

Succinctly Verifiable Computation over Additively-Homomorphically Encrypted Data: Making Privacy-Preserving Blueprints Practical

Scott Griffy¹, Markulf Kohlweiss², Anna Lysyanskaya¹, and Meghna Sengupta³

¹ Brown University, {anna_lysyanskaya,scott_griffy}@brown.edu

² University of Edinburgh and IOG, Edinburgh, markulf.kohlweiss@ed.ac.uk

³ University of Edinburgh, M.Sengupta-1@ed.ac.uk

Abstract. With additively homomorphic encryption (AHE), one can compute, from input ciphertexts $\text{Enc}(x_1), \dots, \text{Enc}(x_n)$, and additional inputs y_1, \dots, y_k , a ciphertext $c_f = \text{Enc}(f(x_1, \dots, x_n, y_1, \dots, y_k))$ for any polynomial f in which each monomial has total degree at most 1 in the x -variables (but with arbitrary degree in the known y -variables). For AHE that satisfies a set of natural requirements, we give a zero-knowledge proof system for showing that a ciphertext c_f is the result of homomorphically evaluating f on ciphertexts $(c_1, \dots, c_n) = (\text{Enc}(x_1), \dots, \text{Enc}(x_n))$ and private inputs y_1, \dots, y_k that correspond to commitments C_1, \dots, C_k where the encrypted values, x_1, \dots, x_n are unknown to the prover. Our proofs are *succinct*, i.e., their size is independent of the number of ciphertexts n , and is instead $O(k \log d)$ where k is the number of private inputs, and d is the maximum degree of any variable in f .

We give two ways of instantiating this framework: with ElGamal-based encryption (under the DDH assumption) and with a variant of the Camenisch-Shoup cryptosystem (under the DCR and Strong RSA assumptions). Both yield proof systems where computing and verifying the proof takes a comparable amount of time to homomorphically evaluating f .

Next, we show that our framework yields a dramatically improved privacy-preserving blueprint (PPB) system. Introduced by Kohlweiss, Lysyanskaya, and Nguyen (Eurocrypt'23), an f -PPB system allows an auditor with secret input x to create a public encoding pk of the function $f(x, \cdot)$ that reveals nothing about x . Yet, it allows a user to compute an encoding, or escrow Z , of the value $f(x, y)$ on input the user's private data y corresponding to a commitment C_y ; Z will verifiably correspond to the commitment C_y . The auditor will be able to recover $f(x, y)$ from Z , but will learn no other information about y . For example, if f is the watchlist function where $f(x, y)$ outputs y only in the event that y is on the list x , then an f -PPB allows the auditor to trace watchlisted users in an otherwise anonymous system.

Using our succinct zero-knowledge proof system for additively homomorphic computation we achieve the following results: (1) We provide efficient schemes for a bigger class of functions f ; for example, we show how to realize f that would allow the auditor to trace private payment

transactions of a criminal suspect which was previously not efficient. (2)
For the watchlist and related functions, we reduce the size of the escrow Z from linear in the size of the auditor's input x , to logarithmic. Additionally, we define and satisfy a stronger notion of security for f -PPBs, where a malicious auditor cannot frame a user in a transaction in which the user was not involved in.

Table of Contents

1	Introduction.....	4
1.1	Our Framework for Verifiable Computation	7
1.2	Non-Frameability and Why It Matters	10
1.3	Related Work and Efficiency Analysis	11
2	Preliminaries	12
2.1	Zero-knowledge Proofs of Knowledge	13
2.2	Additively Homomorphic Encryption	14
2.3	Privacy Preserving f -Blueprint Schemes (PPBs).....	15
3	Succinct Proofs for Verifiable Secure Computation on Additively- Homomorphic Ciphertexts	16
3.1	Basic Building Blocks	17
3.2	Efficient Proof System for R_f for $k = 1$	19
3.3	Proof System for Multivariate Polynomials	22
4	Instantiations of Commitments to Additively-Homomorphic Ciphertexts	22
4.1	Encryption Schemes	23
4.2	Commitments to $ QR_{n^2} $ and Camenisch-Shoup Ciphertexts.....	24
5	Applications of our Framework to Privacy-Preserving Blueprints	27
5.1	Non-Frameable Privacy-Preserving Blueprints	28
5.2	Instantiation of Consistent HEC Schemes.....	32
5.3	Efficient Instantiation of HEC Evaluation Proof Ψ_2	34
A	Discussion on Non-frameability vs. Deniability	41
B	Full Definitions for Privacy Preserving f -Blueprint Schemes	42
C	Additional preliminaries	46
C.1	Motivation for BB-PSL	46
C.2	Useful Lemmas for Composite-Order Groups.....	47
C.3	More <i>eqrep</i> relations and constructions	48
D	Additional HEC definitions, constructions, and proofs	54
D.1	Security Properties of HEC Scheme	54
D.2	Multi-attribute HEC Scheme	55
D.3	Constructions of HEC Schemes.....	56
E	Additional Constructions and Proofs of Security for R_P , R_f , and Ψ_2 ..	60
E.1	Proof of security for the R_P proof system	60
E.2	Proofs for the Ψ_2 proof system	62
E.3	Construction of NIZKs in Ψ_2 Proof Scheme	63
E.4	Construction and proof of security for the R_f proof system	65
F	Commitments to Ciphertexts	66
F.1	Additional Notes and Proofs for Simplified Camenisch-Shoup ...	66
F.2	Security Proof of Damgård-Fujisaki Commitments for $\mathcal{G} = \mathbb{Z}_{n^2}$..	68
F.3	Commitments to \mathbb{G}_p Elements and ElGamal Ciphertexts	71
F.4	Proofs of Hiding and Binding for $ QR_{n^2} $ -commitments in Fig. 4.3	75

F.5	Construction of Commitments to Camenisch-Shoup Ciphertexts	76
F.6	Proofs for Commitments to Camenisch-Shoup Ciphertexts	76

1 Introduction

The need for privacy preserving blueprints. Not all citizens are lawful and not all governments democratic, thus cryptographers have developed powerful tools to trade off our fundamental need to protect our personal privacy with the legitimate needs of systems and governments to enforce rules and laws and to regulate finance. Among these tools, anonymous credentials [Cha90,LRSW99,CL01], [Lys02,CL02,CV02,CL04,BCL04,BL13,HS21,RWGM23,TZ23,HSS23] and related technologies such as e-cash [CFN90] are prominent examples: such systems allow a user with a cryptographic commitment C_y to his data y to prove that y is somehow certified by some authority or authorities; in the case of anonymous payments, they further allow to prove that a payment transaction based on the user's private data y was executed correctly.

In a recent paper, Kohlweiss, Lysyanskaya and Nguyen (KLN) [KLN23] added *privacy-preserving blueprints* (PPBs) to the repertoire of cryptographic algorithms to depolarise privacy and accountability. In an f -PPB system, the goal is to allow an authorized auditor to learn $f(x, y)$ where x is the auditor's secret input that's fixed once and for all, and y is a user's secret input to a transaction; if a PPB system is used in tandem with an anonymous credential system, y can include meaningful information about the user's identity. Via an appropriate choice of f , an f -PPB system makes it possible to perform audits of the system while leaking no information other than what's leaked by f . For example, for x representing a watchlist of suspected criminals, let $f_{watchlist}$ be defined as follows: $f_{watchlist}(x, y) = y$ if y is on the list x , and \perp otherwise. An $f_{watchlist}$ -PPB would allow the auditor to trace all of the suspects' transactions, but none of the transactions of other people. A PPB further requires that the secret x corresponds to a publicly known commitment C_x that can be further certified by an external party, so that a malicious auditor cannot make up x at will.

In a PPB system, first, the auditor sets up his public key \mathbf{pk} and secret key \mathbf{sk} on input his secret x and a commitment C_x to x for which the auditor knows the opening (and which may be signed by an external validator who certifies that x is a correct input). A PPB includes a public verification procedure $\text{VerPK}(\mathbf{pk}, C_x)$ for ensuring that \mathbf{pk} corresponds to the commitment C_x . Now the system is ready for blueprinting transactions; there is no limit on the number of such transactions. In a transaction, a user with secret input y and a commitment C_y to y to which the user knows the opening r (and which meaningfully corresponds to some information about this user, for example validated via an anonymous credential system), computes the escrow $Z = \text{Escrow}(\mathbf{pk}, y, r)$ of y under \mathbf{pk} . A PPB includes a public verification procedure $\text{VerEscrow}(\mathbf{pk}, C_y, Z)$ for ensuring that Z corresponds to \mathbf{pk} and C_y . Finally, using \mathbf{sk} , the auditor runs the decryption algorithm to recover $z = f(x, y)$ from Z . The reason that it is

called a privacy-preserving *blueprint* is that we can think of pk as a “blueprint” of the function $f(x, \cdot)$ of the user’s y .

An f -PPB is realizable for any efficiently computable function f from either fully homomorphic encryption (FHE) or non-interactive secure computation (NISC) [KLN23] by representing the function as a circuit. However, this general approach is not suitable for practical use. KLN additionally gave a more practical construction of $f_{\text{watchlist}}$ -PPB from the ElGamal cryptosystem and proof systems about discrete logarithm relations in the random-oracle model, though even with this more practical construction, the size of their escrow is linear in the size of the watchlist.

As we argue below, this linear size is not sufficient to be useful in practice. To bridge this gap, we develop a commit-and-prove framework for working with additively-homomorphically encrypted data. Additively homomorphic encryption (Definition 4) allows one to compute, on input ciphertexts c_1, \dots, c_n that encrypt x_1, \dots, x_n , and additional inputs y_1, \dots, y_k , the value $f(x_1, \dots, x_n, y_1, \dots, y_k)$ for any polynomial f in which each monomial has total degree at most 1 in the x -variables (but can be arbitrary in the y -variables).

Our first contribution: A modular framework for succinct verifiable secure computation on additively-homomorphically encrypted data.

In this paper, we give a non-interactive zero-knowledge proof system (in the random-oracle model) for showing that a ciphertext c_{out} is the result of homomorphically evaluating f on c_1, \dots, c_n and private inputs y_1, \dots, y_k that correspond to commitments C_1, \dots, C_k . Our proof system, described in Section 3 outputs *succinct* proofs, i.e. their size is $O(k \log d)$ where k is the number of private inputs, and d is an upper bound on the degree of any variable in f ; note that the size of the proof is independent on the number n of the x -variables. For the proof to be efficient, we must include “intermediate” ciphertexts in the proof that allows the verifier to follow along to be convinced of the final evaluation. Thus, to protect these intermediate ciphertexts from being decrypted, only commitments to them will be created as part of the proof; this is the essence of the “commit-and-prove” approach which allows our proofs to be very modular. We begin (in Section 3.1) with the “commit” part of this approach by defining *commitments to ciphertexts*; then Sections 3.2 and 3.3 contain the proof system itself. We give two different practical instantiations of this framework: one under the DDH assumption (using the ElGamal cryptosystem) and the other under the Strong RSA and Decisional Composite Residuosity assumptions (using the Camenisch-Shoup cryptosystem): Section 4.1 is dedicated to the description of these additively homomorphic cryptosystems, while Section 4.2 constructs commitments to their respective ciphertexts and the basic proof systems that serve as building blocks for the framework.

Our second contribution: Realizing PPBs for central bank digital currencies (f_{CBDC} -PPBs). Since the KLN paper first appeared, privacy-preserving blueprints received some attention in the civil liberties discourse [Sta23] because (among other things) of the following motivating application to central bank digital currencies (CBDCs): suppose that the auditor’s input x is a list of

suspected financial criminals’ unique identifiers. Suppose a user’s input y contains this user’s unique identifier y_{id} as well as seed y_{seed} from which all of this user’s e-coins’ serial numbers are generated. This is consistent with, e.g., compact e-cash [CHL05] and related schemes [CHL06,CHK⁺06,KKS22,TBA⁺22], including those proposed specifically for the CBDC application [KKS22,TBA⁺22]. The function f is as follows: $f(x, y) = y$ if $y_{id} \in x$, and \perp otherwise. A PPB with these properties will allow the auditor to not only identify that a transaction was carried out by a suspect, but also to recover the seed y_{seed} and trace all of the user’s transactions, even as the rest of the users of the systems’ privacy is protected. This application to anonymous payments is attractive to those who advocate that a CBDC can be privacy-preserving even while enabling lawful investigations. Unfortunately, the alternative to yielding ground on this to law enforcement is that central banks throughout the world would adopt a CBDC that provides no privacy — even from third-party observers — to individuals, in the name of compliance with law enforcement.⁴

Recently, [DEF⁺24] proposed an updatable PPB scheme with a scoring-based watchlist and dynamic user-risk scores, similar to banks’ anti-money laundering due diligence. However, for the CBDC application, $f_{watchlist}$ is not the right function. Instead, we need $f_{CBDC}(x, y) = y$ if $y = (y_{id}, y_{seed})$, and $y_{id} \in x$. KLN give a practical construction that works for $f_{watchlist}$ but not for f_{CBDC} , because of their use of ElGamal encryption. Instead of recovering y , the auditor in their construction can only recover g^y where g is a generator of a group in which the discrete logarithm problem is hard. From g^y it is possible to recover y by brute-force search if only a small number of bits of y are still unknown; but it wouldn’t be possible to recover y_{seed} , since the size of a pseudorandom seed must be too large to allow brute-force search. Here, we give a construction for the correct f . Let $f(x, y) = y$ if $y = (y_1, y_2)$, and $y_1 \in x$, and \perp otherwise. We give a practical instantiation of a f -PPB construction (Section 5.2). By “practical”, we mean that it can be instantiated efficiently using proof systems for discrete logarithm relations in the random-oracle model.

The existence of a practical cryptographic system that can provide this tracing capability in a way that is transparent to citizens who, even if they shouldn’t know who is on the suspect list, can still see the size of the list and the fact that there was a lawfully obtained warrant for placing a person on it, would strike a reasonable balance, and, as a result, may sway the policy conversation (in which law enforcement voices are often louder than those of privacy advocates) in favor of using more anonymous systems for CBDCs.

Our third contribution: Exponential improvement in the size of the escrow Z . The KLN approach is also not good enough for either f_{CBDC} -PPBs or even $f_{watchlist}$ -PPBs because we expect the watchlist x to be quite large. In the KLN construction, the size of the escrow Z was linear in the size of the

⁴ For example, the analysis of CBDC design choices provided by the White House [Gov22] is lukewarm on using anonymous systems employing zero-knowledge proofs for that reason. See page 14 and page 17 of [Gov22].

watchlist x . Using the fact that our framework produces succinct proofs, we give a substantial improvement:

We give practical constructions of a f_{CBDC} -PPB and a $f_{watchlist}$ -PPB where the size of Z is logarithmic in the size of x ; this is achieved because the proof system we use in the construction (in Section 5.3) uses our succinct approach (i.e. our first contribution described above).

Our fourth contribution: Stronger security. The KLN definition of security [KLN23] does not rule out that a malicious auditor would be able to produce pk , sk , C_y and Z such that the decryption algorithm will output $z \neq f(x, y)$. In Sect. 1.2, we discuss how the KLN construction of $f_{watchlist}$ -PPB allowed for a “framing” attack: a malicious auditor causing an escrow to decrypt to the identity of an honest user y who is not a party to the transaction. Addressing these security issues using our new proof framework from our first contribution and the reworked functionality is our fourth and final contribution.

We improve the definition of security of PPB to that of *non-frameable* PPB: we add the requirement that the decryption algorithm’s output be publicly verifiable. In Section 5.1, we present this improved definition. Our constructions (which are also presented in Section 5.1) achieve non-frameability.

Summary of how this paper is organized. In Section 1.1 below we give a more detailed overview of our techniques for achieving verifiable computation over additively homomorphically encrypted data, and why they lead to an efficient construction of f_{CBDC} -PPBs. In Section 1.2 we explain why and how we improved the definition of privacy-preserving blueprints to incorporate non-frameability. To conclude the introduction, in Section 1.3, we review related work and provide efficiency analysis and comparison with KLN.

After going over the preliminaries in Section 2, we dive into our commit-and-prove framework in Section 3: In Section 3.1, we go over the “commit” part, and in Sections 3.2 and 3.3, over the “prove” part. Section 4 explains how to adapt the ElGamal and the Camenisch-Shoup cryptosystems; the (adapted) additively homomorphic encryption schemes are given in Section 4.1; while the commitment schemes for committing to ciphertexts and related proof systems that fit the requirements of our framework are given in Sections 4.2 (for Camenisch-Shoup) and Appendix F.3 (for ElGamal). Armed with these tools, in Section 5 we define and realize non-frameable privacy-preserving blueprints.

1.1 Our Framework for Verifiable Computation

Let us focus on a concrete example. At a high level, a f_{CBDC} -PPB scheme will work as follows: The auditor will first find the coefficients of the polynomial $P(\chi) = a_0 + a_1\chi + \dots + a_n\chi^n$ of degree n whose roots are values on the list x , and it will output a public key pk of an encryption scheme, as well as the encryptions of the coefficients of P ; i.e. $X = (\text{pk}, \overline{a_0}_{\text{pk}}, \dots, \overline{a_n}_{\text{pk}})$, where \overline{m}_{pk} denotes an encryption of a message m under the public key pk (and we drop the subscript when clear from the context). Let $f(a_0, \dots, a_n, y_{id}, y, s) = (s \sum_{i=0}^n a_i y_{id}^i) + y$.

Note that if $f_{CBDC}(x, y_{id}, y) \neq \perp$, then $f(\mathbf{a}, y_{id}, y, s) = y$; else, if the user picks s uniformly at random, then $f(\mathbf{a}, y_{id}, y, s)$ is also random. Thus, the goal is for the user to compute c_f , an encryption of $f(a_0, \dots, a_n, y_{id}, y, s)$, from X .

If the underlying encryption scheme is additively homomorphic, then $c_f = \overline{f(a_0, \dots, a_n, y_{id}, y, s)}$ can be computed using homomorphic addition: Let the symbol ‘ \oplus ’ denote the homomorphic operation on ciphertexts, and let \odot denote multiplying a ciphertext by a scalar. Then $c_f = (\bigoplus_{i=0}^n (s y_{id}^i) \odot \overline{a_i}) \oplus \overline{y}$. We also need the user to compute a zero-knowledge proof that c_f was computed correctly from X and the user’s secret inputs s , y and y_{id} that correspond to commitments C_s , C_y and $C_{y_{id}}$. While general-purpose ZK proof systems can be used here, a proof system designed hand-in-hand with the underlying encryption scheme can take advantage of efficient Σ -protocols and impose only a minimal overhead over encryption; the classical results on efficient multi-party computation of Cramer, Damgård and Nielsen [CDN01] serve as the inspiration for this approach.

We suggest a modular, commit-and-prove [BCF⁺] approach for constructing a proof that a given ciphertext is the result of computing on additively-homomorphically encrypted data. For example, here the output ciphertext c_f is the result of applying a series of homomorphic operations, starting with the input ciphertexts $\{\overline{a_i}\}$ and the user’s inputs. In order to prove correctness of c_f in our framework, one forms commitments to the intermediate steps of this computation (for example, the intermediate ciphertexts $\overline{a_i} \odot y_{id}^i$) and proves that each of these intermediate steps was carried out correctly.

Thus, our main new building block is an additively homomorphic encryption scheme equipped with (1) a cryptographic commitment scheme for committing to ciphertexts; and (2) proof systems for proving properties of committed ciphertexts, such as the property that a committed ciphertext c was obtained from committed ciphertexts c_1 and c_2 , along with a committed scalar a , as follows: $c = c_1 \oplus (c_2 \odot a)$. (See Sect. 3.1 for the more formal treatment.)

Next, let us explain how to instantiate this framework with the ElGamal cryptosystem. Let G be a group of prime order q with generator g_1 ; an ElGamal public key is a group element g_2 ; an encryption of $M \in G$ is $(g_1^r, g_2^r M)$ where for random $r \in \mathbb{Z}_q$. ElGamal is not, strictly speaking, an additively homomorphic encryption scheme, but a multiplicatively homomorphic one: $(g_1^r, g_2^r M) \oplus (g_1^{r'}, g_2^{r'} M') = (g_1^{r+r'}, g_2^{r+r'} M M')$. However, we can define a “lifted” ElGamal cryptosystem: to encrypt the message m , use the ElGamal cryptosystem to encrypt g_1^m ; i.e. $\overline{m} = (g_1^r, g_2^r g_1^m)$. The problem is that, instead of outputting m , the decryption algorithm outputs g_1^m ; converting it to m requires that m come from a small space, so that it can be found via brute-force search; we call this flavor of encryption “semi”-encryption. Still, for some applications (such as realizing $f_{watchlist}$ -PPBs), this is good enough.

Our techniques for achieving succinct proofs. The naïve way for computing a proof π of correctness of c_f is to form a commitment to the ciphertext that is the result of each intermediate step in the computation (for example, the values $\overline{a_i} \odot y_{id}^i$ in the example above), meaning that the size of the proof will need to be linear

in the degree d of the polynomial f (and in the description of the polynomial altogether). To reduce the dependence on the degree from d to $O(\log d)$, we use a degree reduction technique inspired by the sum-check protocol of Lund, Fortnow, Karloff and Nisan [LFKN92]. The sum-check protocol was used more recently in cryptography by Goldwasser, Kalai and Rothblum [GKR08] and follow-up work on “proofs for Muggles” [XZZ⁺19,ZLW⁺21]. Pietrzak [Pie19,HHKP23] was the first to use it to halve the degree of a polynomial (as we do) rather than to eliminate a linear variable as in the other cited work. As far as we know, our paper is the first time that this technique is used in order to prove correctness of commit-and-prove-style computation on encrypted data.⁵ We compare our technique to more works in Sec. 1.3.

The overall idea, described in Section 3.2 (and generalized to the multivariate case in Section 3.3), is to recursively halve the degree of the polynomial. Suppose that we need to prove that a ciphertext $c_f = \boxed{f(x_1, \dots, x_n, y_1, \dots, y_k)}$; the prover and verifier both know $\{\boxed{x_i}\}_{i \in [n]}$; further, the prover knows y_1, \dots, y_k (and thus can compute c_f) while the verifier knows just the corresponding commitments $\{C_{y_i} = \text{Com}(y_i; r_i)\}$. Suppose the degree of y_1 in f is d . The recursive step is to reduce the proof of this statement to the proof that another ciphertext $c_{f'}$ is an encryption of $f'(x_1, \dots, x_n, y_1, \dots, y_k)$, where in f' the degree of y_1 is $d/2$. This can be accomplished using the Schwartz-Zippel lemma: we obtain f' from f by replacing each occurrence of $y_1^{d/2}$ with a random scalar α ; in the interactive version of the sum-check protocol α would be chosen by the verifier, but here it is chosen by the random oracle. It is important that the ciphertext $c_{f'}$ used in the recursive step not be given to the verifier in the clear; otherwise, it will leak information to the adversary who knows the decryption key. Instead, our proof system works for *committed* ciphertexts.

