



HAL
open science

Rational Modular Encoding in the DCR Setting: Non-Interactive Range Proofs and Paillier-Based Naor-Yung in the Standard Model

Julien Devevey, Benoît Libert, Thomas Peters

► **To cite this version:**

Julien Devevey, Benoît Libert, Thomas Peters. Rational Modular Encoding in the DCR Setting: Non-Interactive Range Proofs and Paillier-Based Naor-Yung in the Standard Model. Public-Key Cryptography (PKC 2022) - 25th IACR International Conference on Practice and Theory of Public-Key Cryptography, Mar 2022, Yokohama (devenu virtuel pour cause de COVID), Japan. hal-03807457

HAL Id: hal-03807457

<https://inria.hal.science/hal-03807457>

Submitted on 9 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Rational Modular Encoding in the DCR Setting: Non-Interactive Range Proofs and Paillier-Based Naor-Yung in the Standard Model

Julien Devevey¹, Benoît Libert^{2,1}, and Thomas Peters³

¹ ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, Inria, UCBL), France

² CNRS, Laboratoire LIP, France

³ FNRS and UCLouvain, ICTEAM, Belgium

Abstract. Range proofs allow a sender to convince a verifier that committed integers belong to an interval without revealing anything else. So far, all known non-interactive range proofs in the standard model rely on groups endowed with a bilinear map. Moreover, they either require the group order to be larger than the range of any proven statement or they suffer from a wasteful rate. Recently (Eurocrypt’21), Couteau *et al.* introduced a new approach to efficiently prove range membership by encoding integers as a modular ratio between small integers. We show that their technique can be transposed in the standard model under the Composite Residuosity (DCR) assumption. Interestingly, with this modification, the size of ranges is not a priori restricted by the common reference string. It also gives a constant ratio between the size of ranges and proofs. Moreover, we show that their technique of encoding messages as bounded rationals provides a secure standard model instantiation of the Naor-Yung CCA2 encryption paradigm under the DCR assumption.

Keywords. Range proofs, NIZK, standard model, Naor-Yung.

1 Introduction

Zero-knowledge proofs [39] make it possible for a prover to convince a verifier about the truth of a statement while revealing nothing else. Since their introduction, they have been used in countless cryptographic protocols to protect users’ privacy or to hedge against malicious adversaries. In many situations, it is desirable to have non-interactive zero-knowledge (NIZK) proofs comprised of a single message from the prover to the verifier. In the non-interactive setting, NIZK proofs necessarily rely on a common reference string generated by some trusted party. While the Fiat-Shamir paradigm [34] allows for non-interactive proofs without a trusted setup in the random oracle model, it is known to only provide heuristic arguments in terms of security.

In the standard model, NIZK proofs are known to exist for all NP languages under well-studied assumptions [7,6,43,61]. For specific languages, however, much more efficient constructions are often possible, by dispensing with the need for an expensive Karp reduction.

Efficient NIZK constructions exist in the context of range proofs [11], where a prover convinces a verifier that a committed value belongs to a specific interval. Range proofs served as a building block of a number of cryptographic protocols, including anonymous credentials or e-cash [15], auction protocols [55], e-voting [41], and many more. Recently, they also served as crucial components of cryptocurrencies [57,12], where transaction amounts are private and only appear in committed [57] or encrypted [12] form. Range proofs then come into play to ensure that the committed/encrypted value lives in the correct range instead of being, e.g., slightly larger than the order of the message space.

A widely used approach [10,53,41] proceeds by committing to integers [37,28], rather than finite field elements. By withholding the order $|\mathcal{M}|$ of the message space, it forces the prover argue over the integers in order to demonstrate that a committed integer fits in a range $[0, B]$, where $B \in \mathbb{Z}$ may be larger than $|\mathcal{M}|$.

Recently, Couteau *et al.* [24] suggested an elegant technique that surprisingly emulates the properties of integer commitments in the discrete logarithm setting over public-order groups. The core idea of their construction is to view each Pedersen commitment [60] $C = g^m \cdot h^r$ as committing to the rounded rational $\lfloor x/c \rfloor \in \mathbb{Z}$, where x and c are small-magnitude integers $x, c \in \mathbb{Z}$ such that $m = x \cdot c^{-1} \bmod q$, where q is the group order. This approach yields instantiations in class groups and under lattice assumptions. In the discrete-log setting, it outperforms the BulletProof technique [13] for a wide range of parameters. It also enables either computational or statistical soundness (whereas integers commitments only offer computational soundness).

In this paper, we consider their approach in the Composite Residuosity setting [58], where we highlight several advantages when proving range membership of Paillier-encrypted values.

1.1 Our Contribution

RANGE PROOFS. We provide the first *unbounded* non-interactive range proof with constant rate in the standard model. The rate is defined in the standard way, as the ratio between the length of the witness and the total length of commitments and proofs. By “unbounded”, we mean that a fixed-size common reference string makes it possible to commit to arbitrarily large integers.⁴ In the standard model, it is also the first non-interactive candidate that does not rely on pairing-friendly groups. Instead, we can prove security under the standard Composite Residuosity (DCR) and Learning-with-Errors (LWE) [62] assumptions. While our construction provides statistical soundness (and computational zero-knowledge), it can be turned into a dual-mode NIZK system – where soundness/zero-knowledge can be either statistical or computational depending on the configuration of the CRS – at the cost of sacrificing unboundedness.

In either case, we obtain space-efficient proofs consisting of a constant number of Damgård-Jurik [29] ciphertexts. Asymptotically, the communication cost

⁴ It is tempting to believe that Groth-Sahai proofs achieve unboundedness. In Supplementary Material D, we explain why it is not the case.

is dominated by $O(\lambda^{3-O(1)} + \log B)$ bits, where B is the range size, which is on par with constructions based on integer commitments [53,41,25] in the random oracle model. In comparison, standard-model solutions based on Groth-Sahai proofs [44] cost $O(\lambda \cdot \log B)$ per proof.

Our unbounded range proof makes it possible to prove that a Paillier ciphertext decrypts to a modular ratio $M = x \cdot c^{-1} \bmod N^\zeta$, for some $\zeta \in \mathbb{N}$ and bounded integers $x, c \in \mathbb{Z}$ such that $\lfloor x/c \rfloor \in \mathbb{Z}$ belongs to a range $[0, B]$. As a second contribution, we show that this encoding technique can be used to instantiate the Naor-Yung CCA2-secure encryption paradigm [56].

DCR-BASED INSTANTIATION OF NAOR-YUNG IN THE STANDARD MODEL. We give a Σ -protocol proving plaintext equalities between Paillier ciphertexts encrypted under distinct moduli, which restores the soundness of a Σ -protocol used by Fouque and Pointcheval [35]. Recently, Devevey *et al.* [31, Appendix E] showed that the Σ -protocol of [35, Section 4.2] does not provide soundness as a cheating prover can exploit the distinct moduli to prove false statements. This invalidates the proof⁵ that the DCR-based threshold cryptosystem of [35] provides IND-CCA2 security in the random oracle model. Devevey *et al.* [31] suggested to fix the problem by additionally proving that the plaintext is smaller than both Paillier moduli. While efficient range proofs (e.g., [41,25,13]) can solve this problem in the random oracle model, we do not know how to instantiate them in the standard model via the Fiat-Shamir paradigm. To achieve standard-model security by exploiting correlation-intractable hash functions as in [19,61], we show that no range proof is actually necessary if the decryption algorithm is modified and “undoes” the rational modular encoding of Couteau *et al.* [24].

We show that the modified decryption algorithm can be combined with the correlation-intractable hash functions of [19,61] so as to instantiate the scheme in the standard model. As a result, we obtain a new construction of a non-interactive threshold CCA2-secure cryptosystem without pairings. Devevey *et al.* [31] recently proposed such a construction under the DCR and LWE assumptions. Our scheme provides several advantages over their construction. It notably inherits a property of the Damgård-Jurik system [29], which makes it possible to encrypt very long messages⁶ for a fixed size public key comprised of an RSA modulus N . Variable-length plaintexts can even be encrypted by flexibly choosing an integer $\zeta > 1$, depending on the message length, and working over $\mathbb{Z}_{N^{\zeta+1}}^*$. In the threshold setting, the key generation phase requires to set a bound on the maximal value of ζ . However, this constraint disappears in the centralized (i.e., non-threshold) case, where we can CCA2-encrypt variable-length messages using a fixed-size public key without using hybrid encryption. To our knowledge,

⁵ We are not aware of any effective attack. Only the proof of IND-CCA2 security in the ROM is affected.

⁶ A common approach to encrypt long messages is to use hybrid encryption. However, it makes it harder to prove properties about encrypted data in zero-knowledge. It also destroys the additive homomorphic properties that we retain when we discard ciphertext components that ensure chosen-ciphertext security. The latter property is useful in the context of voting protocols [5].

this useful property of the Damgård-Jurik cryptosystem was never preserved in the chosen-ciphertext setting (at least in the standard model).

We believe that, even in the random oracle model, properly instantiating Naor-Yung under the DCR assumption is important. For example, it provides a convenient way to encrypt arbitrarily long messages with a fixed-size public key while preserving the possibility of efficiently proving properties (e.g., range membership) about encrypted data, which would be difficult using hybrid encryption. It also provides a “voting-friendly” encryption scheme – in the terminology of [5] – in the sense that the keys/ciphertexts of the threshold CCA2-secure system can be publicly mapped to the keys/ciphertexts of an embedded additively homomorphic encryption scheme.

1.2 Technical Overview

Our range proofs depart from all known standard-model candidates [14,63], which are based on Groth-Sahai proofs [44] and proceed by breaking the committed integers into bits. To our knowledge, this approach either restricts committed integers to be smaller than the group order, or they are inherently stuck with a somewhat wasteful rate $O(1/\lambda)$ caused by bit-by-bit comparisons (as discussed in Supplementary Material D). In the discrete-log setting, the construction of Couteau *et al.* [24] also requires the group order to be sufficiently larger than the maximal magnitude of committed integers.

To avoid this a priori bound on the range of committed values, we leverage a property of the Damgård-Jurik cryptosystem in that the CRS only consists of an RSA modulus $N = pq$. The prover commits to an integer in a range $[0, B]$ by having the prover first choose a sufficiently large $\zeta \geq 1$ such that $B < N^\zeta$ exactly as in the Damgård-Jurik encryption scheme. Following the approach of Kiayias *et al.* [47], we can obtain a constant rate as the ratio between the size of the proof and that of witnesses becomes constant (i.e., about 12) for a large $\zeta \in \text{poly}(\lambda)$. Unlike our main construction, our dual-mode variant requires a CRS that fixes an integer $\zeta \geq 1$ once-and-for-all.

In order to prove security in the standard model, we build on recent progress on instantiations of the Fiat-Shamir paradigm. Canetti *et al.* [16] and Peikert and Shiehian [61] showed that Fiat-Shamir can provide soundness in the standard model under the Learning-With-Errors (LWE) assumption [62], which yields *correlation intractable* (CI) hash functions [17] for efficiently searchable relations.

Correlation intractability for a relation R requires the infeasibility of finding x such that $(x, H_k(x)) \in R$ given a random hashing key k . It guarantees soundness by preventing a cheating prover’s first message a from being hashed into a challenge $\text{Chall} = H_k(a)$ admitting a valid response z . Canetti *et al.* [19] showed that CI hash functions for efficiently searchable relations suffice when Fiat-Shamir is applied to *trapdoor* Σ -protocols. These are Σ -protocols that assume a CRS and where an efficiently computable function `BadChallenge` can identify (on input of a trapdoor τ_Σ , the false statement x and the prover’s first message a) the only challenge Chall such that an accepting transcript (a, Chall, z)

exists for some z . Libert *et al.* [51] (based on earlier observations from [21,54]) showed that the group structure of Paillier allows **BadChallenge** to identify bad challenges within an exponentially large challenge space, thus eliminating the need for parallel repetitions to ensure soundness.

Here, we also achieve soundness without parallel repetitions by exploiting the group structure of $\mathbb{Z}_{N^{\zeta+1}}^*$. However, our **BadChallenge** functions additionally solve integer linear programming instances with a constant number of variables. They also apply the technique of Fouque, Stern and Wackers [36], which decodes Paillier-decrypted values into rational numbers. In our variant of Couteau *et al.*'s range proof [24], the prover first sends DCR-based commitments to integers $\{x_i\}_{i=0}^3$ such that $1 + 4x_0(B - x_0) = \sum_{i=1}^3 x_i^2$ over \mathbb{Z} (recall that, for any positive integer y , there exist $\{x_i \in \mathbb{Z}\}_{i=1}^3$ such that $1 + 4y = \sum_{i=1}^3 x_i^2$, as observed in [41]). Our **BadChallenge** function first computes $\{\tilde{x}_i\}_{i=0}^3$ by decrypting Paillier ciphertexts. Following Fouque *et al.* [36], it then runs Gauss' algorithm to compute pairs $(x_i, c_i) \in [-B^*, B^*] \times [0, C]$ such that $\tilde{x}_i = x_i \cdot c_i^{-1} \bmod N^{\zeta}$ for each i . If no such decomposition exists for a given index $i \in [0, 3]$, the corresponding \tilde{x}_i determines the only bad challenge that can admit a valid response element z_i . We show that this bad challenge is computable by solving an integer linear programming instance $\mathbf{A} \cdot \mathbf{t} \leq \mathbf{b}$ with 3 variables and 8 constraints. By the definition of the language, we know that the solution \mathbf{t} is unique if the statement is false. Moreover, Lenstra's algorithm [48] allows computing it in polynomial time as the number of variables is fixed.

If all decrypted elements $\{\tilde{x}_i\}_{i=0}^3$ can be represented as pairs of integers $(x_i, c_i) \in [-B^*, B^*] \times [0, C]$ such that $\tilde{x}_i = x_i \cdot c_i^{-1} \bmod N^{\zeta}$, our **BadChallenge** function determines if such representations exist for a common denominator $c = c_i$ for each i . If not all \tilde{x}_i have a such a representation with $x_i \in [-B^*, B^*]$, then we know that no response elements $\{z_i\}_{i=0}^3$ will simultaneously satisfy all verification equations for the *same* challenge. In this case, the language definition implies that at most one challenge can satisfy all these verification equations and we can identify this bad challenge by solving an integer linear program with 9 variables. In the last case, the prover's first message commitments decrypt to elements $\{\tilde{x}_i \in \mathbb{Z}_{N^{\zeta}}\}_{i=0}^3$ that all admit a representation $(x'_i, c) \in [-B^*, B^*] \times [0, C]$ such that $\tilde{x}_i = x'_i \cdot c^{-1} \bmod N^{\zeta}$. In this case, if the statement is false, the unique bad challenge is determined by the last verification equation and it is computable by solving a simple modular equation.

Our Paillier-based instantiation of Naor-Yung uses exactly the same Σ -protocol as in [35, Section 4.2]. We prove that its soundness is restored if we introduce a post-processing step in the (distributed) decryption mechanism. Each decryption server computes its partial decryption exactly as in the threshold variant of Damgård-Jurik [29] (as in [35], this is done without interaction among servers). When partial decryptions are combined together, we first compute a Paillier/Damgård-Jurik plaintext $M \in \mathbb{Z}_{N^{\zeta}}$. Using Gauss' algorithm as suggested by Fouque *et al.* [36], we then decode M as a modular ratio $M = x \cdot c^{-1} \bmod N^{\zeta}$ for small-magnitude $x, c \in \mathbb{Z}$ before outputting the rounded rational $\lfloor x/c \rfloor \in \mathbb{Z}$ as a plaintext. We show that this modified decryption algo-

rithm can be safely combined with the Σ -protocol in [35] as it ensures that both Paillier ciphertexts lead to the same plaintext $\lfloor x/c \rfloor \in \mathbb{Z}$. In the case $\zeta = 1$, given two Paillier ciphertexts $\text{ct}_1 = (1 + N_1)^{\text{Msg}} \cdot r_1^{N_1} \bmod N_1^2$ and $\text{ct}_2 = (1 + N_2)^{\text{Msg}} \cdot r_2^{N_2} \bmod N_2^2$, the protocol of [35] guarantees the existence of $\bar{c} \in [0, C]$ and $\bar{m} \in [-R, R]$ such that $\text{ct}_1^{\bar{c}} = (1 + N_1)^{\bar{m}} \cdot w_1^{N_1} \bmod N_1^2$ and $\text{ct}_2^{\bar{c}} = (1 + N_2)^{\bar{m}} \cdot w_2^{N_2} \bmod N_2^2$, for some $w_1 \in \mathbb{Z}_{N_1}^*$, $w_2 \in \mathbb{Z}_{N_2}^*$. While there is no guarantee that $m \cdot \bar{c}^{-1} \bmod N_1$ equals $m \cdot \bar{c}^{-1} \bmod N_2$, we know from [36] that they both decode to the same pair $(m, \bar{c}) \in [-R, R] \times [0, C]$ as long as $2RC < N$ when we run Gauss’ algorithm. This ensures plaintext equality when the decryption algorithm outputs $\lfloor m/\bar{c} \rfloor$.

In order to obtain a trapdoor Σ -protocol, our `BadChallenge` function appeals again to Lenstra’s algorithm and solves an integer linear programming instance with a constant number of variables/constraints. When it comes to proving CCA2-security in the standard model, we need to turn the Σ -protocol into a one-time simulation-sound⁷ NIZK proof system [64]. For this purpose, we could use a construction put forth by Devevey *et al.* [31] but it would unfortunately ruin the length-flexible property of the scheme. If we were to combine it with our trapdoor Σ -protocol showing plaintext equalities, the public key would inherently bound the size of the message space. To avoid this problem, we build a new DCR-based construction that compiles any trapdoor Σ -protocol into a one-time simulation-sound NIZK argument. Unlike the solution of [31, Section 3], simulation-soundness is achieved by augmenting the CRS with a number of bits that does not depend on the underlying trapdoor Σ -protocol.

1.3 Related Work

Range proofs were introduced by Brickell *et al.* [11] and receive continuous attention [22,14,10,53,42,20,25,40] since then. So far, known solutions have been following two main approaches.

The first approach proceeds by breaking integers into bits or small digits [11,3,29,14,42,40,13], which allows communicating a logarithmic (in the range size) number of group elements [14,42,40,13]. This technique is usually implemented using homomorphic commitment schemes over groups of public prime order, while the optimized versions of [14,42,40] require pairings. Within this line of work, Bulletproof [13] obtains the best communication complexity via a clever recursive proof technique and can be realized over standard (i.e., non-pairing-friendly) discrete-logarithm-hard groups. Unfortunately, it is not known to be instantiable in the standard model without interaction.

The second approach [10,53,41,25] relies on integer commitments over groups of hidden order. This approach is often preferred for very large ranges (which arise in applications like anonymous credentials [15], where range elements may be comprised of thousands of bits) where it tends to be more efficient. Also, it does not require the maximal range length to be known ahead of time, when the

⁷ In short, one-time simulation-soundness means that seeing a simulated proof for a false statement of its choice does not help the adversary prove a new false statement.

commitment key is set up. Using homomorphic integer commitments, any range $[\alpha, \beta]$ can be proven by exploiting the homomorphic properties of the commitment scheme and demonstrating that $X - \alpha \in [0, \beta - \alpha]$. Indeed, working over the integers allows showing that $X - \alpha$ and $\beta - X$ are both positive by expressing them as a sum of squares. The idea to rely on square decompositions over the integers dates back to [11]. The square decomposition method was improved by Lipmaa [53] by relying on the Lagrange decomposition of any positive integer as a sum of four squares. Groth [41] observed any positive integer of the form $4Y + 1$, for some $Y \in \mathbb{Z}$, can be more efficiently expressed as a sum of three squares. Further efficiency and security improvements were given in [25]. In this second approach, the underlying integer commitment scheme builds on [37,28] and is usually instantiated using RSA groups. Couteau *et al.* [25] showed that its security relates to a slight variant of the RSA assumption rather than the less standard Strong RSA assumption.

Very recently, Couteau *et al.* [24] managed to reconcile the advantages of both approaches. Their core technique converts any (homomorphic) commitment scheme over groups of (public) prime order into a *bounded* integer commitment scheme. While the conversion does not completely preserve the homomorphic property, it allows committing to bounded-range integers by interpreting them as rounded rationals. It also allows reviving the square decomposition method so as to prove integer relations holding over public ranges. As a result, their range proof consists of a public-coin 3-move interactive protocol that only communicates a constant number of elements. It can be instantiated using standard Pedersen commitments [60] in prime-order groups as long as the group order is large enough to represent the bounded integers. Their technique also applies under lattice assumptions and in class groups. In the latter instantiation, it also inherits the unbounded property of solutions based on hidden-order groups.

We note that a generic transformation due to Ciampi *et al.* [23, Section 4.2] can be used to turn a slight modification (where the first-message group elements are not hashed) of Couteau *et al.*'s discrete-log-based range proof [24] into a trapdoor Σ -protocol, and thus obtain a non-interactive variant in the standard model. However, since the transformation of [23] only applies to Σ -protocols with small challenge space, it has to be repeated $O(\lambda)$ times in parallel to achieve negligible soundness error. In contrast, we achieve soundness without parallel repetitions as in [51]. Moreover, applying [23] to build a non-interactive variant of [24] would still require to fix the maximal cardinality of ranges ahead of time. As it turns out, none of the existing range proofs (even in the bounded case where the CRS depends on $\log(\beta - \alpha)$) in the standard model features proofs comprised of a constant number of element of the base ring/group.

The first non-interactive CCA-secure threshold cryptosystems date back to the work of Shoup and Gennaro [65] who gave DDH-based realizations in the random oracle model. Fouque and Pointcheval [35] gave a generic construction and a DDH-based instantiation using the Naor-Yung paradigm. Until the recent years, all non-interactive solutions in the standard model were pairing-based [8,52]. Boneh *et al.* gave a generic technique [9] to transform any IND-CCA

secure encryption scheme into a non-interactive threshold system using fully homomorphic encryption. Using correlation-intractable hash functions, Devevey *et al.* [31] recently obtained constructions under the DCR and LWE assumptions in the adaptive corruption setting. Back in 1999, Canetti and Goldwasser [18] showed that chosen-ciphertext security was achievable in the standard model by allowing decryption servers to interact with one another. Their approach was subsequently extended to handle adaptive adversaries [46,1].

2 Background

Let S be a finite set. Then, $a \leftarrow U(S)$ means that a is sampled according to the uniform distribution over S . $|a|$ is the bit-length of a .

2.1 Hardness Assumptions

We first recall Paillier’s Composite Residuosity assumption and its variant considered by Damgård and Jurik.

