parameter λ, generates a strong RSA modulus n and secret parameters (p, q) where $n = pq$, such that the best known algorithm for factoring n takes time 2^λ. In the following, we write $(n, (p, q)) \xleftarrow{\$} \mathsf{Gen}(1^\lambda)$.

2.1 Encryption Schemes

The formal definition of an IND-CPA-secure public-key encryption scheme is recalled in the full version [16], but in short, a public-key encryption scheme S is a triple of PPT algorithms $(S.\mathsf{KeyGen}, S.\mathsf{Enc}, S.\mathsf{Dec})$, where $S.\mathsf{KeyGen}$ generates a pair $(\mathsf{ek}, \mathsf{dk})$ with an encryption key and a decryption key, decryption (with dk, deterministically) is the reverse operation of encryption (with ek, randomized), and no adversary can distinguish encryptions of one of two messages of its choice (IND-CPA security).

In this work, we will focus on additively homomorphic encryption schemes, which are homomorphic for both the message and the random coin. More formally, we require that the message space $\mathcal{M}$ and the random source $\mathcal{R}$ are integer sets $(\mathbb{Z}_M, \mathbb{Z}_R)$ for some integers (M, R), and that there exists an efficient operation $\oplus$ such that for any $(\mathsf{ek}, \mathsf{sk}) \xleftarrow{\$} \mathsf{KeyGen}(1^\lambda)$, any $(m_1, m_2) \in \mathbb{Z}_M^2$ and $(r_1, r_2) \in \mathbb{Z}_R^2$, denoting $(C_i)_{i\leq 2} \leftarrow (S.\mathsf{Enc}_{\mathsf{ek}}(m_i; r_i))_{i\leq 2}$, it holds that $C_1 \oplus C_2 = S.\mathsf{Enc}_{\mathsf{ek}}(m_1 + m_2 \bmod M; r_1 + r_2 \bmod R)$. We say an encryption scheme is *strongly additive* if it satisfies these requirements. Note that the existence of $\oplus$ implies (via a standard square-and-multiply method) the existence of an algorithm that, on input a ciphertext $C = S.\mathsf{Enc}_{\mathsf{ek}}(m; r)$ and an integer $\rho \in \mathbb{Z}$, outputs a ciphertext $C' = S.\mathsf{Enc}_{\mathsf{ek}}(\rho m \bmod M; \rho r \bmod R)$. We denote by $\rho \odot C$ the external multiplication of a ciphertext C by an integer ρ, and by $\ominus$ the operation $C \oplus (-1) \odot C'$ for two ciphertexts (C, C'). We will sometimes slightly abuse these notations, and write $C \oplus m$ (resp. $C \ominus m$) for a plaintext m to denote $C \oplus S.\mathsf{Enc}_{\mathsf{ek}}(m; 0)$ (resp. $C \ominus S.\mathsf{Enc}_{\mathsf{ek}}(m; 0)$).

A simple observation on strongly additively homomorphic encryption schemes is that IND-CPA security implies that R must either be equal to 0 mod M, or unknown given ek. Otherwise, an IND-CPA adversary would set $(m_0, m_1) = (0, 1)$ and check if $R \odot C$ equals $S.\mathsf{Enc}_{\mathsf{ek}}(0; 0)$ or $S.\mathsf{Enc}_{\mathsf{ek}}(R; 0)$.

The Paillier Encryption Scheme. The Paillier encryption scheme [45] is a well-known additively homomorphic encryption scheme over $\mathbb{Z}_n$ for an

RSA modulus n. We describe here a standard variant [25, 43], where the random coin is an exponent over $\mathbb{J}_n$ rather than a group element. Note that the exponent space of $\mathbb{J}_n$ is $\mathbb{Z}_{\varphi(n)/2}$, which is a group of unknown order; however, it suffices to draw exponents at random from $\mathbb{Z}_{n/2}$ to get a distribution statistically close from uniform over $\mathbb{Z}_{\varphi(n)/2}$.

- $\mathsf{KeyGen}(1^\lambda)$: run $(n,(p,q)) \xleftarrow{\$} \mathsf{Gen}(1^\lambda)$, pick $g \xleftarrow{\$} \mathbb{J}_n$, set $h \leftarrow g^n \bmod n^2$, and compute $\delta \leftarrow n^{-1} \bmod \varphi(n)$ (n and $\varphi(n)$ are relatively prime). Return $\mathsf{ek} = (n, h)$ and $\mathsf{dk} = \delta$;
- $\mathsf{Enc}(\mathsf{ek}, m; r)$: given $m \in \mathbb{Z}_n$, for a random $r \xleftarrow{\$} \mathbb{Z}_{n/2}$, compute and output $c \leftarrow (1+n)^m \cdot h^r \bmod n^2$;
- $\mathsf{Dec}(\mathsf{dk}, c)$: compute $x \leftarrow c^{\mathsf{dk}} \bmod n$ and $c_0 \leftarrow [c \cdot x^{-n} \bmod n^2]$. Return $m \leftarrow (c_0 - 1)/n$.

Note that knowing dk is equivalent to knowing the factorization of n. The IND-CPA security of the Paillier encryption scheme reduces to the decisional composite residuosity (DCR) assumption, which states that it is computationally infeasible to distinguish random n'th powers over $\mathbb{Z}_{n^2}^*$ from random elements of $\mathbb{Z}_{n^2}^*$.[3] It is also strongly additive, where the homomorphic addition of ciphertexts is the multiplication over $\mathbb{Z}_{n^2}^*$.

The ElGamal Encryption Scheme. We recall the additive variant of the famous ElGamal cryptosystem [28], over an abelian group $(\mathbb{G}, \mathbf{+})$ of order k.

- $\mathsf{KeyGen}(1^\lambda)$: pick $G \xleftarrow{\$} \mathbb{G}$, pick $s \xleftarrow{\$} \mathbb{Z}_k$, set $G \leftarrow s \bullet G$, and return $\mathsf{ek} = (G, H)$ and $\mathsf{dk} = s$;
- $\mathsf{Enc}(\mathsf{ek}, m; r)$: given $m \in \mathbb{Z}_k$, for a random $r \xleftarrow{\$} \mathbb{Z}_k$, output $\boldsymbol{C} \leftarrow (r \bullet G, (m \bullet G) \mathbf{+} (r \bullet H))$;
- $\mathsf{Dec}(\mathsf{dk}, \boldsymbol{C})$: parse $\boldsymbol{C}$ as (C_0, C_1), and compute $M \leftarrow C_1 \mathbf{-} (\mathsf{dk} \bullet C_0)$. Compute the discrete logarithm m of M in base G, and return m.

The IND-CPA security of the ElGamal encryption scheme reduces to the decisional Diffie-Hellman (DDH) assumption over $\mathbb{G}$, which states that it is computationally infeasible to distinguish tuples of the form $(G, H, x \bullet G, x \bullet H)$ for random x from uniformly random 4-tuples over $\mathbb{G}$. It is also strongly additive (and the homomorphic operation is the vector addition over $\mathbb{G}$). However, the decryption procedure is not efficient in

[3] In the variant we consider here, we must restrict our attention to elements of $\mathbb{Z}_{n^2}^*$ which have Jacobi symbol 1 when reduced modulo n as $g \in \mathbb{J}_n$, but this can be checked in polynomial time anyway.

general, as it requires to compute a discrete logarithm. For the decryption process to be efficient, the message m must be restricted to come from a subset of $\mathbb{Z}_k$ of polynomial size.

DVNIZK-Friendly Encryption Scheme. We say that a strongly additive encryption scheme is DVNIZK-*friendly*, when it satisfies the following additional properties:

- Coprimality Property: we require that the size M of the plaintext space and the size R of the random source are coprime[4], *i.e.*, $\gcd(M, R) = 1$;
- Decodable: for any $(\mathsf{ek}, \mathsf{sk}) \xleftarrow{\$} \mathsf{KeyGen}(1^\lambda)$, the function $f_{\mathsf{ek}} : m \mapsto \mathsf{Enc}_{\mathsf{ek}}(m; 0)$ must be efficiently invertible (*i.e.*, there is a PPT algorithm, which is given ek, computing f_{ek}^{-1} on any value from the image of f_{ek}).

One can observe that the Paillier cryptosystem is DVNIZK-friendly ($\gcd(n, \varphi(n)) = 1$, and any message m can be efficiently recovered from $\mathsf{Enc}_{\mathsf{ek}}(m; 0) = (1 + n)^m \bmod n^2$), while the ElGamal cryptosystem is not (it satisfies none of the above properties). Other DVNIZK-friendly cryptosystems include variants of the Paillier cryptosystem [12, 22, 24–26], and the more recent Castagnos-Laguillaumie cryptosystem [15], with prime-order plaintext space. For simplicity, we will also assume that all prime factors of the size M of the plaintext space of a DVNIZK-friendly cryptosystem are of superpolynomial size; our results can be extended to cryptosystems with a small plaintext space (or a plaintext space with small prime factors), but at a cost in efficiency. Note that by the homomorphic property, the decodability property implies that a plaintext can always be recovered from a ciphertext if the random coin is known.

2.2 Non-Interactive Zero-Knowledge Proof Systems

In the definitions below, we focus on proof systems for NP-languages that admit an efficient (polynomial-time) prover. For an NP-language $\mathscr{L}$, we denote $R_{\mathscr{L}}$ its associated relation, *i.e.*, a polynomial-time algorithm which satisfies $\mathscr{L} = \{x \mid \exists w, |w| = \mathsf{poly}(|x|) \wedge R_{\mathscr{L}}(x, w) = 1\}$. It is well known that non-interactive proof systems cannot exist for non-trivial languages in the plain model [44]; our constructions will be described in the common reference string model. For conciseness, the common reference string is always implictly given as input to all algorithms. We note that all of our

[4] In view of our previous observation on IND-CPA security for strongly additive cryptosystems, this implies that R is secret.

constructions can be readily adapted to work in the registered public-key model as well, a relaxation of the common reference string model introduced by Barak et al in [2].

While languages are naturally associated to *statements of membership*, the constructions of this paper will mainly consider *statements of knowledge*. We write $\mathsf{St}(x) = \nrightarrow\{w : R(x,w) = 1\}$ to denote the statement "I know a witness w such that $R(x,w) = 1$" for a word x and a polytime relation R. Similarly, we write $\mathsf{St}(x) = \exists\{w : R(x,w) = 1\}$ to denote the existential statement "there exists a witness w such that $R(x,w) = 1$".