To obtain a commitment to an ElGamal ciphertext $\boxed{a} = (A, A')$, we first extend Pedersen commitments (with generators g and h) to commit to group elements. To commit to A , we sample $s_A, r_A \leftarrow \mathbb{Z}_q$ and the commitment is $C_A = (C_{A,1}, C_{A,2}) = (Ag^{s_A}, g^{s_A}h^{r_A})$; similarly, we can form a commitment $C_{A'} = (C_{A',1}, C_{A',2})$. Thus, a commitment to \boxed{a} is $C_{\boxed{a}} = (C_A, C_{A'})$. It is easy to see that this commitment scheme has convenient homomorphic properties: if ‘ $*$ ’ denotes applying the group operation componentwise, then $C_{\boxed{a}} * C_{\boxed{b}} = C_{\boxed{a+b}}$. As shown in Sect. 4, this allows for efficient proof systems for properties of committed ciphertexts needed for our framework. Additionally, we show in Sect. 4 that our framework can also be instantiated, under the Paillier assumption, with a semantically secure variant of the Camenisch-Shoup cryptosystem [CS03].

Why f_{CBDC} -PPB was not achievable in KLN. KLN’s limitation was that it used lifted ElGamal, and thus, in the event that the user was on the watchlist, the decryption algorithm was only able to recover g^y from the escrow, rather than y in the clear. As explained earlier, this is not good enough if y comes

⁵ Previous work [BG13] used a completely different technique to give a succinct proof that a committed value corresponds to the evaluation of a polynomial, but with the important distinction that the polynomial was known to both Prover and Verifier.

from a large enough domain (for example if it contains a seed for a PRF) and cannot be brute-force-searched. The Camenisch-Shoup based instantiation of the framework we just discussed allows the decryption algorithm to recover y , which yields f_{CBDC} -blueprints. It turns out that the ElGamal-based instantiation can work as well (with some efficiency limitations), if we split the payment tracing seed into sufficiently small chunks, see Appx. D.2.

1.2 Non-Frameability and Why It Matters

Our additional contribution to privacy-preserving blueprints is an additional property — non-frameability, — and our constructions satisfy it. The concept of non-frameability was first introduced in the work of Camenisch [Cam97]. The paper introduced it for the group signature scheme setting as the property that the manager (even if they collude with a group member) cannot falsely accuse group members. Subsequently, Bellare, Shi and Zhang [BSZ05] formalized the property and called it Non-frameability - again for group signature schemes.

At a high-level there are similarities with the property of non-frameability as we define it and as defined by [BSZ05]. Both properties require that if some authority (the opener in the case of [BSZ05] and the auditor in our case) wants to prove that a user took some action (signing a message in the case of [BSZ05] and authenticating themselves in an anonymous credential scheme in the case of blueprints) they must provide verifiable proof. One difference between the schemes is that in [BSZ05] the opener traces any user indiscriminately. In our case, the auditor’s functionality is not "trace" but the function f . (In the case of watchlists, that means the auditor can trace iff the user is on the watchlist.) Also, a group signature scheme provides tracing for group members who are signing messages, whereas in blueprints, the functionality is to trace users who are using an anonymous credential scheme, which does not imply that these traceable users sign any messages. Thus, it is not trivial to construct blueprints from the group signature scheme in [BSZ05].

The watchlist PPB scheme of [KLN23] is frameable, i.e., a malicious auditor can collude with a malicious user to produce Z that will decrypt to the identity of an honest user who was not a party to the transaction (and who may or may not be on the watchlist). The gist of their scheme is that pk includes encrypted coefficients of a polynomial P such that $P(y) = 0$ if and only if y is on the watchlist x . The escrow $Z = (\hat{Z}, \pi)$ produced by the user whose identity is y consists of the encryption \hat{Z} of $rP(y) + y$ for a random r chosen by the user, as well as a proof π that indeed \hat{Z} was computed correctly. In order to frame the user with identity y^* , a malicious user whose identity is y and to whom the coefficients of the polynomial P are known (as would be the case if the auditor is malicious) needs to solve for r^* in the $r^*P(y) + y = y^*$, and will produce an escrow $Z = (\hat{Z}, \pi)$ by following the original algorithm, but just using $r = r^*$.

This attack is outside the KLN security model, and therefore does not contradict their security analysis (which is correct). One could also argue that frameability, also known as deniability, can be a feature and not a bug. We discuss this at greater length in Appx. A.

In Sect. 5.1, we improve the KLN definition of privacy-preserving blueprints by incorporating non-frameability. The decryption algorithm must now produce a proof π_z of correct decryption, and a new algorithm `Judge` verifies this proof. The proof π_z is important when the auditor’s output is used as evidence in legal proceedings⁶ or as input in a smart contract, e.g., an Ethereum Eigenlayer slashing operation or crime restitution.

In order to obtain a practical non-frameable f -PPB for the watchlist function, we modify the KLN construction as follows: our `Escrow` algorithm will output (\hat{Z}, \hat{Z}', π) , where \hat{Z} is an encryption of $rP(y) + y$ (just as before), and the additional value \hat{Z}' is an encryption of $r'P(y)$, while, as before, the proof π is to ensure that \hat{Z} and \hat{Z}' were computed correctly. If π verifies, the decryption algorithm will decrypt \hat{Z} iff \hat{Z}' decrypts to 0; it will output \perp otherwise. Our succinct proofs are compatible with this non-framing construction.

1.3 Related Work and Efficiency Analysis

Freedman, Nissim, and Pinkas (FNP) [FNP04] were the first to give a protocol for the evaluating an encrypted polynomial. Unlike here, the evaluator in their work was not committed to a particular input y on which to evaluate it; it only needed to ensure that some y exists that makes the evaluation correct. In our scheme, the user commits to a y before the protocol starts and must use this y throughout the protocol, making our proof system much more involved. Further, because of the hash functions needed in their construction [FNP04], proving the correct y was used would be expensive. Despite recent work in making hash functions verifiable [GKR⁺21] it is still best to avoid proving hash functions in zero knowledge if possible. FNP initiated the study on secure set intersection (PSI) which is by now an extremely well-studied [CMdG⁺21, CM20, RS21, GPR⁺21] [CRR21, RR22] special case of secure two-party computation. Our framework can be seen as a building block for verifiable PSI [KMRS14, ATD16, JWP22], since verifiable evaluation of encrypted polynomials is a subroutine in many of these protocols.

Recent years have seen an explosion of techniques for zero-knowledge proof systems [BMM⁺21, CBBZ23, GLS⁺23, WHV24, BFK⁺24]; many of these are for general circuits, but especially worthy of comparison to our work are those of them that, like us, take advantage of efficient Σ -protocols for algebraic relations over committed values and, like us, also achieve succinctness [BBB⁺18, ACC⁺22]. In general, the main difference of our work from these is that our framework is suitable for verifiable computation on encrypted data, which is a scenario to which these cited works do not directly apply.

For example, bulletproofs [BCC⁺16, BBB⁺18, BMM⁺21, LMR19] are an especially promising way to achieve a succinct proof of correct computation over

⁶ Interestingly, this is currently rarely the case for existing investigations employing mass or targeted surveillance. Instead, law enforcement follow a complicated process of parallel construction where not always lawfully attained evidence is used to inform a lawful investigation [Boy].

linear data. To the best of our knowledge, Bünz, Maller, Mishra, Tyagi, and Vesely [BMM⁺21] currently prove the relation that resembles ours most closely: their relations are of the form $\prod_{i \in [d]} A_i^{a_i}$ given vector commitments to group elements $\{A_i\}_{i \in [d]}$ and scalars $\{a_i\}_{i \in [d]}$. At a high-level, if we set $A_i = \text{Enc}(x_i)$ and $a_i = y^i$, then bulletproofs could be used to prove the homomorphic operation on these encryptions, $\prod_{i \in [d]} \text{Enc}(x_i)^{y^i}$ which could be useful for blueprints. Unfortunately, their adaptation of bulletproofs requires the use of pairings. Additionally, using bulletproofs would require a large proof to verify a vector commitment to $\{y^i\}_{i \in [d]}$. Our solution only requires verifying commitments to $\{y^{2^i}\}_{i \in [\log(d)]}$ which is much easier to prove. While the technique in [BMM⁺21] achieves a logarithmic verifier complexity, this desirable verifier complexity requires a trusted setup and the CRS cannot be sampled in the ROM. Our solution achieves linear verifier complexity (similar to [LMR19]) and does not require a trusted setup. Their technique also requires the inversion of scalars, which is not possible in \mathbb{Z}_n without knowing the factorization of n (for example, computing x^{-1} such that $(g^{x*x^{-1}} = g \in \mathbb{Z}_n)$). Thus, bulletproofs do not apply directly to our results when using the Camenisch-Shoup cryptosystem and require more expensive operations in the ElGamal setting (since our proofs do not require pairings).⁷

Bhadauria, Hazay, Venkitasubramaniam, Wu, and Zhang [BHV⁺23] provide a way for a prover to compute and prove the encryption of an evaluation of a polynomial without knowing the polynomial. Where our work differs is that their proof system achieves zero-knowledge only in the event that the secret key of the encryption scheme is unknown to the adversary. Bartusek, Garg, Jain and Policharla’s work [BGJP23] is related in spirit to privacy-preserving blueprints: they show a scheme that makes it possible to identify an originator of harmful content (relative to a database of harmful content) while protecting privacy in all other circumstances.

Comparing our construction to that of KLN, our escrows contain 58 group elements and 52 scalar elements per recursion. Thus, our ElGamal escrows contain 2.6kB * log(n) (where n is the size of the watch list) and our Camenisch-Shoup escrows should take roughly 10 times that at 28kB * log(n) (where the modulus is 2048 bits). This means our ElGamal escrow becomes more efficient than KLN when there are more than 60 suspects on the watchlist. The computation time of each escrow is roughly the same between our construction and KLN.

2 Preliminaries

Notation We will write the set of integers from 1 to m as $[m]$. We use bold font to represent vectors, e.g.: $\mathbf{a} = (a_1, \dots, a_n)$.

⁷ Like much other work our techniques are certainly inspired by Bulletproofs and its predecessors.

2.1 Zero-knowledge Proofs of Knowledge

Black-box partially straight-line (BB-PSL) NIZK. Non-interactive zero-knowledge (NIZK) proofs are an important building block for us. We follow the KLN notation and definitions (Sec. 2.1 of [KLN23]) of the completeness and ZK properties of NIZK proof system, provided in abbreviated form in Def. 1 below.

Definition 1 (Completeness and ZK of NIZK [KLN23]). *Let \mathcal{R} be a relation. Let \mathcal{S} be a setup model (e.g., the CRS model or the random oracle model). Let $\mathsf{P}^{\mathcal{S}}$ and $\mathsf{V}^{\mathcal{S}}$ be (non-interactive) algorithms for the prover and the verifier in the \mathcal{S} -setup model. $(\mathsf{P}^{\mathcal{S}}, \mathsf{V}^{\mathcal{S}})$ constitute a complete proof system if for all $(\mathbb{x}, \mathbb{w}) \in \mathcal{R}$, $\Pr[\pi \leftarrow \mathsf{P}^{\mathcal{S}}(\mathbb{x}, \mathbb{w}) : \mathsf{V}^{\mathcal{S}}(\mathbb{x}, \pi) = 0] = 0$.*

They satisfy the zero-knowledge property if for any PPT adversary Adv in the experiment of Fig. 2.1, the advantage function $\nu(\lambda)$ defined below is negligible: $\text{Adv}_{\text{Adv}}^{\text{NIZK}} = |\Pr[\text{NIZK}^{\text{Adv},0}(1^\lambda) = 0] - \Pr[\text{NIZK}^{\text{Adv},1}(1^\lambda) = 0]| = \nu(\lambda)$

$\text{NIZK}^{\text{Adv},0}(1^\lambda)$	$\mathcal{O}_{\mathcal{S}}(m)$	$\mathcal{O}_{\mathcal{P}}(\mathbb{x}, \mathbb{w})$
return $\text{Adv}^{\mathcal{S}(\cdot), \mathsf{P}^{\mathcal{S}}(\cdot, \cdot)}(1^\lambda)$	st, h, $\tau_{\text{Ext}} \leftarrow \text{SimS}(\text{st}, m)$	if $(\mathbb{x}, \mathbb{w}) \notin \mathcal{R}$: return \perp
$\text{NIZK}^{\text{Adv},1}(1^\lambda)$	return h	st, $\pi \leftarrow \text{Sim}(\text{st}, \mathbb{x})$
return $\text{Adv}^{\mathcal{O}_{\mathcal{S}}(\cdot), \mathcal{O}_{\mathcal{P}}(\cdot, \cdot)}(1^\lambda)$		return π

Fig. 2.1: NIZK game

Let us review BB-PSL simulation extractable proof systems [KLN23] (Def. 2). The straight-line extractor here does not extract the entire witness, but just some function of it; simultaneously, a black-box extractor (that's allowed to rewind the adversary) can extract the entire witness. In Appx. C.1, we motivate this definition further.

Definition 2 (Black-box partial straight-line simulation extractability). *A proof system (as defined in Def. 1) is BB-PSL simulation extractable if the advantage (defined below) of any PPT adversary is negligible: $\text{Adv}_{\text{Adv},f}^{\text{NISimBBPSLExt}} = \Pr[f\text{-NISimBBPSLExt}^{\text{Adv}}(1^\lambda) = 1] = \nu(\lambda)$ for some negligible function ν .*

Proofs of Equivalent Representations of Discrete Logarithms. Using known techniques, we can construct a Σ -protocol that proves the relation in Def. 3 in both cyclic groups of prime order where the DDH and CDH problems are hard as well as \mathbb{Z}_{n^2} . We describe such a protocol in Appx. C.3. When using this protocol, we employ Camenisch-Stadler notation to denote witnesses and relations [CS97], e.g. $\text{NIZK}[\text{witness} : \text{statement}]$.

f -NISimBBPSLExt ^{Adv} (1 ^λ)	
1 : $\mathcal{Q}, \mathcal{Q}_S \leftarrow []; (\mathbb{x}, \pi) \leftarrow \text{Adv}^{\tilde{\mathcal{O}}_S(\cdot), \mathcal{O}_{\text{Sim}}(\cdot)}(1^\lambda)$	
2 : $\mathbb{w} \leftarrow \text{Ext}^{\text{BB}(\text{Adv})}(\mathcal{Q}_S, \mathbb{x}, \pi); \mathbb{w}' \leftarrow \text{ExtSL}(\mathcal{Q}_S, \mathbb{x}, \pi)$	
3 : return $\mathbb{V}^{\mathcal{O}_S}(\mathbb{x}, \pi) \wedge (\mathbb{x}, \pi) \notin \mathcal{Q} \wedge ((\mathbb{x}, \mathbb{w}) \notin \mathcal{R} \vee \mathbb{w}' \neq f(\mathbb{w}))$	
$\mathcal{O}_S(m)$ $\tilde{\mathcal{O}}_S(m)$	$\mathcal{O}_{\text{Sim}}(\mathbb{x})$
1 : $\text{st}, h, \tau_{\text{Ext}} \leftarrow \text{SimS}(\text{st}, m)$	1 : $\text{st}, \pi \leftarrow \text{Sim}(\text{st}, \mathbb{x})$
2 : $\mathcal{Q}_S.\text{add}((m, h, \tau_{\text{Ext}}))$	2 : $\mathcal{Q}.\text{add}((\mathbb{x}, \pi))$
3 : return h, τ_{Ext}	3 : return π

Fig. 2.2: f -NISimBBPSLExt game

Definition 3 (Relation for proof of multiplication of witnesses over bases in cryptographic groups). Let $R_{\text{eqrep-}\mathcal{G}^*}(\mathbb{x}, \mathbb{w})$ be the relation that accepts if the following two conditions hold:

- (1) $\mathbb{x} = (\mu, \{x_i, \{g_{i,1}, \dots, g_{i,m}\}_{i=1}^k\})$, all the x_i s and $g_{i,j}$ s are elements of \mathcal{G} , and witness $\mathbb{w} = (\{b_i\}_{i=0}^k, \{w_j\}_{j=1}^m)$ is such that $x_i = b_i \prod_{j=1}^m g_{i,j}^{w_j}$ where $b_i \in \mathcal{B}$.
- (2) If $\forall i \in [m], w_i = \prod_{j \in \mu(i)} w_j$ where μ is a map $\mu : [m] \rightarrow \mathcal{P}([m])$ and $\mathcal{P}([m])$ is the set of all subsets of $[m]$.

We instantiate this definition with $\mathcal{G} = \mathbb{G}_p$, a cyclic group of order p , and with $\mathcal{G} = \mathbb{Z}_{n^2}$. In the former case, $\mathcal{B} = \{1\}$, yielding exact equality of representation. In the latter, $\mathcal{B} = \{-1, 1\}$, so the relation only holds for the absolute values.

2.2 Additively Homomorphic Encryption

Additively homomorphic \mathfrak{g} -semi-encryption scheme. We need an appropriate additively homomorphic (AH) semantically secure public-key encryption scheme. Our application can tolerate a relaxed version of encryption, in which the decryption algorithm need not recover the original plaintext m , but just some function $\mathfrak{g}(m)$, where \mathfrak{g} is a (not necessarily efficiently) invertible function. This relaxation allows us to view the ElGamal cryptosystem as additively homomorphic. Let us define it formally.

Definition 4 (Additively homomorphic \mathfrak{g} -semi-encryption scheme). A set of three polynomial-time algorithms $AH = (\text{KeyGen}_{AH}, \text{Enc}_{AH}, \text{Dec}_{AH})$ constitutes a semantically secure homomorphic \mathfrak{g} -semi-encryption scheme if it satisfies the following input-output specification as well as correctness, security, and homomorphic properties:

Input-output specification KeyGen_{AH} and Enc_{AH} have the same input-output specifications as those for key generation and encryption algorithms, respectively, for a public-key encryption scheme. $\text{Dec}_{AH}(\text{sk}_{AH}, c)$ takes as input a

secret key sk_{AH} and a ciphertext, and outputs a value $m' = \mathbf{g}(m)$ for some $m \in \mathcal{M}$ where \mathcal{M} is the message space of the encryption scheme⁸.

Correctness For all $(\text{pk}, \text{sk}) \in \text{KeyGen}_{AH}$, for all $m \in \mathcal{M}$, for all $c \in \text{Enc}_{AH}(\text{pk}, m)$, $\text{Dec}_{AH}(\text{sk}, c) = \mathbf{g}(m)$. I.e., the decryption algorithm correctly recovers $\mathbf{g}(m)$ from an encryption of m .

Security A semantically secure \mathbf{g} -semi-encryption scheme must satisfy the same definition of semantic security as a regular semantically secure encryption scheme [GM82].

Additively homomorphic properties (1) \mathcal{M} is an algebraic ring and (2) there is an efficient deterministic algorithm Op_{AH} that takes as input the public key pk_{AH} and two ciphertexts, c_1 and c_2 and outputs a ciphertext c' such that for all $\text{pk}_{AH} \in \text{KeyGen}_{AH}$, for all $m_1, m_2 \in \mathcal{M}$, for all ciphertexts $c_1 \in \text{Enc}(\text{pk}_{AH}, m_1)$ and $c_2 \in \text{Enc}(\text{pk}_{AH}, m_2)$, if $c' = \text{Op}_{AH}(\text{pk}_{AH}, c_1, c_2)$, then $c' \in \text{Enc}(\text{pk}_{AH}, m_1 + m_2)$.

For our constructions in Sec. 4.1 we define \mathcal{M} as \mathbb{Z}_p for a prime p for ElGamal or \mathbb{Z}_N for an RSA modulus N for Camenisch-Shoup.

Further (inspired by Cramer, Damgård and Nielsen's [CDN01] formalization of an additively homomorphic cryptosystem), we also need a way to sample new encryptions of messages, i.e., compute $c' \leftarrow \text{Enc}(\text{pk}_{AH}, m)$ given any $c \in \text{Enc}(\text{pk}_{AH}, m)$. I.e. we require that this be achieved by forming a fresh encryption of 0, $c_0 \leftarrow \text{Enc}(\text{pk}_{AH}, 0)$ and then adding to c , resulting in $c' = c \oplus c_0$. Further, we need AH to include efficient algorithms for obtaining $c' \in \text{Enc}(\text{pk}_{AH}, am)$ from $c \in \text{Enc}(\text{pk}_{AH}, m)$ and $a \in \mathcal{M}$.

Notation for additively-homomorphic encryption. We will generally use the lowercase c label to refer to ciphertexts (while uppercase C refers to commitments). If c_1 and c_2 are ciphertexts, will use $c_1 \oplus c_2$ to denote the output of $\text{Op}(\text{pk}, c_1, c_2)$. We use $\underline{a}_{\text{pk}}$ to represent an encryption of a under the public key pk using the scheme AH ; we will drop the subscript and denote it \underline{a} when pk is clear from the context. By $\underline{a} = \underline{c} \oplus \underline{d}$ we denote that the ciphertext \underline{a} was generated by running the algorithm $\text{Op}(\text{pk}, \underline{c}, \underline{d})$; thus $\underline{a} = \underline{c} + \underline{d}$. $y \odot \underline{a}$ denotes applying this operation y times; in our instantiations this will yield $\underline{y a}$ and is efficient for large y with repeated squaring; $\bigoplus_{i=0}^n \underline{a_i}$ denotes applying Op n times on the set $\{\underline{a_i} : i \in [0..n]\}$.

2.3 Privacy Preserving f -Blueprint Schemes (PPBs)

[KLN23] defines a blueprint scheme as in Def. 2.3. We will modify their definition to serve our new use-case of non-frameable privacy preserving blueprints in Sect. 5.1. A blueprint scheme has three parties - an auditor, a set of users and a set of recipients.

KLN define a *secure* f -blueprint scheme as one that possesses the following properties - (a) Correctness of VerPK and VerEscrow , (b) Correctness of Dec (c)

⁸ In general, \mathcal{M} and \mathbf{g} may depend on the public key, pk_{AH} , but in our constructions, this is not the case. Our scalar commitments share this same message space.

$\text{Setup}(1^\lambda, cpar) \rightarrow \Lambda$: Outputs public parameters Λ including 1^λ and commitment scheme, $cpar$. $\text{KeyGen}(\Lambda, x, r_x) \rightarrow (\text{pk}_\Lambda, \text{sk}_\Lambda)$: The key generation algorithm for auditor A. $\text{VerPK}(\Lambda, \text{pk}_\Lambda, C_x) \rightarrow 1$ or 0 : Takes the auditor's public key pk_Λ and a commitment C_x as input, verifies that the auditor's public key was computed correctly for the commitment C_x . $\text{Escrow}(\Lambda, \text{pk}_\Lambda, y, r_y) \rightarrow Z$: Takes Λ , pk_Λ , and commitment value and opening (y, r_y) as input and outputs an escrow Z for commitment $C = \text{Com}(y; r_y)$. $\text{VerEscrow}(\Lambda, \text{pk}_\Lambda, C_y, Z) \rightarrow 1$ or 0 : Takes the auditor's public key pk_Λ , a commitment C_y , and an escrow Z as input and verifies that the escrow was computed correctly for the commitment C_y . $\text{Dec}(\Lambda, \text{sk}_\Lambda, C_y, Z) \rightarrow f(x, y)$ or \perp : Takes the auditor's secret key sk_Λ , a commitment C_y and an escrow Z as input. It decrypts the escrow and returns the output $f(x, y)$ if C_y is a commitment to y and $\text{VerEscrow}(\Lambda, \text{pk}_\Lambda, C_y, Z) = 1$.
--

Fig. 2.3: An f -blueprint scheme

Soundness (d) Blueprint Hiding (e) Privacy against Dishonest Auditor and (f) Privacy with Honest Auditor. We recall these definitions in Appx. B.