Definition 2.1 ([58,29]). *Let integers $N = pq$ and $\zeta > 1$ for primes p, q . The ζ -Decision Composite Residuosity (ζ -DCR) assumption states that the distributions $\{x = w^{N^\zeta} \bmod N^{\zeta+1} \mid w \leftarrow U(\mathbb{Z}_N^*)\}$ and $\{x \mid x \leftarrow U(\mathbb{Z}_{N^{\zeta+1}}^*)\}$ are computationally indistinguishable.*

Lemma 2.2 (Adapted from [29]). *Let $\zeta = \text{poly}(\lambda)$. Then ζ -DCR is equivalent to 1-DCR with a security loss $\leq \zeta$. (The proof is in Supplementary Material B.)*

We now recall the definition of the Learning-With-Errors (LWE) assumption.

Definition 2.3 ([62]). *Let $m \geq n \geq 1$, $q \geq 2$ and $\alpha \in (0, 1)$ be functions of a security parameter λ . The LWE problem consists in distinguishing between the distributions $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ and $U(\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m)$, where $\mathbf{A} \sim U(\mathbb{Z}_q^{m \times n})$, $\mathbf{s} \sim U(\mathbb{Z}_q^n)$ and $\mathbf{e} \sim D_{\mathbb{Z}^m, \alpha q}$.*

2.2 Correlation Intractable Hash Functions

We consider unique-output efficiently searchable relations [16].

Definition 2.4. *A relation $R \subseteq \mathcal{X} \times \mathcal{Y}$ is **searchable** in time T if there exists a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ which is computable in time T and such that, if there exists y such that $(x, y) \in R$, then $f(x) = y$.*

Let $\lambda \in \mathbb{N}$ a security parameter. A hash family with input length $n(\lambda)$ and output length $m(\lambda)$ is a collection $\mathcal{H} = \{h_\lambda : \{0, 1\}^{s(\lambda)} \times \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)}\}$ of keyed functions induced by efficient algorithms $(\text{Gen}, \text{Hash})$, where $\text{Gen}(1^\lambda)$ outputs a key $k \in \{0, 1\}^{s(\lambda)}$ and $\text{Hash}(k, x)$ computes $h_\lambda(k, x) \in \{0, 1\}^{m(\lambda)}$.

Definition 2.5. For a relation ensemble $\{R_\lambda \subseteq \{0, 1\}^{n(\lambda)} \times \{0, 1\}^{m(\lambda)}\}$, a hash function family $\mathcal{H} = \{h_\lambda : \{0, 1\}^{s(\lambda)} \times \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)}\}$ is **R-correlation intractable** if, for any probabilistic polynomial time (PPT) adversary \mathbb{A} , we have $\Pr [k \leftarrow \text{Gen}(1^\lambda), x \leftarrow \mathcal{A}(k) : (x, h_\lambda(k, x)) \in R] = \text{negl}(\lambda)$.

Peikert and Shiehian [61] described a correlation-intractable hash family for any searchable relation (in the sense of Definition 2.4) defined by functions f of bounded depth. When f is computable by a branching program, their construction relies on the standard SIS assumption with polynomial approximation factors. Under the LWE assumption with polynomial approximation factors, their bootstrapping theorem allows handling arbitrary bounded-depth functions.

2.3 Trapdoor Σ -protocols

Canetti *et al.* [19] considered a definition of Σ -protocols that slightly differs from the usual formulation [26].

Definition 2.6 (Adapted from [19,2]). Let a language $\mathcal{L} = (\mathcal{L}_{\text{zk}}, \mathcal{L}_{\text{sound}})$ associated with two NP relations $\mathcal{R}_{\text{zk}}, \mathcal{R}_{\text{sound}}$. A 3-move interactive proof system $\Pi = (\text{Gen}_{\text{par}}, \text{Gen}_{\mathcal{L}}, \text{P}, \text{V})$ in the common reference string model is a Gap Σ -protocol for \mathcal{L} if it satisfies the following conditions:

- **3-Move Form:** P and V both take as input $\text{crs} = (\text{par}, \text{crs}_{\mathcal{L}})$, with $\text{par} \leftarrow \text{Gen}_{\text{par}}(1^\lambda)$ and $\text{crs}_{\mathcal{L}} \leftarrow \text{Gen}_{\mathcal{L}}(\text{par}, \mathcal{L})$, and a statement x and proceed as follows: (i) P takes in $w \in \mathcal{R}_{\text{zk}}(x)$, computes $(\mathbf{a}, st) \leftarrow \text{P}(\text{crs}, x, w)$ and sends \mathbf{a} to the verifier; (ii) V sends back a random challenge Chall from the challenge space \mathcal{C} ; (iii) P finally sends a response $\mathbf{z} = \text{P}(\text{crs}, x, w, \mathbf{a}, \text{Chall}, st)$ to V ; (iv) On input of $(\mathbf{a}, \text{Chall}, \mathbf{z})$, V outputs 1 or 0.
- **Completeness:** If $(x, w) \in \mathcal{R}_{\text{zk}}$ and P honestly computes (\mathbf{a}, \mathbf{z}) for a challenge Chall , $\text{V}(\text{crs}, x, (\mathbf{a}, \text{Chall}, \mathbf{z}))$ outputs 1 with probability $1 - \text{negl}(\lambda)$.
- **Special zero-knowledge:** There is a PPT simulator ZKSim that inputs crs , $x \in \mathcal{L}_{\text{zk}}$ and a challenge $\text{Chall} \in \mathcal{C}$. It outputs $(\mathbf{a}, \mathbf{z}) \leftarrow \text{ZKSim}(\text{crs}, x, \text{Chall})$ such that $(\mathbf{a}, \text{Chall}, \mathbf{z})$ is computationally indistinguishable from a real transcript with challenge Chall (for $w \in \mathcal{R}_{\text{zk}}(x)$).
- **Special soundness:** For any CRS $\text{crs} = (\text{par}, \text{crs}_{\mathcal{L}})$ obtained as $\text{par} \leftarrow \text{Gen}_{\text{par}}(1^\lambda)$, $\text{crs}_{\mathcal{L}} \leftarrow \text{Gen}_{\mathcal{L}}(\text{par}, \mathcal{L})$, any $x \notin \mathcal{L}_{\text{sound}}$, and any first message \mathbf{a} sent by P , there is at most one challenge $\text{Chall} = f(\text{crs}, x, \mathbf{a})$ for which an accepting transcript $(\text{crs}, x, \mathbf{a}, \text{Chall}, \mathbf{z})$ exists for some third message \mathbf{z} . The function f is called the “bad challenge function” of Π . That is, if $x \notin \mathcal{L}_{\text{sound}}$ and the challenge differs from the bad challenge, the verifier never accepts.

Definition 2.6 is taken from [19] and relaxes the standard special soundness property in that extractability is not required. Instead, it considers a bad challenge function f , which may not be efficiently computable. Canetti *et al.* [19] define trapdoor Σ -protocols as Σ -protocols where the bad challenge function is efficiently computable using a trapdoor. Here, we use a definition where the CRS and the trapdoor may depend on the language.

The common reference string $\text{crs} = (\text{par}, \text{crs}_{\mathcal{L}})$ consists of a fixed part par and a language-dependent part $\text{crs}_{\mathcal{L}}$ which is generated as a function of par and a language parameter $\mathcal{L} = (\mathcal{L}_{\text{zk}}, \mathcal{L}_{\text{sound}})$.

Definition 2.7 (Adapted from [19]). A Σ -protocol $\Pi = (\text{Gen}_{\text{par}}, \text{Gen}_{\mathcal{L}}, \text{P}, \text{V})$ with bad challenge function f for a trapdoor language $\mathcal{L} = (\mathcal{L}_{\text{zk}}, \mathcal{L}_{\text{sound}})$ is a **trapdoor Σ -protocol** if it satisfies the properties of Definition 2.6 and there exist PPT algorithms $(\text{TrapGen}, \text{BadChallenge})$ with the following properties.

- Gen_{par} inputs $\lambda \in \mathbb{N}$ and outputs public parameters $\text{par} \leftarrow \text{Gen}_{\text{par}}(1^\lambda)$.
- $\text{Gen}_{\mathcal{L}}$ is a randomized algorithm that, on input of public parameters par , outputs the language-dependent part $\text{crs}_{\mathcal{L}} \leftarrow \text{Gen}_{\mathcal{L}}(\text{par}, \mathcal{L})$ of $\text{crs} = (\text{par}, \text{crs}_{\mathcal{L}})$.
- $\text{TrapGen}(\text{par}, \mathcal{L}, \tau_{\mathcal{L}})$ takes as input public parameters par and a membership-testing trapdoor $\tau_{\mathcal{L}}$ for the language $\mathcal{L}_{\text{sound}}$. It outputs a common reference string $\text{crs}_{\mathcal{L}}$ and a trapdoor $\tau_{\Sigma} \in \{0, 1\}^{\ell_{\tau}}$, for some $\ell_{\tau}(\lambda)$.
- $\text{BadChallenge}(\tau_{\Sigma}, \text{crs}, x, \mathbf{a})$ takes in a trapdoor τ_{Σ} , a CRS $\text{crs} = (\text{par}, \text{crs}_{\mathcal{L}})$, an instance x , and a first prover message \mathbf{a} . It outputs a challenge Chall .

In addition, the following properties are required.

- **CRS indistinguishability:** For any $\text{par} \leftarrow \text{Gen}_{\text{par}}(1^\lambda)$, and any trapdoor $\tau_{\mathcal{L}}$ for the language \mathcal{L} , an honestly generated $\text{crs}_{\mathcal{L}}$ is computationally indistinguishable from a CRS produced by $\text{TrapGen}(\text{par}, \mathcal{L}, \tau_{\mathcal{L}})$. Namely, for any aux and any PPT distinguisher \mathcal{A} , we have

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{indist-}\Sigma}(\lambda) &:= |\Pr[\text{crs}_{\mathcal{L}} \leftarrow \text{Gen}_{\mathcal{L}}(\text{par}, \mathcal{L}) : \mathcal{A}(\text{par}, \text{crs}_{\mathcal{L}}) = 1] \\ &\quad - \Pr[(\text{crs}_{\mathcal{L}}, \tau_{\Sigma}) \leftarrow \text{TrapGen}(\text{par}, \mathcal{L}, \tau_{\mathcal{L}}) : \mathcal{A}(\text{par}, \text{crs}_{\mathcal{L}}) = 1]| \leq \text{negl}(\lambda). \end{aligned}$$

- **Correctness:** There exists a language-specific trapdoor $\tau_{\mathcal{L}}$ such that, for any instance $x \notin \mathcal{L}_{\text{sound}}$ and all pairs $(\text{crs}_{\mathcal{L}}, \tau_{\Sigma}) \leftarrow \text{TrapGen}(\text{par}, \mathcal{L}, \tau_{\mathcal{L}})$, we have $\text{BadChallenge}(\tau_{\Sigma}, \text{crs}, x, \mathbf{a}) = f(\text{crs}, x, \mathbf{a})$.

Note that the TrapGen algorithm does not take a specific statement x as input, but only a trapdoor $\tau_{\mathcal{L}}$ allowing to recognize elements of $\mathcal{L}_{\text{sound}}$.

2.4 Trapdoor Σ -Protocol Showing Composite Residuosity

We recall a standard Σ -protocol that allows proving that an element of $\mathbb{Z}_{N^{\zeta+1}}^*$ is a N^{ζ} -th residue. In [51], it was shown that the latter protocol is a trapdoor Σ -protocol showing that an element of $\mathbb{Z}_{N^2}^*$ is a composite residue.

Namely, let $\mathcal{L}^{\text{DCR}} := \{x \in \mathbb{Z}_{N^{\zeta+1}}^* \mid \exists w \in \mathbb{Z}_N^* : x = w^{N^{\zeta}} \bmod N^{\zeta+1}\}$, the language of N^{ζ} -th residues, for some integer $\zeta > 1$, where $N = pq$ is an RSA modulus. We assume that the challenge space is $\{0, \dots, 2^\lambda - 1\}$ and that $p, q > 2^{l(\lambda)}$, for some polynomial $l : \mathbb{N} \rightarrow \mathbb{N}$ such that $l(\lambda) > \lambda$ for any sufficiently large $\lambda \in \mathbb{N}$. The condition $p, q > 2^\lambda$ will ensure that the difference between any two challenges be co-prime with N .

In order to obtain a BadChallenge function that identifies bad challenges for elements $x \notin \mathcal{L}^{\text{DCR}}$, [51] uses an observation from Lipmaa [54], which shows that the factorization of N allows computing bad challenges even if $\text{gcd}(x, N) > 1$.

- Gen_{par}**(1^λ) : Given the security parameter λ , define $\mathbf{par} = \{\lambda\}$.
- Gen_L**($\mathbf{par}, \mathcal{L}^{\text{DCR}}$) : Given public parameters \mathbf{par} and the description of a language \mathcal{L}^{DCR} , consisting of an RSA modulus $N = pq$ with primes p and q such that $p, q > 2^{l(\lambda)}$, for some polynomial $l : \mathbb{N} \rightarrow \mathbb{N}$ such that $l(\lambda) > \lambda$, define the language-dependent $\mathbf{crs}_L = \{N\}$. The global CRS is $\mathbf{crs} = (\{\lambda\}, \mathbf{crs}_L)$.
- TrapGen**($\mathbf{par}, \mathcal{L}^{\text{DCR}}, \tau_L$) : Given \mathbf{par} , the description of a language \mathcal{L}^{DCR} that specifies an RSA modulus N and a membership-testing trapdoor $\tau_L = (p, q)$ consisting of the factorization of $N = pq$, output the language-dependent $\mathbf{crs}_L = \{N\}$ which defines $\mathbf{crs} = (\{\lambda\}, \mathbf{crs}_L)$ and the trapdoor $\tau_\Sigma = (p, q)$.
- P**(\mathbf{crs}, x, w) \leftrightarrow **V**(\mathbf{crs}, x) : Given a \mathbf{crs} , a statement $x = w^{N^\zeta} \bmod N^{\zeta+1}$, P (who has the witness $w \in \mathbb{Z}_N^*$) and V interact as follows:
1. P chooses a random $r \leftarrow U(\mathbb{Z}_N^*)$ and sends $a = r^{N^\zeta} \bmod N^{\zeta+1}$ to V .
 2. V sends a random challenge $\mathbf{Chall} \leftarrow U(\{0, \dots, 2^\lambda - 1\})$ to P .
 3. P computes the response $z = r \cdot w^{\mathbf{Chall}} \bmod N$ and sends it to V .
 4. V checks if $a \cdot x^{\mathbf{Chall}} \equiv z^{N^\zeta} \pmod{N^{\zeta+1}}$ and returns 0 otherwise.
- BadChallenge**($\mathbf{par}, \tau_\Sigma, \mathbf{crs}, x, a$) : Given $\tau_\Sigma = (p, q)$, decrypt x and a to obtain $\alpha_x = \mathcal{D}_{\tau_\Sigma}(x) \in \mathbb{Z}_{N^\zeta}$, $\alpha_a = \mathcal{D}_{\tau_\Sigma}(a) \in \mathbb{Z}_{N^\zeta}$.
1. If $\alpha_a = 0$, return $\mathbf{Chall} = 0$.
 2. If $\alpha_a \neq 0$, let $d_x = \gcd(\alpha_x, N^\zeta)$, which lives in the set $\{p^i q^j \mid 0 \leq i < \zeta, 0 \leq j < \zeta\} \cup \{p^i q^\zeta \mid 0 \leq i < \zeta\} \cup \{p^\zeta q^j \mid 0 \leq j < \zeta\}$. Then,
 - a. If $1 < d_x < N^\zeta$, return \perp if d_x does not divide $N^\zeta - \alpha_a$.
 - b. Otherwise, the congruence $\alpha_a + \mathbf{Chall} \cdot \alpha_x \equiv 0 \pmod{\frac{N^\zeta}{d_x}}$ has a unique solution $\mathbf{Chall}' = -\alpha_x^{-1} \cdot \alpha_a \in \mathbb{Z}_{N^\zeta/d_x}$ since $\gcd(\alpha_x, N^\zeta/d_x) = 1$. If $\mathbf{Chall}' \in \mathbb{Z}_{N^\zeta/d_x} \setminus \{0, \dots, 2^\lambda - 1\}$, return \perp . Else, return $\mathbf{Chall} = \mathbf{Chall}'$.

In [51], it is shown that the above construction is a trapdoor Σ -protocol with large challenge space. By applying [61], this implies compact NIZK arguments (i.e., without using parallel repetitions to achieve negligible soundness error) for the language \mathcal{L}^{DCR} assuming that the LWE assumption holds.

Lemma 2.8 ([51]). *The above protocol is a trapdoor Σ -protocol for \mathcal{L}^{DCR} .*

2.5 Encoding and Decoding Bounded Rationals in \mathbb{Z}_N

In [36], Fouque *et al.* suggested a technique that allows computing over rational numbers when they are encrypted using Paillier. The idea is to encode a rational r/s , for co-prime integers $(r, s) \in [-R, R] \times [0, S]$, as the modular ratio $r \cdot s^{-1} \bmod N$. They showed that, as long as, $2RS < N$, it is possible to recover (r, s) from $t = r \cdot s^{-1} \bmod N$ using Gauss' lattice reduction algorithm in dimension 2.

Let an RSA modulus and bounds R, S . Let $r, s \in \mathbb{Z}$ such that $-R \leq r \leq R$, $0 < s \leq S$, $\gcd(r, s) = 1$ and $\gcd(s, N) = 1$. Let the rational $t = r/s \in \mathbb{Q}$

Define the encoding $\mathcal{E}(t) := t' = r \cdot s^{-1} \bmod N$. To decode it and recover $t \in \mathbb{Q}$ from t' , consider the lattice

$$A := \{(x, y) \in \mathbb{Z}^2 : x = y \cdot t' \bmod N\} = \{(x, y) \in \mathbb{Z}^2 : s \cdot x = y \cdot r \bmod N\}.$$

A particular basis of Λ is formed by the vectors $(N, 0)$ and $(t', 1)$. Since s is invertible over \mathbb{Z}_N , the vector $(r, s) \in \mathbb{Z}^2$ also lives in Λ . To recover co-prime integers $(r, s) \in \mathbb{Z}^2$ such that $t' = r \cdot s^{-1} \pmod N$, one can run Gauss' algorithm on input of the initial basis $\vec{u} = (N, 0)$, $\vec{v} = (t', 1)$ to compute a minimal vector of Λ . A result of Vallée [66] ensures that the number of iterations is at most $3 + \log_{1+\sqrt{2}} \max(\|\vec{u}\|, \|\vec{v}\|)$ in the worst case.

Fouque *et al.* proved that the decoding procedure is correct and pointed out that it carries over when computations take place modulo N^ζ for $\zeta > 1$.

Lemma 2.9 ([36, Theorem 1]). *If $t' = r \cdot s^{-1} \pmod N$, $-R \leq r \leq R$, and $0 < s \leq S$, then Gauss' algorithm uniquely recovers r and s if $2RS < N$.*

2.6 Paillier Decryption of (Rounded) Rationals

We first describe a variant of Paillier's cryptosystem used by Fouque, Stern and Wackers [36] to perform homomorphic operations over rational numbers. While the encryption algorithm is identical to that of Paillier/Damgård-Jurik [58,29], the message space is restricted to a specific interval and the decryption algorithm runs Gauss' lattice reduction algorithm in dimension 2. In fact, we modify the decryption algorithm of [36] to make sure that it outputs an integer instead of a rational. In addition, we follow a suggestion of Damgård and Jurik [29] and assume that the message space is *not* a priori bounded by the public key. Instead, it can be flexibly adjusted by the encryption algorithm.

In the following, we let $\ell_M \in \text{poly}(\lambda)$ denote the message length, which can be dynamically determined at encryption time. We also denote by $\text{abs} : \mathbb{Z} \rightarrow \mathbb{N}$ the absolute value function defined as $\text{abs}(x) = x \cdot (x \geq 0) + (-x) \cdot (x < 0)$. Letting $C = 2^\lambda - 1$, the encryptor will fix $R > 2^\lambda \cdot (M + 1)$, where $M = 2^{\ell_M} - 1$ is the largest possible message, and choose ζ in such a way that $2RC < N^\zeta$. After having obtained $\widetilde{\text{Msg}} \in \mathbb{Z}_{N^\zeta}$ from the decryption algorithm of Damgård-Jurik, the receiver will be able to apply Lemma 2.9 so as to decode $\widetilde{\text{Msg}}$ as the ratio $m \cdot c^{-1} \pmod{N^\zeta}$ between bounded rationals $-R \leq m \leq R$ and $0 \leq c \leq C$.

Keygen(1^λ) : Given a security parameter, choose an RSA modulus $N = pq$ such that $p, q > 2^{l(\lambda)}$, for some polynomial $l : \mathbb{N} \rightarrow \mathbb{N}$, and an integer $\zeta \geq 1$. The public key is $\text{pk} = N$ and the secret key is $\text{sk} = (p, q)$.

Encrypt(pk, Msg) : To encrypt $\text{Msg} \in \{0, 1\}^{\ell_M}$, interpret it as a positive integer in $[0, M]$, where $M = 2^{\ell_M} - 1$. Set $\zeta > 1$ as the smallest integer such that $N^\zeta \geq 2^{2\lambda+1}M$. Then, choose $r \leftarrow U(\mathbb{Z}_N^*)$ and compute

$$(\text{ct}, \ell_M) = \left((1 + N)^{\text{Msg}} \cdot r^{N^\zeta} \pmod{N^{\zeta+1}}, \ell_M \right).$$

Decrypt($\text{sk}, (\text{ct}, \ell_M)$) : Given $(\text{ct}, \ell_M) \in \mathbb{Z}_{N^{\zeta+1}}^* \times \mathbb{N}$ and $\text{sk} = (p, q)$. Compute $\widetilde{\text{Msg}} \in \mathbb{Z}_{N^\zeta}$ by running the Damgård-Jurik decryption algorithm, denoted $\mathcal{D}_{\text{sk}}(\text{ct})$. Then, using Gauss' algorithm, find the unique $(m, c) \in \mathbb{Z}^2$ such that $-R \leq m \leq R$, $0 \leq c \leq C$ and $\widetilde{\text{Msg}} = m \cdot c^{-1} \pmod{N^\zeta}$. If no such pair exists, return \perp . Otherwise, return $\text{Msg} = \text{abs}(\lfloor m/c \rfloor)$, where $m/c \in \mathbb{Q}$.