Definition 1. *(Non-Interactive Zero-Knowledge Proof System) A non-interactive zero-knowledge (*NIZK*) proof system* Π *between for a family of languages* $\mathscr{L} = \{\mathscr{L}_{\mathsf{crs}}\}_{\mathsf{crs}}$ *is a quadruple of probabilistic polynomial-time algorithms* $(\Pi.\mathsf{Setup}, \Pi.\mathsf{KeyGen}, \Pi.\mathsf{Prove}, \Pi.\mathsf{Verify})$ *such that*

- $\Pi.\mathsf{Setup}(1^\lambda)$*, outputs a common reference string* crs *(which specifies the language* $\mathscr{L}_{\mathsf{crs}}$*),*
- $\Pi.\mathsf{KeyGen}(1^\lambda)$*, outputs a public key* pk *and a verification key* vk*,*
- $\Pi.\mathsf{Prove}(\mathsf{pk}, x, w)$*, on input the public key* pk*, a word* $x \in \mathscr{L}_{\mathsf{crs}}$*, and a witness* w*, outputs a proof* π*,*
- $\Pi.\mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, x, \pi)$*, on input the public key* pk*, the verification key* vk*, a word* x*, and a proof* π*, outputs* $b \in \{0,1\}$*,*

which satisfies the completeness, zero-knowledge, and soundness properties defined below.

We assume for simplicity that once it is generated, the common reference string crs is implicitly passed as an argument to the algorithms $(\Pi.\mathsf{KeyGen}, \Pi.\mathsf{Prove}, \Pi.\mathsf{Verify})$. In the above definition of NIZK proof systems, we let the key generation algorithm generate a verification key vk which is used by the verifier to check the proofs. We call *publicly verifiable non-interactive zero-knowledge proof system* a NIZK proof system in which vk is set to the empty string (or, equivalently, in which vk is made part of the public key). Otherwise, we call it a *designated-verifier non-interactive zero-knowledge proof system.*

Definition 2. *(Completeness) A* NIZK *proof system* $\Pi = (\Pi.\mathsf{Setup}, \Pi.\mathsf{KeyGen}, \Pi.\mathsf{Prove}, \Pi.\mathsf{Verify})$ *for a family of languages* $\mathscr{L} = \{\mathscr{L}_{\mathsf{crs}}\}_{\mathsf{crs}}$ *with relations* R_{crs} *satisfies the (perfect,statistical) completeness property if for* $\mathsf{crs} \xleftarrow{\$} \Pi.\mathsf{Setup}(1^\lambda)$*, for every* $x \in \mathscr{L}_{\mathsf{crs}}$ *and every witness* w *such that* $R_{\mathsf{crs}}(x,w) = 1$*,*

$$\Pr\left[\begin{matrix}(\mathsf{pk},\mathsf{vk}) \xleftarrow{\$} \Pi.\mathsf{KeyGen}(1^\lambda), \\ \pi \leftarrow \Pi.\mathsf{Prove}(\mathsf{pk}, x, w)\end{matrix} : \Pi.\mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, x, \pi) = 1\right] = 1 - \mu(\lambda)$$

where $\mu(\lambda) = 0$ *for perfect completeness, and* $\mu(\lambda) = \mathrm{negl}(\lambda)$ *for statistical completeness.*

We now define the zero-knowledge property.

Definition 3. *(Composable Zero-Knowledge) A* NIZK *proof system* $\Pi = (\Pi.\mathsf{Setup}, \Pi.\mathsf{KeyGen}, \Pi.\mathsf{Prove}, \Pi.\mathsf{Verify})$ *for a family of languages* $\mathscr{L} = \{\mathscr{L}_{\mathsf{crs}}\}_{\mathsf{crs}}$ *with relations* R_{crs} *satisfies the (perfect, statistical) composable zero-knowledge property if for any* $\mathsf{crs} \xleftarrow{\$} \Pi.\mathsf{Setup}(1^\lambda)$*, there exists a probabilistic polynomial-time simulator* Sim *such that for any stateful adversary* $\mathscr{A}$*,*

$$\left| \Pr\left[\begin{array}{ll} (\mathsf{pk}, \mathsf{vk}) \xleftarrow{\$} \Pi.\mathsf{KeyGen}(1^\lambda), & \\ (x, w) \leftarrow \mathscr{A}(\mathsf{pk}, \mathsf{vk}), & : \ (R_{\mathsf{crs}}(x, w) = 1) \wedge (\mathscr{A}(\pi) = 1) \\ \pi \leftarrow \Pi.\mathsf{Prove}(\mathsf{pk}, x, w) & \end{array} \right] - \right.$$
$$\left. \Pr\left[\begin{array}{ll} (\mathsf{pk}, \mathsf{vk}) \xleftarrow{\$} \Pi.\mathsf{KeyGen}(1^\lambda), & \\ (x, w) \leftarrow \mathscr{A}(\mathsf{pk}, \mathsf{vk}), & : \ (R_{\mathsf{crs}}(x, w) = 1) \wedge (\mathscr{A}(\pi) = 1) \\ \pi \leftarrow \mathsf{Sim}(\mathsf{pk}, \mathsf{vk}, x) & \end{array} \right] \right| \leq \mu(\lambda)$$

where $\mu(\lambda) = 0$ *for perfect composable zero-knowledge, and* $\mu(\lambda) = \mathrm{negl}(\lambda)$ *for statistical composable zero-knowledge. If the composable zero-knowledge property holds against efficient (PPT) verifiers, the proof system satisfies computational composable zero-knowledge.*

The composable zero-knowledge property was first introduced in [36]. It strenghtens the standard zero-knowledge definition, in that it explicitly states that the trapdoor of the simulator is exactly the verification key vk of the verifier. This strong security property guarantees that the same common reference string can be used for many different proofs, as the same trapdoor is used for simulating all proofs, which enhances the proof system with composability properties. We note that [36] additionally required indistinguishability between real and simulated common reference string; in our constructions, this will be trivially satisfied, as the simulated crs will be exactly the real one. We define below the notion of (bounded) adaptive soundness, which allows the input to be adversarially picked after the public key is fixed.

Definition 4. *(Bounded Adaptive Soundness) A* NIZK *proof system* $\Pi = (\Pi.\mathsf{Setup}, \Pi.\mathsf{KeyGen}, \Pi.\mathsf{Prove}, \Pi.\mathsf{Verify})$ *for a family of languages* $\mathscr{L} = \{\mathscr{L}_{\mathsf{crs}}\}_{\mathsf{crs}}$ *with relations* R_{crs} *satisfies the bounded adaptive soundness property if for* $\mathsf{crs} \xleftarrow{\$} \mathsf{Setup}(1^\lambda)$*, for every adversary* $\mathscr{A}$*,*

$$\Pr\left[\begin{array}{ll} (\mathsf{pk}, \mathsf{vk}) \xleftarrow{\$} \Pi.\mathsf{KeyGen}(1^\lambda), & \\ (\pi, x) \leftarrow \mathscr{A}(\mathsf{pk}) & : x \notin \mathscr{L}_{\mathsf{crs}} \wedge \Pi.\mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, x, \pi) \end{array} \right] = \mathrm{negl}(\lambda).$$

Definition 4 is formulated with respect to arbitrary adversaries $\mathscr{A}$, which leads to a statistical notion of soundness. A natural relaxation of this requirement is to consider only *efficient* (PPT) adversarial provers. We denote by *computational soundness* this relaxed notion of soundness. Computationally sound proof systems are called *argument systems*.

Unbounded Soundness. Definition 4 corresponds to a *bounded* notion of soundness, in the sense that soundness is only guaranteed to hold when the prover tries to forge a single proof of a wrong statement, right after the setup phase. However, if the prover is allowed to interact polynomially many times with the verifier before trying to forge a proof, sending proofs and receiving feedback on whether the proof was accepted, the previous definition provides no security guarantees.

Intuitively, in this situation, the distinction between bounded and unbounded soundness is comparable to the distinction between security against chosen plaintext attacks and security against chosen ciphertext attacks for cryptosystems. We define unbounded soundness in a similar fashion, by giving the prover access to a verification oracle $\mathcal{O}_{\mathsf{vk}}[\mathsf{pk}]$ (with crs implicitly given as parameter) which, on input (x, π), returns $b \leftarrow \mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, x, \pi)$.

Definition 5. *(Q-bounded Adaptive Soundness) A* NIZK *proof system* $\Pi = (\Pi.\mathsf{Setup}, \Pi.\mathsf{KeyGen}, \Pi.\mathsf{Prove}, \Pi.\mathsf{Verify})$ *for a family of languages* $\mathscr{L} = \{\mathscr{L}_{\mathsf{crs}}\}_{\mathsf{crs}}$ *with relations* R_{crs} *satisfies the Q-bounded adaptive soundness property if for* $\mathsf{crs} \xleftarrow{\$} \Pi.\mathsf{Setup}(1^\lambda)$*, and every adversary* $\mathscr{A}$ *making at most Q queries to* $\mathcal{O}_{\mathsf{vk}}[\mathsf{pk}]$*, it holds that*

$$\Pr\left[\begin{matrix}(\mathsf{pk}, \mathsf{vk}) \xleftarrow{\$} \Pi.\mathsf{KeyGen}(1^\lambda), \\ (\pi, x) \leftarrow \mathscr{A}^{\mathcal{O}_{\mathsf{vk}}[\mathsf{pk}]}(\mathsf{pk})\end{matrix} : x \notin \mathscr{L}_{\mathsf{crs}} \wedge \Pi.\mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, x, \pi)\right] = \mathrm{negl}(\lambda).$$

Alternatively, the above definition can be formulated with respect to polynomial-time adversarial provers, leading to computational Q-bounded adaptive soundness. Note that the answers of the oracle are bits; therefore, if a NIZK proof system satisfies the bounded adaptive soundness property of Definition 4, it also satisfies the above Q-bounded adaptive soundness property for any $Q = O(\log \lambda)$. Indeed, if Q is logarithmic, one can always guess in advance the answers of the verification oracle with non-negligible (inverse polynomial) probability. We say that a NIZK proof system which is Q-bounded adaptively sound for any $Q = \mathsf{poly}(\lambda)$ satisfies *unbounded adaptive soundness*.