3 Succinct Proofs for Verifiable Secure Computation on Additively-Homomorphic Ciphertexts

In this section, we describe our new succinct proof system for verifiable secure computation, as introduced on Page 5 in our first contribution. Suppose that we have an additively homomorphic \mathfrak{g} -semi cryptosystem $\Gamma^{\text{Enc}} = (\text{Setup}, \text{Enc}, \text{Dec}, \oplus, \odot)$ as discussed in Sec. 2.2. Let pk be a public key for this cryptosystem. Given a set of ciphertexts c_1, \dots, c_n whose plaintexts are x_1, \dots, x_n , and a set of scalars y_1, \dots, y_k , the additively homomorphic property of the cryptosystem allows anyone to compute a ciphertext c_f which is the encryption of $f(x_1, \dots, x_n, y_1, \dots, y_k)$, where f is a polynomial. This polynomial is of max degree 1 in any x_i , or, in other words: the polynomial is made up of n monomials, $\{f_i\}_{i \in [n]}$, where $f_i = a_i x_i^{b_i} \prod_{j=1}^k y_j^{d_{i,j}}$, b_i is a bit ($\{0, 1\}$), a_i is a coefficient of f , and $d_{i,j}$ is an integer. In this section, we provide a framework for efficiently obtaining a proof system, in the random-oracle model, for the following relation, parameterized⁹ by the polynomial f :

$$R_f((r_1, \dots, r_k, y_1, \dots, y_k), (C_1, \dots, C_k, c_1, \dots, c_n, c_f)) = 1 \text{ iff}$$

$$\exists x_1, \dots, x_n \text{ such that}$$

$$C_j = \text{Com}(y_j, r_j) \forall 1 \leq j \leq k$$

$$\wedge c_i \in \text{Enc}(\text{pk}, x_i) \forall 1 \leq i \leq n$$

$$\wedge c_f \in \text{Enc}(\text{pk}, f(x_1, \dots, x_n, y_1, \dots, y_k))$$

Our proof system is complete, zero-knowledge and satisfies the definition of a (not straight-line extractable) proof of knowledge in the random-oracle model. To compile it into a partially straight-line extractable (\mathfrak{g} -BB-PSL) proof system, it will be sufficient to combine it with a \mathfrak{g} -BB-PSL proof of knowledge of the opening of the commitments C_1, \dots, C_k which we do in Sec. 5.3.

⁹ This relation is also parameterized by the parameters of the associated commitment and encryption schemes (for Com and Enc) along with the public key, pk , but we omit this for readability since it is clear from context.

Construction of a proof system for R_f . Using a general NIZK proof system to prove R_f would yield a proof of size $\Omega(kd_{\max})$ where d_{\max} is the largest degree among any $y_i, i \in [k]$. To make this more succinct, our proof system that halves the degree with each step. This reduces the size of the proof from linear in d_{\max} to $O(k \log(d_{\max}))$, which is an exponential improvement. The proof size is independent on the number of monomials in f and ciphertexts, and depends only on k (the number of variables y_1, \dots, y_k) and the degree d_{\max} .

Each step of this proof will reduce the task of proving the correct evaluation of a polynomial f to that of another polynomial, f' . To achieve succinctness, at each step we will pick a variable and ensure that the degree of f' in that variable is at most half that of f . For example, proving that $c_f = \boxed{f(x_1, x_2, y_1, y_2)}$ where $f = x_1 y_1^8 y_2 + x_2 y_1^7 y_2$ will be reduced to proving that $c_{f'} = \boxed{f'(x_1, x_2, y_1, y_2)}$ where $f' = x'_1 y_1^4 y_2 + x'_2 y_1^3 y_2$. We can see that the degree of f' in y_1 has been reduced by 4 here compared to f .¹⁰ In this section, we will explain how the ciphertexts $\boxed{x'_1}$ and $\boxed{x'_2}$ can be computed from $\boxed{x_1}$ and $\boxed{x_2}$ by both the prover and verifier in a way that ensures that proving the correctness of f' implies the correctness of f . Because we want to achieve zero knowledge even when the adversary knows the secret key of the encryption scheme, a zero-knowledge simulator cannot simply make up an arbitrary value for $c_{f'}$: the adversary would be able to decrypt it and detect simulation. Thus, we need to instead commit to this value and perform the proof that the committed value was computed correctly. We call these *commitments to additively homomorphic ciphertexts* and we define them in Sec. 3.1 and construct them in Sec. 4.

3.1 Basic Building Blocks

Commitment to $\{y_1, \dots, y_k\}$. Our relation R_f is defined relative to a non-interactive commitment scheme for scalars (CSetup, Com). Com takes as input an element y from the message space, \mathcal{M} , and a random value r sampled uniformly at random from $[R]$ for some integer R .

Proofs of correct modular addition and multiplication of committed values. In order to construct this proof system, we need to add and multiply the values in our scalar commitments together (using ring arithmetic in the message space, \mathcal{M}). This will be used to realize *egrep* and thus realize our commitments to ciphertexts. Let us define the following relations¹¹:

- $R_{\text{add}}((x_1, r_1, x_2, r_2, x_3, r_3), (C_1, C_2, C_3)) = 1$ iff $\forall i \in [3] : C_i = \text{Com}(x_i; r_i)$, and $x_3 = x_1 + x_2 \in \mathcal{M}$. Let $(\text{Prove}^{\text{add}}, \text{Verify}^{\text{add}})$ be a BB NIZK proof system for R_{add} .
- $R_{\text{mult}}((x_1, r_1, x_2, r_2, x_3, r_3), (C_1, C_2, C_3)) = 1$ iff $\forall i \in [3] : C_i = \text{Com}(x_i; r_i)$, and $x_3 = x_1 x_2 \in \mathcal{M}$. Let $(\text{Prove}^{\text{mult}}, \text{Verify}^{\text{mult}})$ be a BB NIZK proof system for R_{mult} .

¹⁰ As noted in Sec. 1.1 and related work, these techniques resemble [LFKN92] and Bulletproofs but present unique challenges when applied to committed ciphertexts.

¹¹ $R_{\text{Com}_{AH}}$ is trivial with additively homomorphic commitments, but we present this function for generality.

We also need this commitment scheme to have a zero-knowledge proof of knowledge ($\text{Prove}^{\text{Com}}, \text{Verify}^{\text{Com}}$) of opening, i.e. a BB NIZK for the relation $R_{\text{Com}} = ((m, r), (C))$ iff $\text{Com}(m; r) = C$.

Commitment to ciphertexts. In order to prove correctness of an intermediate step in a longer computation over (semi-)encrypted data without revealing the ciphertext obtained in that step itself (which would leak data), we need to be able to commit to ciphertexts and prove properties of committed ciphertexts. Thus, we need a non-interactive statistically hiding, computationally binding commitment scheme Com_{AH} (parameterized by public parameters $params$ generated by Setup_{AH}) for committing to ciphertexts $c \in \text{Enc}_{AH}(\text{pk}, \cdot)$ and we need protocols for proving statements about committed ciphertexts, as described below. We use a subscript notation (i.e. Com_{AH}) to distinguish this scheme from our commitments to scalars which do not have a subscript (the commitment function for scalars is Com). If randomness is not supplied to Com_{AH} , it will sample randomness and output it, e.g.: $(C, r) = \text{Com}_{AH}(\underline{a})$ implies that $C = \text{Com}_{AH}(\underline{a}; r)$.

Proofs of relations between committed ciphertexts. We need BB NIZK proof systems for (1) proving knowledge of a committed ciphertext; (2) proving that a committed ciphertext is the result of applying Op_{AH} to other committed ciphertexts; (3) proving that a committed ciphertext is the result of applying Op_{AH} to another committed ciphertext α times, where α is the opening of a commitment (under the scalar commitment scheme Com) to an element of \mathcal{M} ; and (4) proving that a committed ciphertext is an encryption of a committed scalar. (4) is often called “verifiable encryption” (VE). More precisely, let us define the following relations¹²:

- $R_{\text{Com}_{AH}}((c, r), C) = 1$ iff $C = \text{Com}_{AH}(c; r)$;
- $R_{\oplus}((c_1, r_1, c_2, r_2, c_3, r_3), (C_1, C_2, C_3)) = 1$ iff $\forall i \in [3] : C_i = \text{Com}_{AH}(c_i; r_i)$ and $c_3 = \text{Op}_{AH}(c_1, c_2)$;
- $R_{\odot}((c_1, r_1, c_2, r_2, x, r_3), (C_1, C_2, C_3)) = 1$ iff $\forall i \in [2] : C_i = \text{Com}_{AH}(c_i; r_i)$, $C_3 = \text{Com}(x; r_3)$ and $c_2 = c_1 \odot x$.
- $R_{VE}((c_1, r_1, r_{c_1}, y, r_2), (C_1, C_2)) = 1$ iff $C_1 = \text{Com}_{AH}(c_1; r_1)$, $C_2 = \text{Com}(y; r_2)$ and $c_1 = \text{Enc}_{AH}(\text{pk}_{AH}, y; r_{c_1})$.

Our construction will use as building blocks BB NIZK proof systems ($\text{Prove}^{\text{Com}_{AH}}, \text{Verify}^{\text{Com}_{AH}}$) for the relation $R_{\text{Com}_{AH}}$, ($\text{Prove}^{\oplus}, \text{Verify}^{\oplus}$) for the relation R_{\oplus} , ($\text{Prove}^{\odot}, \text{Verify}^{\odot}$) for the relation R_{\odot} , and ($\text{Prove}^{\text{enc}}, \text{Verify}^{\text{enc}}$) for the relation R_{VE} . As before, we omit the parameters and public keys from these relations when it is clear. These proof systems exist generically for any cryptosystem and any set of commitment schemes by representing the Com function as a circuit; however, for the specific instantiations of semi-encryption and commitment schemes we consider, we also show how to construct them efficiently in Sec. 4.

Notation. We will use the following notation when invoking a proof system (inspired by the Camenisch-Stadler notation): $\pi = \text{NIZK}[X, W : R(X, W)]$ denotes

¹² Similar to our commitments to scalars, R_{\oplus} is not necessary if the commitments are additively homomorphic but we present this here for generality.

that the proof π is computed using the proof system for R on input a statement X and a witness W . When X is clear from the description of the relation R , we may omit it. For example, if we have $A = \text{Com}_{AH}(\underline{a}; r_a)$, $B = \text{Com}_{AH}(\underline{b}; r_b)$, and $C = \text{Com}(c; r_c)$ and want to prove that $a = bc$, we'll denote the output of the prover's computation as $\pi = \text{NIZK}[\underline{a}, \underline{b}, c, r_a, r_b, r_c : A = \text{Com}_{AH}(\underline{a}, r_a) \wedge B = \text{Com}_{AH}(\underline{b}, r_b) \wedge C = \text{Com}(c; r_C, a_C) \wedge \underline{a} = \underline{b} \odot c]$. This π is computed by calling $\text{Prove}^\odot(A, B, C, \underline{a}, r_a, \underline{b}, r_b, c, r_c)$. If π is accepted by the verification algorithm (i.e. $\text{Verify}^\odot(A, B, C, \pi) = 1$) we can extract openings for A , B and C to ciphertexts \underline{a} , \underline{b} and scalar c respectively, such that $\underline{a} = \underline{b} \odot c$.

3.2 Efficient Proof System for R_f for $k = 1$

In this section we show how to efficiently instantiate a NIZK proof for the relation R_f when $k = 1$, i.e. there is a single variable y . Our main result in Appx. E.4 subsumes the result in this section; however, this section makes it easier for the reader to understand the results in Appx. E.4.

Observe that it is sufficient¹³ to provide a proof system for the polynomial $P = \sum_{i=0}^{n-1} x_i y^i$. Thus we give a proof system for the relation R_P . Further, it is sufficient to give a proof system for a slightly more general relation, R_P^* in which the statement contains not the ciphertext c_P but a commitment $C_P = \text{Com}_{AH}(c_P, r_P)$. To get a proof system for R_P , prover and verifier set $C_P = \text{Com}_{AH}(c_P, 0)$ and invoke the proof system for R_P^* . Assume WLOG¹⁴ that n (the number of ciphertexts) is a power of two. More formally,

$$R_P^*((r, y, c_P, r_P), (C_y, c_0, \dots, c_{n-1}, C_P)) = 1 \text{ iff} \\ R_P(r, y, C_y, c_0, \dots, c_{n-1}, c_P) = 1 \wedge C_P = \text{Com}_{AH}(c_P, r_P).$$

Input to the recursive step. Our PoK_P^* algorithm in Algorithm 2 recursively computes a proof until R_P^* is satisfied, i.e., C_P is a commitment to $c_P = \underline{a} = \underline{P}(x_0, \dots, x_{n-1}, y)$. The input to PoK_P^* includes an auxiliary input aux , in addition to the statement and witness for the relation R_P^* . aux consists of commitments to a logarithmic number of powers of y , i.e. commitments $\{C_{y^{2^i}}\}$ to $\{y^{2^i}\} = \{y^2, y^4, y^8, \dots, y^{n/2}\}$ and NIZK proofs that for $i > 2$, each $C_{y^{2^i}}$ is computed correctly from $C_{y^{2^{i-1}}}$ (using the proof system $(\text{Prove}^{\text{mult}}, \text{Verify}^{\text{mult}})$ described above). aux is of size that is logarithmic in n and the verifier need not

¹³ From here, to obtain the proof system for any $f = a_{00} + \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} a_{i,j} x_j y^i$, we use the homomorphic properties of the cryptosystem to compute $c'_i = \underline{\sum_{j=0}^{n-1} a_{i,j} x_j}$ for $0 \leq i < n$, (deterministically, using the all-0 string for encryption) incorporate the term a_{00} by letting $c''_0 = c'_0 \oplus \underline{a_{00}}$ and then invoke the proof system for P on input ciphertexts $c''_0, c'_1, \dots, c'_{n-1}$. Here we assume that the degree of y is n , but this is WLOG as the complexity of the proof system is not dependant on the number of ciphertexts.

¹⁴ This is without loss of generality: to reduce to this case, prover and verifier can both compute the extra ciphertexts c_n, \dots, c_{2^a-1} (so that the total number is a power of two) by encrypting 0 with fixed randomness.

verify any proofs in it more than once. We assume that the prover remembers how it computed aux (so we won't explicitly pass the openings of the commitments in aux to the recursive step). Alg. 1 is a “wrapper” algorithm that, on input the statement-witness pair for relation R_P transforms it into the statement-witness pair for relation R_P^* , initializes aux with $\{C_{y^{2^i}}\}$ and their proofs of correctness, initializes the transcript, τ , and calls PoK_P^* .

Ensuring soundness for the recursive proof. The prover and verifier can both compute encrypted evaluations of the polynomial $P(x_0, \dots, x_{n-1}, \gamma)$ on any input γ using the ciphertexts $\{c_i\}$. They can further break P into two parts such that $P(x_0, \dots, x_{n-1}, \gamma) = P_1(x_0, \dots, x_{n/2-1}, \gamma) + P_2(x_{n/2}, \dots, x_{n-1}, \gamma)$ where P_1 contains the monomials $x_i \gamma^i$ for $i < n/2$, and P_2 contains monomials of higher degree in γ . We can represent P as $P(x_0, \dots, x_{n-1}, \gamma) = P_1(x_0, \dots, x_{n/2-1}, \gamma) + \gamma^{n/2} P_3(x_{n/2}, \dots, x_{n-1}, \gamma)$ where $P_3(\gamma) = P_2(\gamma)/\gamma^{n/2}$.

To recurse, the prover commits to ciphertexts $\boxed{e_1} = \boxed{P_1(x_0, \dots, x_{n/2-1}, y)}$, $\boxed{e_2} = \boxed{P_2(x_{n/2-1}, \dots, x_{n-1}, y)}$, $\boxed{e_3} = \boxed{P_3(x_{n/2}, \dots, x_{n-1}, y)}$, and then proves (using the proof systems for proving properties of committed ciphertexts) that $\boxed{e} = \boxed{e_1} \oplus \boxed{e_2}$ and $\boxed{e_2} = y^{n/2} \odot \boxed{e_3}$ using the commitment $C_{y^{n/2}}$ found in aux . Thus, the prover has reduced the task of proving that C_P is a commitment to $\boxed{e} = \boxed{P(x_0, \dots, x_{n-1}, y)}$ for a polynomial P of degree $n - 1$ to the task of proving that C_{P_1} is a commitment to $\boxed{e_1} = \boxed{P_1(x_0, \dots, x_{n/2-1}, y)}$ and C_{P_3} is a commitment to $\boxed{e_3} = \boxed{P_3(x_0, \dots, x_{n/2-1}, y)}$, where P_1 and P_3 are both polynomials of degree $n/2 - 1$.

To take advantage of recursion, we need to use just one recursive call in order to prove that the openings of C_{P_1} and C_{P_3} (i.e., $\boxed{e_1}$ and $\boxed{e_3}$ respectively) are encrypted evaluations of P_1 and P_3 . To do so, prover and verifier define a new polynomial P' of degree $(n - 1)/2$ by taking a random linear combination of P_1 and P_3 : let α be the output of the random oracle on input the elements of the proof that have been computed so far. Let $P'(x_0, \dots, x_{n-1}, y) = P_1(x_0, \dots, x_{n-1}, y) + \alpha P_3(x_0, \dots, x_{n-1}, y)$. By the Schwartz-Zippel Lemma (Lemma 13), if committed $\boxed{e_1} \neq \boxed{P_1(y)}$ or committed $\boxed{e_3} \neq \boxed{P_3(y)}$, then with overwhelming probability over the choice of α , $\boxed{e_1} \oplus (\alpha \odot \boxed{e_3}) \neq \boxed{P'(x_0, \dots, x_{n-1}, y)}$. Let $C_{P'}$ be a commitment to the ciphertext $\boxed{e'} = \boxed{e_1} \oplus (\alpha \odot \boxed{e_3})$; the prover can provide a proof that indeed $C_{P'}$ is a commitment to $\boxed{e'}$ computed this way based on C_{P_1} and C_{P_3} and α using the proof systems for committed ciphertexts.

Next, we use recursion to prove that $C_{P'}$ corresponds to correctly evaluating the polynomial P' , i.e. it is a commitment to $\boxed{P'(x_0, \dots, x_{n-1}, y)}$. To do so, we call PoK_P^* on input ciphertexts $(c'_0, \dots, c'_{n/2-1})$ where $c'_i = \boxed{x_i} \oplus (\alpha \odot \boxed{x_{n/2+i}})$.

Theorem 1 *Our scheme in Algs. 1 and 2 are complete and ZK (Def. 1).*

Theorem 2 *The PoK_P^* function in Alg. 2 is black-box (BB) simulation extractable with respect to Def. 2 for the relation R_f^* .*

For compactness, here we only present the prover's algorithms; the verifier's algorithms (provided in Appx. E.4) should follow from the prover's algorithms.

Algorithm 1 $\text{PoK}_P(r, y, \underline{C}_y, c_0, \dots, c_{n-1}, c_P) \rightarrow \pi$

- Let $c_i = \{\underline{x}_i\}_{i \in [0..n-1]}$;
 Prover needs to prove that $c_P = \underline{e} = \bigoplus_{i=0}^n (\underline{x}_i \odot y^i)$
 To format c_P for the recursion, we commit to it with known randomness e.g. 0
- 1: $C_P \leftarrow \text{Com}_{AH}(c_P; 0)$
 - 2: For $i = 1$ to $\log n$, let $(C_{y^{2^i}}, r_i) = \text{Com}(y^{2^i})$
 and let $\pi_{y^{2^i}} \leftarrow \text{NIZK}[(z, r_{i-1}, r_i) : C_{y^{2^{i-1}}} = \text{Com}(z; r_{i-1}) \wedge C_{y^{2^i}} = \text{Com}(z^2; r_i)]$.
 - 3: Initialize $\text{aux} = (\{C_{y^{2^i}}, \pi_{y^{2^i}}\}_{i \in [\log(n)]}, \tau = (C_y, c_0, \dots, c_{n-1}, c_P))$.
 - 4: **return** $\text{aux}, \text{PoK}_P^*(r_y, y, c_P, r_P, C_y, c_0, \dots, c_{n-1}, C_P, \text{aux}, \tau)$
-

Algorithm 2 $\text{PoK}_P^*(r_y, y, c_P, r_P, \underline{C}_y, c_0, \dots, c_{n-1}, C_P, \text{aux}, \tau) \rightarrow \pi'$

- Let $c_i = \{\underline{x}_i\}_{i \in [0..n-1]}$; $c_P = \underline{e}$
 Prover needs to prove that $C_P = \text{Com}_{AH}(\underline{e}; r_P)$ where $\underline{e} = \bigoplus_{i=0}^{n-1} \underline{x}_i \odot y^i = \underline{\sum_{i=0}^{n-1} y^i x_i}$ and $C_y = \text{Com}(y; r_y)$
- If the degree of the polynomial is low enough, prove its computation directly:
- 1: **if** $n = 1$, **return** (π_1) where $\pi_1 \leftarrow \text{NIZK}[r : \text{Com}_{AH}(\underline{x}_0, r) = C_P]$
 If not, we will need to reduce the degree needed to prove C and recurse.
 To do so, first, commit to the lower half of the polynomial:
 - 2: $(C_1, \rho_1) = \text{Com}_{AH}(\underline{e}_1)$ where $\underline{e}_1 = \bigoplus_{i=0}^{n/2-1} \underline{x}_i \cdot y^i = \underline{\sum_{i=0}^{n/2-1} y^i x_i}$
 Next, commit to the upper half of the polynomial
 - 3: $(C_2, \rho_2) = \text{Com}_{AH}(\underline{e}_2)$
 where $\underline{e}_2 = \bigoplus_{i=0}^{n/2-1} \underline{x}_{i+n/2} \odot y^{i+n/2} = \underline{\sum_{i=0}^{n/2-1} y^{i+n/2} x_{i+n/2}}$
 Lastly, commit to the upper half of the polynomial with the degree lowered by half
 - 4: $(C_3, \rho_3) = \text{Com}_{AH}(\underline{e}_3)$ where $\underline{e}_3 = \bigoplus_{i=0}^{n/2-1} \underline{x}_{i+n/2} \odot y^i = \underline{\sum_{i=0}^{n/2-1} y^i x_{i+n/2}}$
 Query the random oracle on the current transcript of the proof so far, (C_1, C_2, C_3, τ)
 - 5: $\alpha \leftarrow H(C_1, C_2, C_3, \tau)$
 Compute the encryptions of the new coefficients for a reduced degree polynomial
 - 6: $\forall i \in [n/2 - 1], c'_i = \underline{x}'_i = \underline{x}_i \oplus (\underline{x}_{i+n/2} \odot \alpha)$
 Compute a new evaluation over this reduced degree polynomial:
 - 7: $(C'_P, r') = \text{Com}_{AH}(\underline{e}')$ where $\underline{e}' = \bigoplus_{i=0}^{n/2-1} \underline{x}'_i \odot y^i$
 Prove that this new commitment C'_P is consistent with C_P, C_1, C_2 , and C_3 .
 - 8: $\pi_\alpha \leftarrow \text{NIZK}[r, \rho_1, \rho_2, \rho_3, r', r_y, y, \underline{e}, \underline{e}_1, \underline{e}_2, \underline{e}_3, \underline{e}'] :$
 - 9: $\text{Com}_{AH}(\underline{e}, r) = C_P \wedge \text{Com}_{AH}(\underline{e}', r') = C'_P \wedge \forall 1 \leq i \leq 3 : \text{Com}_{AH}(\underline{e}_i, \rho_i) = C_i$
 - 10: $\wedge \underline{e} = \underline{e}_1 \oplus \underline{e}_2$
 - 11: $\wedge \underline{e}_2 = y^{n/2} \odot \underline{e}_3$ ▷ proven relative to $C_{y^{n/2}}$ in aux
 - 12: $\wedge \underline{e}' = \underline{e}_1 \oplus (\alpha \odot \underline{e}_3)$
 - 13: $\pi = (C_1, C_2, C_3, C'_P, \pi_\alpha)$
 - 14: $\tau' = (\pi, \tau)$ ▷ Append this proof to the transcript
 - 15: **return** $(\pi, \text{PoK}_P^*(r_y, y, \underline{e}', r', C_y, c'_0, \dots, c'_{n/2-1}, C'_P, \text{aux}, \tau'))$
-

For readability, in the list of inputs to the prover, we underline those inputs that are also given to the verifier.