In the decryption algorithm, the absolute value is used to enforce positiveness. The scheme is identical to [36], except that it outputs a positive integer rather than a rational. This decoding method will be applied in our instantiation of Naor-Yung. In our non-interactive range proof of Section 3, we will also use the scheme as a perfectly binding extractable commitment with an extraction algorithm $\text{Decrypt}'$ where $\text{Msg}' = \lfloor m/c \rfloor$, without absolute values.

3 Constant-Rate Unbounded Non-Interactive Range Proofs in the Standard Model

This section presents a range proof where a fixed-size common reference string containing an RSA modulus $N = pq$ allows committing to arbitrarily large integers. We note that, after having committed to an integer, the committer is bound to a specific modulus $N^{\zeta+1}$ and all subsequent proofs related to this commitment are restricted to ranges smaller than a certain bound. Still, the CRS and the underlying algebraic structure do not have to be scaled with the size of the committed integers.

Let positive integers $B, C = 2^\lambda - 1, B^* = 2^\lambda BC$ and $\zeta \geq 1$ satisfying the conditions $2B^*C = 2^{\lambda+1}BC^2 < N^\zeta$. Let $\mathcal{L}_{\text{range}}^{B, B^*, C} = (\mathcal{L}_{\text{zk}}^B, \mathcal{L}_{\text{sound}}^{B, B^*, C})$ be

$$\mathcal{L}_{\text{zk}}^B := \left\{ \text{ct} \in \mathbb{Z}_{N^{\zeta+1}}^* \mid \exists x \in [0, B], w \in \mathbb{Z}_N^* : \text{ct} = (1 + N)^x \cdot w^{N^\zeta} \bmod N^{\zeta+1} \right\}$$

$$\begin{aligned} \mathcal{L}_{\text{sound}}^{B, B^*, C} := \left\{ \text{ct} \in \mathbb{Z}_{N^{\zeta+1}}^* \mid \exists x \in [0, B^*], c \in [1, C], w \in \mathbb{Z}_N^* : \right. \\ \left. \text{ct} = (1 + N)^{x \cdot c^{-1} \bmod N^\zeta} \cdot w^{N^\zeta} \bmod N^{\zeta+1} \wedge \lfloor x/c \rfloor \in [0, B] \right\}. \end{aligned}$$

To prove membership, we will have the prover generate auxiliary commitments $\{C_i\}_{i=1}^3$ and rely on $\bar{\mathcal{L}}_{\text{range}}^{B, B^*, C} = (\bar{\mathcal{L}}_{\text{zk}}^B, \bar{\mathcal{L}}_{\text{sound}}^{B, B^*, C})$ such that

$$\begin{aligned} \bar{\mathcal{L}}_{\text{zk}}^B := \left\{ (\text{ct}, \{C_i\}_{i=1}^3) \in (\mathbb{Z}_{N^{\zeta+1}}^*)^4 \mid \exists x_0, x_1, x_2, x_3 \in [0, B], \right. \\ \exists s_0, s_1, s_2, s_3 \in \mathbb{Z}_N^* : 1 + 4(B - x_0)x_0 = x_1^2 + x_2^2 + x_3^2 \\ \wedge (1 + N)^B \cdot \text{ct}^{-1} = (1 + N)^{x_0} \cdot s_0^{N^\zeta} \bmod N^{\zeta+1} \\ \left. \wedge C_i = (1 + N)^{x_i} \cdot s_i^{N^\zeta} \bmod N^{\zeta+1} \quad \forall i \in [3] \right\}, \end{aligned}$$

$$\begin{aligned} \bar{\mathcal{L}}_{\text{sound}}^{B, B^*, C} := \left\{ (\text{ct}, \{C_i\}_{i=1}^3) \in (\mathbb{Z}_{N^{\zeta+1}}^*)^4 \mid \exists x_0, x_1, x_2, x_3 \in [-B^*, B^*], \right. \\ \exists s_0, s_1, s_2, s_3, \tau \in \mathbb{Z}_N^*, c \in [1, C] : \\ \wedge ((1 + N)^B \cdot \text{ct}^{-1})^c = (1 + N)^{x_0} \cdot s_0^{N^\zeta} \bmod N^{\zeta+1} \\ \wedge C_i^c = (1 + N)^{x_i} \cdot s_i^{N^\zeta} \bmod N^{\zeta+1} \quad \forall i \in [3] \\ \left. \wedge (1 + N)^c = \prod_{i=1}^3 C_i^{x_i} \cdot \text{ct}^{-4x_0} \cdot \tau^{N^\zeta} \bmod N^{\zeta+1} \right\}, \end{aligned} \tag{1}$$

In Lemma 3.2, we show that $(\text{ct}, \{C_i\}_{i=1}^3) \in \bar{\mathcal{L}}_{\text{sound}}^{B, B^*, C}$ implies $\text{ct} \in \mathcal{L}_{\text{sound}}^{B, B^*, C}$, which in turn implies $\text{Decrypt}'(\text{sk}, (\text{ct}, |B^*|)) \in [0, B]$, where $\text{sk} = (p, q)$ and $N = pq$.

Gen_{par}(1^λ): Given the security parameter λ , define $\text{par} = \{\lambda\}$.

Gen_L($\text{par}, \mathcal{L}_{\text{range}}^{B, B^*, C}$): Given public parameters par as well as a description of a language pair $\mathcal{L}_{\text{range}}^{B, B^*, C}$, consisting of an RSA modulus $N = pq$ with primes $p, q > 2^{l(\lambda)}$, for some polynomial $l: \mathbb{N} \rightarrow \mathbb{N}$ such that $l(\lambda) > \lambda$, define the language-dependent CRS $\text{crs}_{\mathcal{L}} = \{N\}$. The global CRS is $\text{crs} = (\{\lambda\}, \text{crs}_{\mathcal{L}})$.

TrapGen($\text{par}, \mathcal{L}_{\text{range}}^{B, B^*, C}, \tau_{\mathcal{L}}$): This algorithm is identical to **Gen_L**($\text{par}, \mathcal{L}_{\text{range}}^{B, B^*, C}$), except that it also outputs the trapdoor $\tau_{\Sigma} = (p, q)$.

P($\text{crs}, \vec{x}, \vec{w}$) \leftrightarrow **V**(crs, \vec{x}): On input of a CRS crs , a statement $\text{ct} \in \mathcal{L}_{\text{zk}}^B$, the prover P (who has $\vec{w} = (x, w) \in [0, B] \times \mathbb{Z}_N^*$) and V interact as follows:

1. P computes $x_1, x_2, x_3 \in [0, B + 1]$ such that $1 + 4x(B - x) = \sum_{i=1}^3 x_i^2$ over \mathbb{Z} . Then, P sets $C_0 = (1 + N)^B \cdot \text{ct}^{-1} \bmod N^{\zeta+1}$, $x_0 = B - x$ and $s_0 = w^{-1} \bmod N$. It randomly picks $s_1, s_2, s_3 \leftarrow U(\mathbb{Z}_{N^\zeta}^*)$ and computes

$$C_i = (1 + N)^{x_i} \cdot s_i^{N^\zeta} \quad \forall i \in [3].$$

Next, to show that $(\text{ct}, \{C_i\}_{i=1}^3) \in \bar{\mathcal{L}}_{\text{zk}}^B$, it chooses $\sigma \leftarrow U(\mathbb{Z}_N^*)$, $r_i \leftarrow U([0, B^*])$ and $\alpha_i \leftarrow U(\mathbb{Z}_{N^\zeta}^*)$ for each $i \in [0, 3]$, to compute

$$R_i = (1 + N)^{r_i} \cdot \alpha_i^{N^\zeta} \bmod N^{\zeta+1} \quad \forall i \in [0, 3]$$

$$R = \sigma^{N^\zeta} \cdot C^{4 \cdot r_0} \cdot \prod_{i=1}^3 C_i^{-r_i} \bmod N^{\zeta+1}.$$

and send $(R, \{R_i\}_{i=0}^3, \{C_i\}_{i=1}^3)$ to V .

2. V sends a random challenge $\text{Chall} \leftarrow U(\{0, \dots, 2^\lambda - 1\})$ to P .
3. P computes the response

$$\tau = \sigma \cdot (s_0^{4 \cdot x_0} \cdot \prod_{i=1}^3 s_i^{x_i})^{\text{Chall}} \bmod N$$

$$z_i = r_i + \text{Chall} \cdot x_i, \quad t_i = \alpha_i \cdot s_i^{\text{Chall}} \bmod N \quad \forall i \in [0, 3]$$

and fails if there exists $i \in [0, 3]$ such that $z_i \notin [0, B^*]$. Otherwise, it sends $(\tau, \{(z_i, t_i)\}_{i=0}^3)$ to V .

4. V sets $C_0 = (1 + N)^B \cdot \text{ct}^{-1} \bmod N^{\zeta+1}$. It accepts iff $z_i \in [0, B^*]$ for each $i \in [0, 3]$ and the following equations hold:

$$R_i \equiv (1 + N)^{z_i} \cdot t_i^{N^\zeta} \cdot C_i^{-\text{Chall}} \pmod{N^{\zeta+1}} \quad \forall i \in [0, 3],$$

$$R \equiv \prod_{i=1}^3 C_i^{-z_i} \cdot \text{ct}^{4 \cdot z_0} \cdot \tau^{N^\zeta} \cdot (1 + N)^{\text{Chall}} \pmod{N^{\zeta+1}}. \quad (2)$$

BadChallenge($\text{par}, \tau_\Sigma, \text{crs}, \mathbf{x}, \mathbf{a}$) : Given the statement $\mathbf{x} = \text{ct} \in \mathbb{Z}_{N^\zeta}$, the message $\mathbf{a} = (R, \{R_i\}_{i=0}^3, \{C_i\}_{i=1}^3)$ and the trapdoor $\tau_\Sigma = (p, q)$, return \perp if $\text{Decrypt}'_{\tau_\Sigma}(\text{ct}) \in [0, B]$. Otherwise, do the following.

1. Let $C_0 = (1 + N)^B \cdot \text{ct}^{-1} \bmod N^{\zeta+1}$. For each index $i \in [0, 3]$, compute $\tilde{x}_i = \mathcal{D}_{\tau_\Sigma}(C_i) \in \mathbb{Z}_{N^\zeta}$ using the Damgård-Jurik decryption algorithm. Also, compute $r = \mathcal{D}_{\tau_\Sigma}(R) \in \mathbb{Z}_{N^\zeta}$ and $r_i = \mathcal{D}_{\tau_\Sigma}(R_i) \in \mathbb{Z}_{N^\zeta}$ for each $i \in [0, 3]$. Then, for each $i \in [0, 3]$, run Gauss' algorithm to compute $x_i \in [-B^*, B^*]$ and $c_i \in [0, C]$ such that $\tilde{x}_i = x_i \cdot c_i^{-1} \bmod N^\zeta$.
2. If there exists $i \in [0, 3]$ such that no pair $(x_i, c_i) \in [-B^*, B^*] \times [0, C]$ satisfies $\tilde{x}_i = x_i \cdot c_i^{-1} \bmod N^\zeta$, let $j \in [0, 3]$ the smallest such index. Compute $(z_j, \text{Chall}_j, k_j) \in \mathbb{Z}^3$ such that

$$\begin{aligned} r_j &= z_j - \tilde{x}_j \cdot \text{Chall}_j + k_j \cdot N^\zeta \\ 0 &\leq z_j \leq B^* \\ 0 &\leq \text{Chall}_j \leq 2^\lambda - 1 \\ 0 &\leq k_j \leq 2^\lambda \end{aligned} \tag{3}$$

This can be achieved by replacing the first equality by inequalities

$$z_j - \tilde{x}_j \cdot \text{Chall}_j + k_j \cdot N^\zeta \leq r_j, \quad -z_j + \tilde{x}_j \cdot \text{Chall}_j - k_j \cdot N^\zeta \leq -r_j$$

and solving an integer linear programming instance with 8 constraints and 3 variables $(z_j, \text{Chall}_j, k_j) \in \mathbb{Z}^3$ using Lenstra's algorithm [48]. If a solution is found (in which case, it is unique), return $\text{Chall} = \text{Chall}_j$.

3. For each $i \in [0, 3]$, let $(x_i, c_i) \in [-B^*, B^*] \times [0, C]$ such that $\{\tilde{x}_i\}_{i=0}^3$ satisfy $\tilde{x}_i = x_i \cdot c_i^{-1} \bmod N^\zeta$. Then, let $c \triangleq \text{lcm}(c_0, c_1, c_2, c_3)$. Check if $c \in [0, C]$ and there exist integers $x'_0, x'_1, x'_2, x'_3 \in [-B^*, B^*]$ such that $\tilde{x}_i = x'_i \cdot c^{-1} \bmod N^\zeta$ for each $i \in [0, 3]$. If no such $\{x'_i\}_{i=0}^3$ and c exist, find the (unique) integer vector $(z_0, z_1, z_2, z_3, \text{Chall}, k_0, k_1, k_2, k_3) \in \mathbb{Z}^9$ such that $0 \leq \text{Chall} \leq 2^\lambda - 1$ and

$$\forall j \in [0, 3] : \begin{cases} r_j = z_j - \tilde{x}_j \cdot \text{Chall} + k_j \cdot N^\zeta \\ 0 \leq z_j \leq B^* \\ 0 \leq k_j \leq 2^\lambda \end{cases}$$

This is done by replacing equalities by pairs of inequalities and solving an integer linear programming instance with 9 variables and 26 constraints. If this vector exists, return the corresponding $\text{Chall} \in [0, 2^\lambda - 1]$.

4. Let $c \in [0, C]$ and $\{x'_i \in [-B^*, B^*]\}_{i=0}^3$ such that $\tilde{x}_i = x'_i \cdot c^{-1} \bmod N^\zeta$. Let $d_x = \text{gcd}(4\tilde{x}\tilde{x}_0 - \sum_{i=1}^3 \tilde{x}_i^2 + 1, N^\zeta)$, where $\tilde{x} = B - \tilde{x}_0 \bmod N^\zeta$, and compute

$$\text{Chall}_0 \triangleq (r + \sum_{i=1}^3 \tilde{x}_i \cdot r_i - 4\tilde{x} \cdot r_0) \cdot (4\tilde{x} \cdot \tilde{x}_0 - \sum_{i=1}^3 \tilde{x}_i^2 + 1)^{-1} \bmod \frac{N^\zeta}{d_x}.$$

If $\text{Chall}_0 \in \{0, \dots, 2^\lambda - 1\}$, return $\text{Chall} = \text{Chall}_0$. Otherwise, return $\text{Chall} = \perp$.

The `BadChallenge` function computes the bad challenge (which is unique when the statement is false) using Lenstra's algorithm [48] that runs in polynomial time since the number of variables is fixed. For an instance with t constraints, each of binary encoding length $O(s)$, the algorithm requires $O(st + s^2)$ arithmetic operations on s -bit numbers.

COMPLETENESS. As long as $z_i \in [0, B^*]$ for all $i \in [0, 3]$ when P computes its response at step 3, i.e., P does not abort, we have

$$\begin{aligned}
& \prod_{i=1}^3 C_i^{-z_i} \cdot \text{ct}^{4 \cdot z_0} \cdot \tau^{N^\zeta} \cdot (1 + N)^{\text{Chall}} \\
&= \prod_{i=1}^3 C_i^{-r_i} \cdot \text{ct}^{4 \cdot r_0} \cdot \left(\prod_{i=1}^3 (1 + N)^{x_i} s_i^{N^\zeta} \right)^{-x_i \text{Chall}} \cdot ((1 + N)^x w^{N^\zeta})^{4x_0 \text{Chall}} \\
&\quad \cdot \sigma^{N^\zeta} \cdot (w^{-4 \cdot x_0} \cdot \prod_{i=1}^3 s_i^{x_i})^{N^\zeta \cdot \text{Chall}} \cdot (1 + N)^{\text{Chall}} \pmod{N^{\zeta+1}} \\
&= \prod_{i=1}^3 C_i^{-r_i} \cdot \text{ct}^{4 \cdot r_0} \cdot (1 + N)^{-\text{Chall} \cdot \sum_{i=1}^3 x_i^2} \cdot (1 + N)^{4 \cdot x_0 \cdot \text{Chall} \cdot x} \\
&\quad \cdot \sigma^{N^\zeta} \cdot (1 + N)^{\text{Chall}} \pmod{N^{\zeta+1}} \\
&= \prod_{i=1}^3 C_i^{-r_i} \cdot \text{ct}^{4 \cdot r_0} \cdot \sigma^{N^\zeta} \pmod{N^{\zeta+1}} = R
\end{aligned}$$

$$(1 + N)^{z_i} \cdot t_i^{N^\zeta} \equiv (1 + N)^{r_i + \text{Chall} \cdot x_i} \cdot \alpha_i^{N^\zeta} \cdot s_i^{\text{Chall} \cdot N^\zeta} \equiv R_i \cdot C_i^{\text{Chall}} \pmod{N^{\zeta+1}},$$

Finally, P only aborts with probability at most $4 \cdot 2^{-\lambda}$.

SPECIAL ZERO-KNOWLEDGE. We first describe a simulator $\text{ZKSim}_B^{\text{range}}$ before showing that a simulated transcript produced by $\text{ZKSim}_B^{\text{range}}(\text{crs}, \vec{x}, \text{Chall})$ is *computationally* indistinguishable from a real transcript generated from a statement-witness pair $(\vec{x}, \vec{w}) \in \mathcal{R}_B^{\text{range}}$ when the challenge is Chall .

Given $\text{crs} = (\{\lambda\}, \text{crs}_{\mathcal{L}})$, an element $\vec{x} = \text{ct} \in \mathbb{Z}_{N^{\zeta+1}}^*$ of the language $\mathcal{L}^{B, B^*, C}$ and a challenge $\text{Chall} \in [0, C]$, $\text{ZKSim}_B^{\text{range}}(\text{crs}, \vec{x}, \text{Chall})$ proceeds as follows: First, it sets $C_0 = (1 + N)^B \cdot \text{ct}^{-1} \pmod{N^{\zeta+1}}$ and randomly picks $s_1, s_2, s_3 \leftarrow U(\mathbb{Z}_{N^\zeta}^*)$ in order to compute an encryption $C_i = s_i^{N^\zeta} \pmod{N^{\zeta+1}}$ of 0 for each $i \in [3]$. Then, the simulator uniformly picks elements of the response \mathbf{z} as $z_i \leftarrow [0, B^*]$, $t_i \leftarrow \mathbb{Z}_N^*$, for all $i \in [0, 3]$, and $\tau \leftarrow \mathbb{Z}_N^*$. Finally, it computes the remaining components $(R, \{R_i\}_{i=0}^3)$ of the first prover message \mathbf{a} in such a way that satisfy the verification equations (2).

We now prove the computational indistinguishability between the transcripts generated by $\text{ZKSim}_B^{\text{range}}$ and real transcripts, which are faithfully computed from $\vec{w} \in \mathcal{R}_B^{\text{range}}(\vec{x})$. We first observe that a simulated transcript $(\mathbf{a}, \text{Chall}, \mathbf{z})$

is computationally indistinguishable from an hybrid transcript where, instead of encrypting 0 in the computation of $\{C_i\}_{i=1}^3$, we encrypt $\{x_i\}_{i=1}^3$ such that $1 + 4x(B - x) = \sum_{i=1}^3 x_i^2$ and $x_0 = B - x$ over \mathbb{Z} , as in the real protocol. In this hybrid transcript, however, we still compute $(R, \{R_i\}_{i=0}^3)$ and \mathbf{z} as in the simulation. A simple reduction shows that the probability to distinguish between simulated transcripts and hybrid transcripts is at most 3 times the advantage of an adversary against the semantic security of Damgård-Jurik (and thus the ζ -DCR assumption). Finally, we show that the distributions of hybrid and real transcripts for $(\vec{x}, \vec{w}) \in \mathcal{R}_B^{\text{range}}$ and the challenge Chall are statistically close (assuming that we use a deterministic algorithm to compute the Lagrange decomposition of $1 + 4x(B - x) \geq 0$) into the sum of 3 squares). This follows from standard arguments. By relying on the generalized Paillier isomorphism we can and split the analysis. Over the “randomness” modulo N , the distributions are the same because each (t_i, α_i) are in one-to-one relation for $i \in [0, 3]$, as well as (τ, σ) . Since the x_i ’s are constant, the distributions “over the plaintext” modulo N^ζ are statistically close because the statistical distance between the z_i -variables is negligible.

More precisely, the ciphertexts $\{C_i\}_{i=0}^3$ have exactly the same distribution in the hybrid and the real transcripts. Now, let $\psi: \mathbb{Z}_{N^\zeta} \times \mathbb{Z}_N^* \mapsto \mathbb{Z}_{N^{\zeta+1}}^*$ denote the generalized Paillier isomorphism. Let also $(r_i, \alpha_i) := \psi^{-1}(R_i)$, for all $i \in [0, 4]$, and $(r, \alpha) := \psi^{-1}(R)$ of an hybrid transcript. We thus have, for all $i \in [0, 3]$,

$$r_i \equiv z_i - \text{Chall} \cdot x_i \pmod{N^\zeta} \quad \alpha_i \equiv t_i \cdot s_i^{-\text{Chall}} \pmod{N},$$

where $x_0 = B - x \pmod{N^\zeta}$ and $s_0 = w^{-1} \pmod{N}$, as well as

$$r \equiv 4z_0(B - x_0) - \sum_{i \in [3]} z_i x_i + \text{Chall} \pmod{N^\zeta},$$

and $\alpha \equiv w^{4z_0} \cdot \prod_{i \in [3]} s_i^{-z_i} \cdot \tau \pmod{N}$. For α and $\{\alpha_i\}_{i \in [0, 3]}$, The congruences in the multiplicative group \mathbb{Z}_N^* show that, given w and $\{(z_i, s_i)\}_{i \in [0, 3]}$, there is a one-to-one relation between α and τ , and between α_i and t_i , for all $i \in [0, 3]$. Then, their distributions are the same as those of the real distributions. (Note that α in the real distribution is also random due to σ .) We are thus left with analyzing the distributions over the additive group \mathbb{Z}_{N^ζ} . For all $i \in [0, 3]$, the congruences on the r_i ensure that, unless $z_i \in [0, CB]$ (which occurs with negligible probability $2^{-\lambda}$), we have $0 \leq r_i = z_i - \text{Chall} \cdot x_i \leq B^*$. That means that, over the integers, we have to show that the statistical distance between $U([0, B^*])$ (which is the distribution of the hybrid z_i) and $\text{Chall} \cdot x_i + U([0, B^*])$ (which is the distribution of the real z) is negligible. Since $x_i \cdot \text{Chall} \leq BC < 2^{-\lambda} B^*$, it is actually bounded by $2^{-\lambda}$. Finally, since $1 + 4x(B - x) = \sum_{i=1}^3 x_i^2$ and $x_0 = B - x$ in both transcripts, we can rewrite the hybrid r as $r = 4r_0(B - x_0) - \sum_{i \in [3]} r_i x_i \pmod{N^\zeta}$, which, given the x_i , is a deterministic function evaluated on independent statistically-closed distributions.