Eventually, we define (unbounded) *knowledge-extractability*, a strenghtening of the soundness property which guarantees that if the prover produces an accepting proof, then the simulator can actually *extract* a witness for the statement. To this aim, we extend the syntax of the Setup algorithm to also output a trapdoor τ, used by the extractor. The knowledge-extractibility guarantee is stronger than soundness, in that the proof guarantees not only that there exists a witness, but also that the prover must know that witness. A NIZK satisfying knowledge-extractability is called a NIZK proof of knowledge.

Definition 6. *(Q-bounded Knowledge-Extractability) A* NIZK *proof system* $\Pi = (\Pi.\mathsf{Setup}, \Pi.\mathsf{KeyGen}, \Pi.\mathsf{Prove}, \Pi.\mathsf{Verify})$ *for a family of languages* $\mathscr{L} = \{\mathscr{L}_{\mathsf{crs}}\}_{\mathsf{crs}}$ *with relations* R_{crs} *satisfies the Q-bounded knowledge-extractability property if for* $(\mathsf{crs}, \tau) \xleftarrow{\$} \Pi.\mathsf{Setup}(1^\lambda)$*, and every adversary* $\mathscr{A}$ *making at most* Q *queries to* $\mathcal{O}_{\mathsf{vk}}[\mathsf{pk}]$*, there is an efficient extractor* Ext *such that*

$$\Pr\left[\begin{array}{l}(\mathsf{pk},\mathsf{vk}) \xleftarrow{\$} \Pi.\mathsf{KeyGen}(1^\lambda),\\ (\pi, x) \leftarrow \mathscr{A}^{\mathcal{O}_{\mathsf{vk}}[\mathsf{pk}]}(\mathsf{pk}),\\ w \leftarrow \mathsf{Ext}(\pi, x, \tau),\end{array} : R_{\mathsf{crs}}(x, w) \text{ iff } \Pi.\mathsf{Verify}(\mathsf{pk},\mathsf{vk}, x, \pi)\right] \approx 1.$$

3 A Framework for Designated-Verifier Non-Interactive Zero-Knowledge Proofs of Knowledge

In this section, we let k be an integer, $(\mathbb{G}, \boldsymbol{+})$ be an abelian group of order k, and (α, β, γ) be three integers. We will describe a framework for proving statements of knowledge over a wide variety of algebraic relations over $\mathbb{G}$, in the spirit of the Groth-Sahai framework for NIZK proofs over bilinear groups. To describe the relations handled by our framework, we describe languages of algebraic relations via linear maps. While this system was previously used to describe membership statements [7–9], we adapt it to statements of knowledge. As previously observed in [7], this system encompasses a wider class of languages than the Groth-Sahai framework.

3.1 Statements Defined by a Linear Map over $\mathbb{G}$

Let $\boldsymbol{G} \in \mathbb{G}^\alpha$ denote a vector of *public parameters*, and let $\boldsymbol{C} \in \mathbb{G}^\beta$ denote a public *word*. We will consider statements $\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C})$ defined by a linear map $\Gamma : (\mathbb{G}^\alpha, \mathbb{G}^\beta) \mapsto \mathbb{G}^{\gamma\times\beta}$ as follows:

$$\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C}) = \daleth\{\boldsymbol{x} \in \mathbb{Z}_k^\gamma \mid \boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = \boldsymbol{C}\}$$

That is, the prover knows a witness-vector $\boldsymbol{x} \in \mathbb{Z}_k^\gamma$ such that the equation $\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = \boldsymbol{C}$ holds. This abstraction captures a wide class of statements. Below, we describe two examples of statements that can be handled by our framework. They aim at clarifying the way the framework can be used, illustrating its power, as well as providing useful concrete instantiations. The examples focus on the most standard primitives (Pedersen commitments, ElGamal ciphertexts), but the reader will easily recognize they can be naturaly generalized to all standard variants of these primitives (e.g., variants of ElGamal secure under t-linear assumptions [11], or under assumptions from the matrix Diffie-Hellman family of assumptions [29]).

Example 1: Knowledge of Opening to a Pedersen Commitment. We consider statements of knowledge of an opening (m, r) to a Pedersen commitment C.

- Public Parameters: $(G, H) \in \mathbb{G}^2$;
- Word: $C \in \mathbb{G}$;
- Witness: a pair $(m, r) \in \mathbb{Z}_k^2$ such that $C = m \bullet G + r \bullet H$;
- Linear Map: $\Gamma_{\mathsf{Ped}} : (G, H, C) \mapsto (G, H)^\intercal$;
- Statement: $\mathsf{St}_{\Gamma_{\mathsf{Ped}}}(G, H, C) = \aleph\{(m, r) \in \mathbb{Z}_k^2 \mid (m, r) \bullet (G, H)^\intercal = C\}$.

Example 2: Multiplicative Relationship Between ElGamal Ciphertexts. This type of statement is of particular interest, as it can be generalized to arbitrary (polynomial) relationships between plaintexts.

- Public Parameters: $(G, H) \in \mathbb{G}^2$;
- Word: $\boldsymbol{C} = ((U_i, V_i)_{0 \le i \le 2}) \in \mathbb{G}^6$;
- Witness: a 5-tuple $\boldsymbol{x} = (m_0, r_0, m_1, r_1, r_2) \in \mathbb{Z}_k^5$ such that $U_i = r_i \bullet G$ and $V_i = m_i \bullet G + r \bullet H$ for $i = 0, 1$, and $U_2 = m_1 \bullet U_0 + r_2 \bullet G$, $V_2 = m_1 \bullet V_0 + r_2 \bullet H$;
- Linear Map:

$$\Gamma_{\mathsf{EM}} : (G, H, \boldsymbol{C}) \mapsto \begin{pmatrix} 0 & G & 0 & 0 & 0 & 0 \\ G & H & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & G & U_0 & V_0 \\ 0 & 0 & G & H & 0 & 0 \\ 0 & 0 & 0 & 0 & G & H \end{pmatrix};$$

- Statement: $\mathsf{St}_{\Gamma_{\mathsf{EM}}}(G, H, \boldsymbol{C}) = \aleph\{\boldsymbol{x} \in \mathbb{Z}_k^5 \mid \boldsymbol{x} \bullet \Gamma_{\mathsf{EM}}(G, H, \boldsymbol{C}) = \boldsymbol{C}\}$.

Conjunction of Statements. The above framework naturally handles conjuctions. Consider two statements $(\mathsf{St}_{\Gamma_0}(\boldsymbol{G_0}, \boldsymbol{C_0}), \mathsf{St}_{\Gamma_1}(\boldsymbol{G_1}, \boldsymbol{C_1}))$, defined by linear maps (Γ_0, Γ_1), with public parameters $(\boldsymbol{G_1}, \boldsymbol{G_1})$, words $(\boldsymbol{C_0}, \boldsymbol{C_1})$, and witnesses $(\boldsymbol{x_0}, \boldsymbol{x_1})$. Let $\boldsymbol{G} \leftarrow (\boldsymbol{G_1}, \boldsymbol{G_1})$, $\boldsymbol{C} \leftarrow (\boldsymbol{C_0}, \boldsymbol{C_1})$, and $\boldsymbol{x} \leftarrow (\boldsymbol{x_0}, \boldsymbol{x_1})$. We construct the linear map Γ associated to $\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C})$ as $\Gamma \leftarrow ((\Gamma_0, 0)^\intercal, (0, \Gamma_1)^\intercal)$. One can immediatly observe that $\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C}) = \mathsf{St}_{\Gamma_0}(\boldsymbol{G_0}, \boldsymbol{C_0}) \wedge \mathsf{St}_{\Gamma_1}(\boldsymbol{G_1}, \boldsymbol{C_1})$. The framework handles disjunction of statements as well, as observed in [1]; we omit the details.

3.2 A Framework for DVNIZK Proofs of Knowledge

We now introduce our framework for constructing designated-verifier non-interactive zero-knowledge proofs of knowledge for statements defined by a linear map over $\mathbb{G}$. Let $S = (S.\mathsf{KeyGen}, S.\mathsf{Enc}, S.\mathsf{Dec})$ denote a DVNIZK-friendly encryption scheme with plaintext space $\mathbb{Z}_k$. We construct a DVNIZK of knowledge $\Pi_\mathsf{K} = (\Pi_\mathsf{K}.\mathsf{Setup}, \Pi_\mathsf{K}.\mathsf{KeyGen}, \Pi_\mathsf{K}.\mathsf{Prove}, \Pi_\mathsf{K}.\mathsf{Verify})$ for a statement $\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C})$ over a word $\boldsymbol{C} \in \mathbb{G}^\beta$, with public parameters $\boldsymbol{G} \in \mathbb{G}^\alpha$, defined by a linear map $\Gamma : (\mathbb{G}^\alpha, \mathbb{G}^\beta) \mapsto \mathbb{G}^{\gamma \times \beta}$. Our construction proceeds as follows:

- $\Pi_\mathsf{K}.\mathsf{Setup}(1^\lambda)$: compute $(\mathsf{ek}, \mathsf{dk}) \xleftarrow{\$} S.\mathsf{KeyGen}(1^\lambda)$. Output $\mathsf{crs} \leftarrow \mathsf{ek}$. Note that ek defines a plaintext space $\mathbb{Z}_k$ and a random source $\mathbb{Z}_R$. As the IND-CPA and strong additive properties of S require R to be unknown, we assume that a bound B on R is publicly available. We denote $\ell \leftarrow 2^\lambda k B$.
- $\Pi_\mathsf{K}.\mathsf{KeyGen}(1^\lambda)$: pick $e \leftarrow \mathbb{Z}_\ell$, set $\mathsf{pk} \leftarrow S.\mathsf{Enc}_\mathsf{ek}(0; e)$ and $\mathsf{vk} \leftarrow e$.
- $\Pi_\mathsf{K}.\mathsf{Prove}(\mathsf{pk}, \boldsymbol{C}, \boldsymbol{x})$: on a word $\boldsymbol{C} \in \mathbb{Z}_k^\beta$, with witness $\boldsymbol{x}$ for the statement $\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C})$, pick $\boldsymbol{x'} \xleftarrow{\$} \mathbb{Z}_k^\gamma$, $\boldsymbol{r} \xleftarrow{\$} \mathbb{Z}_{2^\lambda B}^\gamma$, compute

  $$\boldsymbol{X} \leftarrow S.\mathsf{Enc}_\mathsf{ek}(\boldsymbol{x}, \boldsymbol{r}), \quad \boldsymbol{X'} \leftarrow S.\mathsf{Enc}_\mathsf{ek}(\boldsymbol{x'}, 0) \ominus (\boldsymbol{r} \odot \mathsf{pk}), \quad \boldsymbol{C'} \leftarrow \boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}),$$

  and output $\boldsymbol{\pi} \leftarrow (\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C'})$.
- $\Pi_\mathsf{K}.\mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, \boldsymbol{C}, \boldsymbol{\pi})$: parse $\boldsymbol{\pi}$ as $(\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C'})$. Check that $e \odot \boldsymbol{X} \oplus \boldsymbol{X'}$ is decodable, and decode it to a vector $\boldsymbol{d} \in \mathbb{Z}_k^\gamma$. Check that

  $$\boldsymbol{d} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = e \bullet \boldsymbol{C} \boldsymbol{+} \boldsymbol{C'}.$$

  If all checks succeeded, accept. Otherwise, reject.