We prove Thms. 1 and 2 in Appx. E.1. We show how to instantiate our NIZK proofs from *eqrep* in Appx. E.3.

3.3 Proof System for Multivariate Polynomials

We present our algorithm for polynomials with multiple y_i values in Appendix E.4 in Alg. 8. This algorithm proves the relation R_y described at the start of this section. In essence, the algorithm will perform the same recursive step as Alg. 2 until it has reduced the degree of a y_i variable to 0. The algorithm then recurses on the remaining $k - 1$ y_i variables until none are left. At this point, the evaluation has been fully proven.

For intuition, we describe how our protocol would prove the correct computation of an example polynomial: $f(x_1, x_2, y_1, y_2) = a_1x_1y_1y_2 + a_2x_2y_1^2y_2$. Our proof function will first focus on y_1 , finding that the maximum degree of this variable, $d_{\max} = 2$. It will then compute $f_1(x_1, x_2, y_1, y_2) = a_1x_1y_1y_2$ and $f_2(x_1, x_2, y_1, y_2) = a_2x_2y_1^2y_2$ as well as $f_3(x_1, x_2, y_1, y_2) = (a_2x_2y_1^2y_2)/y_1 = a_2x_2y_1y_2$. The proof function will then commit to encryptions of these polynomials, and hash the transcript to receive the challenge, α . It will then prove the relation $f(\dots) = f_1(\dots) + f_2(\dots)$ and $f_2(\dots) = y * f_3(\dots)$ (for compactness, we've replaced the input to functions, x_1, x_2, y_1, y_2 , with ellipses). It will then compute $f'(\dots) = f_1(\dots) + \alpha f_3(\dots) = a_1x_1y_1y_2 + a'_2x_2y_1y_2$ where $a'_2 = a_2 * \alpha$. This process will repeat for f' , with the prover recomputing d_{\max} for y_1 to be 1. The prover will then compute f'_1, f'_2 , and f'_3 (similar to how they computed f_1, f_2 , and f_3) but this time, the prover will find that $f'_1(\dots) = 0$ (since no monomial has degree of y_1 less than $d_{\max}/2 = 1/2$) and $f'_3(\dots) = (a_1x_1y_1y_2 + a'_2x_2y_1y_2)/y_1 = a_1x_1y_2 + a'_2x_2y_2$. Thus, $f''(\dots) = f'_1(\dots) + \alpha f'_3(\dots) = 0 + \alpha'(a_1x_1y_2 + a'_2x_2y_2)$ (for the challenge, α'). Thus, the prover has removed y_1 from the polynomial to be proven. Once this repeats to remove y_2 , the prover is left with $f^*(\dots) = a_1^*x_1 + a_2^*x_2$ where a_1^* and a_2^* are some combination of the coefficients of f and the challenges from the previous recursive steps. This is a linear function in the x_i 's where the α 's are known by the verifier so the verifier can simply compute the encryption of $f^*(\dots)$ at this point and the prover can prove that they've committed to this encryption. We give our results for this relation in Theorems 4 and 3 which we prove in Appx. E.4.

Theorem 3 *Our scheme in Alg. 8 is complete and ZK (Def. 1).*

Theorem 4 *The function in Alg. 8 is black-box (BB) simulation extractable with respect to Def. 2 for the relation R_f defined in Sect. 3.*

4 Instantiations of Commitments to Additively-Homomorphic Ciphertexts

As explained in Page 5 in our first contribution, for our proof system from Sec. 3 to be efficient, we require schemes to compute over commitments to additively

homomorphic ciphertexts that we define in Sec. 3.1. In this section, we introduce some efficient instantiations of these schemes.

We first define variants of ElGamal and Camenisch-Shoup encryption, in Sec. 4.1. Specifically, we define “lifted” ElGamal and Camenisch-Shoup in a “commitment-friendly” group. We then construct commitments to ciphertexts and associated proof systems for adding and multiplying ElGamal ciphertexts and Camenisch-Shoup ciphertexts. We use (Lifted) ElGamal which is a \mathfrak{g} -semi-encryption as defined in Sec. 3 with message space $\mathcal{M} = \mathbb{Z}_p$ and $\mathfrak{g}(x) = h^x \bmod p$. Camenisch-Shoup encryption has the advantage that it allows for the efficient computation of discrete logarithms in a subgroup of size n where n is an RSA modulus. Thus, with Camenisch-Shoup encryption, we can efficiently decrypt ciphertexts when the message space has exponential size. Thus, our Camenisch-Shoup construction is a \mathfrak{g} -semi-encryption where \mathfrak{g} is the identity function (i.e. a standard encryption scheme). In our Camenisch-Shoup construction, the message space is $\mathcal{M} = \mathbb{Z}_n$. In Sec. 4.2 we construct commitments to Camenisch-Shoup ciphertexts. As the setting of ElGamal is simpler, we defer it to Appx. F.3 though we review some elements of ElGamal encryption which explain why it is similar to Camenisch-Shoup.

4.1 Encryption Schemes

We review (Lifted) ElGamal encryption in Fig. 4.1a and a modified Camenisch-Shoup encryption in Fig. 4.1b. We include an extra generator (h) for lifting to exponents in ElGamal so that we can draw parallels between ElGamal and Camenisch-Shoup (ElGamal encryption generally uses the default generator, $h = g$). We see that both ElGamal and Camenisch-Shoup have similar homomorphic properties. Specifically for two encryptions, $(g^r, k^r h^m)$ and $(g^{r'}, k^{r'} h^{m'})$, $(g^r \cdot g^{r'}, k^r h^m \cdot k^{r'} h^{m'})$ is a valid encryption of $\mathfrak{g}(m + m')$ in both encryption schemes. Also, exponentiation is similar, i.e. $((g^r)^y, (k^r h^m)^y)$ is a valid encryption of $\mathfrak{g}(ym)$ in both encryption schemes.

Our modification to Camenisch-Shoup encryption includes replacing some values (parameter $g \in \mathbb{Z}_n$ and ciphertext c) with their absolute values. Modifying Camenisch-Shoup in this way ensures that the elements of honest Camenisch-Shoup ciphertexts lie in a “commitment-friendly” sub-group $|QR_{n^2}|$ where $|QR_{n^2}| = \{|x| : x \in QR_{n^2}\}$, QR_{n^2} is the quadratic residues in \mathbb{Z}_{n^2} (elements that have a square root in \mathbb{Z}_{n^2}), and $|\cdot|$ is the absolute value function such that $|x|$ of an element $x \in \mathbb{Z}_{n^2}$ is $n^2 - x$ if $x > \lfloor n^2/2 \rfloor$ and x otherwise. One nice property of $|QR_{n^2}|$ is that it is cyclic (which helps with our hiding and ZK proofs) and also $|QR_{n^2}|$ is efficiently sampleable by sampling a random element of \mathbb{Z}_{n^2} , squaring it, and taking its absolute value.

Critically, elements of $|QR_{n^2}|$ are always equal to their absolute value i.e. $|x| = x$ for elements in $|QR_{n^2}|$. This is important because it means that using *eqrep*- \mathbb{Z}_{n^2} (as defined in Sec. 2.1) to prove relations between $|QR_{n^2}|$ elements works perfectly, where-as for \mathbb{Z}_{n^2} it only holds for the absolute values of these elements. As an example, if we wanted to prove that we know a such that $c = g^a$ in $\mathbb{Z}_{n^2}^*$, we could only prove that $c = bg^a$ where $b \in \{-1, 1\}$. Ultimately, we

Fig. 4.1: Encryption schemes

<div style="border: 1px solid black; padding: 5px;"> Setup(1^λ) \rightarrow $params_{ElG}$ <hr/> 1 : Generate cyclic group of prime order p, \mathbb{G}_p 2 : $g, h \leftarrow_{\\$} \mathbb{G}_p$ 3 : return (g, h, \mathbb{G}_p) </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> KeyGen($params_{ElG}$) \rightarrow (pk, sk) <hr/> 1 : $x \leftarrow_{\\$} \mathbb{Z}_p$; 2 : return $pk \leftarrow g^x, sk \leftarrow x$; Enc($pk = k, m$) \rightarrow c <hr/> 1 : $r \leftarrow_{\\$} \mathbb{Z}_p$; 2 : return $c = (g^r, k^r h^m)$ </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> Dec($sk, c = (c_0, c_1)$) \rightarrow M <hr/> 1 : $z = c_0^{sk} = k^r$ 2 : return $c_1/z = M = h^m$ </div>	<div style="border: 1px solid black; padding: 5px;"> Setup(1^λ) \rightarrow $params_{CS}$ <hr/> 1 : Sample a safe RSA modulus, $n = pq = (2p' + 1)(2q' + 1)$ 2 : $g' \leftarrow_{\\$} QR_{n^2} , g = (g')^n , h = (1 + n),$ 3 : return $params_{CS} = (n, g, h)$ </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> KeyGen($params_{CS}$) \rightarrow (pk, sk) <hr/> 1 : $sk = x \leftarrow_{\\$} [n^2/4], pk = k = g^x \quad // \text{ in } QR_{n^2}$ 2 : return pk, sk </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> Enc($pk, m \in [n]$) \rightarrow c <hr/> 1 : $r \leftarrow_{\\$} [n/4],$ 2 : return $c = (g^r , k^r h^m) \quad // \text{ in } QR_{n^2}$ </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> Dec($sk, c = (c_0, c_1)$) \rightarrow m <hr/> 1 : $t = 2^{-1} \pmod n$ 2 : $M = c_1/c_0^x \quad // \text{ in } \mathbb{Z}_{n^2}$ 3 : return $m = ((M^{2t} \pmod{n^2}) - 1)/n$ </div>
--	--

(a) Lifted ElGamal

(b) Simplified Camenisch-Shoup

use $|QR_{n^2}|$ since we want to ensure that after performing exponentiation and multiplication proofs over commitments to ciphertexts, the ciphertext decrypts to the expected value. We can see in Fig. 4.1b that the encryption scheme decrypts the absolute value of a ciphertext exactly the same as the original ciphertext. This is clear from rewriting the decryption process as $m = (((c_1^2/(c_0^2)^x)^t \pmod{n^2}) - 1)/n$. The first operation the decryptor does is square both elements of the ciphertext, and our claim follows from the fact that $|x|^2 = x^2 \in \mathbb{Z}_{n^2}$. The proofs of our ElGamal commitments use the *eqrep*- \mathbb{G}_p protocol (in Sec. 2.1) which does not have the same limitations as *eqrep*- \mathbb{Z}_{n^2} . Unfortunately, $|QR_{n^2}|$ is not efficiently recognizable, but if every ciphertext comes with a proof of correct encryption (starting from honest parameters) we can be assured that the resulting ciphertext lives in $|QR_{n^2}|$. We show in Appx. F.1 that in our modified Camenisch-Shoup scheme, g and h are both in the group $|QR_{n^2}|$. Thus, if we can create commitments to elements of $|QR_{n^2}|$, we can use them to commit to our modified Camenisch-Shoup ciphertexts and construct the associated protocols for multiplication and exponentiation.

4.2 Commitments to $|QR_{n^2}|$ and Camenisch-Shoup Ciphertexts

To construct commitments to Camenisch-Shoup ciphertexts, we need to construct commitments to the elements of the group in which components of a Camenisch-Shoup ciphertexts lie. We accomplish this by using Damgård-Fujisaki

integer commitments [DF02] that are similar to Pedersen commitments. We adapt Damgård-Fujisaki commitments to “live” in \mathbb{Z}_{n^2} in Fig. 4.2 in functions Setup_{DF} and Com_{DF} . We prove them secure in Appx. F.2. To give intuition, these commitments appear very similar to Pedersen commitments, using two generators of a group as the public parameters and exponentiating them in a similar way to commit to integers. In this scheme, B is such that 2^B is larger than the order of $|QR_{n^2}|$ (i.e. $2^B = n^2/4$). Using Damgård-Fujisaki commitments that

Fig. 4.2: Simplified Damgård-Fujisaki commitments in \mathbb{Z}_{n^2}

$\text{Setup}_{\text{DF}}(1^\lambda) \rightarrow \text{params}_{\text{DF}} :$	$\text{Commit}_{\text{DF}}(\text{params}, m) \rightarrow (C, O) :$
1: Sample a safe RSA modulus, $n = pq = (2p' + 1)(2q' + 1)$ 2: Sample random $g, h \in \mathbb{Z}_{n^2}$. 3: return $\text{params} = (g, h)$	1: To commit to integer, m , compute: $C = g^m h^r$ where $r \leftarrow_{\$} [2^{B+\lambda}]$ 2: return (C, O) where $O = r$

live in \mathbb{Z}_{n^2} will allow us to use the *eqrep*- \mathbb{Z}_{n^2} protocol from Def. 3 to complete proofs of multiplication and exponentiation of our $|QR_{n^2}|$ commitments¹⁵.

Next, by employing Damgård-Fujisaki commitments that live in \mathbb{Z}_{n^2} , we can construct a scheme for committing to elements of $|QR_{n^2}|$ (and then we can use $|QR_{n^2}|$ commitments to construct commits to Camenisch-Shoup ciphertexts). We show this scheme in Fig. 4.3 in functions Setup_{QR} and Com_{QR} . Similar to in the Damgård-Fujisaki commitments, B is such that 2^B is larger than the order of $|QR_{n^2}|$. We show that such commitments are hiding and binding in Appx. F.4. We can see that these $|QR_{n^2}|$ commitments are multiplicatively homomorphic, i.e. if you take two $|QR_{n^2}|$ commitments $c = (c_1, c_2)$ committing to element M and $d = (d_1, d_2)$ committing to element N , then if you compute their pair-wise multiplication: $e = (c_1 * d_1, c_2 * d_2)$, this results in a commitment to $M * N$ with opening information $s_c + s_d, r_c + r_d$, computed pair-wise, where s_c, r_c is the opening information for c and s_d, r_d is the opening information for d .

Auxiliary proofs for commitments to $|QR_{n^2}|$ We now describe protocols that we can use to create proofs of opening, multiplication, and exponentiation of elements in $|QR_{n^2}|$ which can be verified using only their commitments.

¹⁵ We could use Damgård-Fujisaki commitments as-is (such that they live in \mathbb{Z}_n), but our $|QR_{n^2}|$ commitments would then consist of elements in \mathbb{Z}_{n^2} and \mathbb{Z}_n , requiring a new *eqrep* protocol that spans both groups. It is not clear if this alternative approach would be more efficient or simpler.

Fig. 4.3: $|QR_{n^2}|$ -Commitments

Setup _{QR} (1^λ) \rightarrow $params_{QR}$	Com _{QR} ($params, M$) \rightarrow (C, O)
1: Sample a safe RSA modulus, $n = pq = (2p' + 1)(2q' + 1)$ 2: Sample random $g \leftarrow QR_{n^2} $ 3: Sample random $(g', h') \leftarrow \mathbb{Z}_{n^2}$ 4: return $params = (n, g, g', h')$	1: $(s, r) \leftarrow_{\S} [2^{B+\lambda}]$ 2: $C = (Mg^s , (g')^s (h')^r)$ 3: return (C, O) where $O = (s, r)$

Proof of knowledge of opening for $|QR_{n^2}|$ -commitments. Comparing Fig. 4.3 to Fig. 4.2, we can see that the second part of a $|QR_{n^2}|$ commitment is simply a Damgård-Fujisaki commitment that lives in \mathbb{Z}_{n^2} . Using the Damgård-Fujisaki commitment opening proof protocol [DF02], we can create a proof of opening of the second part of the commitment which suffices as a proof of opening for a $|QR_{n^2}|$ commitment as we can extract s, r from C_2 and compute: $M = |C_1/(g^s)|$.

Proof of multiplication of $|QR_{n^2}|$ -commitments. We show how to prove knowledge of multiplication of committed $|QR_{n^2}|$ elements by utilizing the homomorphic property of the commitments. Given three commitments, C_1, C_2, C_3 , committing to $|QR_{n^2}|$ elements E_1, E_2, E_3 (where each commitment consists of two elements of $|QR_{n^2}|$, $C_i = (C_{i,1}, C_{i,2})$), we prove that a fourth commitment $C_4 = (C_{4,1}, C_{4,2})$ is a commitment to $1 \in |QR_{n^2}|$, where $C_{4,1} = C_{1,1}/(C_{2,1}C_{3,1})$ and $C_{4,2} = C_{1,2}/(C_{2,2}C_{3,2})$ – the verifier can compute C_4 using $(C_i)_{i \in [3]}$. This is equivalent to proving multiplication because of the homomorphic properties of the relation and can be proven using *eqrep*- \mathbb{Z}_{n^2} from Sect. 2.1 using relation $R((\gamma_1, \gamma_2, \beta_1, \beta_2), (C_{4,1}, C_{4,2})) = 1$ iff $C_{4,1} = \beta_1 g^{\gamma_1} \wedge C_{4,2} = \beta_2 (g')^{\gamma_1} (h')^{\gamma_2} \wedge \beta_1 \in \{-1, 1\} \wedge \beta_2 \in \{-1, 1\}$. This proves that $|C_4| = (|g^{\gamma_1}|, (g')^{\gamma_1} (h')^{\gamma_2})$ which is a commitment to 1. The prover uses $\gamma_1 = s_1 - s_2 - s_3$ and $\gamma_2 = r_1 - r_2 - r_3$ to satisfy this relation, where (s_i, r_i) is the opening of C_i .

Proof of exponentiation of $|QR_{n^2}|$ -commitments with Damgård-Fujisaki commitments. We prove this with *eqrep*- \mathbb{Z}_{n^2} from Sec. 2.1. This proof operates over two commitments C_1, C_2 to $|QR_{n^2}|$ elements E_1, E_2 (resp.) and one commitment C_y to scalar y and proves that $E_1 = E_2^y$. First let $C_1 = (C_{1,1}, C_{1,2})$, $C_2 = (C_{2,1}, C_{2,2})$ and $C_y = (g')^y (h')^{r_y}$. This can be proven with relation $R((\gamma_1, \gamma_2, \beta_1, \beta_2), (C_1, C_2, C_y)) = 1$ iff $C_{1,1} = \beta_1 C_{2,1}^y g^{\gamma_1} \wedge C_{1,2} = \beta_2 C_{2,2}^y g^{\gamma_1} h^{\gamma_2} \wedge \beta_1 \in \{-1, 1\} \wedge \beta_2 \in \{-1, 1\}$. The Prover uses $\gamma_1 = s_1 - y s_2$ and $\gamma_2 = r_1 - y r_2$ to satisfy this relation. If the prover can open C_2 then, $C_{1,1} = \beta_1 E_2^y g^{y s_2 + \gamma_1}$ and $C_{1,2} = \beta_2 (g')^{y s_2 + \gamma_1} (h')^{y r_2 + \gamma_2}$ which is exactly a commitment to $|E_2^y|$.

Commitments to Camenisch-Shoup encryptions Since we constructed commitments to elements of $|QR_{n^2}|$ along with their associated proof protocols,

we can use these commitments with Camenisch-Shoup ciphertexts. We present the full construction in Appx. F.5 and describe it in Fig. 4.4. The Prove functions in Fig. 4.4 satisfy the relations between committed ciphertexts described in Sec. 3.1. Similar to Sect. 3.2, we underline elements known to the verifier so that the reader can intuit how the proof is verified. We prove Theorem 5 in Appx. F.6.

- $\text{Setup}_{CS}(1^\lambda) \rightarrow \text{params}$: Generates the parameters for a Camenisch-Shoup encryption scheme, params_{CS} from Fig. 4.1b, the parameters for a $|QR_{n,2}|$ commitment scheme, params_{QR} from Fig. 4.3, and for a scalar commitment scheme, params_{DF} .
- $\text{Com}_{CS}(\text{params}, c) \rightarrow (C, O)$: Takes in a Camenisch-Shoup encryption, $c = (c_1, c_2)$ and uses the $|QR_{n,2}|$ commitment scheme to commit to each element, $(C_1, O_1) = \text{Com}_{QR}(c_1)$, $(C_2, O_2) = \text{Com}_{QR}(c_2)$ with the opening being $O = (O_1, O_2)$.
- $\text{Prove}_{CS}^{\text{ComAH}}(\text{params}, c, M, O, \underline{C}) \rightarrow \pi$: Parse the encryption into two $|QR_{n,2}|$ elements, $c = (c_1, c_2)$. Perform a proof of opening of both $|QR_{n,2}|$ commitments (C_1 and C_2) to c_1 and c_2 .
- $\text{Prove}_{CS}^{\oplus}(\text{params}, c_a, c_b, c_c, O_a, O_b, O_c, \underline{C_a}, \underline{C_b}, \underline{C_c}) \rightarrow \pi$: Parse the encryption, $c_a = (c_{a,1}, c_{a,2})$ and similarly parse c_b and c_c . Compute two $|QR_{n,2}|$ multiplication proofs over C_a, C_b , and C_c proving that $c_{a,1} * c_{b,1} = c_{c,1}$ and $c_{a,2} * c_{b,2} = c_{c,2}$. Output these proofs as π .
- $\text{Prove}_{CS}^{\ominus}(\text{params}, c_a, c_b, O_a, O_b, r_y, y, \underline{C_a}, \underline{C_b}, \underline{C_y}) \rightarrow \pi$: Let $C_y = \text{Com}(y; r_y)$. Similar to the $\text{Prove}_{CS}^{\oplus}$ function, parse c_a and c_b into their $|QR_{n,2}|$ elements. Perform two proofs of exponentiation of $|QR_{n,2}|$ commitments for $c_{b,1}^y = c_{a,1}$ and $c_{b,2}^y = c_{a,2}$. Output these proofs as π .
- $\text{Prove}_{CS}^{\text{enc}}(\text{params}, \text{pk}_{AH}, c_a, r_a, O_a, y, r_y, \underline{C_a}, \underline{C_y}) \rightarrow \pi$: Parse the encryption as $c_a = (g^y h^{r_a}, g^{r_a})$ where g, h are the parameters for the encryption scheme. Parse the commitments similar to previous functions. Commit to the generator (g) and public key (h) of the Camenisch-Shoup scheme, yielding C_g and C_h . Use these $|QR_{n,2}|$ commitments to prove that $c_{a,1} = g^y h^{r_a}$ and that $c_{a,2} = g^{r_a}$. Output these proofs as π .