Lemma 3.1 ([24]). *Let integers $n, d \in \mathbb{Z}$, $B \geq 2$ and $x = \lfloor \frac{n}{d} \rfloor$. If there exist $x_1, x_2, x_3 \in \mathbb{Q}$ such that $1 + 4 \cdot \frac{n}{d} \cdot \left(B - \frac{n}{d}\right) = \sum_{i=1}^3 x_i^2$, then we have $x \in [0, B]$.*

Lemma 3.2. *The above construction is a trapdoor Σ -protocol for $\bar{\mathcal{L}}^{B, B^*, C}$ assuming that $2^{2\lambda+3}B^2C^2 < N^\zeta$.*

Proof. We first prove that $(\text{ct}, \{C_i\}_{i=1}^3) \in \bar{\mathcal{L}}^{B, B^*, C}$ ensures that $\text{ct} \in \mathcal{L}^{B, B^*, C}$.

Indeed, letting $\gamma = c^{-1} \bmod N^\zeta$ and $k \in \mathbb{Z}$ such that $\gamma \cdot c + k \cdot N^\zeta = 1$, the first four equations of (1) imply

$$\begin{aligned} C_i &= (1 + N)^{x_i \cdot \gamma} \cdot (s_i^\gamma \cdot C_i^k)^{N^\zeta} \bmod N^{\zeta+1}, & \forall i \in [0, 3] \\ &= (1 + N)^{x_i \cdot (c^{-1} \bmod N^\zeta)} \cdot \tilde{s}_i^{N^\zeta} \bmod N^{\zeta+1} \end{aligned}$$

for some $\tilde{s}_i \in \mathbb{Z}_N^*$, and thus $\text{ct} = (1 + N)^{B - x_0 \cdot (c^{-1} \bmod N^\zeta)} \cdot \tilde{s}_0^{-N^\zeta} \bmod N^{\zeta+1}$. Hence, the ciphertexts $(\text{ct}, \{C_i\}_{i=1}^3)$ decrypt to $\tilde{x} = B - x_0 \cdot c^{-1} \bmod N^\zeta$ and $\{\tilde{x}_i = x_i \cdot c^{-1} \bmod N^\zeta\}_{i=1}^3$. Then, decrypting the last equation of (1) implies

$$c = \sum_{i=1}^3 \left(\frac{x_i}{c} \right) \cdot x_i - 4x_0 \cdot \left(B - \frac{x_0}{c} \right) \bmod N^\zeta.$$

If we multiply both members of the latter equation by c , we obtain

$$c^2 + 4(BC - x_0)x_0 = \sum_{i=1}^3 x_i^2 \bmod N^\zeta. \quad (4)$$

The latter equality holds over \mathbb{Z} since the left-hand-side member is bounded by $C^2 + 4(BC + B^*)B^* = C^2 + 4(BC)^2(1 + 2^\lambda)2^\lambda \leq 2^{\lambda+3}B^2C^2 < N^\zeta$ and the right-hand-side member is bounded by $3B^{*2} = 3 \cdot 2^{2\lambda}B^2C^2 < N^\zeta$. If we divide both members by c^2 over the rationals, we obtain

$$1 + 4\left(B - \frac{x_0}{c}\right) \cdot \frac{x_0}{c} = \sum_{i=1}^3 \left(\frac{x_i}{c} \right)^2 \quad \text{over } \mathbb{Q}.$$

By Lemma 3.1, this in turn implies $\lfloor x_0/\bar{c} \rfloor \in [0, B]$ and thus $B - \lfloor x_0/\bar{c} \rfloor \in [0, B]$.

We now prove that **BadChallenge** output the correct result when the prover sends commitments $\{C_i\}_{i=1}^3$ such that $(\text{ct}, \{C_i\}_{i=1}^3) \notin \bar{\mathcal{L}}^{B, B^*, C}$. For a given first message $\mathbf{a} = (R, \{R_i\}_{i=0}^3, \{C_i\}_{i=1}^3)$ sent by the prover, **BadChallenge** obtains $r, \{r_i\}_{i=0}^3 \in \mathbb{Z}_{N^\zeta}$ and $\{x_i\}_{i=0}^3 \in \mathbb{Z}_{N^\zeta}$ at step 1. It only stops at step 2 if there exists $i \in [0, 3]$ such that C_i decrypts to a value $\tilde{x}_i \in \mathbb{Z}_{N^\zeta}$ which has no representation $\tilde{x}_i = x_i \cdot c^{-1} \bmod N^\zeta$ with $(x_i, c_i) \in [-B^*, B^*] \times [0, C]$. In this case, only one pair $(\text{Chall}_i, z_i) \in [0, C] \times [0, B^*]$ can satisfy the first verification equation of (2). Indeed, if we had distinct such pairs $(\text{Chall}_i, z_i), (\text{Chall}'_i, z'_i) \in [0, C] \times [0, B^*]$ with $\text{Chall}'_i \neq \text{Chall}_i$, we would have $C_i^{\text{Chall}_i - \text{Chall}'_i} = (1 + N)^{z_i - z'_i} \cdot (t_i/t'_i)^{N^\zeta} \bmod N^{\zeta+1}$ and thus $\tilde{x}_i = (z_i - z'_i) \cdot (\text{Chall}_i - \text{Chall}'_i)^{-1} \bmod N^\zeta$. Hence, the unique valid pair $(\text{Chall}_i, z_i) \in [0, C] \times [0, B^*]$ that can satisfy the first equation (2) can be found by applying Gauss' algorithm. Note that **BadChallenge** might output $\text{Chall} \neq \perp$ when

no bad challenge exists at all.⁸ However, **BadChallenge** only needs to find the bad challenge when it exists. When there is no bad challenge, the Fiat-Shamir hash function can output arbitrary values without hurting soundness.

If step 3 is reached, each plaintext in $\{\tilde{x}_i \in \mathbb{Z}_{N^\zeta}^*\}_{i=0}^3$ is decoded as a pair $(x_i, c_i) \in [-B^*, B^*] \times [0, C]$ such that $\tilde{x}_i = x_i \cdot c_i^{-1} \bmod N^\zeta$. We then define $c \triangleq \text{lcm}(c_0, c_1, c_2, c_3)$ and distinguish two cases:

- (a) $c \notin [0, C]$ or $c \in [0, C]$ but there exist no integers $x'_0, x'_1, x'_2, x'_3 \in [-B^*, B^*]$ such that $\tilde{x}_i = x'_i \cdot c^{-1} \bmod N^\zeta$ for each $i \in [0, 3]$.
- (b) $c \in [0, C]$ and there exist integers $x'_0, x'_1, x'_2, x'_3 \in [-B^*, B^*]$ such that we have $\tilde{x}_i = x'_i \cdot c^{-1} \bmod N^\zeta$ for each $i \in [0, 3]$.

In case (a), we observe from the first four verification equations (2) that a valid response $(\tau, \{(z_i, t_i)\}_{i=0}^3)$ can exist for at most one $\text{Chall} \in [0, 2^\lambda - 1]$. This unique challenge value can be determined by solving an integer linear program and finding $(z_0, z_1, z_2, z_3, \text{Chall}, k_0, k_1, k_2, k_3) \in \mathbb{Z}^9$ satisfying (4).

We are left with case (b). In order to satisfy the verification equations (2), the challenge-response pair $(\text{Chall}, (\tau, \{(z_i, t_i)\}_{i=0}^3))$ must satisfy

$$z_i = r_i + \tilde{x}_i \cdot \text{Chall} \bmod N^\zeta \quad r = - \sum_{i=1}^3 \tilde{x}_i \cdot z_i + 4\tilde{x}z_0 + \text{Chall} \bmod N^\zeta.$$

Letting $\tilde{x} = B - \tilde{x}_0 \bmod N$, the above implies

$$\text{Chall} \cdot (4\tilde{x} \cdot \tilde{x}_0 - \sum_{i=1}^3 \tilde{x}_i^2 + 1) = r + \sum_{i=1}^3 \tilde{x}_i \cdot r_i - 4\tilde{x} \cdot r_0 \bmod N^\zeta, \quad (5)$$

Observe that we cannot have $4\tilde{x} \cdot \tilde{x}_0 - \sum_{i=1}^3 \tilde{x}_i^2 + 1 = 0 \bmod N^\zeta$ as this would imply $4\tilde{x} \cdot x'_0 - \sum_{i=1}^3 \tilde{x}_i \cdot x'_i + c = 0 \bmod N^\zeta$, which would mean that

$$(1 + N)^c \cdot \prod_{i=1}^3 C_i^{-x'_i} \cdot \text{ct}^{4x'_0} \bmod N^{\zeta+1}$$

is an N^ζ -th residue in $\mathbb{Z}_{N^{\zeta+1}}^*$. Since we are in case (b), this would contradict the hypothesis $(\text{ct}, \{C_i\}_{i=1}^3) \notin \tilde{\mathcal{L}}^{B, B^*, C}$.

From the inequality $4\tilde{x} \cdot \tilde{x}_0 - \sum_{i=1}^3 \tilde{x}_i^2 + 1 \neq 0 \bmod N^\zeta$, we are guaranteed that $d_x = \text{gcd}(4\tilde{x}\tilde{x}_0 - \sum_{i=1}^3 \tilde{x}_i^2 + 1, N^\zeta) < N^\zeta$ and (5) then yields

$$\text{Chall} \cdot (4\tilde{x} \cdot \tilde{x}_0 - \sum_{i=1}^3 \tilde{x}_i^2 + 1) = r + \sum_{i=1}^3 \tilde{x}_i \cdot r_i - 4\tilde{x} \cdot r_0 \bmod \frac{N^\zeta}{d_x}. \quad (6)$$

Since $\text{gcd}(4\tilde{x} \cdot \tilde{x}_0 - \sum_{i=1}^3 \tilde{x}_i^2 + 1, N^\zeta/d_x) = 1$, equation (6) has a unique solution $\text{Chall}_0 \in \mathbb{Z}_{N^\zeta/d_x}$. Since $N^\zeta/d_x > \min(p, q) > 2^\lambda$, we have $\text{Chall} = \text{Chall}_0 \bmod N^\zeta$.

⁸ This can happen when more than one $\{\tilde{x}_i\}_{i=0}^3$ has no valid representation $(x_i, c_i) \in [-B^*, B^*] \times [0, C]$, in which case they can possibly determine incompatible bad challenges.

N^ζ/d_x for any $\text{Chall} \in \{0, 1, \dots, 2^\lambda - 1\}$, meaning that `BadChallenge` returns the correct result by outputting Chall_0 whenever $\text{Chall}_0 \in \{0, \dots, 2^\lambda - 1\}$. \square

COMPILING THE Σ -PROTOCOL INTO MULTI-THEOREM NIZK. The trapdoor Σ -protocol immediately implies a single-theorem NIZK construction via the Fiat-Shamir transform when we apply the CI hash function of [61]. In order to obtain NIZK proofs in the multi-theorem setting, we could apply the compiler of [50, Appendix B]. One issue is that the latter proceeds by encrypting the Σ -protocol’s first prover message using an equivocable lossy encryption system [4]. Unfortunately, while Paillier can serve as an equivocable lossy encryption scheme (as observed in [45]), we would lose the unbounded property of the range proof if we were to use it. The reason is that the CRS should contain a lossy/injective Paillier public key component that should be longer than messages to be encrypted.

Fortunately, multi-theorem NIZK proofs can be achieved (with computational zero-knowledge and statistical soundness) by adapting the Feige-Lapidot-Shamir compiler using correlation intractable hash functions. The OR trick of [32] builds multi-theorem NIZK proofs by showing OR statements of the form “either the statement is true OR some component of the CRS is in the image of a pseudorandom generator.” Here, we can instantiate their approach using a DCR-based PRG. Recall that the DCR assumption immediately implies a length-doubling PRG that maps a seed $s \in \mathbb{Z}_N^*$ to $y = s^N \bmod N^2$. Here, we can apply the trapdoor Σ -protocol of [51] (which is recalled in Section 2.4) together with the OR Σ -protocols of [26] to prove that “either the range statement is true OR the CRS component $y \in \mathbb{Z}_{N^2}^*$ is an N -th residue.” In the real construction, the CRS contains a uniformly random $y \sim U(\mathbb{Z}_{N^2}^*)$ so as to obtain statistical soundness. In the simulation, y is sampled as a composite residue and its N -th root allows simulating proofs. Using this approach, since the zero-knowledge property is only computational, we can obtain adaptive soundness by hashing the statement together with the prover’s first message when the Fiat-Shamir transform is applied (as observed in [23, Theorem 4]).

In Section 4.2, we will apply a similar instantiation of the FLS paradigm to obtain one-time simulation-soundness in our DCR-based variant of Naor-Yung.

DUAL-MODE RANGE PROOFS/ARGUMENTS. If we give up unboundedness, we can obtain statistically zero-knowledge or even dual-mode range arguments as follows. The CRS initially chooses $\zeta > 1$ and a modulus N such that committed integers always live in a range $[0, B]$ for which $2^{3\lambda+1}B < N^\zeta$. The CRS is augmented with an element $g \in \mathbb{Z}_{N^{\zeta+1}}^*$ that is chosen as an N^ζ -th residue in the zero-knowledge setting (and uniformly over $\mathbb{Z}_{N^{\zeta+1}}^*$ in the soundness setting).

Then, each occurrence of $1 + N$ is replaced by g in the Σ -protocol. The DCR assumption immediately implies the indistinguishability of CRS distributions for the soundness and zero-knowledge settings. Moreover, our simulator $\text{ZKSim}_B^{\text{range}}$ produces statistically indistinguishable transcripts as it computes $\{C_i\}_{i=1}^3$ as dual-mode (or lossy) encryption of 0 instead of random elements modulo $N^{\zeta+1}$.

ACHIEVING CONSTANT RATE. Let $x \in [0, B]$ and $N^{\zeta-1} \leq B \leq N^\zeta$, for some integer ζ , and where only N is fixed by the CRS. We now assess the ratio between

the input size and the proof size assuming that $n := |N|$. We see the witness x as a $|B|$ -bit string since the zero-knowledge property requires a commitment whose message space can contain $[0, B]$. For simplicity we assume that $\zeta = 2\zeta' + 1$ since our proof system requires $2^{2\lambda+3}B^2C^2 < N^\zeta$.

Since the commitment ct to x is a ciphertext over $\mathbb{Z}_{N^{\zeta+1}}^*$, we have

$$\frac{|\text{ct}|}{|B|} \leq \frac{(\zeta + 1)n}{m} \leq \frac{(2\zeta' + 2)n}{(\zeta' - 1)n} = 2 + \frac{4}{(\zeta' - 1)} \downarrow 2.$$

The range proof π for x consists of $\{C_i\}_{i=1}^3, \{R_i\}_{i=0}^3, R$, each of size $(\zeta + 1)n$, and of $\tau, \{(z_i, t_i)\}_{i=0}^3$, where $|\tau| = n$ and $|(z_i, t_i)| = (m + 3\lambda + 1) + n \leq (\zeta + 1)n$, for each $i = 0$ to 3. The total proof size amounts to $12(\zeta + 1)n + n$ and

$$\frac{|\pi|}{|\text{ct}|} \leq \frac{12(\zeta + 1)n + n}{(\zeta + 1)n} = 12 + \frac{1}{2\zeta' + 2} \downarrow 12,$$

leading to a total rate of $|\pi|/|\text{ct}| \leq (24(\zeta' + 1) + 1)/(\zeta' - 1) \leq 73$ for $\zeta' > 1$, which goes down to 24 when ζ grows. If the OR trick is used in the multi-theorem case, it is easy to see that the asymptotic rate remains unchanged as the OR-branch involving the N -th residue only adds a component of size at most $4n$.

4 Instantiating Naor-Yung under the DCR Assumption

In this section, we show that decoding Paillier plaintexts as a rounded rationals provides a secure instantiation of Naor-Yung under the DCR assumption. We first give a trapdoor Σ -protocol showing plaintext equalities before upgrading it into a one-time simulation-sound NIZK argument.

4.1 A Trapdoor Σ -Protocol Showing Plaintext Equalities Between Paillier Ciphertexts for Distinct Moduli

We now give a trapdoor Σ -protocol showing that two ciphertexts decrypt to the same plaintext in the encryption scheme of Section 2.6. Let $N_1 = p_1q_1$ and $N_2 = p_2q_2$ be RSA moduli. Let $C = 2^\lambda - 1$ and let also the languages

$$\begin{aligned} \mathcal{L}_{\text{zk}}^{\text{eq-dcr}} &:= \{(\text{ct}_1, \text{ct}_2, \ell_M) \in \mathbb{Z}_{N_1^\zeta}^* \times \mathbb{Z}_{N_2^\zeta}^* \times \mathbb{N} \mid \exists m \in [0, M], \\ &\quad w_1 \in \mathbb{Z}_{N_1}^*, w_2 \in \mathbb{Z}_{N_2}^* : \text{ct}_1 = (1 + N_1)^m \cdot w_1^{N_1^\zeta} \bmod N_1^{\zeta+1} \\ &\quad \wedge \text{ct}_2 = (1 + N_2)^m \cdot w_2^{N_2^\zeta} \bmod N_2^{\zeta+1}\}, \end{aligned}$$

$$\begin{aligned} \mathcal{L}_{\text{sound}}^{\text{eq-dcr}} &:= \{(\text{ct}_1, \text{ct}_2, \ell_M) \in \mathbb{Z}_{N_1^{\zeta+1}}^* \times \mathbb{Z}_{N_2^{\zeta+1}}^* \times \mathbb{N} \mid \exists m \in [-R, R], \bar{c} \in [0, C], \\ &\quad w_1 \in \mathbb{Z}_{N_1}^*, w_2 \in \mathbb{Z}_{N_2}^* : \text{ct}_1^{\bar{c}} = (1 + N_1)^m \cdot w_1^{N_1^\zeta} \bmod N_1^{\zeta+1} \\ &\quad \wedge \text{ct}_2^{\bar{c}} = (1 + N_2)^m \cdot w_2^{N_2^\zeta} \bmod N_2^{\zeta+1}\}, \end{aligned}$$

where $M = 2^{\ell_M} - 1$ and $\zeta \geq 1$ is the smallest integer such that

$$2RC < 2^{\lambda+1}R < \min(N_1^\zeta, N_2^\zeta)$$

with $R > 2^\lambda(C+1)(M+1)$. Note that $\mathcal{L}_{\text{zk}}^{\text{eq-dcr}} \subset \mathcal{L}_{\text{sound}}^{\text{eq-dcr}}$ since $M < R$.

We note that, for any pair of ciphertexts $((\text{ct}_1, \ell_M), (\text{ct}_2, \ell_M))$ such that $(\text{ct}_1, \text{ct}_2, \ell_M) \in \mathcal{L}_{\text{sound}}^{\text{eq-dcr}}$, the decryption algorithms of Section 2.6 for N_1 and N_2 output the same $\text{Msg} = \text{abs}(\lfloor m/\bar{c} \rfloor)$. Indeed, there exist $u_1, v_1, u_2, v_2 \in \mathbb{Z}$ with $|u_1| < N_1^\zeta$ and $|u_2| < N_2^\zeta$ such that $u_1 \cdot \bar{c} + v_1 \cdot N_1^\zeta = 1$ and $u_2 \cdot \bar{c} + v_2 \cdot N_2^\zeta = 1$, which implies

$$\begin{aligned} \text{ct}_1 &= (1 + N_1)^{u_1 \cdot m} \cdot (w_1^{u_1} \cdot \text{ct}_1^{v_1})^{N_1^\zeta} \bmod N_1^{\zeta+1}, \\ \text{ct}_2 &= (1 + N_2)^{u_2 \cdot m} \cdot (w_2^{u_2} \cdot \text{ct}_2^{v_2})^{N_2^\zeta} \bmod N_2^{\zeta+1}. \end{aligned}$$

Since $u_1 = \bar{c}^{-1} \bmod N_1^\zeta$ and $u_2 = \bar{c}^{-1} \bmod N_2^\zeta$, the decryption algorithm necessarily outputs $\text{Msg} = \lfloor m/\bar{c} \rfloor$ in both cases.

We assume that the challenge space is $\{0, \dots, C\}$, where $C = 2^\lambda - 1$, and that $p, q > 2^{l(\lambda)}$, for some polynomial $l : \mathbb{N} \rightarrow \mathbb{N}$ such that $l(\lambda) > \lambda$ for any sufficiently large $\lambda \in \mathbb{N}$. We now give a trapdoor Σ -protocol proving membership of $\mathcal{L}_{\text{sound}}^{\text{eq-dcr}}$.

Gen_{par}(1^λ): Given the security parameter λ , define $\text{par} = \{\lambda\}$.

Gen_L($\text{par}, \mathcal{L}^{\text{eq-dcr}}$): Given public parameters par and a language description $\mathcal{L}^{\text{eq-dcr}}$, consisting of RSA moduli $N_1 = p_1 q_1$ and $N_2 = p_2 q_2$ with primes $p_1, q_1, p_2, q_2 > 2^{l(\lambda)}$, for some polynomial $l : \mathbb{N} \rightarrow \mathbb{N}$ such that $l(\lambda) > \lambda$, define the language-dependent CRS $\text{crs}_{\mathcal{L}} = \{N_1, N_2\}$. The global CRS is $\text{crs} = (\{\lambda\}, \text{crs}_{\mathcal{L}})$.

TrapGen($\text{par}, \mathcal{L}^{\text{eq-dcr}}, \tau_{\mathcal{L}}$): This algorithm is identical to **Gen_L**($\text{par}, \mathcal{L}^{\text{eq-dcr}}$), except that it also outputs the trapdoor $\tau_{\Sigma} = (p_1, q_1, p_2, q_2)$.