The proof $\boldsymbol{\pi}$ consists of 2γ ciphertexts of S, and β elements of $\mathbb{G}$. Below, we illustrate our construction of DVNIZK on the examples of statements given in the previous section. For the sake of concreteness, we instantiate the

DVNIZK-friendly encryption scheme S with Paillier (hence the operation $\boldsymbol{+}$ is instantiated as the multiplication modulo n^2), so that the message space is $\mathbb{Z}_n$ and the randomizer space is $\mathbb{Z}_{\varphi(n)/2}$ for an RSA modulus n. In the examples, we use a bound $B = n$ and draw Paillier random coins from $\mathbb{Z}_{2^\lambda B}$, following our generic framework. However, observe that in the case of Paillier, we can also draw the coins from $\mathbb{Z}_{n/2}$ to get a distribution statistically close to uniform over $\mathbb{Z}_{\varphi(n)/2}$, which is more efficient.

Example 1: Knowledge of Opening to a Pedersen Commitment.

- $\Pi_{\mathsf{Ped}}.\mathsf{Setup}(1^\lambda)$: Compute $((n,h),\delta) = (\mathsf{ek},\mathsf{dk}) \xleftarrow{\$} S.\mathsf{KeyGen}(1^\lambda)$. Output $\mathsf{crs} \leftarrow \mathsf{ek}$. Let $\ell \leftarrow 2^\lambda n^2$. Let $\mathbb{G} \xleftarrow{\$} \mathsf{GGen}(1^\lambda, n)$, $(G,H) \xleftarrow{\$} \mathbb{G}^2$.
- $\Pi_{\mathsf{Ped}}.\mathsf{KeyGen}(1^\lambda)$: pick $e \xleftarrow{\$} \mathbb{Z}_\ell$, set $\mathsf{pk} \leftarrow h^e \bmod n^2$ and $\mathsf{vk} \leftarrow e$.
- $\Pi_{\mathsf{Ped}}.\mathsf{Prove}(\mathsf{pk}, C, \boldsymbol{x})$: on a word $C \in \mathbb{G}$, with witness $\boldsymbol{x} = (m,r) \in \mathbb{Z}_n^2$ for the statement $\mathsf{St}_{\Gamma_{\mathsf{Ped}}}(\boldsymbol{G},\boldsymbol{C})$, pick $\boldsymbol{x'} \xleftarrow{\$} \mathbb{Z}_n^2$, $\boldsymbol{\rho} \xleftarrow{\$} \mathbb{Z}_{2^\lambda B}^2$, compute $\boldsymbol{X} \leftarrow (1+n)^{\boldsymbol{x}} h^{\boldsymbol{\rho}} \bmod n^2, \boldsymbol{X'} \leftarrow (1+n)^{\boldsymbol{x'}} \mathsf{pk}^{-\boldsymbol{\rho}} \bmod n^2, \boldsymbol{C'} \leftarrow \boldsymbol{x'} \bullet (G,H)^\intercal$, and output $\boldsymbol{\pi} \leftarrow (\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C'})$.
- $\Pi_{\mathsf{Ped}}.\mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, \boldsymbol{C}, \boldsymbol{\pi})$: parse $\boldsymbol{\pi}$ as $(\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C'})$. Check that $\boldsymbol{X}^e \boldsymbol{X'}$ is of the form $(1+n)^{\boldsymbol{d}}$, and recover the vector $\boldsymbol{d} \in \mathbb{Z}_n^2$. Check that $\boldsymbol{d} \bullet (G,H)^\intercal = e \bullet \boldsymbol{C} \boldsymbol{+} \boldsymbol{C'}$.

Example 2: Multiplicative Relationship Between ElGamal Ciphertexts.

- $\Pi_{\mathsf{EM}}.\mathsf{Setup}(1^\lambda)$ as $\Pi_{\mathsf{Ped}}.\mathsf{Setup}(1^\lambda)$.
- $\Pi_{\mathsf{EM}}.\mathsf{KeyGen}(1^\lambda)$ as $\Pi_{\mathsf{Ped}}.\mathsf{KeyGen}(1^\lambda)$.
- $\Pi_{\mathsf{EM}}.\mathsf{Prove}(\mathsf{pk}, \boldsymbol{C}, \boldsymbol{x})$: on a word $\boldsymbol{C} \in \mathbb{G}^6$, with witness $\boldsymbol{x} = (m_0, r_0, m_1, r_1, r_2) \in \mathbb{Z}_n^5$ for the statement $\mathsf{St}_{\Gamma_{\mathsf{EM}}}(\boldsymbol{G},\boldsymbol{C})$, pick $\boldsymbol{x'} \xleftarrow{\$} \mathbb{Z}_n^5$, $\boldsymbol{\rho} \xleftarrow{\$} \mathbb{Z}_{2^\lambda B}^5$, compute $\boldsymbol{X} \leftarrow (1+n)^{\boldsymbol{x}} h^{\boldsymbol{\rho}} \bmod n^2, \boldsymbol{X'} \leftarrow (1+n)^{\boldsymbol{x}} \mathsf{pk}^{-\boldsymbol{\rho}} \bmod n^2, \boldsymbol{C'} \leftarrow \boldsymbol{x'} \bullet \Gamma_{\mathsf{EM}}(\boldsymbol{G},\boldsymbol{C})$, and output $\boldsymbol{\pi} \leftarrow (\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C'})$.
- $\Pi_{\mathsf{EM}}.\mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, \boldsymbol{C}, \boldsymbol{\pi})$: parse $\boldsymbol{\pi}$ as $(\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C'})$. Check that $\boldsymbol{X}^e \boldsymbol{X'}$ is of the form $(1+n)^{\boldsymbol{d}}$, and recover the vector $\boldsymbol{d} \in \mathbb{Z}_n^5$. Check that $\boldsymbol{d} \bullet \Gamma_{\mathsf{EM}}(\boldsymbol{G},\boldsymbol{C}) = e \bullet \boldsymbol{C} \boldsymbol{+} \boldsymbol{C'}$.

3.3 Security Proof

We now prove the generic DVNIZK construction from Section 3.2 is secure.

Perfect Completeness. It follows from straighforward calculations: $e \odot \boldsymbol{X} \oplus \boldsymbol{X'} = S.\mathsf{Enc}_{\mathsf{ek}}(e \cdot \boldsymbol{x} + \boldsymbol{x'}; e \cdot \boldsymbol{r} - e \cdot \boldsymbol{r}) = S.\mathsf{Enc}_{\mathsf{ek}}(e \cdot \boldsymbol{x} + \boldsymbol{x'}; 0)$ is decodable and decodes to $\boldsymbol{d} = e \cdot \boldsymbol{x} + \boldsymbol{x'} \bmod k$. Then, $\boldsymbol{d} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = e \bullet (\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C})) \boldsymbol{+} \boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = e \bullet \boldsymbol{C} \boldsymbol{+} \boldsymbol{C'}$ by the correctness of the statement $(\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = \boldsymbol{C})$ and by construction of $\boldsymbol{C'}$.

Composable Zero-Knowledge. We prove the following theorem:

Theorem 7 (Zero-Knowledge of Π_{K}). *If the encryption scheme S is* IND-CPA *secure, the* DVNIZK *scheme Π_{K} is composable zero-knowledge.*

We describe a simulator $\mathsf{Sim}(\boldsymbol{C}, \mathsf{pk}, \mathsf{vk})$ producing proofs computationally indistinguishable from those produced by an honest prover on true statements. The simulator operates as follows: let $\boldsymbol{d} \xleftarrow{\$} \mathbb{Z}_k^\gamma$, and $\boldsymbol{C'} \leftarrow \boldsymbol{d} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \boldsymbol{-} e \bullet \boldsymbol{C}$. Sample $\boldsymbol{x} \xleftarrow{\$} \mathbb{Z}_k^\gamma$, $\boldsymbol{r} \xleftarrow{\$} \mathbb{Z}_{2^\lambda B}^\gamma$, and compute $\boldsymbol{X} \leftarrow S.\mathsf{Enc}_{\mathsf{ek}}(\boldsymbol{x}, \boldsymbol{r}), \boldsymbol{X'} \leftarrow S.\mathsf{Enc}_{\mathsf{ek}}(\boldsymbol{d} \boldsymbol{-} \boldsymbol{e} \cdot \boldsymbol{x}, \boldsymbol{-e} \cdot \boldsymbol{r})$. Output $\pi_s = (\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C'})$.

Let $\mathscr{A}$ be an adversary that can distinguish Sim from Prove. We will build a reduction against the IND-CPA security of S. The reduction obtains $\boldsymbol{C}, \boldsymbol{x}$ from $\mathscr{A}$, samples $\tilde{\boldsymbol{x}} \leftarrow \mathbb{Z}_k^\gamma$, sends $(\boldsymbol{x}, \tilde{\boldsymbol{x}})$ to the IND-CPA game and sets $\boldsymbol{X}$ to be the challenge from the IND-CPA game. Now, the reduction samples $\boldsymbol{d} \leftarrow \mathbb{Z}_k^\gamma$ and sets $\boldsymbol{X'} := S.\mathsf{Enc}_{\mathsf{ek}}(\boldsymbol{d}; 0) \ominus \boldsymbol{X} \odot e$. Finally, the reduction sets $\boldsymbol{C'} := \boldsymbol{d} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \boldsymbol{-} e \bullet \boldsymbol{C}$. Send $\pi^* = (\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C})$ to $\mathscr{A}$.