Fig. 4.4: Summary of Camenisch-Shoup commitments

Theorem 5 (Security of Camenisch-Shoup commitments) *The construction in Fig. F.4 in Appx. F.5 (partially described in Fig. 4.4) satisfies four properties: (1) statistically hiding; (2) computationally binding; (3) our protocols in Fig. F.4 ($\text{Prove}_{CS}^{\text{ComAH}}$, $\text{Prove}_{CS}^{\text{enc}}$, $\text{Prove}_{CS}^{\ominus}$, and $\text{Prove}_{CS}^{\oplus}$) are computationally zero-knowledge; and (4) computationally black-box knowledge extractable assuming that the Strong RSA and Decisional Composite Residuosity assumptions hold.*

5 Applications of our Framework to Privacy-Preserving Blueprints

Given this new efficient framework for verifiable computation on ciphertexts, we are now equipped to build a PPB scheme that supports the f_{CBDC} functionality, has more succinct escrow proofs, and can withstand the framing attack in

Sec. 1.2. We present the non-frameability result first as this allows us to recall the generic construction of PPB from homomorphic-enough encryption (HEC), which one can think of as a passively secure PPB. In Sec. 5.2 we instantiate HEC using additively homomorphic encryption in such a way that Camenisch-Shoup gives us a HEC for f_{CBDC} , and in Sec. 5.3 we show a succinct proof system which ensures escrows are created honestly.

5.1 Non-Frameable Privacy-Preserving Blueprints

To systematically prevent framing attacks and formally define and prove non-frameability, we extend the formal definition of a blueprint scheme of [KLN23], see Sect. 2.3. We change the Dec algorithm to additionally output a proof of correct decryption and introduce a Judge algorithm for verifying this proof.

Definition 5 (A non-frameable f -blueprint scheme). *For a non-interactive commitment scheme $(\text{CSetup}, \text{Com})$, a non-frameable f -blueprint scheme consists of all the algorithms of a basic f -blueprint scheme with an adapted Dec algorithm and an additional Judge algorithm:*

- $\text{Dec}(\Lambda, \text{sk}_A, C_y, Z) \rightarrow (f(x, y), \pi_z)$ or \perp : Takes the auditor’s secret key sk_A , commitment C_y and escrow Z such that $\text{VerEscrow}(\Lambda, \text{pk}_A, C_y, Z) = 1$ as input. Decrypts the escrow and returns the output $f(x, y)$ if C_y is a commitment to y . Additionally it returns a proof, π_z , that proves to the Judge algorithm that $f(x, y)$ was decrypted correctly from Z .
- $\text{Judge}(\Lambda, \text{pk}_A, C_x, C_y, Z, z, \pi_z) \rightarrow 0$ or 1 : Takes as input the public key of the auditor, pk_A , the commitment to the watchlist and user data, C_x and C_y , the escrow, Z , the decrypted value, z , a proof of correct decryption, π_z and verifies that z was obtained correctly from escrow Z .

Intuitively, non-frameability requires the Judge algorithm to only accept valid results.

Definition 6 (Non-Frameability). *Let C_x and C_y be commitments computed from (x, r_x) and (y, r_y) respectively. Non-frameability guarantees that any pk_A , Z, z, π_z that passes $\text{Judge}(\Lambda, \text{pk}_A, C_x, C_y, Z, z, \pi_z)$ will imply that $f(x, y) = z$ with overwhelming probability. More formally, for all PPT adversaries \mathcal{A} , there exists a negligible function ν such that: $\Pr[\text{NonFraming}_{\text{BU}}^{\text{Adv}}(\lambda) = 1] < \nu(\lambda)$*

KLN uses a “homomorphic-enough” encryption (HEC) scheme to generically construct their PPB scheme. We extend their construction by adapting Dec and adding Judge in Fig. 5.4.

The KLN HEC scheme is parameterized by a function family and is correct if it is possible to compute any function from that family using only the ciphertexts. Existing HEC schemes that are only correct as defined by KLN will not be sufficient to construct non-frameable blueprint schemes.

Our main insight for adapting the generic construction of blueprints from a HEC scheme is that the adversary now controls the randomness r to the

NonFraming ^{Adv} _{Blu} (λ)
1 : $cpar \leftarrow \text{CSetup}(1^\lambda)$
2 : $A \leftarrow \text{Setup}(1^\lambda, cpar)$
3 : $(pk_A, x, r_x, y, r_y, Z, z, \pi_z) \leftarrow \mathcal{A}(1^\lambda, A)$
4 : $C_x = \text{Com}_{cpar}(x, r_x); C_y = \text{Com}_{cpar}(y, r_y)$
5 : return $[(\text{Judge}(A, pk_A, C_x, C_y, Z, z, \pi_z) = 1) \wedge (f(x, y) \neq z)]$

 Fig. 5.1: Experiments NonFraming^{Adv}_{Blu}(λ)

HEC encryption algorithm, in addition to the randomness r_Z , and can thus exercise additional control over the output of HECENC. Thus, we need to create a definition that is stronger than correctness that we refer to as *HEC consistency*. We will construct such HEC schemes using AHE in Sec. 5.2.

Definition 7 (Consistent homomorphic-enough cryptosystem (HEC) for a function family). Let $F = \{f \mid f : \text{domain}_{f,x} \times \text{domain}_{f,y} \mapsto \text{range}_f\}$ be a set of polynomial-time computable functions. We say that algorithms $\text{HEC} = (\text{HECSETUP}, \text{HECENC}, \text{HECEVAL}, \text{HECDEC}, \text{HECDIRECT})$ constitute a HEC for F if they satisfy the input-output, consistency, and security requirements below:

- $\text{HECSETUP}(1^\lambda) \rightarrow \text{hecpa}$ takes the security parameter as input, outputs the parameters hecpa .
- $\text{HECENC}(\text{hecpa}, f, x) \rightarrow (X, d)$ takes parameters hecpa , a function $f \in F$, and a value $x \in \text{domain}_{f,x}$ as input, outputs an encrypted representation X of the function $f(x, \cdot)$, and a decryption key d .
- $\text{HECEVAL}(\text{hecpa}, f, X, y) \rightarrow Z$ takes as input the parameters hecpa , a function $f \in F$, an encrypted representation of $f(x, \cdot)$, and a value $y \in \text{domain}_{f,y}$ and outputs a ciphertext Z , an encryption of $f(x, y)$.
- $\text{HECDEC}(\text{hecpa}, d, Z) \rightarrow z$ takes as input the parameters hecpa , the decryption key d , and a ciphertext Z , decrypts Z to obtain a value z .
- $\text{HECDIRECT}(\text{hecpa}, X, z) \rightarrow Z$ on input hecpa , an encrypted representation X of some function, and a value z , outputs a ciphertext Z .

 Fig. 5.2: Algorithms of HEC scheme for F

HEC consistency. For a given adversary Adv and HEC, let $\text{Adv}_{\text{HEC}, \text{Adv}}$ be the probability that the experiment HECCONSISTENT in Fig. 5.3 accepts. HEC is *consistent* if $\text{Adv}_{\text{HEC}, \text{Adv}}$ is negligible for all PPT algorithms Adv .

HEC security. We provide the formal definitions for the Security of x , security of x and y from third parties, and security of DIRECTZ in Appx. D.1.

Relying on HEC consistency we reprove Theorem 2 of [KLN23].

$\text{HECCONSISTENT}^{\text{Adv}}(\lambda)$	$\text{HECCORRECT}^{\text{Adv}}(\lambda)$
1 : $\text{hecpa} \leftarrow \text{HECSETUP}(\lambda)$	1 : $\text{hecpa} \leftarrow \text{HECSETUP}(\lambda)$
2 : $(f, x, r, y, rz) \leftarrow \text{Adv}(1^\lambda, \text{hecpa})$	2 : $(f, x, \text{st}) \leftarrow \text{Adv}(1^\lambda, \text{hecpa})$
3 : if $f \notin F \vee x \notin \text{domain}_{f,x} \vee y \notin \text{domain}_{f,y}$	3 : if $f \in F, x \in \text{domain}_{f,x}$
4 : return 0	4 : $(X, d) \leftarrow \text{HECENC}(\text{hecpa}, f, x)$
5 : $(X, d) \leftarrow \text{HECENC}(\text{hecpa}, f, x; r)$	5 : $(y, rz) \leftarrow \text{Adv}(\text{st}, X)$
6 : $Z \leftarrow \text{HECEVAL}(\text{hecpa}, f, X, y, rz)$	6 : if $y \in \text{domain}_{f,y}$
7 : if $\text{HECDEC}(\text{hecpa}, d, Z) \neq f(x, y)$	7 : $Z \leftarrow \text{HECEVAL}(\text{hecpa}, f, X, y, rz)$
8 : return 1	8 : if $\text{HECDEC}(\text{hecpa}, d, Z) \neq f(x, y)$
9 : return 0	9 : return 1
	10 : return 0

Fig. 5.3: The HEC consistency game and for comparison the original correctness game of KLN: In the HEC consistency game, the adversary outputs x, y , and the randomness for the HEC scheme (r, rz) , and the encryption and evaluation algorithms cannot produce a ciphertext that decrypts to a plaintext other than $f(x, y)$.

Theorem 6 *If HEC is a consistent and secure homomorphic-enough cryptosystem, the commitment scheme is binding, and the NIZK PoKs Ψ_1, Ψ_2 and Ψ_3 are zero-knowledge and BB-PSL simulation extractable then our generic blueprint scheme is a secure, non-frameable f -blueprint scheme.*

Proof. Since the property of HEC consistency implies HEC correctness, the proofs of correctness of VerEscrow, VerPK and Dec from the original PPB proof of KLN remains unchanged. Similarly, the *Soundness* of the generic f -blueprint scheme is also proven using the BB-Extractability of the NIZK Ψ_2 using the same reduction technique as KLN. The proofs for the properties of *Blueprint Hiding*, *Privacy against dishonest auditor* and *Privacy against honest auditor* do not involve the decryption proof and are unchanged from KLN. We thus focus on the proof of non-frameability.

Lemma 1 *Let Ψ_1, Ψ_2 , and Ψ_3 be a BB extractable NIZK schemes, let $(\text{CSetup}, \text{Com})$ be a computationally binding commitment scheme, and HEC be consistent, then our proposed scheme achieves Non-frameability.*

Proof. Consider Fig. 5.1. Suppose, for the sake of contradiction, that there exists a PPT adversary \mathcal{A} such that $\text{Adv}_{\mathcal{A}, \text{Blu}}^{\text{NonFraming}} = \nu(\lambda)$ is non negligible. Let $(\text{pk}_A, x, r_x, y, r_y, Z, z, \pi_z)$, be the adversary's output in the experiment for which $C_x = \text{Com}_{\text{cpar}}(x, r_x)$ and $C_y = \text{Com}_{\text{cpar}}(y, r_y)$.

We first extend the experiment by running the extractors for the proofs π_A for Ψ_1 in pk_A , π_U for Ψ_2 in Z and π_z for Ψ_3 . If any of the extractors fail with non-negligible probability we break the BB extractability of the corresponding

<p>Setup($\lambda, cpar, S_1, S_2, S_3$)</p> <pre> 1 : $hecpair \leftarrow \text{HECSETUP}(1^\lambda)$ 2 : return $\Lambda = (\lambda, cpar, hecpair, S_1, S_2, S_3)$ </pre> <hr/> <p>KeyGen(Λ, x, r_x)</p> <pre> 1 : parse $\Lambda = (\lambda, cpar, hecpair, S_1, S_2, S_3)$ 2 : $(X, d) \xleftarrow{r} \text{HECENC}(hecpair, f, x)$ 3 : $C_x \leftarrow \text{Com}_{cpar}(x; r_x); C_d \xleftarrow{r_d} \text{Com}_{cpar}(d)$ 4 : $\pi_A \leftarrow \text{PoK}_{\Psi_1}^{S_1} \{ (x, d, r, r_x, r_d) :$ 5 : $(X, d) = \text{HECENC}(hecpair, f, x; r)$ 6 : $\wedge C_x = \text{Com}_{cpar}(x; r_x)$ 7 : $\wedge C_d = \text{Com}_{cpar}(d; r_d) \}$ 8 : $\text{pk}_A \leftarrow (X, C_x, C_d, \pi_A); \text{sk}_A \leftarrow (\text{pk}_A, d, r_d)$ 9 : return $(\text{pk}_A, \text{sk}_A)$ </pre> <hr/> <p>VerPK($\Lambda, \text{pk}_A, C_x$)</p> <pre> 1 : parse $\Lambda = (\lambda, cpar, hecpair, S_1, S_2, S_3)$ 2 : parse $\text{pk}_A = (X, C'_x, \pi_A)$ 3 : return $\text{V}_1^{S_1}((X, hecpair, f, C_x, cpar), \pi_A)$ 4 : $\wedge (C'_x = C_x)$ </pre> <hr/> <p>Judge($\Lambda, \text{pk}_A, C_x, C_y, Z = (\hat{Z}, \pi_U), z, \pi_Z$)</p> <pre> 1 : parse $\Lambda = (\lambda, cpar, hecpair, S_1, S_2, S_3)$ 2 : parse $\text{pk}_A = (_, _, C_d, _)$ 3 : return $\text{V}_3^{S_3}((z, hecpair, \hat{Z}, C_d), \pi_Z)$ 4 : $\wedge \text{VerPK}(\Lambda, \text{pk}_A, C_x)$ 5 : $\wedge \text{VerEscrow}(\Lambda, \text{pk}_A, C_y, Z)$ </pre>	<p>Escrow($\Lambda, \text{pk}_A, y, r_y$)</p> <pre> 1 : parse $\Lambda = (\lambda, cpar, hecpair, S_1, S_2, S_3)$ 2 : parse $\text{pk}_A = (X, C_x, _)$ 3 : if $\text{VerPK}(\Lambda, \text{pk}_A, C_x) = 0$ 4 : return 0 5 : $\hat{Z} \xleftarrow{r_{\hat{Z}}} \text{HECEVAL}(hecpair, f, X, y)$ 6 : $C_y \leftarrow \text{Com}_{cpar}(y; r_y)$ 7 : $\pi_U \leftarrow \text{PoK}_{\Psi_2}^{S_2} \{ (y, r_y, r_{\hat{Z}}) :$ 8 : $\hat{Z} = \text{HECEVAL}(hecpair, f, X, y; r_{\hat{Z}})$ 9 : $\wedge C_y = \text{Com}_{cpar}(y; r_y) \}$ 10 : return (\hat{Z}, π_U) </pre> <hr/> <p>VerEscrow($\Lambda, \text{pk}_A, C_y, Z = (\hat{Z}, \pi_U)$)</p> <pre> 1 : parse $\Lambda = (\lambda, cpar, hecpair, S_1, S_2, S_3)$ 2 : parse $\text{pk}_A = (_, C_x, _, _)$ 3 : return $\text{VerPK}(\Lambda, \text{pk}_A, C_x)$ 4 : $\wedge \text{V}_2^{S_2}((\hat{Z}, hecpair, f, X, C_y, cpar), \pi_U)$ </pre> <hr/> <p>Dec($\Lambda, \text{sk}_A, C_y, Z = (\hat{Z}, \pi_U)$)</p> <pre> 1 : parse $\Lambda = (\lambda, cpar, hecpair, S_1, S_2, S_3)$ 2 : parse $\text{sk}_A = (\text{pk}_A, d, r_d)$ 3 : parse $\text{pk}_A = (_, _, C_d, _)$ 4 : if $\text{VerEscrow}(\Lambda, \text{pk}_A, C_y, Z) = 0$ 5 : return \perp 6 : $z \leftarrow \text{HECDEC}(hecpair, d, \hat{Z})$ 7 : $\pi_Z \leftarrow \text{PoK}_{\Psi_3}^{S_3} \{ d, r_d :$ 8 : $z = \text{HECDEC}(hecpair, d, \hat{Z})$ 9 : $\wedge C_d = \text{Com}_{cpar}(d; r_d) \}$ 10 : return (z, π_Z) </pre>
--	---

Fig. 5.4: Construction of generic f -blueprint scheme from HEC and NIZK PoKs Ψ_1, Ψ_2 and Ψ_3 with setup S_1, S_2 , and S_3 respectively.

NIZK scheme. Otherwise, we obtain witnesses $(x', d, r, r'_x, r_d), (y', r'_y, r_{\hat{Z}})$, and (d', r'_d) for which $(X, d) = \text{HECENC}(hecpair, f, x'; r), C_x = \text{Com}_{cpar}(x'; r'_x), C_d = \text{Com}_{cpar}(d; r_d), \hat{Z} = \text{HECEVAL}(hecpair, f, X, y'; r_{\hat{Z}}), C_y = \text{Com}_{cpar}(y'; r'_y), z = \text{HECDEC}(hecpair, d', \hat{Z})$ and $C_d = \text{Com}_{cpar}(d'; r'_d)$.

The events where \mathcal{A} outputs 1 can be divided into four cases: (i) when $x \neq x'$ (ii) when $x = x'$ but $y \neq y'$, (iii) the case where $x = x'$ and $y = y'$ but $d \neq d'$ (iv) $x = x', y = y'$ and $d = d'$. We express the probabilities of these events with the

functions $\nu_0(\lambda)$, $\nu_1(\lambda)$, $\nu_2(\lambda)$, and $\nu_3(\lambda)$ respectively. Since $\nu(\lambda)$ is non negligible and these three events covers all cases where Adv would output 1, at least one of $\nu_0(\lambda)$, $\nu_1(\lambda)$, $\nu_2(\lambda)$ or $\nu_3(\lambda)$ must be non negligible.

Suppose $\nu_1(\lambda)$ is non negligible. By the structure of the experiment and the validity of the witness we know that $C_x = \text{Com}_{\text{par}}(x; r_x)$ and $C_x = \text{Com}_{\text{par}}(x'; r'_x)$, and we can thus construct an adversary against the binding property of the commitment by outputting $(x, r_x, x'; r'_x)$. By the same argument we show that $\nu_2(\lambda)$ and $\nu_3(\lambda)$ are negligible. Finally, in the fourth case, we have that $(X, d) = \text{HECENC}(\text{heccpar}, f, x; r)$ and $\hat{Z} = \text{HECEVAL}(\text{heccpar}, f, X, y; r_{\hat{Z}})$, and $f(x, y) \neq \text{HECDEC}(\text{heccpar}, d, \hat{Z})$. We can thus construct an adversary against the consistency property of the HEC scheme by outputting $(f, x, r, y, r_{\hat{Z}})$.

5.2 Instantiation of Consistent HEC Schemes

It turns out that the FHE based HEC scheme of [KLN23] is already consistent, see Appx. D.3.

In this section, we provide a consistent HEC scheme for realizing a non-frameable CBDC and watchlist PPB scheme. In Fig. 5.5 we construct the algorithms HECEVAL and HECDEC for the function family $\{f_{n,k}\}_{n,k \in \mathbb{Z}}$, where n is the length of the auditor's list $x = \{x_1, \dots, x_n\}$ and k is the bit length of the user's attribute y_{at} , where the user's input consists of the user's identifier y_{id} and an attribute: $y = (y_{id}, y_{at})$. $f_{n,k}$ is defined as $f_{n,k}(x, y) = y$ if $y_{id} \in x$ and $f_{n,k} = \emptyset$ otherwise. We discuss why this watchlist function is useful for the watchlist/CBDC application in Sec. 1. y_{id} uniquely identifies a user and y_{at} could be any useful data about the user such as a seed for tracing the user's anonymous payments. We construct a HECEVAL algorithm for multiple attributes in Appx. D.2 and show the full construction of our HEC scheme in Appx. D.3.

Overview of the construction. The HECENC algorithm (Fig. 5.5) takes as input the list x of n watchlisted identities, and computes a polynomial $P(\chi) = \sum_{i \in [n]} a_i \chi^i$ such that $P(y_{id}) = 0$ if and only if $y_{id} \in x$. Then, it samples a key pair $(\text{pk}_{AH}, \text{sk}_{AH})$ for a semantically secure \mathfrak{g} -semi-encryption scheme (Def. 4), and outputs the public key $X = (\text{pk}_{AH}, \{A_i = \text{Enc}(\text{pk}_{AH}, a_i)\}_{i \in [0..n]})$ where the a_i 's are coefficients of P , and the decryption key $d = (\text{sk}_{AH}, x)$.

On input the public key X and the value $y = (y_{id}, y_{at})$, HECEVAL will output the escrow $Z = (Z_{id}, Z_{at}, Z_{nf})$ which consists of three ciphertexts under the key pk_{AH} ; these will decrypt to the values $(y_{id}, y_{at}, 0)$ if and only if $y_{id} \in x$; otherwise they will decrypt to uniformly random elements of the message space, independent of y . As we show in more detail in Fig. 5.5, additively homomorphic properties of the underlying (semi-)encryption scheme allow the evaluator to form the ciphertext E so that it will be an encryption of $P(y_{id})$. The evaluator also encrypts the identity y_{id} and attribute y_{at} , yielding ciphertexts Y_{id} and Y_{at} . The escrow of y_{id} is then formed as $Z_{id} = (r_1 \odot E) \oplus Y_{id} = ((r_1 \odot \overline{P(y_{id})}) \oplus \overline{y_{id}}) = \overline{r_1 P(y_{id}) + y_{id}}$, which is an encryption of y_{id} if E is an encryption of 0 (i.e. whenever $y_{id} \in x$), and an encryption of a random value otherwise,

thanks to the randomizer r_1 . Similarly, the escrow of y_{at} is $Z_{at} = (r_2 \odot E) \oplus Y_{at} = \boxed{r_2 P(y_{id}) + y_{at}}$. To make the HEC consistent, we include $Z_{nf} = r_3 \odot E = \boxed{r_3 P(y_{id})}$, which will decrypt to 0 if and only if $y_{id} \in x$.

HECDEC takes as input the HEC decryption key $d = (\text{sk}_{AH}, x)$ and the escrow Z . It recovers y'_{id}, y'_{at} , and y' by decrypting the escrows (Z_{id}, Z_{at}, Z_{nf}) using the secret key, sk_{AH} . By the correctness property the decryption algorithm for \mathbf{g} -semi-encryption, we know that for $Z \in \text{HECENC}(X, y)$, $y' = \mathbf{g}(r_3 P(y_{id})) = \mathbf{g}(0)$ if and only if $y_{id} \in x$; so if $y' \neq \mathbf{g}(0)$, HECDEC outputs \perp . Else, we know that $y_{id} \in x$, so HECDEC must somehow determine (1) y_{id} from $y'_{id} = \mathbf{g}(y_{id})$, and (2) y_{at} from $y'_{at} = \mathbf{g}(y_{at})$. Let us explain how HECDEC can do so.