P($\text{crs}, \vec{x}, \vec{w}$) \leftrightarrow **V**(crs, \vec{x}): On input of a common reference string crs , a statement $\vec{x} = (\text{ct}_1, \text{ct}_2, \ell_M) \in \mathbb{Z}_{N_1^{\zeta+1}}^* \times \mathbb{Z}_{N_2^{\zeta+1}}^* \times \mathbb{N}$, the prover P (who has $\vec{w} = (m, w_1, w_2) \in [0, M] \times \mathbb{Z}_{N_1}^* \times \mathbb{Z}_{N_2}^*$) and the verifier V interact as follows:

1. P chooses $a \leftarrow U([0, R])$, $r_1 \leftarrow U(\mathbb{Z}_{N_1}^*)$, $r_2 \leftarrow U(\mathbb{Z}_{N_2}^*)$ and sends

$$A_1 = (1 + N_1)^a \cdot r_1^{N_1^\zeta} \bmod N_1^{\zeta+1}, \quad A_2 = (1 + N_2)^a \cdot r_2^{N_2^\zeta} \bmod N_2^{\zeta+1}.$$

2. V sends back a random challenge $\text{Chall} \leftarrow U(\{0, \dots, 2^\lambda - 1\})$.
3. P aborts if $a + \text{Chall} \cdot m \notin [0, R]$. Otherwise, it sends V the response

$$z = a + \text{Chall} \cdot m, \quad z_1 = r_1 \cdot w_1^{\text{Chall}} \bmod N_1^\zeta, \quad z_2 = r_2 \cdot w_2^{\text{Chall}} \bmod N_2^\zeta$$

4. V checks if $z \in [0, R]$ and accepts iff the following conditions hold:

$$\begin{aligned} A_1 \cdot \text{ct}_1^{\text{Chall}} &\equiv z_1^{N_1^\zeta} \cdot (1 + N_1)^z \pmod{N_1^{\zeta+1}}, \\ A_2 \cdot \text{ct}_2^{\text{Chall}} &\equiv z_2^{N_2^\zeta} \cdot (1 + N_2)^z \pmod{N_2^{\zeta+1}} \end{aligned}$$

BadChallenge(par, τ_Σ , crs, \mathbf{x} , \mathbf{a}) : Given $\mathbf{x} = (\text{ct}_1, \text{ct}_2, \ell_M) \in (\mathbb{Z}_{N_1^{\zeta+1}}^*)^2 \times \mathbb{N}$, the message $\mathbf{a} = (A_1, A_2) \in (\mathbb{Z}_{N_1^{\zeta+1}}^*)^2$ and the trapdoor $\tau_\Sigma = (p_1, q_1, p_2, q_2)$,

1. Using $\text{sk}_1 = (p_1, q_1)$, decrypt ct_1 and A_1 using Paillier's decryption algorithm to obtain $m_1 \in \mathbb{Z}_{N_1^\zeta}$ and $a_1 \in \mathbb{Z}_{N_1^\zeta}$. Likewise, use $\text{sk}_2 = (p_2, q_2)$ to compute $m_2 \in \mathbb{Z}_{N_2^\zeta}$ and $a_2 \in \mathbb{Z}_{N_2^\zeta}$ by decrypting ct_2 and A_2 .
2. Find an integer vector $(z, \text{Chall}, k_1, k_2) \in \mathbb{Z}^4$ satisfying

$$\begin{aligned}
a_1 &= z - m_1 \cdot \text{Chall} + k_1 \cdot N_1^\zeta \\
a_2 &= z - m_2 \cdot \text{Chall} + k_2 \cdot N_2^\zeta, \\
0 &\leq \text{Chall} \leq 2^\lambda - 1 \\
0 &\leq k_1 \leq 2^\lambda \\
0 &\leq k_2 \leq 2^\lambda
\end{aligned} \tag{7}$$

This can be achieved by replacing the equalities by inequality pairs

$$\forall b \in \{1, 2\} : \begin{cases} z - m_b \cdot \text{Chall} + k_b \cdot N_b^\zeta \leq a_b, \\ -z + m_b \cdot \text{Chall} - k_b \cdot N_b^\zeta \leq -a_b \end{cases}$$

and running Lenstra's algorithm [48] to solve an integer linear programming instance with 10 constraints and 4 variables.

If a suitable $(z, \text{Chall}, k_1, k_2) \in \mathbb{Z}^4$ is found (in which case, Chall is uniquely determined), output the corresponding Chall . Otherwise, return \perp .

Again, Lenstra's algorithm [48] allows computing the unique bad challenge (when it exists) in polynomial time since the number of variables is fixed.

Lemma 4.1. *The construction is a trapdoor Σ -protocol for $(\mathcal{L}_{\text{zk}}^{\text{eq-dcr}}, \mathcal{L}_{\text{sound}}^{\text{eq-dcr}})$.*

Proof. We first show the completeness and special zero-knowledge properties.

COMPLETENESS. Given $\vec{w} \in \mathcal{R}_{\text{zk}}^{\text{eq-dcr}}(\vec{x})$, P computes (\mathbf{a}, \mathbf{z}) for a challenge Chall such that $\text{V}(\text{crs}, \vec{x}, (\mathbf{a}, \text{Chall}, \mathbf{z})) = 1$ as long as P does not abort at step 3 of the interactive protocol. Therefore, an honest run of the protocol always leads to a valid transcript except if $a + \text{Chall} \cdot m \notin [0, R]$ which occurs with probability at most $2^{-\lambda}$ since $\text{Chall} \cdot m \leq CM < 2^{\lambda+\ell_M}$ and $R > 2^{2\lambda+\ell_M}$.

SPECIAL ZERO-KNOWLEDGE. The simulator ZKSim proceeds in a standard way. It that inputs $\text{crs} = (\{\lambda\}, \text{crs}_{\mathcal{L}})$, a statement $\vec{x} = (\text{ct}_1, \text{ct}_2, \ell_M) \in \mathcal{L}_{\text{zk}}^{\text{eq-dcr}}$ and a challenge $\text{Chall} \in \{0, \dots, 2^\lambda - 1\}$. First, the simulator $\text{ZKSim}(\text{crs}, \vec{x}, \text{Chall})$ picks $z \leftarrow U([0, R])$ as well as $z_1 \leftarrow U(\mathbb{Z}_{N_1}^*)$ and $z_2 \leftarrow U(\mathbb{Z}_{N_2}^*)$. Then, it computes $A_1 = z_1^{N_1^\zeta} \cdot (1 + N_1)^z \cdot \text{ct}_1^{-\text{Chall}} \bmod N_1^{\zeta+1}$, as well as

$$A_2 = z_2^{N_2^\zeta} \cdot (1 + N_2)^z \cdot \text{ct}_2^{-\text{Chall}} \bmod N_2^{\zeta+1},$$

and outputs (\mathbf{a}, \mathbf{z}) , where $\mathbf{a} = (A_1, A_2)$ and $\mathbf{z} = (z, z_1, z_2)$. We turn to showing that $(\mathbf{a}, \text{Chall}, \mathbf{z})$ is statistically indistinguishable from a real transcript computed using the witness $\vec{w} = (m, w_1, w_2) \in [0, M] \times \mathbb{Z}_{N_1}^* \times \mathbb{Z}_{N_2}^*$ (i.e., $\vec{w} \in \mathcal{R}_{zk}^{\text{eq-dcr}}(\vec{x})$) and with challenge Chall . For each $i \in \{1, 2\}$, let $\psi_i: \mathbb{Z}_{N_i}^\zeta \times \mathbb{Z}_{N_i}^* \mapsto \mathbb{Z}_{N_i}^{\zeta+1}$ denote the generalized Paillier isomorphism. By applying $\{\psi_i^{-1}\}_{i=1}^2$ to compute $(a_1, r_1) := \psi_1^{-1}(A_1)$ and $(a_2, r_2) := \psi_2^{-1}(A_2)$ for a simulated transcript $((A_1, A_2), \text{Chall}, (z, z_1, z_2))$, we find

$$\begin{aligned} a_1 &\equiv z - \text{Chall} \cdot m \pmod{N_1^\zeta} & r_1 &\equiv z_1 \cdot w_1^{-\text{Chall}} \pmod{N_1}, \\ a_2 &\equiv z - \text{Chall} \cdot m \pmod{N_2^\zeta} & r_2 &\equiv z_2 \cdot w_2^{-\text{Chall}} \pmod{N_2}. \end{aligned}$$

The congruences on the left ensure that, unless $z \in [0, CM]$ (which occurs with negligible probability $2^{-\lambda}$), we have $0 \leq a_1 = z - \text{Chall} \cdot m = a_2 \leq R$. Given Chall , the distributions of $\{(z_i, r_i)\}_{i=1}^2$ over the multiplicative rings are exactly the same between the real and the simulated transcripts. Finally, we show that, over the integers, the statistical distance between $U([0, R])$ (which is the distribution of the simulated z) and $\text{Chall} \cdot m + U([0, R])$ (in the real z) is negligible. Since $m \cdot \text{Chall} \leq MC < 2^{\lambda+\ell_M} < 2^{-\lambda}R$, it is actually bounded by $2^{-\lambda}$.

SPECIAL SOUNDNESS. Let us assume two transcripts $((A_1, A_2), \text{Chall}, (z, z_1, z_2))$ and $((A_1, A_2), \text{Chall}', (z', z'_1, z'_2))$ that both satisfy the verification equations with $z, z' \in [0, R]$ and $\text{Chall} \neq \text{Chall}'$ for a given first message (A_1, A_2) sent by the prover. We assume w.l.o.g. that $0 \leq \text{Chall}' < \text{Chall} \leq 2^\lambda - 1$. This implies that $\bar{c} = \text{Chall} - \text{Chall}' \in [0, 2^\lambda - 1]$ and $\bar{z} = z - z' \in [-R, R]$ satisfy the congruences $\text{ct}_1^{\bar{c}} \equiv (z_1/z'_1)^{N_1^\zeta} (1 + N_1)^{\bar{z}} \pmod{N_1^{\zeta+1}}$ and

$$\text{ct}_2^{\bar{c}} \equiv (z_2/z'_2)^{N_2^\zeta} (1 + N_2)^{\bar{z}} \pmod{N_2^{\zeta+1}},$$

which implies $(\text{ct}_1, \text{ct}_2) \in \mathcal{L}_{\text{sound}}^{\text{eq-dcr}}$. This shows that, for any first message (A_1, A_2) sent by the prover, only one bad challenge can exist if $(\text{ct}_1, \text{ct}_2) \notin \mathcal{L}_{\text{sound}}^{\text{eq-dcr}}$.

CRS INDISTINGUISHABILITY. The distribution of the CRS output by TrapGen is exactly the same as the distribution of the CRS output by $\text{Gen}_{\mathcal{L}}$.

BADCHALLENGE CORRECTNESS. Let a false statement $\vec{x} \notin \mathcal{L}_{\text{sound}}^{\text{eq-dcr}}$. Special soundness ensures the existence of at most one bad challenge for any given \mathbf{a} . Lenstra's algorithm can efficiently determine if the bad challenge exists since it can solve the integer feasibility problem in polynomial time when the number of variables is fixed. Moreover, whenever an admissible integer solution $(z, \text{Chall}, k_1, k_2) \in \mathbb{Z}^4$ exists (in which case it is unique), it is efficiently computable from the decrypted values (m_1, m_1, a_1, a_2) . \square

4.2 New Construction of One-Time Simulation-Sound NIZK Arguments from Trapdoor Σ -Protocols

In this section, we aim at one-time simulation soundness without imposing a bound on the plaintext space in the centralized version our scheme of Section

4.3. To this end, we cannot use the constructions of [51,31] because they follow an idea from [27] and encrypt the prover’s first message using a DCR-based lossy encryption scheme [4]. Unfortunately, the latter’s public key should be larger than the first prover message in the underlying trapdoor Σ -protocol.

We describe a new one-time simulation-sound argument which departs from [50,51,31] in that it does not proceed by encrypting the first prover message of the trapdoor Σ -protocol. Instead, it uses an OR proof [26] inspired by the FLS technique [32]. In order to achieve one-time simulation-soundness, we introduce a twist and program the CRS $(u, v) \in (\mathbb{Z}_{N^2}^*)^2$ in such a way that $u^{\text{VK}} \cdot v$ is a composite residue for exactly one VK .

- A trapdoor Σ -protocol $\Pi^{(1)} = (\text{Gen}_{\text{par}}^{(1)}, \text{Gen}_{\mathcal{L}}^{(1)}, \text{P}^{(1)}, \text{V}^{(1)})$ for an NP language \mathcal{L} . This protocol should satisfy the properties of Definition 2.7. We assume that $\Pi^{(1)}$ has challenge space $\mathcal{C} = \{0, 1\}^\lambda$, where λ is the security parameter. In addition, the function $\text{BadChallenge}^{(1)}$ should be computable within time $T_1 \in \text{poly}(\lambda)$ for any input $(\tau, \text{crs}^{(1)}, x, a_1)$.
- A strongly unforgeable one-time signature scheme $\text{OTS} = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ with verification keys in $\{0, 1\}^L$, where $L \in \text{poly}(\lambda)$.
- An RSA modulus $N = pq$, for large primes $p, q > 2^L$.
- A trapdoor Σ -protocol $\Pi^{(0)} = (\text{Gen}_{\text{par}}^{(0)}, \text{Gen}_{\mathcal{L}}^{(0)}, \text{P}^{(0)}, \text{V}^{(0)})$ for the language $\mathcal{L}^{\text{DCR}} := \{x \in \mathbb{Z}_{N^2}^* \mid \exists w \in \mathbb{Z}_N^* : x = w^N \bmod N^2\}$. We assume that the function $\text{BadChallenge}^{(0)}$ is computable within time $T_0 \in \text{poly}(\lambda)$ for any input $(\tau, \text{crs}^{(0)}, x, a_0)$. This protocol can be instantiated as in Section 2.4
- A correlation intractable hash family $\mathcal{H} = (\text{Gen}, \text{Hash})$ for the class \mathcal{R}_{CI} of relations that are efficiently searchable within time T .

Gen_{par}(1^λ): Run $\text{par} \leftarrow \text{Gen}_{\text{par}}^{(1)}(1^\lambda)$ and output par .

Gen _{\mathcal{L}} (par, \mathcal{L}): Given public parameters par and a language \mathcal{L} , the CRS is generated as follows.

1. Generate a CRS $\text{crs}_{\mathcal{L}}^{(1)} \leftarrow \text{Gen}_{\mathcal{L}}^{(1)}(\text{par}, \mathcal{L})$ for the trapdoor Σ -protocol $\Pi^{(1)}$.
2. Choose the description of a one-time signature scheme $\text{OTS} = (\mathcal{G}, \mathcal{S}, \mathcal{V})$ with verification keys in $\{0, 1\}^L$, where $L \in \text{poly}(\lambda)$.
3. Choose an RSA modulus $N = pq$, for primes $p, q > 2^L$, where $L \in \text{poly}(\lambda)$ is the verification key length of OTS . Then, choose $u_0, v_0 \leftarrow \mathbb{Z}_N^*$ and compute $u = u_0^N \bmod N^2$, $v = v_0^N \bmod N^2$.
4. Generate a CRS $\text{crs}^{(0)} \leftarrow \text{Gen}_{\mathcal{L}}^{(0)}(\text{par}, \mathcal{L}^{\text{DCR}})$ for $\Pi^{(0)}$, where \mathcal{L}^{DCR} is associated with $N = pq$.
5. Generate a key $k \leftarrow \text{Gen}(1^\lambda)$ for a correlation intractable hash function with output length λ .

Output the language-dependent CRS $\text{crs}_{\mathcal{L}} := (N, u, v, \text{crs}^{(0)}, \text{crs}_{\mathcal{L}}^{(1)}, k)$ and the simulation trapdoor $\tau_{\text{zk}} := (u_0, v_0)$. The global common reference string consists of $\text{crs} = (\text{par}, \text{crs}_{\mathcal{L}}, \text{OTS})$.

P($\text{crs}, x, w, \text{lbl}$): To prove a statement $x \in \mathcal{L}$ for a label $\text{lbl} \in \{0, 1\}^*$ using the witness w , generate a one-time signature key pair $(\text{VK}, \text{SK}) \leftarrow \mathcal{G}(1^\lambda)$. Then,

1. Compute $(a_1, st) \leftarrow \mathbf{P}^{(1)}(\text{crs}_{\mathcal{L}}^{(1)}, x, w)$. Then, generate a simulated proof $(a_0, \text{Chall}_0, z_0) \in \mathbb{Z}_{N^2}^* \times \{0, 1\}^\lambda \times \mathbb{Z}_N^*$ that $(u^{\text{VK}} \cdot v) \in \mathcal{L}^{\text{DCR}}$. Namely, choose random elements $z_0 \leftarrow U(\mathbb{Z}_N^*)$, $\text{Chall}_0 \leftarrow U(\{0, 1\}^\lambda)$ and compute $a_0 = z_0^N \cdot (u^{\text{VK}} \cdot v)^{-\text{Chall}_0} \bmod N^2$.
2. Compute $\text{Chall} = \text{Hash}(k, (x, a, \text{VK})) \in \{0, 1\}^\lambda$, where $a = (a_0, a_1)$, and set $\text{Chall}_1 = \text{Chall} \oplus \text{Chall}_0$.
3. Compute $z_1 = \mathbf{P}^{(1)}(\text{crs}_{\mathcal{L}}^{(1)}, x, w, a_1, \text{Chall}_1, st)$ by executing the prover of $\Pi^{(1)}$. Define $z = (z_0, z_1, \text{Chall}_0, \text{Chall}_1)$.
4. Generate $sig \leftarrow \mathcal{S}(\text{SK}, (x, a, z, \text{lbl}))$ and output $\vec{\pi} = (\text{VK}, (a, z), sig)$.

$\mathbf{V}(\text{crs}, x, \vec{\pi}, \text{lbl})$: Given a statement x , a label lbl as well as a purported proof $\vec{\pi} = (\text{VK}, (a, z), sig)$, return 0 if $\mathcal{V}(\text{VK}, (x, a, z, \text{lbl}), sig) = 0$. Otherwise,

1. Write $z = (z_0, z_1, \text{Chall}_0, \text{Chall}_1)$ and return 0 if any of these does not parse properly or if $\text{Hash}(k, (x, a, \text{VK})) \neq \text{Chall}_0 \oplus \text{Chall}_1$.
2. If $\mathbf{V}^{(1)}(\text{crs}_{\mathcal{L}}^{(1)}, x, a_1, \text{Chall}_1, z_1) = 1$ and $a_0 \cdot (u^{\text{VK}} \cdot v)^{\text{Chall}_0} = z_0^N \bmod N^2$, return 1. Otherwise, return 0.

Theorem 4.2. *The above construction is a one-time simulation-sound NIZK argument if: (i) OTS is a strongly unforgeable one-time signature; (ii) The DCR assumption holds; (iii) The hash function is correlation-intractable for efficiently searchable relations. (The proof is given in Supplementary Material C.1.)*

4.3 A DCR-Based CCA2-Secure Threshold Cryptosystem from the Naor-Yung Paradigm

The syntax and the security definitions of threshold PKE schemes is recalled in Supplementary Material A.2. Using the tools of Section 4.1 and Section 4.2, we obtain the following variant of the threshold encryption scheme in [35].

We assume that the key generation step chooses a value ζ' that determines a maximal length of encrypted messages (note that this is only necessary in the threshold setting and not in the centralized version of the scheme). However, the encryptor can still choose $\zeta \leq \zeta'$ in a flexible way at encryption time.

For simplicity, we first describe the non-robust version of the scheme, where decryption servers do not provide a proof that partial decryptions are correctly generated. However, robustness can be achieved in a modular way as in [31].

Keygen $(1^\lambda, 1^B, 1^t, 1^n)$: On input of a security parameter λ , a maximal bitlength $B \in \text{poly}(\lambda)$ of encrypted messages, a number of servers $n \in \text{poly}(\lambda)$, and a threshold $t \in \text{poly}(\lambda)$, conduct the following steps.

1. Generate two safe prime products $N_1 = p_1 q_1$ and $N_2 = p_2 q_2$ such that $p_i, q_i > 2^{l(\lambda)}$, for some polynomial $l : \mathbb{N} \rightarrow \mathbb{N}$, and primes $p_i = 2p'_i + 1$, $q_i = 2q'_i + 1$ for which p'_i, q'_i are also prime for each $i \in \{1, 2\}$.
2. Choose an integer $\zeta' > 0$ such that $2^{B+2\lambda+1} < \min(N_1^{\zeta'}, N_2^{\zeta'})$.
3. Choose an integer d such that $d = 1 \bmod N_1^{\zeta'}$ and $d = 0 \bmod \lambda(N_1)$.

4. Choose a random polynomial $f[X] = \sum_{i=0}^{t-1} a_i X^i \in \mathbb{Z}_{N_1^{\zeta'} p_1' q_1'}[X]$ such that $a_0 = d \bmod N_1^{\zeta'} p_1' q_1'$.
5. Generate the CRS $\text{crs}_{\mathcal{L}} := (N, u, v, \text{crs}^{(0)}, \text{crs}_{\mathcal{L}}^{(1)}, k)$ of a simulation-sound NIZK argument for the language $(\mathcal{L}_{\text{zk}}^{\text{eq-dcr}}, \mathcal{L}_{\text{sound}}^{\text{eq-dcr}})$ of Section 4.1, which is induced by the moduli N_1 and N_2 .

The public key is $\text{pk} = (N_1, N_2, \text{crs}_{\mathcal{L}})$ and the secret key shares $\{\text{sk}_i\}_{i=1}^n$ are defined as $\text{sk}_i = f(i) \bmod N_1^{\zeta'} p_1' q_1'$ for each $i \in [n]$.

Encrypt(pk, Msg): To encrypt $\text{Msg} \in \{0, 1\}^{\ell_M}$, return \perp if $\ell_M > B$. Otherwise, interpret Msg as a positive integer in $[0, M]$, where $M = 2^{\ell_M} - 1$. Set $\zeta > 1$ as the smallest integer such that $\min(N_1^{\zeta}, N_2^{\zeta}) \geq 2^{2\lambda+1}M$. Then, choose $(r_1, r_2) \leftarrow U(\mathbb{Z}_{N_1}^* \times \mathbb{Z}_{N_2}^*)$ and compute

$$\text{ct}_1 = (1 + N_1)^{\text{Msg}} \cdot r_1^{N_1^{\zeta}} \bmod N_1^{\zeta+1}, \quad \text{ct}_2 = (1 + N_2)^{\text{Msg}} \cdot r_2^{N_2^{\zeta}} \bmod N_2^{\zeta+1}.$$

Then, using the empty label $\text{lbl} = \varepsilon$, generate a simulation-sound NIZK argument $\vec{\pi} \leftarrow \text{P}(\text{crs}, (\text{ct}_1, \text{ct}_2, \ell_M), (\text{Msg}, r_1, r_2), \text{lbl})$ that $(\text{ct}_1, \text{ct}_2, \ell_M) \in \mathcal{L}_{\text{zk}}^{\text{eq-dcr}}$. Finally, output $\text{ct} = (\text{ct}_1, \text{ct}_2, \ell_M, \vec{\pi})$.