Direct calculation shows that if the IND-CPA game outputs an encryption of $\tilde{\boldsymbol{X}}$, then $\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C}$ are distributed as those produced by Sim, whereas when it outputs an encryption of $\boldsymbol{X}$ then π^* is distributed identical to a real proof. Thus, whatever advantage $\mathscr{A}$ has in distinguishing Sim from Prove is also achieved by the reduction against IND-CPA. Note that for simplicity, our proof assume that the IND-CPA game is directly played over vectors, but standard methods allow to reduce this to the classical IND-CPA game with a single challenge ciphertext.

Adaptive Unbounded Knowledge-Extractability. We start by showing that Π_{K} satisfies statistical adaptive unbounded knowledge-extractability. More precisely, we prove the following theorem:

Theorem 8 (Soundness of Π_{K}). *There is an efficient simulator* Sim *such that for any (possibly unbounded) adversary $\mathscr{A}$ that outputs an accepting proof $\boldsymbol{\pi}$ with probability ε on an arbitrary word $\boldsymbol{C}$ after making at most Q queries to the oracle $\mathcal{O}_{\mathsf{vk}}[\mathsf{pk}]$,* Sim *extracts a valid witness for the statement $\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C})$ with probability at least $\varepsilon - (Q+1)\beta/p_k$, where p_k is the smallest prime factor of k.*

The proof describes an efficient simulator Sim that correctly emulates the verifier, without knowing vk mod k. The simulation is done as follows:

- $\mathsf{Sim.Setup}(1^\lambda)$: compute $(\mathsf{ek}, \mathsf{dk}) \xleftarrow{\$} S.\mathsf{KeyGen}(1^\lambda)$. Output $\mathsf{crs} \leftarrow \mathsf{ek}$. The encryption key ek defines a plaintext space $\mathbb{Z}_k$ and a random source $\mathbb{Z}_R$ with bound B. Let $\ell \leftarrow 2^\lambda kB$.
- $\mathsf{Sim.KeyGen}(1^\lambda)$: compute $(\mathsf{pk}, \mathsf{vk}) \xleftarrow{\$} \Pi_\mathsf{K}.\mathsf{KeyGen}(1^\lambda)$, output pk, store $e_R \leftarrow \mathsf{vk} \bmod R$, and erase vk.
- $\mathsf{Sim.Verify}(\mathsf{pk}, \mathsf{dk}, e_R, \boldsymbol{C}, \boldsymbol{\pi})$: parse $\boldsymbol{\pi}$ as $(\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{C'})$. Using the secret key dk of S, decrypt $\boldsymbol{X}$ to a vector $\boldsymbol{x}$, and $\boldsymbol{X'}$ to a vector $\boldsymbol{x'}$. Check that $(-e_R) \odot (\boldsymbol{X} \ominus \boldsymbol{x}) = \boldsymbol{X'} \ominus \boldsymbol{x'}$. Check that $\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = \boldsymbol{C}$, and that $\boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = \boldsymbol{C'}$. If all checks succeeded, accept. Otherwise, reject.

The simulator Sim first calls $\mathsf{Sim.Setup}(1^\lambda)$ to generate the common reference string (note that our simulator generates the common reference string honestly, hence the simulation of Setup cannot be distinguished from an honest run of Setup), and stores dk. Each time the adversary $\mathscr{A}$ sends a query $(\boldsymbol{C}, \boldsymbol{\pi})$ to the oracle $\mathcal{O}_\mathsf{vk}[\mathsf{pk}]$, Sim simulates $\mathcal{O}_\mathsf{vk}[\mathsf{pk}]$ (without knowing vk mod k) by running $\mathsf{Sim.Verify}(\mathsf{pk}, \mathsf{dk}, e_R, \boldsymbol{C}, \boldsymbol{\pi})$, and accepts or rejects accordingly. When $\mathscr{A}$ outputs a final answer $(\boldsymbol{C}, \boldsymbol{\pi})$, Sim computes a witness $\boldsymbol{x}$ for $\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C})$ by decrypting $\boldsymbol{C}$ with dk.

Observe that the distribution $\{(\mathsf{pk}, \mathsf{vk}) \xleftarrow{\$} \Pi_\mathsf{K}.\mathsf{KeyGen}(1^\lambda), e_k \leftarrow \mathsf{vk} \bmod k : (\mathsf{pk}, e_k)\}$ is statistically indistinguishable from the distribution $\{(\mathsf{pk}, \mathsf{vk}) \xleftarrow{\$} \Pi_\mathsf{K}.\mathsf{KeyGen}(1^\lambda), e_k \xleftarrow{\$} \mathbb{Z}_k : (\mathsf{pk}, e_k)\}$. Put otherwise, the distribution of vk mod k is statistically indistinguishable from random, even given pk. Indeed, as S is a DVNIZK-friendly encryption scheme, it holds by definition that $\gcd(k, R) = 1$. As $\ell = 2^\lambda Bk \geq 2^\lambda Rk$, the distribution $\{e \xleftarrow{\$} \mathbb{Z}_\ell, e_k \leftarrow e \bmod k, e_R \leftarrow e \bmod R : (e_k, e_R)\}$ is statistically indistinguishable from the uniform distribution over $\mathbb{Z}_k \times \mathbb{Z}_R$, and the value pk only leaks e_R, even to an unbounded adversary (as $S.\mathsf{Enc}_\mathsf{ek}(0; e) = S.\mathsf{Enc}_\mathsf{ek}(0; e \bmod R)$). We now prove the following claim:

Claim. For any public parameters $\boldsymbol{G}$ and word $\boldsymbol{C}$, it holds that

$$\Pr\left[\begin{array}{l}(\mathsf{pk}, \mathsf{vk}) \xleftarrow{\$} \Pi_\mathsf{K}.\mathsf{KeyGen}(1^\lambda),\\ b \leftarrow \mathsf{Sim.Verify}(\mathsf{pk}, \mathsf{dk}, \boldsymbol{C}, \boldsymbol{\pi}), : b' = b\\ b' \leftarrow \Pi_\mathsf{K}.\mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, \boldsymbol{C}, \boldsymbol{\pi})\end{array}\right] \geq 1 - \beta/p_k,$$

where p_k is one of the prime factors of k.

Proof. First, we show that if $b = 1$, then $b' = 1$. Indeed, let us denote $(\boldsymbol{x}, \boldsymbol{x'})$ the plaintexts associated to $(\boldsymbol{X}, \boldsymbol{X'})$. Let $(\boldsymbol{r}, \boldsymbol{r'})$ be the random

coins of the ciphertexts $(\boldsymbol{X}, \boldsymbol{X'})$. Observe that, by the homomorphic properties of S, the equation $(-e_R) \odot (\boldsymbol{X} \ominus \boldsymbol{x}) = \boldsymbol{X'} \ominus \boldsymbol{x'}$ is equivalent to $S.\mathsf{Enc}_{\mathsf{ek}}(0; -e_R \cdot \boldsymbol{r}) = S.\mathsf{Enc}_{\mathsf{ek}}(0; \boldsymbol{r'})$, which is equivalent to $e \odot \boldsymbol{X} \oplus \boldsymbol{X'} = S.\mathsf{Enc}(e \cdot \boldsymbol{x} + \boldsymbol{x'} \bmod k; e \cdot \boldsymbol{r} + \boldsymbol{r'} \bmod R) = S.\mathsf{Enc}(e \cdot \boldsymbol{x} + \boldsymbol{x'} \bmod k; 0)$ as $e = e_R \bmod R$. Therefore, the verifier's check that $e \odot \boldsymbol{X} \oplus \boldsymbol{X'}$ is decodable succeeds if and only if Sim's first check succeeds, and the decoded value $\boldsymbol{d} \in \mathbb{Z}_k^\gamma$ satisfies $\boldsymbol{d} = e \cdot \boldsymbol{x} + \boldsymbol{x'} \bmod k$. Moreover, if the equations $\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = \boldsymbol{C}$ and $\boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = \boldsymbol{C'}$ are both satisfied (*i.e.* Sim's other checks succeed), then it necessarily holds that $\boldsymbol{d} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = (e \cdot \boldsymbol{x} + \boldsymbol{x'}) \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = e \bullet (\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C})) \boldsymbol{+} \boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = e \bullet \boldsymbol{C} \boldsymbol{+} \boldsymbol{C'}$. This concludes the proof that, conditioned on Sim's checks succeeding, the verifier's checks necessarily succeed.

Now, assume for the sake of contradiction that the converse is not true: suppose that Sim rejected the proof, while the verifier accepted. We already showed that the equation $(-e_R) \odot (\boldsymbol{X} \ominus \boldsymbol{x}) = \boldsymbol{X'} \ominus \boldsymbol{x'}$ is equivalent to the equation $e \odot \boldsymbol{X} \oplus \boldsymbol{X'} = S.\mathsf{Enc}(e \cdot \boldsymbol{x} + \boldsymbol{x'} \bmod k; 0)$; therefore, if $e \odot \boldsymbol{X} \oplus \boldsymbol{X'}$ is decodable (it has random coin 0), then Sim's check that $(-e_R) \odot (\boldsymbol{X} \ominus \boldsymbol{x}) = \boldsymbol{X'} \ominus \boldsymbol{x'}$ succeeds. As we assumed that Sim rejects the proof, this means that at least one of Sim's last checks must fail: either $\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \neq \boldsymbol{C}$, or $\boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \neq \boldsymbol{C'}$. By the first check of the verifier, it holds that $e \odot \boldsymbol{X} \oplus \boldsymbol{X'}$ is decodable; denoting $(\boldsymbol{x}, \boldsymbol{x'})$ the plaintexts associated to $(\boldsymbol{X}, \boldsymbol{X'})$, it therefore decodes to $\boldsymbol{d} = e \cdot \boldsymbol{x} + \boldsymbol{x'} \bmod k$. By the second check of the verifier, it holds that $\boldsymbol{d} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = e \bullet \boldsymbol{C} \boldsymbol{+} \boldsymbol{C'}$, which implies $e \bullet (\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C})) \boldsymbol{+} \boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = e \bullet \boldsymbol{C} \boldsymbol{+} \boldsymbol{C'}$. This last equation rewrites to

$$e \bullet (\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \boldsymbol{-} \boldsymbol{C}) = \boldsymbol{C'} \boldsymbol{-} \boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \tag{1}$$