If \mathbf{g} is the identity function then this step is trivial; we will show in Sec. 4 that we can achieve an additively homomorphic \mathbf{g} -semi-encryption scheme where \mathbf{g} is the identity function under the decisional composite residuosity assumption using the Camenisch-Shoup cryptosystem.

If, however, \mathbf{g} is a one-way injective function, then (1) can be done by looking for $\mathbf{g}(y_{id})$ on the list $\mathbf{g}(x_1), \dots, \mathbf{g}(x_n)$ where $x_i \in x$ and (2) can only be done by exhaustive search, which is only possible if y_{at} comes from a small space. This is the approach that was (implicitly) taken by the original PPB paper of Kohlweiss et al.: since the ElGamal cryptosystem is only additively homomorphic when viewed as a \mathbf{g} -semi-encryption scheme, and \mathbf{g} is a one-way function, they could only achieve attributes from a small space.

HECEVAL($hecp\text{ar}, f_{n,k}, X, y; r_{\hat{Z}}$)	HECDEC($hecp\text{ar}, d, Z$)
1 : parse $X = (\text{pk}_{AH}, A_0, \dots, A_n)$, $y = (y_{id}, y_{at}), \quad // y_{at} = k$ $r_{\hat{Z}} = (r_{id}, r_{at}, r_1, r_2, r_3)$	1 : parse $d = (\text{sk}_E, f_{n,k}, x)$, $Z = (Z_{id}, Z_{at}, Z_{nf})$
2 : $E \leftarrow \bigoplus_{i=0}^n (A_i \odot y_{id}^i)$	2 : $y'_{id} \leftarrow \text{Dec}(\text{sk}_E, Z_{id})$
3 : $Y_{id} \leftarrow \text{Enc}(\text{pk}_{AH}, y_{id}; r_{id})$	3 : $y'_{at} \leftarrow \text{Dec}(\text{sk}_E, Z_{at})$
4 : $Y_{at} \leftarrow \text{Enc}(\text{pk}_{AH}, y_{at}; r_{at})$	4 : $y' \leftarrow \text{Dec}(\text{sk}_E, Z_{nf})$
5 : $Z_{id} \leftarrow (r_1 \odot E) \oplus Y_{id}$	5 : if $y' \neq \mathbf{g}(0)$, return \emptyset
6 : $Z_{at} \leftarrow (r_2 \odot E) \oplus Y_{at}; Z_{nf} = r_3 \odot E$	6 : return (y_{id}, y_{at}) where $\mathbf{g}(y_{at}) = y'_{at} \wedge \mathbf{g}(y_{id}) = y'_{id}$ $\wedge y_{id} \in x \wedge y_{at} \in \text{domain}_{f,y,at}$
7 : return $Z = (Z_{id}, Z_{at}, Z_{nf})$	

Fig. 5.5: HEC algorithms: we omit the functions HECENC and HECDIRECT as they are unchanged from KLN and instead present them in Appx. D.3.

Theorem 7 (Security of the construction in Fig. 5.5) *Assuming our AHE scheme is secure, our construction in Fig. D.3 in Appx. D.3 (partially described in Fig 5.5) achieves HEC consistency as defined in Def. 5.3; and security of y ,*

security of x and y from third parties, and security of DIRECTZ, as defined in Def. 14.

Proof of Thm. 7 We focus on consistency and prove the remaining properties as Lemmas 10, 11, and 12 in Appx. D.3 following KLN. Because we include $Z_{nf} = E \odot r_3$ in the escrow, an auditor can prove that this is an encryption of 0. This ensures that the y_{id} is actually on the watchlist as the polynomial has roots at each entry of the watchlist. Formally, if an adversary were to be able to produce a (f, x, st, r, y, r_Z) such that $Z \leftarrow \text{HECEVAL}(hecp, f, X, y; r_Z)$ but $\text{HECDEC}(hecp, d, Z) \neq f(x, y)$, we see that $E \odot r_3 = 0$ in this case, which implies that $r_3 P(y) = 0$. This is only true if $y \in x$ since $r_3 > 0$. In this case, because HECEVAL is proven to be correctly computed, $E \odot r_1$ decrypts to 0. Thus, $y' = \text{Dec}(Y)$. Thus, this decrypts to the correct value.

5.3 Efficient Instantiation of HEC Evaluation Proof Ψ_2

In this section we show how to use the techniques introduced in Sect. 3.2 to efficiently instantiate a NIZK proof used in the Escrow algorithm in Fig. 5.4 to compute π_U . This proof is for the following relation: $R_{\Psi_2}((y, r_y, r_{\hat{Z}}), (\hat{Z}, X, f_{n,k}, C_y)) = 1$ iff $\hat{Z} = \text{HECEVAL}(hecp, f_{n,k}, X, y; r_{\hat{Z}}) \wedge C_y = \text{Com}(y; r_y)$ where $f_{n,k}$ is the watchlist blueprinting function described at the start of this section. We can use similar techniques to that of Ψ_2 to construct proofs for Ψ_1 and Ψ_3 .

In Alg. 3, we give the construction of Ψ_2 for HECEVAL. This function calls the proof function for R_f from Sec. 3 on lines 11, 12, and 13 in order to prove correct computation of Z_{id}, Z_{at} , and Z_{nf} . Because of the succinctness of our proof for R_f , the complexity of our proof will be $O(\log(n))$ since we evaluate with a constant number of variables.

Our proof system must have the zero-knowledge and extractability properties needed for the proofs of both blueprint hiding (Def. 9) and user privacy (Def. 10 and 11) for our construction in Fig. 5.4. The zero-knowledge property is standard; for extractability recall that we require both the usual black-box proof of knowledge property, as well as partial straight-line extraction of $\mathbf{g}(y)$; \mathbf{g} is some function such that $\mathbf{g}(y)$, jointly with x is sufficient to compute $f(x, y)$ because there is some efficiently computable function f^* such that $f^*(x, \mathbf{g}(y)) = f(x, y)$. In order to achieve straight-line extractability of $\mathbf{g}(y)$, our proof system requires that the prover \mathbf{g} -semi-encrypt y under a public key “in the sky”, i.e. a public key that’s part of the parameters generated during setup; the knowledge extractor’s trapdoor will be the decryption key. To that end, we use a public-key \mathbf{g} -semi-encryption scheme $(\Gamma_{sky} = \{\text{KeyGen}_{sky}, \text{Enc}_{sky}, \text{Dec}_{sky}\})$. (Using our notation from Def. 2, the prover retrieves the public key in the sky by querying setup S_2 .)

We present the verification functions for PoK_{Ψ_2} and PoK_P^* in Appx. E.

We prove Thms. 8 and 9 in Appx. E.2.

Theorem 8 *Our scheme in Alg. 3 is complete and ZK (Def. 1).*

Algorithm 3 $\text{PoK}_{\Psi_2}^{S_2}(\text{hecpa}, f, X, y, r_y, r_{\hat{z}}) \rightarrow \pi$

parse $X = (\text{pk}_{AH}, \{\overline{a_i}\}_{i \in [0 \dots n]}); r_{\hat{z}} = (r_1, r_2, r_3, r_{id}, r_{attr})$
1: $(y_{id}, y_{at}) \leftarrow y; (C_{id}, r_{id}) = \text{Com}(y_{id}); (C_{at}, r_{at}) = \text{Com}(y_{at}); C_y \leftarrow \text{Com}(y; r_y)$
2: $Z_{id} = \text{Enc}_{AH}(\text{pk}_{AH}, y_{id}; r_{id}) \oplus (r_1 \odot \overline{a}); Z_{at} = \text{Enc}_{AH}(\text{pk}_{AH}, y_{at}; r_{attr}) \oplus (r_2 \odot \overline{a})$
3: $Z_{nf} = r_3 \odot \overline{a}; Z = (Z_{id}, Z_{at}, Z_{nf})$
4: $\text{pk}_{sky} \leftarrow S_2(1^\lambda); C_{sky} = \text{Enc}_{sky}(\text{pk}_{sky}, y; r_{sky});$
5: $\pi_{sky} = \text{NIZK}[y, r_y, r_{sky} : C_{sky} = \text{Enc}(\text{pk}_{sky}, y; r_{sky}) \wedge C_y = \text{Com}(y; r_y)]$
6: $\forall i \in [3], (C_{r_i}, \rho_i) = \text{Com}(r_i)$
7: $f_{id}(a_0, \dots, a_n, y_{id}, r_1) = y_{id} + a_0 r_1 y_{id}^0 + a_1 r_1 y_{id}^1 + \dots + a_n r_1 y_{id}^n$
8: $f_{at}(a_0, \dots, a_n, y_{id}, y_{at}, r_2) = y_{at} + a_0 r_2 y_{id}^0 + a_1 r_2 y_{id}^1 + \dots + a_n r_2 y_{id}^n$
9: $f_{nf}(a_0, \dots, a_n, y_{id}, r_3) = a_0 r_3 y_{id}^0 + a_1 r_3 y_{id}^1 + \dots + a_n r_3 y_{id}^n$
10: $\pi_y = \text{NIZK}[y_{id}, r_{id}, y_{at}, r_{at} : C_{id} = \text{Com}(y_{id}; r_{id}) \wedge C_{at} = \text{Com}(y_{at}; r_{at}) \wedge (y_{id}, y_{at}) = y \wedge C_y = \text{Com}(y; r_y)]$
11: $\pi_{id} = \text{NIZK}[y_{id}, r_{id}, r_1, \rho_{id} : Z_{id} = \text{Enc}_{AH}(f_{id}(a_0, \dots, a_n, y_{id}, r_1)) \wedge C_{id} = \text{Com}(y_{id}; r_{id}) \wedge C_{r_1} = \text{Com}(r_1; \rho_1)]$
12: $\pi_{at} = \text{NIZK}[y_{id}, r_{id}, y_{at}, r_{at}, r_2, \rho_2 : Z_{at} = \text{Enc}_{AH}(f_{at}(a_0, \dots, a_n, y_{id}, y_{at}, r_2)) \wedge C_{id} = \text{Com}(y_{id}; r_{id}) \wedge C_{at} = \text{Com}(y_{at}; r_{at}) \wedge C_{r_2} = \text{Com}(r_2; \rho_2)]$
13: $\pi_{nf} = \text{NIZK}[y_{id}, r_3 : Z_{nf} = \text{Enc}_{AH}(f_{nf}(a_0, \dots, a_n, y_{id}, r_3)) \wedge C_{id} = \text{Com}(y_{id}; r_{id}) \wedge C_{r_1} = \text{Com}(r_1; \rho_1)]$
14: **return** $(\pi_{id}, \pi_{at}, \pi_{nf}, \pi_{sky}, C_{sky}, \{C_{r_i}\}_{i \in [3]}, \pi_y)$

Theorem 9 (*g^* -BB-PSL for Ψ_2*) *If PoK_P^* is a BB NIZK for the relation R_P (where R_P is defined as $R_P((C, C_{y_{id}}, X, n), (O, O_{y_{id}}, y_{id})) = 1$ iff $C = \text{Com}_{AH}(\text{Enc}_{AH}(\text{pk}_{AH}, \bigoplus_{i=0}^n (\overline{a_i} \odot y_{id}^i); O)) \wedge C_{y_{id}} = \text{Com}(y, O_{y_{id}})$) and if $\Gamma_{sky} = \{\text{KeyGen}_{sky}, \text{Enc}_{sky}, \text{Dec}_{sky}\}$ is a semantically secure g -semi-encryption scheme, our Ψ_2 proof is a g^* -BB-PSL protocol, where $g^*(y, r_{\hat{z}}) = g(y)$.*

Acknowledgements

Anna Lysyanskaya and Scott Griffy were supported by NSF Grants 2312241, 2154170, and 2247305 as well as the Peter G. Peterson Foundation and the Ethereum Foundation. Markulf Kohlweiss and Meghna Sengupta were supported by Input Output (iohk.io) through their funding of the University of Edinburgh ZK Lab. We'd also like to acknowledge Victor Youdom Kemmoe and Eileen Nolan for their helpful discussions.

References

- ACC⁺22. Thomas Attema, Ignacio Cascudo, Ronald Cramer, Ivan Damgård, and Daniel Escudero. Vector commitments over rings and compressed Σ -protocols. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 173–202. Springer, Cham, November 2022.
- AHH⁺23. Nuttapong Attrapadung, Goichiro Hanaoka, Ryo Hiromasa, Takahiro Matsuda, and Jacob C. N. Schuldt. Maliciously circuit-private multi-key FHE and MPC based on LWE. *DCC*, 91(5):1645–1684, 2023.

- ATD16. Aydin Abadi, Sotirios Terzis, and Changyu Dong. VD-PSI: Verifiable delegated private set intersection on outsourced private datasets. In Jens Grossklags and Bart Preneel, editors, *FC 2016*, volume 9603 of *LNCS*, pages 149–168. Springer, Berlin, Heidelberg, February 2016.
- BBB⁺18. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018.
- BCC⁺16. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 327–357. Springer, Berlin, Heidelberg, May 2016.
- BCF⁺. Daniel Benarroch, Matteo Campanell, Dario Fiore, Jihye Kim, Jiwon Lee, Hyunok Oh, and Anaïs Querol. Proposal: Commit-and-prove zero-knowledge proof systems and extensions. <https://docs.zkproof.org/pages/standards/accepted-workshop4/proposal-commit.pdf>.
- BCL04. Endre Bangert, Jan Camenisch, and Anna Lysyanskaya. A cryptographic framework for the controlled release of certified data. In *Security Protocols Workshop*, volume 3957 of *Lecture Notes in Computer Science*, pages 20–42. Springer, 2004.
- BdMW16. Florian Bourse, Rafaël del Pino, Michele Minelli, and Hoeteck Wee. FHE circuit privacy almost for free. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 62–89. Springer, Berlin, Heidelberg, August 2016.
- BFK⁺24. Alexander R. Block, Zhiyong Fang, Jonathan Katz, Justin Thaler, Hendrik Waldner, and Yupeng Zhang. Field-agnostic SNARKs from expand-accumulate codes. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part X*, volume 14929 of *LNCS*, pages 276–307. Springer, Cham, August 2024.
- BG13. Stephanie Bayer and Jens Groth. Zero-knowledge argument for polynomial evaluation with application to blacklists. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 646–663. Springer, Berlin, Heidelberg, May 2013.
- BGJP23. James Bartusek, Sanjam Garg, Abhishek Jain, and Guru-Vamsi Policharla. End-to-end secure messaging with traceability only for illegal content. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 35–66. Springer, Cham, April 2023.
- BGV12. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012*, pages 309–325. ACM, January 2012.
- BHV⁺23. Rishabh Bhadauria, Carmit Hazay, Muthuramakrishnan Venkatasubramanian, Wenxuan Wu, and Yupeng Zhang. Private polynomial commitments and applications to MPC. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part II*, volume 13941 of *LNCS*, pages 127–158. Springer, Cham, May 2023.
- BL13. Foteini Baldimtsi and Anna Lysyanskaya. Anonymous credentials light. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 1087–1098. ACM Press, November 2013.

- BMM⁺21. Benedikt Bünz, Mary Maller, Pratyush Mishra, Nirvan Tyagi, and Psi Vesely. Proofs for inner pairing products and applications. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 65–97. Springer, Cham, December 2021.
- Boy. Dennis Boyle. The problem of “parallel construction” in criminal investigations. <https://www.boylejasari.com/the-problem-of-parallel-construction-in-criminal-investigations/>. Accessed: 2024-02-13.
- BSZ05. Mihir Bellare, Haixia Shi, and Chong Zhang. Foundations of group signatures: The case of dynamic groups. In Alfred Menezes, editor, *CT-RSA 2005*, volume 3376 of *LNCS*, pages 136–153. Springer, Berlin, Heidelberg, February 2005.
- BV11. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011.
- Cam97. Jan Camenisch. Efficient and generalized group signatures. In Walter Fumy, editor, *EUROCRYPT’97*, volume 1233 of *LNCS*, pages 465–479. Springer, Berlin, Heidelberg, May 1997.
- CBBZ23. Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang. HyperPlonk: Plonk with linear-time prover and high-degree custom gates. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part II*, volume 14005 of *LNCS*, pages 499–530. Springer, Cham, April 2023.
- CDN01. Ronald Cramer, Ivan Damgård, and Jesper Buus Nielsen. Multiparty computation from threshold homomorphic encryption. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 280–299. Springer, Berlin, Heidelberg, May 2001.
- CFN90. David Chaum, Amos Fiat, and Moni Naor. Untraceable electronic cash. In Shafi Goldwasser, editor, *CRYPTO’88*, volume 403 of *LNCS*, pages 319–327. Springer, New York, August 1990.
- Cha90. David Chaum. Showing credentials without identification transferring signatures between unconditionally unlinkable pseudonyms. In Jennifer Seberry and Josef Pieprzyk, editors, *AUSCRYPT’90*, volume 453 of *LNCS*, pages 246–264. Springer, Berlin, Heidelberg, January 1990.
- CHK⁺06. Jan Camenisch, Susan Hohenberger, Markulf Kohlweiss, Anna Lysyanskaya, and Mira Meyerovich. How to win the clonewars: Efficient periodic n-times anonymous authentication. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 201–210. ACM Press, October / November 2006.
- CHL05. Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. Compact e-cash. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 302–321. Springer, Berlin, Heidelberg, May 2005.
- CHL06. Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. Balancing accountability and privacy using e-cash (extended abstract). In Roberto De Prisco and Moti Yung, editors, *Proceedings of the 5th International Conference on Security and Cryptography for Networks (SCN)*, volume 4116 of *Lecture Notes in Computer Science*, pages 141–155. Springer, 2006.
- CL01. Jan Camenisch and Anna Lysyanskaya. An identity escrow scheme with appointed verifiers. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 388–407. Springer, Berlin, Heidelberg, August 2001.

- CL02. Jan Camenisch and Anna Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 61–76. Springer, Berlin, Heidelberg, August 2002.
- CL04. Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 56–72. Springer, Berlin, Heidelberg, August 2004.
- CM20. Melissa Chase and Peihan Miao. Private set intersection in the internet setting from lightweight oblivious PRF. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 34–63. Springer, Cham, August 2020.
- CMdG⁺21. Kelong Cong, Radames Cruz Moreno, Mariana Botelho da Gama, Wei Dai, Iliia Iliashenko, Kim Laine, and Michael Rosenberg. Labeled PSI from homomorphic encryption with reduced computation and communication. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1135–1150. ACM Press, November 2021.
- CRR21. Geoffroy Couteau, Peter Rindal, and Srinivasan Raghuraman. Silver: Silent VOLE and oblivious transfer from hardness of decoding structured LDPC codes. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part III*, volume 12827 of *LNCS*, pages 502–534, Virtual Event, August 2021. Springer, Cham.
- CS97. Jan Camenisch and Markus Stadler. Proof systems for general statements about discrete logarithms. Technical Report TR 260, Institute for Theoretical Computer Science, ETH Zürich, March 1997.
- CS03. Jan Camenisch and Victor Shoup. Practical verifiable encryption and decryption of discrete logarithms. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 126–144. Springer, Berlin, Heidelberg, August 2003.
- CV02. Jan Camenisch and Els Van Herreweghen. Design and implementation of the idemix anonymous credential system. In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 21–30. ACM Press, November 2002.
- DD22. Nico Döttling and Jesko Dujmovic. Maliciously circuit-private FHE from information-theoretic principles. *Cryptology ePrint Archive*, Report 2022/495, 2022.
- DEF⁺24. Bernardo David, Felix Engemann, Tore Kasper Frederiksen, Markulf Kohlweiss, Elena Pagnin, and Mikhail Volkhov. Updatable privacy-preserving blueprints. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part I*, volume 15484 of *Lecture Notes in Computer Science*, pages 105–139. Springer, 2024.
- DF02. Ivan Damgård and Eiichiro Fujisaki. An integer commitment scheme based on groups with hidden order. In *ASIACRYPT 2002*, volume 2501 of *LNCS*, 2002.
- FP04. Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. Efficient private matching and set intersection. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 1–19. Springer, Berlin, Heidelberg, May 2004.
- Gen09. Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of STOC 2009*, pages 169–178, 2009.

- GKL21. Matthew Green, Gabriel Kaptchuk, and Gijs Van Laer. Abuse resistant law enforcement access systems. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 553–583. Springer, Cham, October 2021.
- GKR08. Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 113–122. ACM Press, May 2008.
- GKR⁺21. Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. Poseidon: A new hash function for zero-knowledge proof systems. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 519–535. USENIX Association, August 2021.
- GLS⁺23. Alexander Golovnev, Jonathan Lee, Srinath T. V. Setty, Justin Thaler, and Riad S. Wahby. Brakedown: Linear-time and field-agnostic SNARKs for R1CS. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 193–226. Springer, Cham, August 2023.
- GM82. Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In *14th ACM STOC*, pages 365–377. ACM Press, May 1982.
- Gov22. United States Government. Technical design choices for a U.S. central bank digital currency system. <https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/09/09-2022-Technical-Design-Choices-US-CBDC-System.pdf>, 2022.
- GPR⁺21. Gayathri Garimella, Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. Oblivious key-value stores and amplification for private set intersection. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 395–425, Virtual Event, August 2021. Springer, Cham.
- GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Berlin, Heidelberg, August 2013.
- HHKP23. Charlotte Hoffmann, Pavel Hubáček, Chethan Kamath, and Krzysztof Pietrzak. Certifying giant nonprimes. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part I*, volume 13940 of *LNCS*, pages 530–553. Springer, Cham, May 2023.
- HS21. Lucjan Hanzlik and Daniel Slamanig. With a little help from my friends: Constructing practical anonymous credentials. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 2004–2023. ACM Press, November 2021.
- HSS23. Julia Hesse, Nitin Singh, and Alessandro Sorniotti. How to bind anonymous credentials to humans. In Joseph A. Calandrino and Carmela Troncoso, editors, *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9-11, 2023*, pages 3047–3064. USENIX Association, 2023.
- IR90. K. Ireland and M.I. Rosen. *A Classical Introduction to Modern Number Theory*. Graduate Texts in Mathematics. Springer, 1990.