PartDec(sk_i, ct): Given a ciphertext $\text{ct} = (\text{ct}_1, \text{ct}_2, \ell_M, \vec{\pi})$ and $\text{sk}_i \in \mathbb{Z}_{N_1^{\zeta'} p_1' q_1'}$, the i -th server proceeds as follows.

1. If $\text{V}(\text{crs}, (\text{ct}_1, \text{ct}_2, \ell_M), \vec{\pi}, \text{lbl}) = 0$, return \perp .
2. Compute $\mu_i = \text{ct}_1^{2\Delta \cdot \text{sk}_i} \bmod N_1^{\zeta+1}$, where $\Delta = n!$, and return (i, μ_i) .

Combine($\text{pk}, \mathcal{S}, \{\mu_i\}_{i \in \mathcal{S}}, \text{ct}$): Let $R = 2^\lambda \cdot (M + 1)$ and $C = 2^\lambda - 1$. If \mathcal{S} contains less than t shares in $\mathbb{Z}_{N_1^{\zeta+1}}^*$, return \perp . Otherwise, do the following.

1. Define scaled Lagrange coefficients $\lambda_{0,i}^{\mathcal{S}} = \Delta \cdot \prod_{i' \in \mathcal{S} \setminus \{i\}} \frac{-i}{i-i'}$ in \mathbb{Z} for each $i \in \mathcal{S}$ and compute $\mu_0 = \prod_{i \in \mathcal{S}} \mu_i^{2 \cdot \lambda_{0,i}^{\mathcal{S}}} \bmod N_1^{\zeta+1}$, which should be of the form $\mu_0 = \text{ct}_1^{4\Delta^2 f(0)} = \text{ct}_1^{4\Delta^2 d} \bmod N_1^{\zeta+1}$.
2. Compute $\tilde{\mu} = L(\mu_0, N_1^{\zeta}) \cdot 4^{-1} \cdot (\Delta)^{-2} \bmod N_1^{\zeta}$, where $L(\cdot, N_1^{\zeta})$ extracts the discrete logarithm w.r.t. base $(1 + N_1)$ of the elements modulo $N_1^{\zeta+1}$ that are congruent to 1 modulo N_1 as in [29]. Then, using Gauss' algorithm, find the unique $(m, c) \in \mathbb{Z}^2$ such that $-R \leq m \leq R$, $0 \leq c \leq C$ and $\tilde{\mu} = m \cdot c^{-1} \bmod N_1^{\zeta}$. If no such pair exists, return \perp . Otherwise, return $\text{Msg} = \text{abs}(\lfloor m/c \rfloor)$, where the division is computed over \mathbb{Q} .

Theorem 4.3. *The scheme provides IND-CCA security under static corruptions if: (i) The DCR assumption holds; (ii) Π^{OTSS} is a one-time simulation-sound argument. (The proof is in Supplementary Material C.2.)*

COMPARISONS. Devevey *et al.* [31, Section 4] gave a non-interactive threshold CCA2-secure scheme based on DCR and LWE in the standard model. While they can prove security under adaptive corruptions, our scheme provides several

advantages over [31] although we only prove static security.⁹ In the robust version of the scheme, if we do not consider commitments to the secret key shares as being part of the public key (which is reasonable as the encryptor does not need them), the public key size grows with $|N|$ instead of $|N^\zeta|$. Also, the scheme of [31] requires larger secret key shares, which grow super-linearly with the number of servers. Finally, our scheme allows the sender to adjust the block length by choosing ζ according to the message length.

We note that a non-interactive threshold cryptosystem can also be obtained from the DCR/LWE assumption by applying the generic construction of Boneh *et al.* [9]. However, their use of a threshold FHE scheme requires a stronger LWE assumption than our approach¹⁰ as our use of correlation-intractable hash functions just incurs an LWE assumption with polynomial approximation factor as in [61]. Moreover, our scheme features constant-size secret key shares whereas the generic compiler of [9] has shares that grow with the number n of servers. In particular, we obtain a much more efficient voting-friendly system [5, Section 3.1] when we extract an additively homomorphic¹¹ scheme by throwing away the CCA-related ciphertext components (see Supplementary Material C.3).

In the random oracle model, our security proof carries over when we replace the CI hash function by a random oracle. In this case, the scheme remains as efficient as the original Fouque-Pointcheval construction [35]. In addition, it allows encrypting long messages with a $O(|N|)$ -size public key without affecting the voting-friendly property of the scheme nor the ability to efficiently prove certain properties (e.g., range membership) about the plaintext.

Acknowledgements

Part of this research was funded by the French ANR ALAMBIC project (ANR-16-CE39-0006). This work is also partially supported by Indo French Center for the Promotion of Advanced Research (IFCPAR, project number: 6002-1). Thomas Peters is a research associate of the Belgian Fund for Scientific Research (F.R.S.-FNRS).

References

1. M. Abe and S. Fehr. Adaptively secure Feldman VSS and applications to universally-composable threshold cryptography. In *Crypto*, 2004.

⁹ Adaptive security is non-trivial to achieve when $t, n \in \text{poly}(\lambda)$. In many applications like e-voting, one can expect the number of servers to be small (e.g., logarithmic in λ), in which case adaptive security can be achieved via complexity leveraging.

¹⁰ The reason is that, in [9], each partial threshold FHE decryption should be noisy (so as to not leak secret key shares), which incurs an LWE assumption with super-polynomial modulus.

¹¹ Like the commitment of [24], the scheme is homomorphic as long as ciphertexts are honestly generated and computations take place within some bounds.

2. G. Asharov, A. Jain, and D. Wichs. Multipart computation with low communication, computation and interaction via threshold FHE. Cryptology ePrint Archive: Report 2011/613, 2012.
3. M. Bellare and S. Goldwasser. Verifiable partial key escrow. In *ACM-CCS*, 1997.
4. M. Bellare, D. Hofheinz, and S. Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In *Eurocrypt*, 2009.
5. D. Bernhard, V. Cortier, O. Pereira, B. Smyth, and B. Warinschi. Adapting helios for provable ballot privacy. In *ESORICS*, 2011.
6. M. Blum, A. De Santis, S. Micali, and G. Persiano. Non-interactive zero-knowledge. *SIAM J. of Computing*, 1991.
7. M. Blum, P. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications. In *STOC*, 1988.
8. D. Boneh, X. Boyen, and S. Halevi. Chosen ciphertext secure public key threshold encryption without random oracles. In *CT-RSA*, 2006.
9. D. Boneh, R. Gennaro, S. Goldfeder, A. Jain, S. Kim, P. Rasmussen, and A. Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. In *Crypto*, 2018.
10. F. Boudot. Efficient proofs that a committed number lies in an interval. In *Eurocrypt*, 2000.
11. E. Brickell, D. Chaum, I. Damgård, and J. van de Graaf. Gradual and verifiable release of a secret. In *Crypto*. Springer, 1988.
12. B. Bünz, S. Agrawal, M. Zamani, and D. Boneh. Zether: Towards privacy in a smart contractworld. In *FC*, 2020.
13. B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *IEEE S&P*, 2018.
14. J. Camenisch, R. Chaabouni, and A. shelat. Efficient protocols for set membership and range proofs. In *Asiacrypt*, 2008.
15. J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *Eurocrypt*, 2001.
16. R. Canetti, Y. Chen, J. Holmgren, A. Lombardi, G. Rothblum, R. Rothblum, and D. Wichs. Fiat-Shamir: From practice to theory. In *STOC*, 2019.
17. R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisited. *J. of the ACM*, 51(4), 2004.
18. R. Canetti and S. Goldwasser. An efficient threshold public key cryptosystem secure against adaptive chosen-ciphertext attacks. In *Eurocrypt*, 1999.
19. R. Canetti, A. Lombardi, and D. Wichs. Fiat-Shamir: From Practice to Theory, Part II (NIZK and Correlation Intractability from Circular-Secure FHE). Cryptology ePrint Archive: Report 2018/1248.
20. R. Chaabouni, H. Lipmaa, and B. Zhang. A non-interactive range proof with constant communication. In *Financial Cryptography*, 2012.
21. P. Chaidos and J. Groth. Making Sigma-protocols non-interactive without random oracles. In *PKC*, 2015.
22. A. Chan, Y. Frankel, and Y. Tsiounis. Easy come – easy go divisible cash. In *Eurocrypt*, 1998.
23. M. Ciampi, R. Parisella, and D. Ventury. On adaptive security of delayed-input Sigma protocols and Fiat-Shamir NIZKs. In *SCN*, 2020.
24. G. Couteau, M. Kloof, H. Lin, and M. Reichle. Efficient range proofs with transparent setup from bounded integer commitments. In *Eurocrypt*, 2021.
25. G. Couteau, T. Peters, and D. Pointcheval. Removing the strong RSA assumption from arguments over the integers. In *Eurocrypt*, 2017.

26. R. Cramer, I. Damgård, and B. Schoenmaekers. Proofs of partial knowledge and simplified design of witness hiding protocols. In *Crypto*, 1994.
27. I. Damgård. Efficient concurrent zero-knowledge in the auxiliary string model. In *Eurocrypt*, 2000.
28. I. Damgård and E. Fujisaki. A statistically-hiding integer commitment scheme based on groups with hidden order. In *Asiacrypt*, 2002.
29. I. Damgård and M. Jurik. A generalisation, a simplification and some applications of Paillier’s probabilistic public-key system. In *PKC*, 2001.
30. A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero-knowledge. In *Crypto*, 2001.
31. J. Devevey, B. Libert, K. Nguyen, T. Peters, and M. Yung. Non-interactive CCA2-secure threshold cryptosystems: Achieving adaptive security in the standard model without pairings. In *PKC*, 2021.
32. U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *FOCS*, 1990.
33. U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero-knowledge under general assumptions. *SIAM J. of Computing*, 29(1), 1999.
34. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Crypto*, 1986.
35. P.-A. Fouque and D. Pointcheval. Threshold Cryptosystems Secure against Chosen-Ciphertext Attacks. In *Asiacrypt*, 2001.
36. P.-A. Fouque, J. Stern, and G. Wacker. Cryptocomputing with rationals. In *FC*, 2002.
37. E. Fujisaki and T. Okamoto. Statistical zero knowledge protocols to prove modular polynomial relations. In *Crypto*, 1997.
38. C. Gentry and S. Halevi. Compressible FHE with applications to PIR. In *TCC*, 2019.
39. S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 1989.
40. A. Gonzalez and C. Ràfols. New techniques for non-interactive shuffle and range arguments. In *ACNS*, 2017.
41. J. Groth. Non-interactive zero-knowledge arguments for voting. In *ACNS*, 2005.
42. J. Groth. Efficient zero-knowledge arguments from two-tiered homomorphic commitments. In *Asiacrypt*, 2011.
43. J. Groth, R. Ostrovsky, and A. Sahai. New techniques for noninteractive zero-knowledge. *J. ACM*, 2012.
44. J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *Eurocrypt*, 2008.
45. B. Hemenway, B. Libert, R. Ostrovsky, and D. Vergnaud. Lossy encryption: Constructions from general assumptions and efficient selective opening chosen ciphertext security. In *Asiacrypt*, 2011.
46. S. Jarecki and A. Lysyanskaya. Adaptively secure threshold cryptography: Introducing concurrency, removing erasures. In *Eurocrypt*, 2000.
47. A. Kiayias, N. Leonardos, H. Lipmaa, K. Pavlyk, and Q. Tang. Near optimal rate homomorphic encryption for branching programs. *Priv. Enhancing Technol.*, 2015.
48. H. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4), 1983.
49. B. Libert, S. Ling, K. Nguyen, and H. Wang. Lattice-based zero-knowledge arguments for integer relations. In *Crypto*, 2018.
50. B. Libert, K. Nguyen, A. Passelègue, and R. Titu. Simulation-sound arguments for LWE and applications to KDM-CCA2 security. In *Asiacrypt*, 2020.

51. B. Libert, K. Nguyen, T. Peters, and M. Yung. One-shot Fiat-Shamir-based NIZK arguments of composite residuosity in the standard model. *Cryptology ePrint Archive: Report 2020/1334*, 2020.
52. B. Libert and M. Yung. Non-interactive cca-secure threshold cryptosystems with adaptive security: New framework and constructions. In *TCC*, 2012.
53. H. Lipmaa. On Diophantine complexity and statistical zero-knowledge arguments. In *Asiacrypt*, 2003.
54. H. Lipmaa. Optimally sound sigma protocols under DCRA. In *FC*, 2017.
55. H. Lipmaa, N. Asokan, and V. Niemi. Secure vickrey auctions without threshold trust. In *Financial Cryptography*, 2002.
56. M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC*, 1990.
57. S. Noether. Ring signature confidential transactions for monero. *Cryptology ePrint Archive Report 2015/1098*, 2015.
58. P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Eurocrypt*, 1999.
59. R. Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In *TCC*, 2013.
60. T. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Crypto*, 1991.
61. C. Peikert and S. Shiehian. Non-interactive zero knowledge for NP from (plain) Learning With Errors. In *Crypto*, 2019.
62. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.
63. A. Rial, M. Kohlweiss, and B. Preneel. Universally composable adaptive priced oblivious transfer. In *Pairing*, 2009.
64. A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *FOCS*, 1999.
65. V. Shoup and R. Gennaro. Securing threshold cryptosystems against chosen ciphertext attack. In *Eurocrypt*, 1998.
66. B. Vallée. Gauss' algorithm, revisited. *J. of Algorithms*, 1991.

Supplementary Material

A Other Definitions for Cryptographic Primitives

A.1 Non-Interactive Zero-Knowledge and Simulation-Sound Arguments

We recall the definitions of NIZK proofs. Since it is sufficient for our applications, we allow the common reference string to be generated as a function of the language \mathcal{L} . We actually give a slightly different definition than the standard ones, defining NIZK for gap languages. That is, a language is defined by a pair of language $\mathcal{L}_{\text{zk}} \subseteq \mathcal{L}_{\text{sound}}$, and completeness is guaranteed for statements in \mathcal{L}_{zk} while soundness is guaranteed for statement outside $\mathcal{L}_{\text{sound}}$. This is sufficient for our purpose.

In addition, we consider NIZK argument systems where each argument comes with a label lbl taken as input by both the prover and the verifier. Labels will only be useful when we consider simulation-soundness, which is necessary in our CCA-secure encryption schemes.

Definition A.1. *A non-interactive zero-knowledge (NIZK) argument system Π for a language $\mathcal{L} = (\mathcal{L}_{\text{zk}}, \mathcal{L}_{\text{sound}})$ associated to two NP relations $(R_{\text{zk}}, R_{\text{sound}})$ consists of four PPT algorithms $(\text{Gen}_{\text{par}}, \text{Gen}_{\mathcal{L}}, \text{P}, \text{V})$ with the following syntax:*

- $\text{Gen}_{\text{par}}(1^\lambda)$ takes as input a security parameter λ and outputs public parameters par .
- $\text{Gen}_{\mathcal{L}}(1^\lambda, \mathcal{L}, \tau_{\mathcal{L}})$ takes as input a security parameter λ , the description of \mathcal{L} which specifies a statement length N , and a membership testing trapdoor $\tau_{\mathcal{L}}$ for \mathcal{L} . It outputs the language-dependent part $\text{crs}_{\mathcal{L}}$ of the common reference string $\text{crs} = (\text{par}, \text{crs}_{\mathcal{L}})$.
- $\text{P}(\text{crs}, x, w, \text{lbl})$ is a proving algorithm taking as input the common reference string crs , a statement $x \in \{0, 1\}^N$, a witness w such that $(x, w) \in R_{\text{zk}}$ and a label lbl . It outputs a proof π .
- $\text{V}(\text{crs}, x, \pi, \text{lbl})$ is a verification algorithm taking as input a common reference string crs , a statement $x \in \{0, 1\}^N$, and a proof π . It outputs 1 or 0.

Moreover, Π should satisfy the following properties. For simplification we denote below by Setup an algorithm that runs successively Gen_{par} and $\text{Gen}_{\mathcal{L}}$ to generate a common reference string.

- **Completeness:** For any $(x, w) \in R_{\text{zk}}$ and any $\text{lbl} \in \{0, 1\}^*$, we have

$$\Pr [\text{crs} \leftarrow \text{Setup}(1^\lambda, \mathcal{L}), \pi \leftarrow \text{P}(\text{crs}, x, w, \text{lbl}) : \text{V}(\text{crs}, x, \pi, \text{lbl}) = 1] \geq 1 - \text{negl}(\lambda).$$

- **Soundness:** For any $x \in \{0, 1\}^N \setminus \mathcal{L}_{\text{sound}}$ and any PPT prover P^* , we have

$$\Pr [\text{crs} \leftarrow \text{Setup}(1^\lambda, \mathcal{L}), (\pi, \text{lbl}) \leftarrow P^*(\text{crs}, x) : \text{V}(\text{crs}, x, \pi, \text{lbl}) = 1] \leq \text{negl}(\lambda).$$

- **Adaptive Soundness:** For any PPT prover P^* , we have

$$\Pr[\text{crs} \leftarrow \text{Setup}(1^\lambda, \mathcal{L}), (\pi, x, \text{lbl}) \leftarrow P^*(\text{crs}) : \\ \forall(\text{crs}, x, \pi, \text{lbl}) = 1 \wedge x \notin \mathcal{L}_{\text{sound}}] \leq \text{negl}(\lambda).$$

- **Multi-theorem Zero-Knowledge:** There is a PPT simulator $(\text{Sim}_0, \text{Sim}_1)$ such that, for any PPT adversary \mathcal{A} , we have

$$|\Pr[\text{crs} \leftarrow \text{Setup}(1^\lambda, \mathcal{L}) : 1 \leftarrow \mathcal{A}^{\text{P}(\text{crs}, \cdot, \cdot)}(\text{crs})] \\ - \Pr[(\text{crs}, \tau_{\text{zk}}) \leftarrow \text{Sim}_0(1^\lambda, \mathcal{L}) : 1 \leftarrow \mathcal{A}^{\mathcal{O}(\text{crs}, \tau_{\text{zk}}, \cdot, \cdot)}(\text{crs})]| \leq \text{negl}(\lambda).$$

Here, $\text{P}(\text{crs}, \cdot, \cdot)$ is an oracle that outputs \perp on input of $(x, w, \text{lbl}) \notin R_{\text{zk}}$ and outputs a valid proof $\pi \leftarrow \text{P}(\text{crs}, x, w, \text{lbl})$ otherwise; $\mathcal{O}(\text{crs}, \tau_{\text{zk}}, \cdot, \cdot)$ is an oracle that outputs \perp on input of (x, w, lbl) such that $(x, w) \notin R_{\text{zk}}$ and outputs a simulated argument $\pi \leftarrow \text{Sim}_1(\text{crs}, \tau_{\text{zk}}, x, \text{lbl})$ on input of (x, w, lbl) such that $(x, w) \in R_{\text{zk}}$. Note that this simulated proof π is generated independently of the witness w provided as input.¹²

The notion of multi-theorem zero-knowledge generalizes the notion of adaptive ZK to the multi-query setting. The first notion of soundness is *non-adaptive* in that the statement is given as input to the dishonest prover and chosen independently of the common reference string. The stronger notion of *adaptive soundness* allows the target statement to be chosen by the adversary after having received the common reference string. Ciampi *et al.* [23] showed how to achieve adaptive soundness by hashing the statement together with the prover’s first message.

It is known [59] that NIZK arguments cannot simultaneously provide adaptive soundness and statistical zero-knowledge under falsifiable assumptions. The reason lies in the impossibility of recognizing when the adversary wins and outputs a proof for a false statement (as the winning condition $x \notin \mathcal{L}_{\text{sound}}$ may not be efficiently checkable). One way to bypass the impossibility results is to either settle for computational zero-knowledge or consider *trapdoor languages*, where a trapdoor can be used to recognize false statements.

In the context of CCA security, we will consider a notion of adaptive soundness for trapdoor languages. Definition A.1 captures a notion of multi-theorem zero-knowledge, which allows the adversary to obtain proofs for multiple statements. Feige *et al.* [33] gave a generic transformation of a multi-theorem NIZK argument system from a single-theorem one (where the adversary can only invoke the oracle once).

SIMULATION-SOUNDNESS. We now recall the definition of simulation-soundness introduced in [64], which informally captures the adversary’s inability to create a new proof for a false statement x^* even after having seen simulated proofs for possibly false statements $\{x_i\}_i$ of its choice.

¹² In particular, Sim_1 can be run on any statement x , even $x \notin \mathcal{L}_{\text{sound}}$.

In the following, in order to allow a challenger to efficiently check the winning condition (ii) in the security experiment, we restrict ourselves to *trapdoor languages*, where a language-specific trapdoor $\tau_{\mathcal{L}}$ makes it possible to determine if a given statement $x^* \in \{0, 1\}^N$ belongs to the language \mathcal{L}_{zk} with overwhelming probability. This restriction has no impact on our applications where we always have a membership testing trapdoor $\tau_{\mathcal{L}}$ at our disposal.

Definition A.2 ([64,30]). *Let a language $\mathcal{L} = (\mathcal{L}_{\text{zk}}, \mathcal{L}_{\text{sound}})$. A NIZK argument system for \mathcal{L} provides **unbounded simulation soundness** if no PPT adversary has non-negligible advantage in this game.*

1. *The challenger chooses a membership testing trapdoor $\tau_{\mathcal{L}}$ that allows recognizing elements of \mathcal{L}_{zk} . Let $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ be an efficient NIZK simulator for \mathcal{L} . The challenger runs $(\text{crs}, \tau_{\text{zk}}) \leftarrow \text{Sim}_0(1^\lambda, \mathcal{L})$ and gives $(\text{crs}, \tau_{\mathcal{L}})$ to the adversary \mathcal{A} .*
2. *The adversary \mathcal{A} is given oracle access to $\text{Sim}_1(\text{crs}, \tau_{\text{zk}}, \cdot, \cdot)$. At each query, \mathcal{A} chooses a statement $x \in \{0, 1\}^N$ and a label $\text{lbl} \in \{0, 1\}^*$. It obtains a simulated argument $\pi \leftarrow \text{Sim}_1(\text{crs}, \tau_{\text{zk}}, x, \text{lbl})$.*
3. *\mathcal{A} outputs $(x^*, \text{lbl}^*, \pi^*)$.*

Let \mathcal{Q} be the set of all simulation queries and responses $(x_i, \text{lbl}_i, \pi_i)$ made by \mathcal{A} to $\text{Sim}_1(\text{crs}, \tau_{\text{zk}}, \cdot, \cdot)$. The adversary \mathcal{A} wins if the following conditions are satisfied: (i) $(x^, \text{lbl}^*, \pi^*) \notin \mathcal{Q}$; (ii) $x^* \notin \mathcal{L}_{\text{sound}}$; and (iii) $\text{V}(\text{crs}, x^*, \pi^*, \text{lbl}^*) = 1$. The adversary's advantage $\text{Adv}_{\mathcal{A}}^{\text{USS}}(\lambda)$ is its probability of success taken over all coin tosses.*

When the adversary is restricted to making only one query in the experiment, the notion is called *one-time simulation-soundness*.