Now, recall that by assumption, either $\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \neq \boldsymbol{C}$, or $\boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \neq \boldsymbol{C'}$. Observe that Equation 1 further implies, as $e \neq 0$ (with overwhelming probability), that $\boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \boldsymbol{-} \boldsymbol{C'} \neq 0$ if and only if $\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \boldsymbol{-} \boldsymbol{C} \neq 0$. Therefore, conditioned on Sim rejecting the proof, it necessarily holds that $\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \boldsymbol{-} \boldsymbol{C} \neq 0$ and $\boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \boldsymbol{-} \boldsymbol{C'} \neq 0$. Let (μ_i, ν_i) be two non-zero entries of the vectors $(\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) \boldsymbol{-} \boldsymbol{C}, \boldsymbol{C'} \boldsymbol{-} \boldsymbol{x'} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}))$ at the same position $i \leq \beta$; by Equation 1, it holds that $e = \nu_i \cdot \mu_i^{-1} \bmod p$ for at least one of the prime factors p of k. However, recall that the value $e \bmod k$ is *statistically hidden* to the prover (and therefore, so is the value $e \bmod p$), hence the probability of this event happening can be upper-bounded by $\beta/p \leq \beta/p_k$. This concludes the proof of the claim. □

Now, consider an adversary $\mathscr{A}$ that outputs an accepting proof $(\boldsymbol{C}, \boldsymbol{\pi})$ with probability at least ε after a polynomial number Q of interactions

with the oracle $\mathcal{O}_{\mathsf{vk}}[\mathsf{pk}]$. By the above claim and a union bound, it necessarily holds that $\mathscr{A}$ outputs an accepting proof $(\boldsymbol{C}, \boldsymbol{\pi})$ with probability at least $\varepsilon - Q\beta/p_k$ after interacting Q times with $\mathsf{Sim.Verify}(\mathsf{pk}, \mathsf{dk}, e_R, \cdot, \cdot)$; moreover, with probability at least $1 - \beta p_k$, this proof is also accepted by Sim's verification algorithm. Overall, Sim obtains a proof accepted by his verification algorithm with probability at least $\varepsilon - (Q+1)\beta/p_k$. In particular, this implies that the vector $\boldsymbol{x}$ extracted by Sim from $\boldsymbol{\pi}$ satisfies $\boldsymbol{x} \bullet \Gamma(\boldsymbol{G}, \boldsymbol{C}) = \boldsymbol{C}$ with probability at least $\varepsilon - (Q+1)\beta/p_k$. Therefore, Sim extracts a valid witness for the knowledge statement $\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C})$ with probability at least $\varepsilon - (Q+1)\beta/p_k$. As the size k of a DVNIZK-friendly cryptosystem has only superpolynomially large prime-factors, it holds that p_k is superpolynomially large. As $(Q+1)\beta$ is polynomial, we conclude that if $\mathscr{A}$ outputs an accepting proof with non-negligible probability, then Sim extracts a valid witness with non-negligible probability.

4 Dual Variant of the Framework

In the previous section, we described a framework for constructing efficient DVNIZKs of knowledge for relations between words defined over an abelian group $(\mathbb{G}, \boldsymbol{+})$, using a cryptosystem with specific properties as the underlying commitment scheme for the proof system. In this section, we show that the framework can also be used in a dual way, by considering languages of relations between the plaintexts of the underlying encryption scheme – we call this variant 'dual variant' of the framework, as the roles of the underlying encryption scheme (which is used as a commitment scheme for the proof) and of the abelian group (which contains the words on which the proof is made) are partially exchanged. This allows for example to handle languages of relations between Paillier ciphertexts. To instantiate the framework, it suffices to have any perfectly binding commitment scheme defined over $\mathbb{G}$. This dual variant leads to efficient DVNIZK proofs for relations between, e.g., Paillier ciphertexts, whose zero-knowledge property reduces to the binding property of the commitment scheme over $\mathbb{G}$ (e.g. the DDH assumption, or its variants), and with statistical (unbounded, adaptive) soundness.

4.1 Perfectly Binding Commitment over $\mathbb{G}$

Suppose that we are given a perfectly binding homomorphic commitment $C = (C.\mathsf{Setup}, C.\mathsf{Com}, C.\mathsf{Verify})$, where $C.\mathsf{Com} : \mathbb{Z}_k \times \mathbb{Z}_k \mapsto \mathbb{G}^*$. Assume further that $C.\mathsf{Setup}$ generates a public vector of parameters $\boldsymbol{G} \in \mathbb{G}^*$, and

that there is a linear map Γ_C associated to this commitment such that for all $(m, r) \in \mathbb{Z}_k^2$, $C.\mathsf{Com}(m, r) = (m, r) \bullet \Gamma_C(\boldsymbol{G})$. Note this implies the commitment scheme is homomorphic over $\mathbb{G}$. ElGamal (Sect. 2.1), can be used as a commitment scheme satisfying these properties, is hiding under the DDH assumption and perfectly binding. We do so by using $\mathsf{KeyGen}(1^\lambda)$ in place of $\mathsf{Setup}(1^\lambda)$ to generate group elements (G, H) (the public key of the encryption scheme), and commit (i.e encrypt) via $\Gamma_C(G, H) = ((0, G)^\intercal, (G, H)^\intercal)$. We generalize this to commitments to length-t vectors as follow: we let $\Gamma_{C,t}$ denote the extended matrix such that $C.\mathsf{Com}(\boldsymbol{m}, \boldsymbol{r}) = (\boldsymbol{m}, \boldsymbol{r}) \bullet \Gamma_{C,t}(\boldsymbol{G})$, where $(\boldsymbol{m}, \boldsymbol{r})$ are vectors of length t ($\Gamma_{C,t}$ is simply the block-diagonal matrix whose t blocks are all equal to Γ_C). Consider now the following statement, where the word is a vector $\boldsymbol{C}$ of commitments:

$$\begin{aligned}\mathsf{St}_{\Gamma_{C,t}}(\boldsymbol{G}, \boldsymbol{C}) &= \mathsf{Ж}\{(\boldsymbol{m}, \boldsymbol{r}) \mid (\boldsymbol{m}, \boldsymbol{r}) \bullet \Gamma_{C,t}(\boldsymbol{G}) = \boldsymbol{C}\} \\ &= \mathsf{Ж}\{(\boldsymbol{m}, \boldsymbol{r}) \mid C.\mathsf{Com}(\boldsymbol{m}, \boldsymbol{r}) = \boldsymbol{C}\}.\end{aligned}$$

One can immediatly observe that this statement (which is a proof of knowledge of openings to a vector of commitments with C) is handled by the framework of Section 3.

4.2 Equality of Plaintexts between C and S

In this section, we describe a simple method to convert a DVNIZK on the statement $\mathsf{St}_{\Gamma_{C,t}}(\boldsymbol{G}, \boldsymbol{C}) = \mathsf{Ж}\{(\boldsymbol{m}, \boldsymbol{r}) \mid C.\mathsf{Com}(\boldsymbol{m}, \boldsymbol{r}) = \boldsymbol{C}\}$ into a DVNIZK on the statement $\mathsf{St}'(\boldsymbol{G}, \boldsymbol{C}, \boldsymbol{X_m}) = \exists\{(\boldsymbol{m}, \boldsymbol{\rho_m}, \boldsymbol{r}) \mid \boldsymbol{X_m} = S.\mathsf{Enc}_{\mathsf{ek}}(\boldsymbol{m}, \boldsymbol{\rho_m}) \wedge \boldsymbol{C} = C.\mathsf{Com}(\boldsymbol{m}, \boldsymbol{r})\}$ for a length-t vector $\boldsymbol{C}$ of commitments with a commitment scheme over $\mathbb{G}$ satisfying the requirements defined in the previous section, and a length-t vector of DVNIZK-friendly ciphertexts $\boldsymbol{X_m}$. Instantiating the framework of Section 3 for the statement $\mathsf{St}_{\Gamma_{C,t}}(\boldsymbol{G}, \boldsymbol{C})$, we get the following DVNIZK Π:

- $\Pi.\mathsf{Setup}(1^\lambda)$: compute $(\mathsf{ek}, \mathsf{dk}) \xleftarrow{\$} S.\mathsf{KeyGen}(1^\lambda)$. Output $\mathsf{crs} \leftarrow \mathsf{ek}$. Note that ek defines the plaintext space $\mathbb{Z}_k$ and the random source $\mathbb{Z}_R$ with bound B. We denote $\ell \leftarrow 2^\lambda k B$.
- $\Pi.\mathsf{KeyGen}(1^\lambda)$: pick $e \leftarrow \mathbb{Z}_\ell$, set $\mathsf{pk} \leftarrow S.\mathsf{Enc}_{\mathsf{ek}}(0; e)$ and $\mathsf{vk} \leftarrow e$.
- $\Pi.\mathsf{Prove}(\mathsf{pk}, \boldsymbol{C}, (\boldsymbol{m}, \boldsymbol{r}))$: on a word $\boldsymbol{C} \in \mathbb{Z}_k^t$, with witness $(\boldsymbol{m}, \boldsymbol{r})$ for the statement $\mathsf{St}_{\Gamma_{C,t}}(\boldsymbol{G}, \boldsymbol{C})$ (where $\boldsymbol{G} \xleftarrow{\$} C.\mathsf{Setup}(1^\lambda)$), pick random $(\boldsymbol{m'}, \boldsymbol{r'})$, random coins $(\boldsymbol{\rho_m}, \boldsymbol{\rho_r})$ for S, and compute

$$\begin{aligned}&\boldsymbol{X_m} \leftarrow S.\mathsf{Enc}_{\mathsf{ek}}(\boldsymbol{m}, \boldsymbol{\rho_m}), && \boldsymbol{X_r} \leftarrow S.\mathsf{Enc}_{\mathsf{ek}}(\boldsymbol{r}, \boldsymbol{\rho_r}), \\ &\boldsymbol{X'_m} \leftarrow S.\mathsf{Enc}_{\mathsf{ek}}(\boldsymbol{m'}, 0) \ominus (\boldsymbol{\rho_m} \odot \mathsf{pk}), && \boldsymbol{X'_r} \leftarrow S.\mathsf{Enc}_{\mathsf{ek}}(\boldsymbol{r'}, 0) \ominus (\boldsymbol{\rho_r} \odot \mathsf{pk}), \\ &\boldsymbol{C'} \leftarrow (\boldsymbol{m'}, \boldsymbol{r'}) \bullet \Gamma_{C,t}(\boldsymbol{G}, \boldsymbol{C}),\end{aligned}$$