- JWP22. Yuting Jiang, Jianghong Wei, and Jing Pan. Publicly verifiable private set intersection from homomorphic encryption. In *Security and Privacy in Social Networks and Big Data - 8th International Symposium, SocialSec 2022, Xi'an, China, October 16-18, 2022, Proceedings*, volume 1663 of *Communications in Computer and Information Science*, pages 117–137. Springer, 2022.
- KKS22. Aggelos Kiayias, Markulf Kohlweiss, and Amirreza Sarencheh. PEReDi: Privacy-enhanced, regulated and distributed central bank digital currencies. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 1739–1752. ACM Press, November 2022.
- KL20. J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC Cryptography and Network Security Series. CRC Press, 2020.
- KLN23. Markulf Kohlweiss, Anna Lysyanskaya, and An Nguyen. Privacy-preserving blueprints. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part II*, volume 14005 of *LNCS*, pages 594–625. Springer, Cham, April 2023.
- KM15. Markulf Kohlweiss and Ian Miers. Accountable metadata-hiding escrow: A group signature case study. *PoPETs*, 2015(2):206–221, April 2015.
- KMRS14. Seny Kamara, Payman Mohassel, Mariana Raykova, and Seyed Saeed Sadeghian. Scaling private set intersection to billion-element sets. In Nicolas Christin and Reihaneh Safavi-Naini, editors, *FC 2014*, volume 8437 of *LNCS*, pages 195–215. Springer, Berlin, Heidelberg, March 2014.
- LFKN92. Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992.
- LMR19. Russell W. F. Lai, Giulio Malavolta, and Viktoria Ronge. Succinct arguments for bilinear group arithmetic: Practical structure-preserving cryptography. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2057–2074. ACM Press, November 2019.
- LRSW99. Anna Lysyanskaya, Ron Rivest, Amit Sahai, and Stefan Wolf. Pseudonym systems. In Howard Heys and Carlisle Adams, editors, *Selected Areas in Cryptography*, volume 1758 of *LNCS*, 1999.
- Lys02. Anna Lysyanskaya. *Signature schemes and applications to cryptographic protocol design*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, September 2002.
- OPCPC14. Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private fhe. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014*, pages 536–553, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- OPP14. Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 536–553. Springer, Berlin, Heidelberg, August 2014.
- PEB21. Charlotte Peale, Saba Eskandarian, and Dan Boneh. Secure complaint-enabled source-tracking for encrypted messaging. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1484–1506. ACM Press, November 2021.
- Pie19. Krzysztof Pietrzak. Simple verifiable delay functions. In Avrim Blum, editor, *ITCS 2019*, volume 124, pages 60:1–60:15. LIPIcs, January 2019.

- RR22. Srinivasan Raghuraman and Peter Rindal. Blazing fast PSI from improved OKVS and subfield VOLE. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 2505–2517. ACM Press, November 2022.
- RS21. Peter Rindal and Philipp Schoppmann. VOLE-PSI: Fast OPRF and circuit-PSI from vector-OLE. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 901–930. Springer, Cham, October 2021.
- RWGM23. Michael Rosenberg, Jacob D. White, Christina Garman, and Ian Miers. zk-creds: Flexible anonymous credentials from zkSNARKs and existing identity infrastructure. In *2023 IEEE Symposium on Security and Privacy*, pages 790–808. IEEE Computer Society Press, May 2023.
- Sch80. J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, oct 1980.
- Sho97. Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT’97*, volume 1233 of *LNCS*, pages 256–266. Springer, Berlin, Heidelberg, May 1997.
- Sta23. Jay Stanley. Paths toward an acceptable public digital currency. ACLU White Paper, 2023. https://www.aclu.org/wp-content/uploads/legal-documents/cbdc_white_paper_-_0882_0.pdf.
- TBA⁺22. Alin Tomescu, Adithya Bhat, Benny Applebaum, Ittai Abraham, Guy Gueta, Benny Pinkas, and Avishay Yanai. UTT: Decentralized ecash with accountable privacy. Cryptology ePrint Archive, Report 2022/452, 2022.
- TZ23. Stefano Tessaro and Chenzhi Zhu. Revisiting BBS signatures. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 691–721. Springer, Cham, April 2023.
- WHV24. Ruihan Wang, Carmit Hazay, and Muthuramakrishnan Venkatasubramanian. Ligetron: Lightweight scalable end-to-end zero-knowledge proofs post-quantum zk-snarks on a browser. In *IEEE Symposium on Security and Privacy, SP 2024, San Francisco, CA, USA, May 19-23, 2024*, pages 1760–1776. IEEE, 2024.
- XZZ⁺19. Tiancheng Xie, Jiaheng Zhang, Yupeng Zhang, Charalampos Papamanthou, and Dawn Song. Libra: Succinct zero-knowledge proofs with optimal prover computation. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 733–764. Springer, Cham, August 2019.
- ZLW⁺21. Jiaheng Zhang, Tianyi Liu, Weijie Wang, Yinuo Zhang, Dawn Song, Xiang Xie, and Yupeng Zhang. Doubly efficient interactive proofs for general arithmetic circuits with linear prover time. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 159–177. ACM Press, November 2021.

A Discussion on Non-frameability vs. Deniability

Non-frameability is a desirable feature, but it is fundamentally at odds with deniability. In a deniable system, data may be authenticated at the moment when it is received, but this authentication information quickly becomes useless. This way, Alice cannot use her authenticated transcript from a conversation with Bob to prove to a third party what Bob did or did not say. Typically, to define deniability, one would explicitly give Alice an algorithm to “frame” Bob, i.e., to

authenticate any transcript on his behalf. That way, a real transcript will not be any more believable than a bogus one, and Bob may convincingly deny ever talking to Alice. Deniability of a ciphertext’s origin, for example, is valuable for encrypted messaging systems, especially when users might face coercion, and in other contexts [PEB21,GKL21]. Kohlweiss and Miers [KM15] attempted to address the question whether the properties of non-frameability and deniability can both be achieved together and reached disappointing conclusions, as did Bartusek et al. [BGJP23].

In a system like PPBs, deniability would allow for an efficient algorithm for creating a convincing-looking escrow that would decrypt to any value the algorithm takes as input. A deniable PPB would give an auditor a meaningful ability to monitor the system only so long as it trusts the escrow recipients that they did not make up the escrows but in fact collected them as part of a legitimate transaction. It may be an interesting direction to pursue in future work if well-motivated in practice.

In this work, however, similarly to Bartusek et al. [BGJP23], we prioritized non-frameability and thus abandoned deniability, because, in our view, systems like ours that are designed to detect illegal activity require not only the ability to identify a watchlisted user’s actions but also the means to only convince a judge of these actions if they have in fact taken place. It is more important to us that innocent users cannot be credibly accused of wrongdoing than that perpetrators be able to deny their activities.

B Full Definitions for Privacy Preserving f -Blueprint Schemes

A blueprint scheme has three parties - an auditor, a set of users and a set of recipients. It is defined as follows:

Definition 8. *For a non-interactive commitment scheme $(\text{CSetup}, \text{Com})$, an f -blueprint scheme consists of the following probabilistic polynomial time algorithms:*

$\text{Setup}(1^\lambda, \text{cpar}) \rightarrow \Lambda$: This algorithm takes as input the security parameter 1^λ and the commitment parameters cpar output by $\text{CSetup}(1^\lambda)$. It outputs the public parameters Λ which includes 1^λ and cpar .

$\text{KeyGen}(\Lambda, x, r_x) \rightarrow (\text{pk}_A, \text{sk}_A)$: The key generation algorithm for auditor A takes 1^λ , Λ , and commitment value and opening (x, r_x) as input, and outputs the key pair $(\text{pk}_A, \text{sk}_A)$. The values (x, r_x) define a commitment C_x .

$\text{VerPK}(\Lambda, \text{pk}_A, C_x) \rightarrow 1$ or 0 : This is the algorithm that, on input the auditor’s public key pk_A and a commitment C_x , verifies that the auditor’s public key was computed correctly for the commitment C_x .

$\text{Escrow}(\Lambda, \text{pk}_A, y, r_y) \rightarrow Z$: This algorithm takes Λ , pk_A , and commitment value and opening (y, r_y) as input and outputs an escrow Z for commitment $C = \text{Com}(y; r_y)$.

$\text{VerEscrow}(\Lambda, \text{pk}_A, C, Z) \rightarrow 1$ or 0 : This algorithm takes the auditor's public key pk_A , a commitment C , and an escrow Z as input and verifies that the escrow was computed correctly for the commitment C .

$\text{Dec}(\Lambda, \text{sk}_A, C, Z) \rightarrow f(x, y)$ or \perp : This algorithm takes the auditor's secret key sk_A , a commitment C and an escrow Z as input. It decrypts the escrow and returns the output $f(x, y)$ if C is a commitment to y and $\text{VerEscrow}(\Lambda, \text{pk}_A, C, Z) = 1$.

[KLN23] also defines a *secure* f -blueprint scheme as one that possesses the following properties:

Correctness of VerPK and VerEscrow: For honestly generated values $(\text{cpar}, \text{pk}_A, C_x, C, Z)$, the algorithms VerEscrow and VerPK should accept with probability 1.

Correctness of Dec: For honestly generated values $(\text{cpar}, \text{pk}_A, \text{sk}_A, C, Z)$, $\text{Dec}(\Lambda, \text{sk}_A, C, Z) = f(x, y)$ should hold with overwhelming probability .

Soundness: For all PPT adversaries \mathcal{A} involved in the experiment in Fig. B.1, there exists a negligible function ν such that:

$$\text{Adv}_{\text{Adv,Blu}}^{\text{Sound}} = \Pr \left[\text{Sound}_{\text{Blu}}^{\text{Adv}}(\lambda) = 1 \right] = \nu(\lambda)$$

$\text{Sound}_{\text{Blu}}^{\text{Adv}}(\lambda)$
1: $\text{cpar} \leftarrow \text{CSetup}(1^\lambda)$
2: $\Lambda \leftarrow \text{Setup}(1^\lambda, \text{cpar})$
3: $x, r_x \leftarrow \text{Adv}(1^\lambda, \Lambda)$
4: $(\text{pk}_A, \text{sk}_A) \leftarrow \text{KeyGen}(\Lambda, x, r_x)$
5: $(C, y, r_y, Z) \leftarrow \text{Adv}(\text{pk}_A)$
6: return $[C = \text{Com}(y; r_y) \wedge$
7: $\text{VerEscrow}(\Lambda, \text{pk}_A, C, Z) \wedge \text{Dec}(\Lambda, \text{sk}_A, C, Z) \neq f(x, y)]$

Fig. B.1: Experiments $\text{Sound}_{\text{Blu}}^{\text{Adv}}(\lambda)$

Definition 9 (Blueprint Hiding). *The blueprint-hiding property makes sure that pk_A just reveals that x is a valid first argument to f . Otherwise, x is hidden even from an adversary who (1) may already know a lot of information about x a-priori; and (2) has oracle access to $\text{Dec}(\Lambda, \text{sk}_A, \cdot, \cdot)$.*

This is formalized by requiring that there exist a simulator $\text{Sim} = (\text{SimSetup}, \text{SimKeygen}, \text{SimDecrypt})$ such that for any PPT adversary the following two games are indistinguishable:

1. **Real Game:** Λ is chosen honestly, the public key pk_A is computed correctly for adversarially chosen x, r_x , and the adversary's decryption queries (C, Z) are answered with $\text{Dec}(\Lambda, \text{sk}_A, C, Z)$.

2. **Ideal Game:** Λ is computed using `SimSetup`, the public key pk_A is computed using `SimKeygen` independently of x (although with access to the commitment C_A), and the adversary's decryption query Z_i is answered by first running `SimDecrypt` to obtain enough information about the user's data y_i to be able to compute $f(x, y_i)$. "Enough information" means that for an efficiently computable f^* and a function g such that $f(x, y) = f^*(x, g(y))$ for all possible inputs (x, y) , `SimDecrypt` obtains $y_i^* = g(y_i)$.

Formally, for all probabilistic poly-time adversaries Adv involved in the game described in Fig. B.2, the advantage function satisfies:

$$\text{Adv}_{\text{Adv, Sim}}^{\text{BH}} = \left| \Pr \left[\text{BHreal}_{\text{Blu}}^{\text{Adv}}(\lambda) = 0 \right] - \Pr \left[\text{BHideal}_{\text{Blu, Sim}}^{\text{Adv}}(\lambda) = 0 \right] \right| = \nu(\lambda)$$

for some negligible ν .

$\text{BHreal}_{\text{Blu}}^{\text{Adv}}(\lambda)$	$\text{BHideal}_{\text{Blu, Sim}}^{\text{Adv}}(\lambda)$
1 : $\text{cpar} \leftarrow \text{CSetup}(1^\lambda)$	1 : $\text{cpar} \leftarrow \text{CSetup}(1^\lambda)$
2 : $\Lambda \leftarrow \text{Setup}(1^\lambda, \text{cpar})$	2 : $(\Lambda, \text{st}) \leftarrow \text{SimSetup}(1^\lambda, \text{cpar})$
3 : $(x, r_x, \text{st}_{\text{Adv}}) \leftarrow \text{Adv}(1^\lambda, \Lambda)$	3 : $(x, r_x, \text{st}_{\text{Adv}}) \leftarrow \text{Adv}(1^\lambda, \Lambda)$
4 :	4 : $\text{dsim} \leftarrow (x , \text{Com}(x; r_x))$
5 : $(\text{pk}_A, \text{sk}_A) \leftarrow \text{KeyGen}(\Lambda, x, r_x)$	5 : $(\text{pk}_A, \text{sk}_A) \leftarrow \text{SimKeygen}(1^\lambda, \text{st}, \text{dsim})$
6 : return $\text{Adv}^{\mathcal{O}_0(\text{pk}_A, \text{sk}_A, \cdot, \cdot)}(\text{pk}_A, \text{st}_{\text{Adv}})$	6 : return $\text{Adv}^{\mathcal{O}_1(\text{pk}_A, \text{st}, x, \cdot, \cdot)}(\text{pk}_A, \text{st}_{\text{Adv}})$
$\mathcal{O}_0(\text{pk}_A, \text{sk}_A, C, Z)$	$\mathcal{O}_1(\text{pk}_A, \text{st}, x, C, Z)$
1 : if $\neg \text{VerEscrow}(\Lambda, \text{pk}_A, C, Z)$	1 : if $\neg \text{VerEscrow}(\Lambda, \text{pk}_A, C, Z)$
2 : return \perp	2 : return \perp
3 : return $\text{Dec}(\Lambda, \text{sk}_A, C, Z)$	3 : $y^* \leftarrow \text{SimDecrypt}(\text{st}, C, Z)$
	4 : return $f(x, y) = f^*(x, y^*)$

Fig. B.2: Experiments $\text{BHreal}_{\text{Blu}}^{\text{Adv}}(\lambda)$ and $\text{BHideal}_{\text{Blu, Sim}}^{\text{Adv}}(\lambda)$

Definition 10 (Privacy against Dishonest Auditor). *There exists a simulator such that the adversary's views in the following two games are indistinguishable:*

1. **Real Game:** *The adversary generates the public key and the data x corresponding to this public key, honest users follow the Escrow protocol using adversarial inputs and openings.*
2. **Privacy-Preserving Game:** *The adversary generates the public key and the data x corresponding to this public key. Next, for adversarially chosen inputs and openings, the users run a simulator algorithm that depends only on the commitment and $f(x, y)$ but is independent of the commitment openings.*

More formally, there exists algorithms $\text{Sim} = (\text{SimSetup}, \text{SimEscrow})$ such that, for any PPT adversary Adv involved in the game described in Fig. B.3, the following equation holds for some negligible function ν :

$$\text{Adv}_{\text{Adv,Blu,Sim}}^{\text{PADA}} = \left| \Pr \left[\text{PADA}_{\text{Blu,Sim}}^{\text{Adv},0}(\lambda) = 1 \right] - \Pr \left[\text{PADA}_{\text{Blu,Sim}}^{\text{Adv},1}(\lambda) = 1 \right] \right| = \nu(\lambda)$$

$\text{PADA}_{\text{Blu,Sim}}^{\text{Adv},b}(\lambda)$	
1 :	$cpar \leftarrow \text{CSetup}(1^\lambda)$
2 :	$\Lambda_0 \leftarrow \text{Setup}(1^\lambda, cpar); (\Lambda_1, \text{st}) \leftarrow \text{SimSetup}(1^\lambda, cpar)$
3 :	$(x, r_A, \text{pk}_A, \text{st}_{\text{Adv}}) \leftarrow \text{Adv}(1^\lambda, \Lambda_b)$
4 :	if $\text{VerPK}(\Lambda_b, \text{pk}_A, \text{Com}(x; r_A)) = 0$: return \perp
5 :	return $\text{Adv}^{\mathcal{O}_b(\cdot, \cdot)}(\text{st}_{\text{Adv}})$
$\mathcal{O}_0(y, r_y)$	$\mathcal{O}_1(y, r_y)$
1 :	1 : return $\text{Escrow}(\Lambda_0, \text{pk}_A, y, r_y)$
	2 : $f(x, y)$

Fig. B.3: Game $\text{PADA}_{\text{Blu}}^{\text{Adv},b}(\lambda)$

Definition 11 (Privacy with Honest Auditor). *There exists a simulator Sim such that the adversary's views in the following two games are indistinguishable:*

1. **Real Game:** *The honest auditor generates the public key on input x provided by the adversary, and honest users follow the **Escrow** protocol on input adversarially chosen openings.*
2. **Privacy-Preserving Game:** *The honest auditor generates the public key on input x provided by the adversary. On input adversary-generated commitments and openings, the users run a simulator that is independent of y (although with access to the commitment C_y) to form their escrows.*

In both of these games, the adversary has oracle access to the decryption algorithm.

We formalize these two games in Fig. B.4. We require that there exists a simulator $\text{Sim} = (\text{SimSetup}, \text{SimEscrow})$ such that, for any PPT adversary Adv involved in the game described in the figure, the following equation holds:

$$\text{Adv}_{\text{Blu,Sim}}^{\text{PWA}} = \left| \Pr \left[\text{PWA}_{\text{Blu,Sim}}^{\text{Adv},0}(\lambda) = 0 \right] - \Pr \left[\text{PWA}_{\text{Blu,Sim}}^{\text{Adv},1}(\lambda) = 0 \right] \right| = \nu(\lambda)$$

for some negligible function ν .

$\text{PWHA}_{\text{Blu,Sim}}^{\text{Adv},b}(\lambda)$	
1 : $cpar \leftarrow \text{CSetup}(1^\lambda)$ 2 : $\Lambda_0 \leftarrow \text{Setup}(1^\lambda, cpar); \Lambda_1 \leftarrow \text{SimSetup}(1^\lambda, cpar)$ 3 : $M \leftarrow []$ 4 : $x, r_x \leftarrow \text{Adv}(1^\lambda, \Lambda_b)$ 5 : $(pk_A, sk_A) \leftarrow \text{KeyGen}(\Lambda_b, x, r_x)$ 6 : return $\text{Adv}_{\mathcal{O}_b^{\text{Escrow}(\cdot, \cdot), \mathcal{O}^{\text{Dec}(\Lambda_b, sk_A, \cdot, \cdot)}}}(\text{pk}_A)$	
$\mathcal{O}_0^{\text{Escrow}}(y, r_y)$	$\mathcal{O}_1^{\text{Escrow}}(y, r_y)$
1 : return $\text{Escrow}(\Lambda_0, pk_A, y, r_y)$	1 : $C = \text{Com}(y; r_y)$ 2 : $Z \leftarrow \text{SimEscrow}(\text{st}, \Lambda_1, pk_A, C)$ 3 : $M[C, Z] \leftarrow f(x, y)$ 4 : return Z
$\mathcal{O}^{\text{Dec}}(\Lambda_1, sk_A, C, Z)$	
1 : if $M[C, Z]$ is defined return $M[C, Z]$ 2 : return $\text{Dec}(\Lambda_1, sk_A, C, Z)$	

Fig. B.4: Game $\text{PWHA}_{\text{Blu,Sim}}^{\text{Adv},b}(\lambda)$

C Additional preliminaries

C.1 Motivation for BB-PSL

For concreteness, let us imagine that π is the NIZK we get by running a Σ -protocol for a proof of knowledge, and making it non-interactive by replacing the message from the verifier with the output of the random oracle. The prover's side of the Σ -protocol consists of two algorithms, P_1 and P_2 . $P_1(\text{pk}, m, r; R)$ generates the first message, a , of the proof of knowledge of how $c = \text{Enc}(\text{pk}, m, r)$ was computed using random coins R ; $P_2(\text{pk}, m, r, e; R)$ generates the prover's response, z , to the challenge e using the same randomness. The verifier's part of the Σ protocol is just the algorithm $V(\text{pk}, c, a, e, z)$. It is well-known that, in the random-oracle model, the following proof system is black-box simulation-extractable: the prover computes $a = P_1(\text{pk}, m, r; R)$, $e = H(\text{pk}, c, a)$, and $z = P_2(\text{pk}, m, r, e; R)$ and outputs the proof $\pi = (a, z)$. To verify π , the verifier computes $e = H(\text{pk}, c, a)$ and runs $V(\text{pk}, c, a, e, z)$.

However, when we plug this proof system into the attempted construction above of a CCA-secure cryptosystem from a semantically secure one, we don't (easily) get a proof of CCA security. This is because the adversary can interleave his decryption queries and his random-oracle queries in such a way that he will force the security reduction to run in exponential time in the number q of queries.

In order to respond to the i^{th} decryption query (c_i, π_i) where $\pi_i = (a_i, z_i)$, the reduction needs to rewind the adversary to the point in time where the adversary queried the random oracle to get $e_i = H(\text{pk}, c_i, a_i)$. By first issuing all the random-oracle queries in reverse order, i.e. obtaining $e_q = H(\text{pk}, c_q, a_q)$, and then e_{q-1}, \dots, e_1 before issuing any decryption queries at all, and then querying for the decryptions of $(c_1, \pi_1), \dots, (c_q, \pi_q)$, the adversary will ensure that the reduction will need to rewind $O(2^q)$ times¹⁶. This is because each time the reduction rewinds the adversary, they also need to rewind for each previous query to ensure the adversary receives the correct decryptions to run normally. Thus, each decryption query doubles the number of required rewinds.

There are two ways of fixing this problem. One is to use a straight-line extractable proof system that does not need to rewind at all; but that can be inefficient. The other way to fix it (implicitly in the spirit of Shoup and Genaro) is to not require the straight-line extraction of the entire witness: the reduction does not need both m and r to proceed, just the message m alone is sufficient. The fact that, with rewinding, it is possible to extract the entire witness is still crucial since it guarantees that the adversary's interaction with the security reduction results in exactly the same view as in its interaction with the decryption oracle: if not, then a separate reduction would break the soundness of the proof system.

C.2 Useful Lemmas for Composite-Order Groups

Lemma 2 $(n + 1) \in QR_{n^2}$

Proof of Lemma 2. In Ireland and Rosen's textbook [IR90] Proposition 5.1.1 gives us that an element, a , in \mathbb{Z}_{n^2} is a quadratic residue iff $a^{(p-1)/2} = 1 \pmod{p}$ and $a^{(q-1)/2} = 1 \pmod{q}$. We can see that $(n + 1)^{(p-1)/2} = \sum_{i=0}^{(p-1)/2} 1^{(p-1)/2-i} \cdot n^{(p-1)/2-i} = 1 + kn$ for some k . Since n is divisible by both p and q , this value is simply $1 \pmod{p}$ and q . Thus, $(n + 1)$ is in QR_{n^2} .

Lemma 3 $(-1) \in QNR_{n^2}$ for RSA modulus, n .

Proof of Lemma 3. Using Proposition 5.1.1 from Ireland and Rosen's textbook [IR90] again we see that $(-1)^{(p-1)/2} \pmod{p}$ is equal to $(-1)^{(4k+2)/2} \pmod{p}$ since we are working with primes that are equal to $3 \pmod{4}$. Thus, this equals $(-1)^{2k+1} \pmod{p}$. Note that $2k+1$ is odd and thus this equals $(-1) \pmod{p}$ thus failing the criteria in Proposition 5.1.1 and thus $(-1) \in QNR_{n^2}$.