A.2 Threshold PKE

In this section, we recall the TPKE syntax defined by Boneh *et al.* [9].

Definition A.3 (Threshold PKE). *Given some message space \mathcal{M} , a Threshold Public Key Encryption scheme (TPKE) is a tuple of efficient PPT algorithms (Keygen, Encrypt, PartDec, PartVerify, Combine) with the following specifications:*

- $\text{Keygen}(1^\lambda, 1^t, 1^n) \rightarrow (\text{pp}, \text{pk}, \text{sk}_1, \text{sk}_2, \dots, \text{sk}_n)$: *On input a security parameter λ , a number of servers n and a threshold t , the algorithm outputs a set of public parameters pp (which are implicit in the inputs of all other algorithms), a public key pk and a set of secret key shares $\text{sk}_1, \text{sk}_2, \dots, \text{sk}_n$.*
- $\text{Encrypt}(\text{pk}, \text{Msg}) \rightarrow \text{ct}$: *On input the public parameters pp , the encryption key pk and a message $\text{Msg} \in \mathcal{M}$, the algorithm outputs a ciphertext ct .*
- $\text{PartDec}(\text{pp}, \text{ct}, \text{sk}_i) \rightarrow \mu_i$: *Given public parameters pp , a ciphertext ct and a secret key share sk_i , this algorithm outputs a partial decryption μ_i .*
- $\text{PartVerify}(\text{pk}, \text{ct}, \mu_i) \rightarrow \text{b} \in \{0, 1\}$: *On input of public parameters pp , a ciphertext ct and a partial decryption μ_i , this algorithm outputs a bit b .*

- **Combine** $(\text{pk}, B = (\mathcal{S}, \{\phi(\mu_i)\}_{i \in \mathcal{S}}), \text{ct}) \rightarrow \text{Msg}'$: Given public parameters and a set of images of ϕ of partial decryptions, the algorithm outputs a message $\text{Msg}' \in \mathcal{M}$. The function ϕ is public and deterministic.¹³

We recall the different properties that our construction will satisfy.

Definition A.4 (Decryption Correctness). A TPKE defined as above provides decryption correctness if the following holds. For any $\lambda \in \mathbb{N}$, any number of servers $n = \text{poly}(\lambda)$, any threshold $t \leq n$, any set $|\mathcal{S}| \geq t$ and any message $\text{Msg} \in \mathcal{M}$, if we run $(\text{pp}, \text{pk}, \text{sk}_1, \text{sk}_2, \dots, \text{sk}_n) \leftarrow \text{Keygen}(1^\lambda, 1^t, 1^n)$, $\text{ct} \leftarrow \text{Encrypt}(\text{pk}, \text{Msg})$ and then $\mu_i \leftarrow \text{PartDec}(\text{pp}, \text{sk}_i, \text{ct}), \forall i \in \mathcal{S}$, we have $\Pr[\text{Combine}(\text{pk}, (\mathcal{S}, \{\phi(\mu_i)\}_{i \in \mathcal{S}}), \text{ct}) = \text{Msg}] = 1 - \text{negl}(\lambda)$.

Definition A.5 (Partial Verification Correctness). A TPKE provides partial verification correctness if the following holds. For any $\lambda \in \mathbb{N}$, any number of servers $n = \text{poly}(\lambda)$, any threshold $t \leq n$ and any message $\text{Msg} \in \mathcal{M}$, if we run $(\text{pp}, \text{pk}, \text{sk}_1, \text{sk}_2, \dots, \text{sk}_n) \leftarrow \text{Keygen}(1^\lambda, 1^t, 1^n)$, $\text{ct} \leftarrow \text{Encrypt}(\text{pk}, \text{Msg})$ and $\mu_i \leftarrow \text{PartDec}(\text{pp}, \text{sk}_i, \text{ct}), \forall i \in \mathcal{S}$, then $\Pr[\text{PartVerify}(\text{pk}, \text{ct}, \mu_i) = 1] = 1 - \text{negl}(\lambda)$.

Various security notions have been considered for threshold PKE schemes. The construction of Section 4.3 is proven IND-CCA2 secure under static corruptions, where an adversary corrupts $t - 1$ servers and obtains their secret key shares at the outset of the IND-CCA2 game.

Definition A.6 (CCA2 security under static corruptions). A TPKE system provides IND-CCA2 security under static corruptions if no PPT adversary \mathcal{A} has non-negligible advantage in the following game.

1. On input the security parameter λ , the adversary \mathcal{A} chooses a number of servers $n = \text{poly}(\lambda)$ and a threshold $t \leq n$, as well as $\mathcal{C} \subseteq [n]$, comprised of exactly $t - 1$ elements.
2. The challenger generates $(\text{pp}, \text{pk}, \text{sk}_1, \text{sk}_2, \dots, \text{sk}_n) \leftarrow \text{Keygen}(1^\lambda, 1^t, 1^n)$. It sends (pp, pk) to \mathcal{A} as well as $\{\text{sk}_i\}_{i \in \mathcal{C}}$.
3. The adversary \mathcal{A} is granted access to a partial decryption oracle: At each query, \mathcal{A} chooses an index $i \in [n]$ and a ciphertext ct and the challenger returns a partial decryption $\mu_i \leftarrow \text{PartDec}(\text{pk}, \text{sk}_i, \text{ct})$.
4. In the challenge phase, \mathcal{A} chooses $\text{Msg}_0^*, \text{Msg}_1^* \in \mathcal{M}$. The challenger replies with a challenge ciphertext $\text{ct}^* \leftarrow \text{Encrypt}(\text{pk}, \text{Msg}_b^*)$, where $b \leftarrow U(\{0, 1\})$ is a random bit.
5. In a post-challenge query phase, \mathcal{A} makes more partial decryption queries for ciphertexts $\text{ct} \neq \text{ct}^*$.
6. The experiment ends with \mathcal{A} outputting a bit $b' \in \{0, 1\}$.

The advantage of \mathcal{A} is defined as $\text{Adv}^{\text{ind-cca}}(\mathcal{A}) := |\Pr[b' = b] - \frac{1}{2}|$.

We now recall the notion of *robustness*, which informally captures that no malicious adversary can prevent a honest majority from decrypting a valid ciphertext.

¹³ It helps defining robustness. For non-robust TPKE, ϕ is the identity function.

Definition A.7 ([9]). A TPKE scheme satisfies **robustness** if no PPT adversary \mathcal{A} can cause the following experiment $\text{Expt}_{\mathcal{A}, \text{TPKE}}^{\text{robust}}(1^\lambda)$ to output 1 with non-negligible probability.

1. On input the security parameter λ , \mathcal{A} chooses a polynomial number of servers $n = \text{poly}(\lambda)$ and a threshold $t \leq n$.
2. The challenger samples $(\text{pp}, \text{pk}, \text{sk}_1, \text{sk}_2, \dots, \text{sk}_n) \leftarrow \text{Keygen}(1^\lambda, 1^t, 1^n)$ and provides $(\text{pp}, \text{pk}, \text{sk}_1, \text{sk}_2, \dots, \text{sk}_n)$ to \mathcal{A} .
3. \mathcal{A} outputs a partial decryption forgery $(\text{ct}^*, \mu_i^*, i)$, where $i \in [n]$.
4. The experiment outputs 1 if we have $\phi(\hat{\mu}_i^*) \neq \phi(\text{PartDec}(\text{pk}, \text{sk}_i, \text{ct}^*))$ while $\text{PartVerify}(\text{pk}, \text{ct}^*, \mu_i^*) = 1$.

We note that the function ϕ allows considering as robust a TPKE such that $\mu_i^* = (\hat{\mu}_i^*, \pi_i^*)$ and where **Combine** only runs on $\hat{\mu}_i^*$ and not on π_i^* . While, given $(\hat{\mu}_i^*, \pi_i^*)$, **Combine** could have simply striped π_i^* , such a formalization would prevent showing as robust a TPKE where $\hat{\mu}_i^*$ is an element of an admissible language and π_i^* is a probabilistic membership argument whose validity does not necessarily guarantee that π_i^* is in the range of honestly computed arguments. Thanks to ϕ , such case will not be artificially discarded.

A.3 Achieving robustness for the DCR-based CCA2-Secure Threshold Cryptosystem from the Naor-Yung Paradigm

Following [35,29], we briefly describe a robust variant of our CCA2-secure TPKE (**KeyGen**, **Enc**, **PartDec**, **Combine**) under static corruptions from Section 4.3.

We assume a Σ -protocol for proving membership of $(a, b, c, d) \in \mathcal{L}^{\log}$ where a, b, c, d are quadratic residues in $\mathbb{Z}_{N_1^{\zeta+1}}^*$, for some $\zeta \leq \zeta'$, such that $a^w = b$ and $c^w = d$, for some integer $w \in \mathbb{Z}$.

Using a generic transformation suggested by Ciampi *et al.* [23, Section 4.2], one can turn this Σ -protocol into a trapdoor Σ -protocol, which can in turn provide a NIZK system $(\text{Setup}_{\log}, \text{P}_{\log}, \text{V}_{\log})$ with the construction from Section 4.2.¹⁴

The construction of [23, Section 4.2] uses a semantically secure PKE, which is used to encrypt responses to the two possible challenges in $\{0, 1\}$ (of course, parallel repetitions are needed to ensure negligible soundness error). This ciphertext is then sent along with the prover's first message.

Then, as in [29], we need to modify the key generation algorithm and introduce commitments $\{\text{vk}_i\}_{i=1}^n$ to the secret key shares.

KeyGen' $(1^\lambda, 1^B, 1^t, 1^n)$: Run $(\text{pk}, \{\text{sk}_i\}_{i=1}^n) \leftarrow \text{KeyGen}(1^\lambda, 1^B, 1^t, 1^n)$, where the public key is $\text{pk} = (N_1, N_2, \text{crs}_{\mathcal{L}})$ and the secret key shares $\{\text{sk}_i\}_{i=1}^n$ are defined as $\text{sk}_i = f(i) \bmod N_1^{\zeta'} p_1' q_1'$ for each $i \in [n]$. Then, pick a random generator vk of the subgroup (which has order $N_1^{\zeta'} p_1' q_1'$) of quadratic residues

¹⁴ For this purpose, we do not need one-time simulation-soundness and we can use the simplified multi-theorem NIZK version suggested in Section 3.

in $\mathbb{Z}_{N_1^{\zeta'+1}}^*$, and set $\Delta = n!$ so as to avoid division in the exponents. For each, $i \in [n]$, compute $\mathbf{vk}_i = \mathbf{vk}^{\Delta \mathbf{sk}_i} \bmod N_1^{\zeta'+1}$. Finally, generate a common reference string crs for a trapdoor Σ -protocol showing the equality of discrete logarithms over $\langle v \rangle$. Update the public key with $\mathbf{pk}' := (\mathbf{pk}, \mathbf{vk}, \{\mathbf{vk}_i\}_{i=1}^n, \text{crs})$ and return

$$(\mathbf{pk}', \{\mathbf{sk}_i\}_{i=1}^n).$$

PartDec'($\mathbf{sk}_i, \text{ct} := (\text{ct}_1, \text{ct}_2, \ell_M, \vec{\pi})$): Compute $\mu_i = \text{PartDec}(\mathbf{sk}_i, \text{ct})$. If $\mu_i = \perp$, return \perp . Otherwise, let $\mu_i = \text{ct}_1^{2\Delta \mathbf{sk}_i} \bmod N_1^{\zeta'+1}$ and compute a NIZK proof π_i^{\log} that $\log_{\text{ct}_1}(\mu_i^2) = \log_{\mathbf{vk}}(\mathbf{vk}_i)$, where \mathbf{vk} and \mathbf{vk}_i can be reduced modulo $N_1^{\zeta'+1}$. Return the modified decryption share

$$\mu'_i = (\mu_i, \pi_i^{\log}).$$

PartVerify(\mathbf{pk}' , $(\text{ct}_1, \text{ct}_2, \ell_M, \vec{\pi}), \mu'_i$): Parse μ'_i as $\mu'_i = (\mu_i, \pi_i^{\log})$ and return 1 if and only if $1 = \text{PartVerify}(\mathbf{pk}, (\text{ct}_1, \text{ct}_2, \ell_M, \vec{\pi}), \mu_i)$ and π_i^{\log} is a valid argument for $\text{ct}_1, \mu_i, \mathbf{vk}, \mathbf{vk}_i$.

This variant is still a CCA2-secure threshold PKE scheme in the static corruptions setting. Indeed, in the key generation phase, we can choose \mathbf{vk} by setting $\mathbf{vk} = (1 + N_1)^v \cdot r_v^{2N_1^{\zeta'}}$ mod $N_1^{\zeta'+1}$ for random $v \leftarrow \mathbb{Z}_{N_1^{\zeta'}}$ and $r_v \leftarrow \mathbb{Z}_{N_1}^*$ to obtain a generator of the subgroup of quadratic residues with overwhelming probability. Knowing $\mathbf{vk}^{f(0)} = (1 + N_1)^v \bmod N_1^{\zeta'+1}$, we can also simulate the verification commitment keys $\{\mathbf{vk}_i\}_{i \in [n] \setminus \mathcal{C}}$, as done in the threshold variant of Damgård-Jurik [29, Section 4.1]. Therefore, this TPKE variant retains CCA2 security under static corruptions thanks to the zero-knowledge property of the arguments of membership of \mathcal{L}^{\log} . Robustness then follows from the soundness of $(\text{Setup}_{\log}, \text{P}_{\log}, \text{V}_{\log})$.

B Equivalence of ζ -DCR and 1-DCR for $\zeta \in \text{poly}(\lambda)$

Damgård and Jurik [29] initially gave their security proof using a recursive argument (rather than a sequence of hybrid experiments) that loses a factor 2 at each step, thus incurring an apparent security loss 2^ζ . However, the semantic security of their scheme under the 1-DCR assumption for any polynomial ζ is a well-known result (see, e.g., [47,38]). In particular, it was crucially used to achieve the first single-server rate-1 PIR [47].

Although the proof of [29] was not explicitly written as a hybrid argument, it is fairly straightforward to re-write so as to have a security loss $\zeta \in \text{poly}(\lambda)$. The proof below is perhaps folklore but we give it for completeness.

Proof. In order to prove that N^ζ -th residues are indistinguishable from random in $\mathbb{Z}_{N^{\zeta+1}}^*$, let us consider a hybrid argument with $\zeta + 1$ hybrids $(H_i)_{0 \leq i \leq \zeta}$. In H_i ,

the adversary obtains a Damgård-Jurik encryption of a message of the form

$$a = \sum_{j=1}^i a_j \cdot N^{\zeta-j} \in \mathbb{Z}_{N^\zeta},$$

where each a_j is uniform over \mathbb{Z}_N . Namely, it gets $c = (1 + N)^a \cdot r^N \pmod{N^2}$ for some $r \sim U(\mathbb{Z}_N)$. So, H_ζ corresponds to an encryption of a random element of \mathbb{Z}_{N^ζ} (so that the ciphertext is uniformly distributed over $\mathbb{Z}_{N^{\zeta+1}}^*$) and H_0 corresponds to an encryption of 0 (i.e., an N^ζ -th residue). If an adversary can distinguish between the cases $i = \zeta$ and $i = 0$, it can distinguish between two consecutive hybrids H_{i-1} and H_i . The reduction just has to guess which one (with probability $1/\zeta$, which is non-negligible if ζ is polynomial ;) in order to embed a DCR instance in the right place. Specifically, the fact any 1-DCR challenge $c = (1 + N)^a \cdot r^N \pmod{N^2}$ lives in $\mathbb{Z}_{N^{\zeta+1}}^*$ implies that it has a representation of the form

$$c = (1 + N)^u \cdot v^{N^\zeta} \pmod{N^{\zeta+1}}$$

for some $v \in \mathbb{Z}_N^*$ and some $u \in \mathbb{Z}_{N^\zeta}$ such that $a = u \pmod{N}$. This means that

$$c^{N^{\zeta-i}} \pmod{N^{\zeta+1}} = (1 + N)^{\sum_{j=1}^{\zeta-i+1} u_j \cdot N^{\zeta-j} + a \cdot N^{\zeta-i}} \cdot w^{N^\zeta} \pmod{N^{\zeta+1}},$$

for some $w \in \mathbb{Z}_N^*$ and $u_1, \dots, u_{\zeta-i+1} \in \mathbb{Z}_N$. Hence, given 1-DCR challenge $c = (1 + N)^a \cdot r^N \pmod{N^2}$ (where a is either uniform over \mathbb{Z}_N or 0), the reduction can interpolate between hybrid $i - 1$ and hybrid i for any $i \in [1, \zeta]$. To do that, the reduction chooses $a_1, \dots, a_{i-1} \leftarrow U(\mathbb{Z}_N)$ and sets

$$C = (1 + N)^{\sum_{j=1}^{i-1} a_j \cdot N^{\zeta-j}} \cdot c^{N^{\zeta-i}} \pmod{N^{\zeta+1}},$$

which is given as a challenge to the distinguisher. Clearly, if $a = 0$, C is distributed as in hybrid $i - 1$. If $a \sim U(\mathbb{Z}_N)$, it is distributed as in hybrid i . \square

C Deferred Material for the Naor-Yung-Based Construction

C.1 Proof of Theorem 4.2

Proof. In the following proof, we denote by $\text{crs} \leftarrow \text{Setup}(1^\lambda, \mathcal{L})$ the act of sampling $\text{crs} \leftarrow \text{Gen}_{\mathcal{L}}(\text{Gen}(1^\lambda), \mathcal{L})$.

COMPLETENESS. Run $\text{crs} \leftarrow \text{Setup}(1^\lambda, \mathcal{L})$ and $(\text{VK}, (a, z), \text{sig}) \leftarrow \text{P}(\text{crs}, x, w, \text{lbl})$ for any $x \in \mathcal{L}$ with label $\text{lbl} \in \{0, 1\}^*$ and witness w . By correctness of OTS, we have $\mathcal{V}(\text{VK}, (x, a, z, \text{lbl}), \text{sig}) = 1$ with probability $1 - \text{negl}(\lambda)$. Moreover, z is correctly parsed as $(z_0, z_1, \text{Chall}_0, \text{Chall}_1)$, where the challenges satisfy $\text{Chall}_0 \oplus \text{Chall}_1 = \text{Hash}(k, (x, a, \text{VK}))$.

The completeness of $\Pi^{(1)}$ implies that $V^{(1)}(\text{crs}_{\mathcal{L}}^{(1)}, x, a_1, \text{Chall}_1, z_1) = 1$ with probability $1 - \text{negl}(\lambda)$. Since $a_0 = z_0^N \cdot (u^{\text{VK}} \cdot v)^{\text{Chall}_0} \bmod N^2$, the verifier will return $V(\text{crs}, x, (\text{VK}, (a, z), \text{sig}), \text{lbl}) = 1$ with probability $1 - \text{negl}(\lambda)$.

ZERO-KNOWLEDGE. We start by describing the simulator $(\text{Sim}_0, \text{Sim}_1)$ before proving that simulated proofs are indistinguishable from genuine ones.

First, $\text{Sim}_0(1^\lambda, \mathcal{L})$ behaves exactly as $\text{Setup}(1^\lambda, \mathcal{L})$ except that it outputs both crs and $\tau_{\text{zk}} = (u_0, v_0)$.

Next, to generate a simulated proof, $\text{Sim}_1(\tau_{\text{zk}}, \text{crs}, x, \text{lbl})$ starts by generating a key pair $(\text{VK}, \text{SK}) \leftarrow \mathcal{G}(1^\lambda)$ for the one-time signature. It then picks $\text{Chall}_1 \leftarrow U(\{0, 1\}^\lambda)$ and runs the simulator of the special ZK property of $\Pi^{(1)}$ to get a simulated transcript $(a_1, z_1) \leftarrow \text{ZKSim}^{(1)}(\text{crs}, x, \text{Chall}_1)$. Since $(u^{\text{VK}} \cdot v) \in \mathcal{L}^{\text{DCR}}$ for any VK , the simulator can use its N -th root $u_0^{\text{VK}} \cdot v_0$ as a witness to generate a real transcript for $\Pi^{(0)}$. Namely, the simulator first samples $a'_0 \leftarrow U(\mathbb{Z}_N^*)$, sets $a_0 := a'_0{}^N \bmod N^2$ and computes $\text{Chall} := \text{Hash}(k, (x, (a_0, a_1), \text{VK}))$ as well as $\text{Chall}_0 := \text{Chall} \oplus \text{Chall}_1$. Finally, it obtains $z_0 = a'_0 \cdot (u_0^{\text{VK}} \cdot v_0)^{\text{Chall}_0} \bmod N$ using the witness (u_0, v_0) .

It then generates a one-time signature $\text{sig} \leftarrow \mathcal{S}(\text{SK}, (x, a, z, \text{lbl}))$ and outputs the simulated proof $\bar{\pi} := (\text{VK}, (a, z), \text{sig})$, where $z = (z_0, z_1, \text{Chall}_0, \text{Chall}_1)$.

Note that $(\text{Chall}_0, \text{Chall}_1)$ are uniformly distributed among pairs in $\{0, 1\}^\lambda$ such that $\text{Chall}_0 \oplus \text{Chall}_1 = \text{Chall}$. Moreover, we replaced the transcript of a real proof for $\Pi^{(1)}$ by that of a simulated proof, and vice versa for $\Pi^{(0)}$. The special zero-knowledge property of $\Pi^{(0)}$ (which holds in the statistical sense) and that of $\Pi^{(1)}$ then imply the zero-knowledge property of the OR proof. Moreover, if $\Pi^{(1)}$ is statistically special ZK, the simulated proof is statistically indistinguishable from a real proof.

ONE-TIME SIMULATION-SOUNDNESS. We consider a sequence of games, where we call W_i the event that the adversary outputs 1 in Game_i .