and output $\boldsymbol{\pi} \leftarrow (\boldsymbol{X_m}, \boldsymbol{X'_m}, \boldsymbol{X_r}, \boldsymbol{X'_r}, \boldsymbol{C'})$.
- $\Pi_{\mathsf{K}}.\mathsf{Verify}(\mathsf{pk}, \mathsf{vk}, \boldsymbol{C}, \boldsymbol{\pi})$: parse $\boldsymbol{\pi}$ as $(\boldsymbol{X_m}, \boldsymbol{X'_m}, \boldsymbol{X_r}, \boldsymbol{X'_r}, \boldsymbol{C'})$. Check that $e \odot \boldsymbol{X_m} \oplus \boldsymbol{X'_m}$ and $e \odot \boldsymbol{X_r} \oplus \boldsymbol{X'_r}$ are decodable, and decode them to vectors $(\boldsymbol{d_m}, \boldsymbol{d_r}) \in (\mathbb{Z}_k^t)^2$. Check that $(\boldsymbol{d_m}, \boldsymbol{d_r}) \bullet \Gamma_{C,t}(\boldsymbol{G}, \boldsymbol{C}) = e \bullet \boldsymbol{C} + \boldsymbol{C'}$.

By the result of Section 3, this is an unbounded statistical adaptive knowledge-extractable DVNIZK of knowledge of an opening for C. Suppose now that we modify the above scheme as follow: we let $\boldsymbol{X_m}$ be part of the *word* on which the proof is executed, rather than being computed as part of the *proof* by the algorithm $\Pi.\mathsf{Prove}$. That is, we consider words of the form $(\boldsymbol{C}, \boldsymbol{X_m})$ with witness $(\boldsymbol{m}, \boldsymbol{r}, \boldsymbol{\rho_m})$ such that $(\boldsymbol{C}, \boldsymbol{X_m}) = (C.\mathsf{Com}(\boldsymbol{m}; \boldsymbol{r}), S.\mathsf{Enc}_{\mathsf{ek}}(\boldsymbol{m}, \boldsymbol{\rho_m}))$. Let Π' denote the modified proof, in which $\boldsymbol{X_m}$ is part of the word and $(\boldsymbol{X'_m}, \boldsymbol{X_r}, \boldsymbol{X'_r}, \boldsymbol{C'})$ are computed as in Π. Observe that the proof of security of our framework immediatly implies that Π' is a secure DVNIZK for *plaintext equality* between commitments with C and encryptions with S: our statistical argument shows that a (possibly unbounded) adversary has negligible probability of outputting a word $\boldsymbol{C}$ together with an accepting proof $\boldsymbol{\pi} = (\boldsymbol{X_m}, \boldsymbol{X'_m}, \boldsymbol{X_r}, \boldsymbol{X'_r}, \boldsymbol{C'})$ where the plaintext extracted by the simulator from $\boldsymbol{X_m}$ is not also the plaintext of $\boldsymbol{C}$. Hence, it is trivial that the probability of outputting a word $(\boldsymbol{C}, \boldsymbol{X_m})$ and an accepting proof $\boldsymbol{\pi'} = (\boldsymbol{X'_m}, \boldsymbol{X_r}, \boldsymbol{X'_r}, \boldsymbol{C'})$ where the plaintext extracted by the simulator from $\boldsymbol{X_m}$ is not also the plaintext of $\boldsymbol{C}$ is also negligible. Thus, we get:

Theorem 9. *The proof system Π' is an adaptive unbounded statistically sound proof for equality between plaintexts of C and plaintexts of S, whose composable zero-knowledge property reduces to the* IND-CPA *security of S.*

Note that the proof Π' is no longer a proof of knowledge: while the simulator can extract $(\boldsymbol{m}, \boldsymbol{r})$ from the prover, he cannot necessarily extract the random coins $\boldsymbol{\rho_m}$ of $\boldsymbol{X_m}$, which are now part of the witness. Therefore, for the protocol to make sense, it is important that C is perfectly binding .

4.3 A Framework for Relations between Plaintexts of S

The observations of the above section suggest a very natural way for designing DVNIZKs for relations between plaintexts $\boldsymbol{m} \in \mathbb{Z}_k^*$ of the encryption scheme S, which intuitively operates in two steps: first, we create commitments to the plaintexts $\boldsymbol{m}$ over $\mathbb{G}$ using C and prove them consistent with the encrypted values using the method described in the previous section. Then, we are able to use the framework of Section 3 to demonstrate the

desired relation holds between the commited values (this is a statement naturally captured by the framework). More formally, on input a vector of ciphertexts $\boldsymbol{X_m}$ encrypting plaintexts $\boldsymbol{m}$ with random coins $\boldsymbol{\rho_m}$,

- Pick $\boldsymbol{r}$ and compute $\boldsymbol{C} \leftarrow C.\mathsf{Com}(\boldsymbol{m}, \boldsymbol{r})$.
- Construct a DVNIZK for the statement $\mathsf{St}'(\boldsymbol{G}, \boldsymbol{C}, \boldsymbol{X_m})$ with witness $(\boldsymbol{m}, \boldsymbol{\rho_m}, \boldsymbol{r})$, using the method described in Section 4.2.
- Construct a DVNIZK for the statement $\mathsf{St}_\Gamma(\boldsymbol{G}, \boldsymbol{C})$ with witness $(\boldsymbol{m}, \boldsymbol{r})$, using the framework of Section 3.

The correctness of this approach is immediate: the second DVNIZK guarantees that the appropriate relation is satisfied between the plaintexts of the commitments, while the first one guarantees that the ciphertexts indeed encrypt the committed values. This leads to a DVNIZK proof of relation between plaintexts of S, with unbounded adaptive statistical soundness. Regarding zero-knowledge, as the proof starts by committing to $\boldsymbol{m}$ with C, we must in addition assume that the commitment scheme is hiding (the security analysis is straightforward).

Theorem 10. *The above system is an adaptive unbounded statistically sound proof for relations of plaintexts of S, whose composable zero-knowledge reduces to the* IND-CPA *security of S and the hiding property of C.*

We note that we can also obtain a variant of Theorem 10, where zero-knowledge only relies on the IND-CPA of S, and hiding of C implies the soundness property, using commitment schemes *a la* Groth-Sahai where the crs can be generated in two indistinguishable ways, one leading to a perfectly hiding scheme, and one leading to a perfectly binding scheme (such commitments are known, e.g., from the DDH assumption).

Example: Multiplicative Relationship Between Paillier Ciphertexts. We focus now on the useful case of multiplicative relationship between plaintexts of Paillier ciphertexts. We instantiate S with the Paillier encryption scheme over an RSA group $\mathbb{Z}_n$, with a public key (n, h) ($h = g^n \bmod n^2$ for a generator g of $\mathbb{J}_n$), and the commitment scheme C with the ElGamal encryption scheme over a group $\mathbb{G}$ of order n, with public key (G, H). Let $(P_0, P_1, P_2) \in (\mathbb{Z}^*_{n^2})^3$ be three Paillier ciphertexts, and let $(m_0, m_1, m_2, \rho_0, \rho_1, \rho_2)$ be such that $m_2 = m_0 m_1 \bmod n$, and $P_0 = (1+n)^{m_0} h^{\rho_0} \bmod n^2, P_1 = (1+n)^{m_1} h^{\rho_1} \bmod n^2, P_2 = (1+n)^{m_2} h^{\rho_2} \bmod n^2$. Let $E = h^e \bmod n^2$ denote the public key of the verifier. The designated-verifier NIZK for proving that P_2 encrypts $m_0 m_1$ proceeds as follows:

- **Committing over $\mathbb{G}$:** pick (r_0, r_1, r_2) and send $(U_i, V_i)_{0\leq i\leq 2} \leftarrow (r_i \bullet G, r_i \bullet H \boldsymbol{+} m_i \bullet G)_{0\leq i\leq 2}$ (which are commitments with ElGamal to (m_0, m_1, m_2) over $\mathbb{G}$).
- **Proof of Plaintext Equality:** pick $(m'_i, r'_i, \rho'_i)_{0\leq i\leq 2} \xleftarrow{\$} (\mathbb{Z}_n \times \mathbb{Z}_n \times \mathbb{Z}_{n/2})^3$, and send for $i = 0$ to 2, $X_i \leftarrow (1+n)^{r_i} h^{\rho'_i} \bmod n^2, X'_i \leftarrow (1+n)^{r'_i} E^{-\rho'_i} \bmod n^2, P'_i \leftarrow (1+n)^{m'_i} E^{-\rho_i} \bmod n^2$, and $(U'_i, V'_i) \leftarrow (r'_i \bullet G, r'_i \bullet H \boldsymbol{+} m'_i \bullet G)$.
- **Proof of Multiplicative Relationship Between the Committed Values:** apply the proof system of Example 2 from Section 3 to the word $(U_i, V_i)_{0\leq i\leq 2}$, with public parameters (G, H), and the witness $\boldsymbol{x} = (m_0, r_0, m_1, r_1, r_2 - r_0 m_1)$ which satisfies $(U_0, V_0) = (r_0 \bullet G, r_0 \bullet H \boldsymbol{+} m_0 \bullet G), (U_1, V_1) = (r_1 \bullet G, r_1 \bullet H \boldsymbol{+} m_1 \bullet G)$, and $(U_2, V_2) = ((r_2 - r_0 m_1) \bullet G \boldsymbol{+} m_1 \bullet U_0, (r_2 - r_0 m_1) \bullet H \boldsymbol{+} m_1 \bullet V_0)$.
- **Proof Verification:** upon receving $(U_i, V_i, X_i, X'_i, P'_i, U'_i, V'_i)_{0\leq i\leq 2}$ together with the proof of multiplicative relationship between the values committed with $(U_i, V_i)_i$, the verifier with verification key $\mathsf{vk} = e$ checks that $e \odot P_i \oplus P'_i$ and $e \odot X_i \oplus X'_i$ successfully decode (respectively) to values p_i, x_i, and that $e \bullet U_i \boldsymbol{+} U'_i = x_i \bullet G$ and $e \bullet V_i \boldsymbol{+} V'_i = x_i \bullet H \boldsymbol{+} p_i \bullet G$, for $i = 0$ to 2. The verifier additionally checks the multiplicative proof, as in Example 4 from Section 3. She accepts iff all checks succeed.