Game₀: This is the real one-time simulation soundness game. Namely, the challenger runs $(\text{crs}, \tau_{\text{zk}}) \leftarrow \text{Sim}_0(1^\lambda, \mathcal{L})$ and gives crs to the adversary \mathcal{A} . The adversary is allowed to make exactly one query to the simulation oracle. It chooses a statement x and a label lbl and the challenger replies with $\bar{\pi} = (\text{VK}, (a, z), \text{sig}) \leftarrow \text{Sim}_1(\text{crs}, \tau_{\text{zk}}, x, \text{lbl})$. Eventually, \mathcal{A} halts and outputs a triple $(x^*, \bar{\pi}^*, \text{lbl}^*)$, where $\bar{\pi}^* := (\text{VK}^*, (a^*, z^*), \text{sig}^*)$. The flag W_0 is set to 1 if and only if $(x^*, \bar{\pi}^*, \text{lbl}^*) \neq (x, \bar{\pi}, \text{lbl})$, $x^* \notin \mathcal{L}$ (we assume that \mathcal{L} is a trapdoor language so that the challenger has a trapdoor allowing to check the latter winning condition) and $V(\text{crs}, x^*, \bar{\pi}^*, \text{lbl}^*) = 1$. By definition, the advantage of \mathcal{A} is $\Pr[W_0]$.

Game₁: This is the same game as Game_0 , except that W_1 is no longer set to 1 if $\text{VK}^* = \text{VK}$. If this event occurs and V still accepts the proof $\bar{\pi}^*$, \mathcal{A} was necessarily able to forge a non-trivial one-time signature. Under the assumption that the one-time signature is strongly unforgeable, we know that $|\Pr[W_0] - \Pr[W_1]| \leq \text{Adv}^{\text{OTS}}(\lambda)$ is negligible.

Game₂: This game is identical to Game_1 except that we change the distribution of v at step 3 of the $\text{Gen}_{\mathcal{L}}$ algorithm. We choose u as a random composite

residue, as before, but we now set $v = u^{-\text{VK}} \cdot v_0^N \bmod N^2$, for a random $v_0 \leftarrow U(\mathbb{Z}_N^*)$. In order to simulate a proof for \mathcal{A} 's unique query, the challenge uses the witness v_0 (which is the N -th root of $u^{\text{VK}} \cdot v \bmod N^2$), so that the N -th root of u is not used anymore. This change is only conceptual since $(u, v) \in (\mathbb{Z}_{N^2}^*)^2$ are still uniformly distributed in the subgroup of N -th residue in $\mathbb{Z}_{N^2}^*$. Hence, we have $\Pr[W_2] = \Pr[W_1]$.

Game₃: This game is like **Game₂** except that we change the distribution of u at step 3 of the $\text{Gen}_{\mathcal{L}}$ algorithm. We now sample u as a random element $u \leftarrow U(\mathbb{Z}_{N^2}^*)$ instead of a composite residue. Under the DCR assumption, this change goes unnoticed and we have $|\Pr[W_2] - \Pr[W_3]| \leq \text{Adv}^{\text{DCR}}(\lambda)$. Note that, for any $\text{VK}^* \neq \text{VK}$, $u^{\text{VK}^*} \cdot v = u^{\text{VK}^* - \text{VK}} \cdot v_0^N \bmod N^2$ is no longer an N -th residue as $|\text{VK}^* - \text{VK}| < \min(p, q)$ implies $\gcd(\text{VK}^* - \text{VK}, N) = 1$.

Game₄: In this game, we change the generation of crs . The challenger now samples $(\text{crs}_0, \tau_0) \leftarrow \text{TrapGen}^{(0)}(\text{Gen}_{\text{par}}^{(0)}(1^\lambda), \mathcal{L}^{\text{DCR}}, \tau_{\mathcal{L}^{\text{DCR}}})$ and $(\text{crs}_1, \tau_1) \leftarrow \text{TrapGen}^{(1)}(\text{Gen}_{\text{par}}^{(1)}(1^\lambda), \mathcal{L}, \tau_{\mathcal{L}})$. Under the CRS indistinguishability of both trapdoor Σ -protocols, the distance $|\Pr[W_3] - \Pr[W_4]| \leq \text{Adv}^{\text{crs},(0)}(\lambda) + \text{Adv}^{\text{crs},(1)}(\lambda)$ is negligible.

However, in this last game, if \mathcal{A} outputs a proof $\bar{\pi}^* = (\text{VK}^*, (a^*, z^*), \text{sig}^*)$ for a false statement x^* such that W_2 is set to 1, we know that neither of the two statements $x^* \in \mathcal{L}$ and $(u^{\text{VK}^*} \cdot v) \in \mathcal{L}^{\text{DCR}}$ is true. This implies that the functions $\text{BadChallenge}^{(0)}(\tau_0, \text{crs}_0, u^{\text{VK}^*}v, a_0^*)$ and $\text{BadChallenge}^{(1)}(\tau_1, \text{crs}_1, x, a_1^*)$ both identify a unique element of $\{0, 1\}^\lambda$. Consequently, if $\bar{\pi}^* = (\text{VK}^*, (a^*, z^*), \text{sig}^*)$ properly verifies, \mathcal{A} was able to come up with $(x^*, a^* = (a_0^*, a_1^*), \text{VK}^*)$ such that

$$\begin{aligned} \text{Hash}(k, (x^*, a^*, \text{VK}^*)) &= \text{BadChallenge}^{(0)}(\tau_0, \text{crs}_0, u^{\text{VK}^*}v, a_0^*) \oplus \\ &\quad \text{BadChallenge}^{(1)}(\tau_1, \text{crs}_1, x^*, a_1^*). \end{aligned}$$

This means that \mathcal{A} broke the correlation intractability of the hash function for the unique-output relation that maps (x^*, a^*, VK^*) to the XOR of both bad challenge functions. Since this relation is efficiently searchable, we can rely on standard assumptions [61] to argue that $\Pr[W_4]$ is negligible, which in turn implies that $\Pr(W_0)$ is negligible.

In particular, the advantage of \mathcal{A} is at most $\text{Adv}^{\text{OTS}}(\lambda) + \text{Adv}^{\text{crs},(0)}(\lambda) + \text{Adv}^{\text{crs},(1)}(\lambda) + \text{Adv}^{\text{DCR}}(\lambda) + \text{Adv}^{\text{CI}}(\lambda)$. \square

C.2 Proof of Theorem 4.3

Proof. To prove the result, we follow the usual strategy of Naor-Yung-based construction. We proceed with a sequence of games, where we call W_i the event that the adversary outputs 1 in **Game_i**.

Game₀: This is the real IND-CCA2 security game with static corruptions. In this game, the adversary \mathcal{A} chooses a bound $B \in \text{poly}(\lambda)$, a number of servers n , a threshold t , a subset $\mathcal{C} \subset [n]$ comprised of $t - 1$ secret players. We assume

w.l.o.g. that $\mathcal{C} = \{1, \dots, t-1\}$. Then, \mathcal{A} obtains the public key pk and secret key shares $\{\text{sk}_i\}_{i \in \mathcal{C}}$. It is then granted access to a partial decryption oracle PartDec . In the challenge phase, it outputs two messages $\text{Msg}_0, \text{Msg}_1$ of the same length ℓ_M^* . The challenger encrypts Msg_0 and sends the challenge ciphertext $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \ell_M^*, \bar{\pi}^*)$ to \mathcal{A} . After another series of partial decryption queries on ciphertexts $\text{ct} \neq \text{ct}^*$, \mathcal{A} outputs a guess b' . We define W_0 to be the event that $b' = 1$.

Game₁: We change the generation of the polynomial $f[X]$. Upon receiving the set of corrupted servers \mathcal{C} , the challenger sets $f(i) \leftarrow U(\mathbb{Z}_{[N_1^{\zeta'+1}/4]})$, for all $i \in \mathcal{C}$. Together with the additional constraint $f(0) = d$, this uniquely determines $f[X]$ (recall that we assumed $|\mathcal{C}| = t-1$). Moreover, the distribution of secret key shares is within distance $t \cdot 2^{-\lambda}$ from the uniform distribution over $(\mathbb{Z}_{N_1^{\zeta'} p_1' q_1'})^{t-1}$, which corresponds to the distribution of $\{\text{sk}_i\}_{i=1}^{t-1}$ in Game_0 . We have $|\Pr[W_1] - \Pr[W_0]| \leq t \cdot 2^{-\lambda}$.

Game₂: We now appeal to the NIZK simulator to generate a common reference string $(\text{crs}_{\mathcal{L}}, \tau_{\mathcal{L}}) \leftarrow \text{Sim}_0(1^\lambda, (\mathcal{L}_{\text{zk}}^{\text{eq-dcr}}, \mathcal{L}_{\text{sound}}^{\text{eq-dcr}}))$ together with its simulation trapdoor $\tau_{\mathcal{L}}$. In the challenge ciphertext $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \ell_M^*, \bar{\pi}^*)$, we now compute $\bar{\pi}^*$ as a simulated proof $\bar{\pi}^* \leftarrow \text{Sim}_1(\tau_{\mathcal{L}}, \text{crs}_{\mathcal{L}}, (\text{ct}_1^*, \text{ct}_2^*, \ell_M^*), |\text{bl}|)$ instead of a real proof. Using the zero-knowledge property of the NIZK, we have $|\Pr[W_1] - \Pr[W_2]| \leq \text{Adv}^{\text{zk}}(\lambda)$, which is negligible.

Game₃: In this game, we modify the partial decryption oracle in the following way. At each query $(j \in [n] \setminus \mathcal{C}, \text{ct} = (\text{ct}_1, \text{ct}_2, \ell_M, \bar{\pi}))$ for which $\bar{\pi}$ properly verifies, the challenger uses the factorization of $N_2 = p_2 q_2$ to decrypt ct_2 and obtain $\widetilde{\text{Msg}}_2 \in \mathbb{Z}_{N_2}^{\zeta}$ using the Damgård-Jurik decryption algorithm. Then, it decodes the resulting $\widetilde{\text{Msg}}_2$ as a pair $(m_2, c_2) \in [-R, R] \times [0, C]$ such that $\widetilde{\text{Msg}}_2 = m_2 \cdot c_2^{-1} \bmod N_2^{\zeta}$ and returns \perp if no such pair exists. Otherwise, it computes $\widetilde{\text{Msg}}_1 = m_2 \cdot c_2^{-1} \bmod N_1^{\zeta}$ and defines

$$\hat{\mu}_0 \triangleq (1 + N_1)^{\widetilde{\text{Msg}}_1} \bmod N_1^{\zeta+1},$$

which allows it to use available shares $\{\text{sk}_i = f(i) \bmod N_1^{\zeta} p_1' q_1'\}_{i \in \mathcal{C}}$ to simulate a partial decryption on behalf of the k -th server as

$$\mu_j = \hat{\mu}_0^{2 \cdot \lambda_{j,0}^{\mathcal{S}}} \cdot \prod_{i \in \mathcal{C}} (\text{ct}_1^{\text{sk}_i})^{2 \cdot \lambda_{j,i}^{\mathcal{S}}} \bmod N_1^{\zeta+1},$$

where $\mathcal{S} \triangleq \mathcal{C} \cup \{0\}$ using the Lagrange coefficients $\lambda_{j,i}^{\mathcal{S}} = \Delta \cdot \prod_{k \in \mathcal{S} \setminus \{i\}} \frac{j-k}{i-k}$ for all $j \in \mathcal{S}$.

We note that, as long as ct_1 decrypts to $\widetilde{\text{Msg}}_1 = m_2 \cdot c_2^{-1} \bmod N_1^{\zeta}$, the simulation is correct since we have $\hat{\mu}_0 = (1 + N_1)^{\widetilde{\text{Msg}}_1} = \text{ct}_1^{f(0)} \bmod N_1^{\zeta+1}$ and thus $\mu_j = \text{ct}_1^{2 \cdot \Delta \cdot f(i)} \bmod N_1^{\zeta+1}$.

By the definition of the language $(\mathcal{L}_{\text{zk}}^{\text{eq-dcr}}, \mathcal{L}_{\text{sound}}^{\text{eq-dcr}})$ in Section 4.1, this modification does not affect \mathcal{A} 's view until the event that \mathcal{A} makes a partial

decryption query $\text{ct} = (\text{ct}_1, \text{ct}_2, \ell_M, \vec{\pi})$ such that $(\text{ct}_1, \text{ct}_2, \ell_M) \notin \mathcal{L}_{\text{sound}}^{\text{eq-dcr}}$. Indeed, $\mathcal{L}_{\text{sound}}^{\text{eq-dcr}}$ is the language of ciphertexts such that ct_1 and ct_2 decrypt to $\widetilde{\text{Msg}}_1 = m \cdot c^{-1} \bmod N_1^\zeta$ and $\widetilde{\text{Msg}}_2 = m \cdot c^{-1} \bmod N_2^\zeta$, for some $m \in [-R, R]$, $c \in [0, C]$. By the soundness of the argument system for $(\mathcal{L}_{\text{zk}}^{\text{eq-dcr}}, \mathcal{L}_{\text{sound}}^{\text{eq-dcr}})$, we have $|\Pr[W_3] - \Pr[W_2]| \leq \mathbf{Adv}^{\text{sound}}(\lambda)$.

Game₄: In this game, we now modify the generation of the challenge ciphertext $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \ell_M^*, \vec{\pi}^*)$ and replace ct_1^* by an encryption of $\text{Msg}_1 \in [0, M]$ (instead of Msg_0 as in the previous games). Recall that, in **Game₃**, the factorization of N_1 was not used and $\vec{\pi}^*$ was simulated without using the random encryption coins. We can thus rely on the semantic security of Damgård-Jurik to argue that $|\Pr[W_4] - \Pr[W_3]| \leq \mathbf{Adv}^{\text{DCR}}(\lambda)$.

Game₅: In this game, we change again the partial decryption oracle and come back to the normal decryption procedure, which uses all secret key shares $\{\text{sk}_i\}_{i \in [n] \setminus \mathcal{S}}$ (we are free to use the factorization of N_1 since we are done with the DCR assumption in $\mathbb{Z}_{N_1^{\zeta+1}}^*$).

As in the transition to **Game₃**, this modification will not affect \mathcal{A} 's view until the event that \mathcal{A} queries the partial decryption of a ciphertext $\text{ct} = (\text{ct}_1, \text{ct}_2, \ell_M, \vec{\pi})$ such that $(\text{ct}_1, \text{ct}_2, \ell_M) \notin \mathcal{L}_{\text{sound}}^{\text{eq-dcr}}$. However, this event would imply that $\vec{\pi}$ breaks the one-time simulation-soundness of the NIZK argument for $(\mathcal{L}_{\text{zk}}^{\text{eq-dcr}}, \mathcal{L}_{\text{sound}}^{\text{eq-dcr}})$. We have $|\Pr[W_5] - \Pr[W_4]| \leq \mathbf{Adv}^{\text{ss}}(\lambda)$.

Game₆: We change again the challenge ciphertext $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \ell_M^*, \vec{\pi}^*)$ and now compute ct_2^* as an encryption of $\text{Msg}_1 \in [0, M]$. Since the factorization of N_2 was not used in **Game₅**, the semantic security of Damgård-Jurik implies that **Game₆** is indistinguishable from **Game₅** under the DCR assumption in $\mathbb{Z}_{N_2^{\zeta+1}}^*$. Concretely, we have $|\Pr[W_6] - \Pr[W_5]| \leq \mathbf{Adv}^{\text{DCR}}(\lambda)$.

Game₇: In this game, we modify again the challenge $\text{ct}^* = (\text{ct}_1^*, \text{ct}_2^*, \ell_M^*, \vec{\pi}^*)$ and now compute $\vec{\pi}^*$ as a real proof (which is computable using the real witnesses since $\vec{\pi}^*$ argues a true statement in **Game₆**) instead of a simulated proof $\vec{\pi}^* \leftarrow \text{Sim}_1(\tau_{\mathcal{L}}, \text{crs}_{\mathcal{L}}, (\text{ct}_1^*, \text{ct}_2^*, \ell_M^*), \text{lbl})$. The NIZK property implies $|\Pr[W_6] - \Pr[W_6]| \leq \mathbf{Adv}^{\text{zk}}(\lambda)$.

We note that **Game₇** (resp. **Game₀**) is the real game, when the challenger's bit is $b = 1$ (resp. $b = 0$). Under the assumptions that DCR holds and that the NIZK argument provides one-time simulation-soundness, we conclude that the threshold system provides IND-CCA2 security under static corruptions. \square

C.3 Application to Voting-Friendly Encryption

When dropping the CCA2 component, the above scheme retains homomorphism for honestly-generated ciphertexts (as long as homomorphic calculations are performed inside some bound), where the denominator c is 1. This is similar to [24], where commitments retain some sort of homomorphic properties for honestly generated commitments. We are in the same setting, except that we consider

decryption. This enables the scheme to be used in voting systems even when ciphertexts are not trusted to be honestly generated, as we now comment.

Our instantiation of Naor-Yung can be used in voting systems with mix-nets or homomorphic tallying. After having checked the validity of a ciphertext, the talliers can only retain the CPA components. In a mix-net, the CPA part can go through shuffles as we preserve re-randomizability and scalar multiplication. Homomorphic tallying often requires more involved proofs before aggregating ciphertexts. In the simplest case, one can just ask that ciphertexts encrypt a 0/1 vote in the standard model. In our case, this additional proof can be appended to ciphertexts, which ensures that there is no “denominator” (i.e., $c = 1$) in the rational encoding in order to preserve the homomorphism.

D On the (Non)-Unboundedness of Groth-Sahai Proofs

At first glance, it is tempting to believe that Groth-Sahai proofs do provide unbounded range proofs. Below, we provide an explanation as to why Groth-Sahai does not give unbounded range proofs, except with a poor rate $O(1/\lambda)$.

A first attempt is to use a simple bit decomposition plus a NIZK proving that all commitments open to bits (and that the sum of the $b_i \cdot 2^i$ is the original value), with the hope of handling integers of arbitrary size (with a CRS just containing the commitment key). However, this approach does not give what we want. Suppose we want to commit to an integer x and later prove that $x \in [0, 2^L]$. With Groth-Sahai, we can only do it when $L < \log p$, where p is the group order, since the prover is only committed to $x \bmod p$ (nothing is guaranteed about x over \mathbb{Z} if $L > \log p$). When we break the initial value x into bits $\{x_i\}_{i=0}^L$, we can only argue about $\sum_{i=1}^L x_i \cdot 2^{i-1} \bmod p$. Here, the CRS allows choosing a larger ζ depending on the range width.”

Groth-Sahai can achieve unboundedness by directly committing to the bits of x . However, this is rather wasteful as the ratio between the sizes of x and the commitment is $O(1/\lambda)$. To address this issue, a possible attempt¹⁵ is as follows: “If the range is super large, say $[a, b]$ such that $a - b$ is about 2^{λ^2} , then one can commit to $x \in [a, b]$ through its decomposition in base p (where p is the order of the group where Groth-Sahai is done, with p about 2^λ). Then, only the “least significant p -digit” of x must be decomposed bitwise. Overall, the commitment size has length $O(\lambda^2)$, which gives a constant rate.”

First, this comes at the cost of a pretty bad rate $O(1/\lambda)$ as the proof requires to proceed bit-by-bit, *even with a p -ary decomposition* with $p > 2$ (as we explain below). Moreover, if the prover wants to prove other statements about committed integers than range membership, the only solution is to argue about bits, which is wasteful in terms of space.

In addition, we contend that Groth-Sahai does not achieve constant rate using the p -ary decomposition of committed integers. In this case, the commitment does have constant rate, but the proof does *not*. The reason is that only proving

¹⁵ We are quoting an argument from a reviewer in an earlier submission.

a statement about one digit does not suffice. With $p = 2$ (for example, but the argument works for any $p > 2$) and $n = 7$, let $B = 1000001$ and let a committed integer $x = 1100001$ which is larger than B . In this case, only considering the most significant non-zero bit does not provide soundness when we want to prove $x < B$. The above approach works when B is a power of 2. However, to our knowledge, going from the case $B = 2^n$ to an arbitrary B requires the commitment to be homomorphic over \mathbb{Z} . Indeed, in order to prove $x \in [A, B]$, we need to set $n \in \mathbb{N}$ such that $2^n > |B - A|$ and prove $(x \in [A, A + 2^n]) \wedge (x \in [B - 2^n, B])$, which reduces to proving $(x - A \in [0, 2^n]) \wedge (B - x \in [0, 2^n])$. When we use GS commitments to the binary decomposition of integers, we do not have the homomorphism over \mathbb{Z} . For a general B , the bit-by-bit approach requires the prover to prove statements about *all* the bits of committed integers.

When using the p -ary decomposition (where p is the group order in Groth-Sahai), the prover has to consider **all** individual digits to make a comparison because there is no additive homomorphism over p^n . If the group order p is smaller than B , the only solution we see is to prove that $B - x$ is positive by emulating a p -ary subtraction **with carries** (this is the approach taken in [49] with $p = 2$). This requires to prove that each carry is done correctly. For each p -ary digit, it basically requires a commitment to a bit that indicates whether a carry takes place or not. In addition, for each p -ary digit, the prover has to perform bit-by-bit comparisons and prove in ZK that these comparisons are done correctly (in the standard model). Hence, if we commit to a λ^2 -bit integer using λ p -ary digits, the commitment has constant rate but we end up with a proof of rate $O(1/\lambda)$. At the end of the day, choosing a p -ary decomposition with $p > 2$ does *not* give constant rate.

More precisely, let us assess the rate with the Groth-Sahai non-interactive approach, where p is the prime order of the bilinear group. Let $x \in [0, B]$ and $p^{k-1} \leq B < p^k$, for some integer k , and where only p is fixed by the CRS. First, we focus on the rate between the input size and the proof size assuming that $n := |p|$. To preserve zero-knowledge, we must encode x using at least $m := |B|$ bits. The commitment com to x is a dual-mode commitment consisting of a ciphertext of two first-source-group elements under SXDH, we thus have the rate

$$\frac{|\text{com}|}{|B|} = \frac{2kn}{m} \geq \frac{2kn}{kn} = 2$$

between the commitment size and the input size. As we do not see any other solution than proving comparison on encrypted inputs than decomposing x (or $B - x$) bit-by-bit, the resulting range proof π_{GS} for x consists of at least m commitments. Even without counting the “proof part” of Groth-Sahai (i.e., we only count the commitment to bits without the longer proofs that committed values are actually bits nor the proofs that carries are correctly computed), the rate between the total proof size and the commitment size already amounts to

$$\frac{|\pi_{GS}|}{|\text{com}|} \geq \frac{m2kn}{2kn} = m,$$

so that we have $|\pi_{GS}|/|B| \geq 2|B| = \Omega(\lambda^e)$ for very large ranges. In our setting, the ratio is almost constant (roughly $1/2 - 1/(\zeta \cdot \omega(1))$ by Lemma 3.2) when ζ is large.