The proof for the multiplicative statement involves 10 Paillier ciphertexts and 3 ElGamal ciphertexts. Overall, the total proof involves 20 Paillier ciphertexts, and 9 ElGamal ciphertexts. However, this size is obtained by applying the framework naively; in this situation, it introduces a lot of redudancy. For instance, instead of computing Paillier encryptions of (m_0, r_0, m_1, r_1) in the third phase, one can simply reuse the word (P_0, P_1) and the ciphertexts (X_0, X_1), as well as reusing $(P'_i, X'_i)_i$ for the corresponding masks $(m'_i, r'_i)_i$, saving 8 Paillier ciphertexts; similar savings can be obtained for the ElGamal ciphertexts, leading to a proof of total size 12 Paillier ciphertexts + 7 ElGamal ciphertexts.

Furthermore, if we eschew unbounded soundness and accept bounds on m_i we are able to produce a much shorter proof, comprising only two Paillier ciphertexts, outperforming even Fiat-Shamir. We detail this in the full version [16].

Acknowledgements

We thank Jens Groth for insightful discussions and contributions to early versions of this work. The first author was supported by EU Horizon 2020

grant 653497 (project PANORAMIX). The second author was supported by ERC grant 339563 (project CryptoCloud) and ERC grant 724307 (project PREP-CRYPTO).

References

1. Abdalla, M., Benhamouda, F., Pointcheval, D.: Disjunctions for hash proof systems: New constructions and applications. pp. 69–100. LNCS, Springer (2015)
2. Barak, B., Canetti, R., Nielsen, J.B., Pass, R.: Universally composable protocols with relaxed set-up assumptions. In: 45th FOCS. pp. 186–195. IEEE Computer Society Press (Oct 2004)
3. Belenkiy, M., Camenisch, J., Chase, M., Kohlweiss, M., Lysyanskaya, A., Shacham, H.: Randomizable proofs and delegatable anonymous credentials. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 108–125. Springer (Aug 2009)
4. Belenkiy, M., Chase, M., Kohlweiss, M., Lysyanskaya, A.: P-signatures and noninteractive anonymous credentials. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 356–374. Springer (Mar 2008)
5. Belenkiy, M., Chase, M., Kohlweiss, M., Lysyanskaya, A.: Compact e-cash and simulatable VRFs revisited. In: Shacham, H., Waters, B. (eds.) PAIRING 2009. LNCS, vol. 5671, pp. 114–131. Springer (Aug 2009)
6. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Ashby, V. (ed.) ACM CCS 93. pp. 62–73. ACM Press (Nov 1993)
7. Benhamouda, F., Blazy, O., Chevalier, C., Pointcheval, D., Vergnaud, D.: New techniques for SPHFs and efficient one-round PAKE protocols. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 449–475. Springer (Aug 2013)
8. Benhamouda, F., Couteau, G., Pointcheval, D., Wee, H.: Implicit zero-knowledge arguments and applications to the malicious setting. In: CRYPTO 2015, Part II. pp. 107–129. LNCS, Springer (Aug 2015)
9. Benhamouda, F., Pointcheval, D.: Trapdoor smooth projective hash functions. Cryptology ePrint Archive, Report 2013/341 (2013), `http://eprint.iacr.org/2013/341`
10. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications (extended abstract). In: 20th ACM STOC. pp. 103–112. ACM Press (May 1988)
11. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer (Aug 2004)
12. Bresson, E., Catalano, D., Pointcheval, D.: A simple public-key cryptosystem with a double trapdoor decryption mechanism and its applications. In: Laih, C.S. (ed.) ASIACRYPT 2003. LNCS, vol. 2894, pp. 37–54. Springer (Nov / Dec 2003)
13. Camenisch, J., Damgård, I.: Verifiable encryption, group encryption, and their applications to separable group signatures and signature sharing schemes. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 331–345. Springer (Dec 2000)
14. Camenisch, J., Neven, G., shelat, a.: Simulatable adaptive oblivious transfer. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 573–590. Springer (May 2007)

15. Castagnos, G., Laguillaumie, F.: Linearly homomorphic encryption from DDH. In: CT-RSA 2015. pp. 487–505. LNCS, Springer (2015)
16. Chaidos, P., Couteau, G.: Efficient designated-verifier non-interactive zero-knowledge proofs of knowledge. Cryptology ePrint Archive, Report 2017/1029 (2017), http://eprint.iacr.org/2017/1029
17. Chaidos, P., Groth, J.: Making sigma-protocols non-interactive without random oracles. In: PKC 2015. pp. 650–670. LNCS, Springer (2015)
18. Chase, M., Meiklejohn, S., Zaverucha, G.: Algebraic MACs and keyed-verification anonymous credentials. In: ACM CCS 14. pp. 1205–1216. ACM Press (2014)
19. Chaum, D., Fiat, A., Naor, M.: Untraceable electronic cash. In: Goldwasser, S. (ed.) CRYPTO'88. LNCS, vol. 403, pp. 319–327. Springer (Aug 1990)
20. Chaum, D., van Heyst, E.: Group signatures. In: Davies, D.W. (ed.) EUROCRYPT'91. LNCS, vol. 547, pp. 257–265. Springer (Apr 1991)
21. Ciampi, M., Persiano, G., Siniscalchi, L., Visconti, I.: A transform for NIZK almost as efficient and general as the Fiat-Shamir transform without programmable random oracles. pp. 83–111. LNCS, Springer (2016)
22. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer (Apr / May 2002)
23. Damgård, I., Fazio, N., Nicolosi, A.: Non-interactive zero-knowledge from homomorphic encryption. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 41–59. Springer (Mar 2006)
24. Damgård, I., Jurik, M.: A length-flexible threshold cryptosystem with applications. In: Safavi-Naini, R., Seberry, J. (eds.) ACISP 03. LNCS, vol. 2727, pp. 350–364. Springer (Jul 2003)
25. Damgård, I., Jurik, M., Nielsen, J.B.: A generalization of paillier's public-key system with applications to electronic voting. International Journal of Information Security 9(6), 371–385 (2010)
26. Damgård, I., Jurik, M.: A generalisation, a simplification and some applications of Paillier's probabilistic public-key system. In: Kim, K. (ed.) PKC 2001. LNCS, vol. 1992, pp. 119–136. Springer (Feb 2001)
27. De Santis, A., Di Crescenzo, G., Ostrovsky, R., Persiano, G., Sahai, A.: Robust non-interactive zero knowledge. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 566–598. Springer (Aug 2001)
28. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. In: Blakley, G.R., Chaum, D. (eds.) CRYPTO'84. LNCS, vol. 196, pp. 10–18. Springer (Aug 1984)
29. Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 129–147. Springer (Aug 2013)
30. Feige, U., Lapidot, D., Shamir, A.: Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In: 31st FOCS. pp. 308–317. IEEE Computer Society Press (Oct 1990)
31. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO'86. LNCS, vol. 263, pp. 186–194. Springer (Aug 1987)
32. Fischlin, M.: Communication-efficient non-interactive proofs of knowledge with online extractors. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 152–168. Springer (Aug 2005)

33. Gay, R., Hofheinz, D., Kiltz, E., Wee, H.: Tightly CCA-secure encryption without pairings. pp. 1–27. LNCS, Springer (2016)
34. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM Journal on Computing 18(1), 186–208 (1989)
35. Green, M., Hohenberger, S.: Universally composable adaptive oblivious transfer. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 179–197. Springer (Dec 2008)
36. Groth, J.: Simulation-sound NIZK proofs for a practical language and constant size group signatures. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 444–459. Springer (Dec 2006)
37. Groth, J., Ostrovsky, R., Sahai, A.: Non-interactive zaps and new techniques for NIZK. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 97–111. Springer (Aug 2006)
38. Groth, J., Ostrovsky, R., Sahai, A.: Perfect non-interactive zero knowledge for NP. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 339–358. Springer (May / Jun 2006)
39. Groth, J., Sahai, A.: Efficient non-interactive proof systems for bilinear groups. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 415–432. Springer (Apr 2008)
40. Kiltz, E., Pan, J., Wee, H.: Structure-preserving signatures from standard assumptions, revisited. In: CRYPTO 2015, Part II. pp. 275–295. LNCS, Springer (Aug 2015)
41. Kiltz, E., Wee, H.: Quasi-adaptive NIZK for linear subspaces revisited. pp. 101–128. LNCS, Springer (2015)
42. Lindell, Y.: An efficient transform from sigma protocols to NIZK with a CRS and non-programmable random oracle. pp. 93–109. LNCS, Springer (2015)
43. Lipmaa, H.: Optimally sound sigma protocols under DCRA. Cryptology ePrint Archive, Report 2017/703 (2017), `http://eprint.iacr.org/2017/703`
44. Oren, Y.: On the cunning power of cheating verifiers: Some observations about zero knowledge proofs (extended abstract). In: 28th FOCS. pp. 462–471. IEEE Computer Society Press (Oct 1987)
45. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT'99. LNCS, vol. 1592, pp. 223–238. Springer (May 1999)
46. Pass, R., shelat, a., Vaikuntanathan, V.: Construction of a non-malleable encryption scheme from any semantically secure one. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 271–289. Springer (Aug 2006)
47. Pointcheval, D., Stern, J.: Security proofs for signature schemes. In: Maurer, U.M. (ed.) EUROCRYPT'96. LNCS, vol. 1070, pp. 387–398. Springer (May 1996)
48. Teranishi, I., Furukawa, J., Sako, K.: k-Times anonymous authentication (extended abstract). In: Lee, P.J. (ed.) ASIACRYPT 2004. LNCS, vol. 3329, pp. 308–322. Springer (Dec 2004)
49. Ventre, C., Visconti, I.: Co-sound zero-knowledge with public keys. In: Preneel, B. (ed.) AFRICACRYPT 09. LNCS, vol. 5580, pp. 287–304. Springer (Jun 2009)