Lemma 4 (Any element to the 2-nd power likely generates QR_{n^2}) *Let $a \neq 1$, $a \in \mathbb{Z}_{n^2}$ be an element output by a polynomial-time adversary on input n . Then the probability that $\langle a^2 \rangle \neq QR_{n^2}$ is negligible. As a corollary, we know that sampling a random element in QR_{n^2} or squaring a random element in \mathbb{Z}_{n^2} results in a generator of QR_{n^2} .*

¹⁶ The adversary must also base the first message of each Σ -protocol on the output of the random oracle from the last query to ensure rewinding is impossible.

Proof of Lemma 4. QR_{n^2} is cyclic and thus every element in QR_{n^2} can be represented as g^i for some g . We see that any g^i doesn't generate QR_{n^2} when $i \nmid \#QR_{n^2}$. The order of QR_{n^2} is $pqq'q'$ and thus, this only occurs when i is a multiple of p, q, p', q' . Thus, there are at most $pqp' + pqq' + pp'q' + qp'q'$ elements that don't generate QR_{n^2} . When we compare this to the total elements, we see: $(pqp' + pqq' + pp'q' + qp'q')/pqq'q' = 1/q' + 1/p' + 1/p + 1/q$ which is negligible if p, q, p', q' are large.

Lemma 5 *If $2^B > \text{ord}(g)$ then no PPT adversary running in time polynomial to λ can distinguish distribution $\{g^s : s \leftarrow \mathbb{S} 2^{B+\lambda}\}$ from $\{u : u \leftarrow \mathbb{S} \langle g \rangle\}$ for any g such that $g \in \mathbb{Z}_{n^2}$ and $\text{ord}(g) > 2$.*

We refer to [DF02] for a proof of Lemma 5.

Lemma 6 *If $x, x' \in QR_p$ and $y, y' \in QNR_p$ then $xy \in QNR_p$, $xx', yy' \in QR_p$.*

Lemma 7 *For $n = pq$ where p, q are safe primes, if $x, x' \in QR_{n^2}$ and $y, y' \in QNR_{n^2}$ then $xy \in QNR_{n^2}$, $xx', yy' \in QR_{n^2}$.*

Lemma 8 $\#QR_{n^2} = \mathbb{Z}_{n^2}/4$

Proofs of Lemmas 7, 6, and 8 are present in [KL20] (deriving Lemma 8 from [KL20] is a trivial exercise and stems from the fact that $QR_{n^2} \cong QR_p \times QR_q \times QR_{p'} \times QR_{q'}$).

C.3 More *eqrep* relations and constructions

In [KLN23], they define the following relation in Def. 12 and use it to construct their protocols.

Definition 12 (Relation for proof of equality of discrete logarithm representations in cyclic groups of prime order). *Let $R_{eqrep-\mathbb{G}_p}$ be the following relation: $R_{eqrep-\mathbb{G}_p}(\mathbb{x}, \mathbb{w})$ accepts if $\mathbb{x} = (\mathcal{G}, \{x_i, \{g_{i,1}, \dots, g_{i,m}\}_{i=1}^k})$ where \mathcal{G} is the description of a group of order q , and all the x_i s and $g_{i,j}$ s are elements of \mathcal{G} , and witness $\mathbb{w} = \{w_j\}_{j=1}^m$ such that $x_i = \prod_{j=1}^m g_{i,j}^{w_j}$.*

We can enhance this protocol to multiply witnesses with the relation in Def. 3 in Sec. 2.1.

Using known techniques, e.g. KLM from which we took the following description, we can construct the protocol in Def. 12 in cyclic groups of prime order where the DDH and CDH assumptions are hard. We do so in Def. 13.

Definition 13 (Σ -protocol for proof of equality of discrete logarithm representations cyclic groups of prime order). *Let $R_{eqrep-\mathbb{G}_p}$ be the following relation: $R_{eqrep-\mathbb{G}_p}(\mathbb{x}, \mathbb{w})$ accepts if $\mathbb{x} = (\mathcal{G}, \{x_i, \{g_{i,1}, \dots, g_{i,m}\}_{i=1}^k})$ where \mathcal{G} is the description of a group of order q , and all the x_i s and $g_{i,j}$ s are elements of \mathcal{G} , and witness $\mathbb{w} = \{w_j\}_{j=1}^m$ such that $x_i = \prod_{j=1}^m g_{i,j}^{w_j}$.*

- P \rightarrow V** On input the $(\mathbb{x}, \mathbb{w}) \in R_{eqrep-\mathbb{G}_p}$, the Prover chooses $e_j \leftarrow \mathbb{Z}_q$ for $1 \leq j \leq m$ and computes $d_i = \prod_{j=1}^m g_{i,j}^{e_j}$ for $1 \leq i \leq k$. Finally, the Prover sends to the Verifier the values $\mathbf{com} = (d_1, \dots, d_n)$.
- P \leftarrow V** On input \mathbb{x} and \mathbf{com} , the Verifier responds with a challenge $\mathbf{chal} = c$ for $c \leftarrow \mathbb{Z}_q$.
- P \rightarrow V** The Prover receives $\mathbf{chal} = c$ and computes $s_i = e_i + cw_i \bmod q$ for $1 \leq i \leq m$, and sends $\mathbf{res} = (s_1, \dots, s_m)$ to the Verifier.
- Verification** The Verifier accepts if for all $1 \leq i \leq n$, $d_i x_i^c = \prod_{j=1}^m g_{i,j}^{s_j}$; rejects otherwise.
- Simulation** On input \mathbb{x} and $\mathbf{chal} = c$, the simulator chooses $s_j \leftarrow \mathbb{Z}_q$ for $1 \leq j \leq m$, and sets $d_i = (\prod_{j=1}^m g_{i,j}^{s_j})/x_i^c$ for $1 \leq i \leq k$. He then sets $\mathbf{com} = (d_1, \dots, d_n)$ and $\mathbf{res} = (s_1, \dots, s_m)$.
- Extraction** On input two accepting transcripts for the same $\mathbf{com} = (d_1, \dots, d_n)$, namely $\mathbf{chal} = c$, $\mathbf{res} = (s_1, \dots, s_m)$, and $\mathbf{chal}' = c'$, $\mathbf{res}' = (s'_1, \dots, s'_m)$, output $w_j = (s_j - s'_j)/(c - c') \bmod q$ for $1 \leq j \leq m$.

It's easy to see how the relation in Def. 3 can be satisfied by the protocol in Def. 13 if the map μ only maps to sets of size 1 (i.e., there are no multiplicative relations between the witnesses, w_i). Though, we want to be able to prove this for arbitrary μ . In Appx. C.3 we will explain how to construct this.

Constructing a protocol to prove the relation in Def. 3 In this section we will refer to the relation in Def. 12 as $eqrep-\mathbb{G}_p$ and the relation in Def. 3 for $\mathcal{G} = \mathbb{G}_p$ and $\beta = \{1\}$ as $eqrep-\mathbb{G}_p^*$. We gave a construction to prove $eqrep-\mathbb{G}_p$ relations in Def. 13, though this is not fully general as it does not allow for arbitrary multiplication of witnesses. In this section, we give a construction of an example relation for the $eqrep-\mathbb{G}_p^*$ protocol for $\mathcal{G} = \mathbb{G}_p$ and $\beta = \{1\}$. In Alg. 4 we show how to implement a $eqrep-\mathbb{G}_p^*$ protocol from an underlying $eqrep-\mathbb{G}_p$ protocol by construction intermediate Pedersen commitments. In this example, we are proving that a Pedersen commitment C_a is committed to the product of the values in three other Pedersen commitments, C_b , C_c , and C_d . Formally, Alg. 4 proves the following relation: $R((C_a, C_b, C_c, C_d), (a, b, c, d, r_a, r_b, r_c, r_d)) = 1$ iff $C_a = g^a h^{r_a} \wedge C_b = g^b h^{r_b} \wedge C_c = g^c h^{r_c} \wedge C_d = g^d h^{r_d} \wedge a = bcd$. Because E is a commitment to bc with fresh randomness, revealing it to the verifier does not affect the zero knowledge of the scheme. The only other communication in this proof for $eqrep-\mathbb{G}_p$ is the proof for an $eqrep$ relation. Thus this scheme is zero knowledge. We can see that the relation proves that $E = g^{bc} h^{c\beta_2}$ which is a valid Pedersen commitment to bc . Thus, because the prover also proves that $C_a = E^d h^{\beta_2}$, the verifiers knows that $C_a = g^{bcd} h^{d\beta_2}$ which is a valid Pedersen commitment to bcd and thus, $a = bcd$. This means we've proven soundness with extraction for this protocol. Using the notation from Def. 3, the map μ would be $\mu(a) = \{b, c, d\}$ (and $\mu(x) = \{x\}$ otherwise). This would ensure that the witness $a = bcd$ with no constraints on the other witnesses. Using the multiplication proof in Algorithm 4, the relation in Def. 3 can be constructed by “flattening” the given relation (i.e. reducing the size of the sets in the image of μ). To do this, take $i \in [m]$ such that $|\mu(w_i)| > 1$. Take two $w_j, w_k \in \mu(w_i)$ and commit to

these values (and their product) using a Pedersen commitment: $c_j = g^{w_j} h^{r_j}$, $c_k = g^{w_k} h^{r_k}$, $c = g^{w_j w_k} h^r$. The prover then uses the protocol in Def. 13 to prove that c_k is the product of these two values, reusing the s_i and e_i values to open c_j and c_k . We can now recurse, but we see that the map for the relationship we need to prove $\mu'(w_i)$ now includes $w' = w_j w_k$ instead of w_j, w_k , thus reducing the size of the set. If we recurse like this for each non-singleton set in the image of μ' , we can see that this eventually reduces to a map with only singletons and can thus be proven directly using 13.

Algorithm 4 Example $eqrep\text{-}\mathbb{G}_p^*$ proof

- 1: $\rho \leftarrow_{\$} \mathbb{Z}_p$; $E = g^{bc} h^\rho$
 - 2: $\beta_1 = \rho - cr_b$; $\beta_2 = r_a - d\rho$
 - 3: Send E to the verifier
 - 4: Prove the following relation via $eqrep\text{-}\mathbb{G}_p$
 - 5: $\text{PoK}_{eqrep\text{-}\mathbb{G}_p}[a, b, c, r_a, r_b, r_c, \beta_1, \beta_2 :$
 $C_a = g^a h^{r_a} \wedge C_b = g^b h^{r_b} \wedge C_c = g^c h^{r_c} \wedge C_d = g^d h^{r_d}$
 $\wedge E = C_b^c h^{\beta_1} \wedge C_a = E^d h^{\beta_2}]$
-

Constructing a protocol for the relation in 3 for $\mathcal{G} = |QR_{n^2}|$ and $\beta = \{-1, 1\}$ Construction 1 shows an example construction of a proof of a relation for $eqrep\text{-}\mathbb{Z}_{n^2}$ defined in Sec. 2.1. We note that to reduce a construction of $eqrep\text{-}\mathbb{Z}_{n^2}$ to the soundness of Damgård-Fujisaki commitments, we need to create Damgård-Fujisaki commitments to each witness in the relation and use a proof of opening in the protocol to ensure we can extract the witnesses. This step is not necessarily required, but is sufficient to realize $eqrep\text{-}\mathbb{Z}_{n^2}$ and allows us to reduce to the auxiliary proofs for Damgård-Fujisaki commitments rather than number theoretic lemmas. In this example, we'll use Damgård-Fujisaki commitments in \mathbb{Z}_{n^2} which we prove are secure in Appx. F.2. In this example, we prove the exponentiation of an element in a $|QR_{n^2}|$ commitment (which we define in Sec. 4.2) by a scalar committed to by a Damgård-Fujisaki commitment. This proof can be seen as proving the relation $R((c_1, c_2, t, d_1, d_2), (x_1, r_1, x_2, r_2, x_3, r_3, M, N, x_1, x_2, x_3)) = 1$ iff $c_2 = g^{x_1} h^{r_1} \wedge t = g^{x_2} h^{r_2} \wedge d_2 = g^{x_3} h^{r_3} \wedge c_1 = M g^{x_1} \wedge d_1 = N g^{x_3} \wedge N = M^{x_2}$.

For this proof, both the prover P and the verifier V have a scalar commitment t to value x_2 along with two $|QR_{n^2}|$ commitments $c = (c_1, c_2)$ and $d = (d_1, d_2)$ to two \mathbb{Z}_{n^2} elements, M , and N . The prover wants to show that $N = M^{x_2}$. Damgård and Fujisaki [DF02] give a multiplication protocol which yields a commitment scheme for integers in any group that satisfies certain properties. We prove in Appx. F.2 that QR_{n^2} and \mathbb{Z}_{n^2} both satisfy these properties. We can see that the second elements of both of our $|QR_{n^2}|$ commitments (c_2 and d_2) are exactly Damgård-Fujisaki commitments. We also note that our commitments to

scalars (the commitment t in this example) are simply Damgård-Fujisaki commitments. The Damgård-Fujisaki exponentiation proof is a Σ -protocol and thus has transcripts a, e, z . If the prover uses the z value from a proof of opening of the scalar commitment (t) and reuses this z value in a relation to the $|QR_{n^2}|$ commitments, the prover can prove this exponentiation property for the c , and d commitments. We construct this exponentiation protocol in Construction 1. This example should give the reader enough intuition to build a proof for any $eqrep\text{-}\mathbb{Z}_{n^2}$ relation by adding more Damgård-Fujisaki commitments to witnesses similar to the extension of $eqrep\text{-}\mathbb{G}_p$ to $eqrep\text{-}\mathbb{G}_p^*$.

The prover must also prove knowledge of the opening of each commitment in addition to running this protocol.

Construction 1 ($|QR_{n^2}|$ -commitments - proof of exponentiation) *The goal is for the prover to prove that the $|QR_{n^2}|$ -commitment d is committed to $N = M^{x_2}$ where c is a $|QR_{n^2}|$ -commitment to M and t is a Damgård-Fujisaki commitment to the integer x_2 . Both prover and verifier have the public values: $\mathbb{x} = (c, d, t)$ where $c = (c_1, c_2)$, $d = (d_1, d_2)$, $c_2 = g^{x_1} h^{r_1}$, $t = g^{x_2} h^{r_2}$, $d_2 = g^{x_3} h^{r_3}$, $c_1 = Mg^{x_1}$, and $d_1 = M^{x_2} g^{x_3}$. The prover has the secret values: $\mathbb{w} = (x_1, x_2, x_3, r_1, r_2, r_3, M)$. First, the prover uses the proof of knowledge of commitment opening from Damgård and Fujisaki [DF02] to prove that $t = g^{x_2} h^{r_2}$. The prover then shows that the prover can open c and d such that $M = \pm c_1/g^{x_1}$ and $N = \pm d_1/g^{x_3}$. The prover then produces the following NIZK:*

Prove(\mathbb{w}) $\rightarrow \pi$

$$\rho_1 \leftarrow_{\$} [CT2^\lambda] \quad (\rho_1 \text{ will hide } ex_2)$$

$$\rho_2 \leftarrow_{\$} [C2^{B+2\lambda}] \quad (\rho_2 \text{ will hide } er_2)$$

$$\rho_3 \leftarrow_{\$} [CT2^{2\lambda}] \quad (\rho_3 \text{ will hide } e(-x_2x_1 + x_3))$$

$$\rho_4 \leftarrow_{\$} [CT2^{B+2\lambda}] \quad (\rho_4 \text{ will hide } e(-r_1x_1 + r_3))$$

$$a_1 = g^{\rho_1} h^{\rho_2}; a_2 = c_1^{\rho_1} g^{\rho_3}; a_3 = c_2^{\rho_1} g^{\rho_3} h^{\rho_4}$$

$$e = H(a_1, a_2, a_3)$$

$$z_1 = \rho_1 + ex_2; z_2 = \rho_2 + er_2; z_3 = \rho_3 + e(-x_1x_2 + x_3); z_4 = \rho_4 + e(-r_1x_2 + r_3)$$

return $\pi = (a_1, a_2, a_3, z_1, z_2, z_3)$

Verify(\mathbb{x}, π) $\rightarrow \{0, 1\}$

$$g^{z_1} h^{z_2} = a_1 t^e$$

$$c_1^{z_1} g^{z_3} = a_2 d_1^e$$

$$c_2^{z_1} g^{z_3} h^{z_4} = a_3 d_2^e$$

return 1 if the above equations hold and 0 otherwise.

Lemma 9 (Strong special soundness property of [DF02]) *If we find $a, e, e', z_1, z'_1, z_2, z'_2$ such that a, e, z_1, z_2 and a, e', z'_1, z'_2 are both valid transcripts for a Damgård-Fujisaki opening protocol. If $g^{z_1} h^{z_2} = ac^e$ and $g^{z'_1} h^{z'_2} = ac^{e'}$, where c is a Damgård-Fujisaki commitment, then we know that $(e - e')|(z_1 - z'_1)$ and $(e - e')|(z_2 - z'_2)$ and we can extract a b such that $bg^{(z_1 - z_1)/(e - e')} h^{(z_2 - z'_2)/(e - e')} = c$*

Proof of Lemma 9 can be found in [DF02]. This is stronger than simple extraction as it ensures that $e - e'$ divides both $z_1 - z'_1$ and $z_2 - z'_2$.

Theorem 10 *Our exponentiation protocol in Construction 1 has special soundness i.e. given two accepting transcripts, there exists an efficient extractor that extracts an opening of d to M^{x_2} , c to M and t to x_2 .*

Special soundness proof overview. Over the course of the proof, we'll extract $\Delta_e = e - e'$ as well as $\forall i \in [4], \Delta_{z_i} = z_i - z'_i, \delta_{z_i} = \Delta_{z_i}/\Delta_e \forall i \in [4]$ along with β_1, β_2 , and β_3 such that: $b_1 g^{\delta_{z_1}} h^{\delta_{z_2}} = t$, $b_2 c_1^{\delta_{z_1}} g^{\delta_{z_3}} = d_1$, and $b_3 c_2^{\delta_{z_1}} g^{\delta_{z_3}} h^{\delta_{z_4}} = d_2$.

Our proof will proceed as follows: First, we'll extract the opening of t , then we'll extract the values from the third equation, $c_2^{\delta_{z_1}} g^{\delta_{z_3}} h^{\delta_{z_4}} = a_3 d_2^e$, and use our knowledge of the opening of t to help us. Lastly, we'll extract values from the second equation ($c_1^{\delta_{z_1}} g^{\delta_{z_3}} = a_2 d_1^e$) using our knowledge of the last two extractions (from the first and third equations). Using these extracted values, we'll be able to prove that the commitments are sound. We need to proceed in this order to ensure we've extracted enough values to compute $(z_3 - z'_3)/(e - e')$ and $(z_4 - z'_4)/(e - e')$. Without knowing previously extracted values, we cannot trivially reduce to the soundness of the proof of knowledge of opening protocol in [DF02] because c_1 and c_2 are used as the bases for verification in the second two equations. We will see that we can carefully craft final messages s_1, s_2 to give to the [DF02] challenger which will allow us to compute $(z_3 - z'_3)/(e - e')$ and $(z_4 - z'_4)/(e - e')$ in the final two equations to prove them secure. In the proof, we'll use Δ and δ to refer to values used in the extraction. For example, Δ_{z_1} will refer to $z_1 - z'_1$ after rewinding a prover and δ_{z_1} will refer to $(z_1 - z'_1)/(e - e')$.

Proof of special soundness. Since we have the prover prove they know the openings of t , c , and d individually, our extractor can compute $c = (M g^{x_1} a_d, g^{x_1} h^{r_1})$, $d = (N g^{x_3} a_d, g^{x_3} h^{r_3})$, and $t = g^{x_2} h^{r_2} b_t$.

Using rewinding, we can extract $\Delta_{z_1} = z_1 - z'_1, \Delta_{z_2} = z_2 - z'_2, \Delta_{z_3} = z_3 - z'_3, \Delta_{z_4} = z_4 - z'_4$, and $\Delta_e = e - e'$. We can see that the first equality, $g^{z_1} h^{z_2} = a_1 t^e$ appears exactly like a proof of opening for Damgård-Fujisaki commitments, and thus, we can extract $\delta_{z_1} = \Delta_{z_1}/\Delta_e, \delta_{z_2} = \Delta_{z_2}/\Delta_e, b_1$, from this due to Lemma 9. To show why we can extract, we can create a reduction to the soundness of proof of opening of [DF02].

Our reduction will take t from our adversary, then claim to the [DF02] opening soundness challenger that we can open this. We can discard all other values from the adversary when doing this. Then, we also pass a_1 to the challenger and we receive the challenge, e from the challenger and pass this to the adversary. The adversary will then produce z_1, z_2 , and we can discard the other z values and simply pass the first two to the challenger. We see that this satisfies $g^{z_1} h^{z_2} = a_1 t^e$ and thus is a valid proof and thus we can rewind and use the same algorithm as the challenger in the knowledge proof of [DF02] to extract $\delta_{z_1}, \delta_{z_2}, b_1$ such that $t = g^{\delta_{z_1}} h^{\delta_{z_2}} b_1$ and $b_1^2 = 1$.

The rest of our proof will create more reductions to the soundness game in [DF02], but the details will be omitted.

Next, we observe that we can continue rewinding until we obtain an even $e - e'$. See that any subset of $[C]$ must be at least half even or odd and the adversary must be able to answer a super polynomial subset of $[C]$. Thus, with probability at least $1/4$ it will be the case that e and e' will both be even or both be odd, thus ensuring that $e - e'$ is even. Let us focus on the case where $e - e'$ is even, knowing that we'll only reduce our chance of breaking soundness in this case by $1/4$ which is still efficient.

Next, we'll prove that because our extractor can open c_2 , if we can't extract $\delta_{z_3} = \Delta_{z_3}/\Delta_e$, $\delta_{z_4} = \Delta_{z_4}/\Delta_e$, and β_3 such that $c_2^{\delta_{z_1}} g^{\delta_{z_3}} h^{\delta_{z_4}} \beta_3 = d_2$, we can reduce to the proof of opening protocol. We can see that this is true with another reduction similar to our reduction for t . We pass d_2, a_3 to the challenger to receive e to pass back to the adversary. After our adversary proves they can open c_2 , we receive x_1, r_1, b_3 such that $g^{x_1} h^{r_1} b_3 = c_2$ and $b_3^2 = 1$. We see that the verifier accepts, so, $c_2^{z_1} g^{z_3} h^{z_4} = a_3 d_2^e$ and thus, $c_2^{\Delta_{z_1}} g^{\Delta_{z_3}} h^{\Delta_{z_4}} = a_3 d_2^{\Delta_e}$. We can replace this with $(\beta_3)^{\Delta_{z_1}} g^{x_1 \Delta_{z_1}} h^{r_1 \Delta_{z_1}} g^{\Delta_{z_3}} h^{\Delta_{z_4}} = a_3 d_2^{\Delta_e}$. Since $e - e'$ is even and we know that $e - e'$ divides Δ_{z_1} , we know that Δ_{z_1} is even. Because $b^2 = 1$ and Δ_{z_1} is even, we see that $g^{x_1 \Delta_{z_1}} h^{r_1 \Delta_{z_1}} g^{\Delta_{z_3}} h^{\Delta_{z_4}} = a_3 d_2^{\Delta_e}$. We then give: $s_1 = x_1 \Delta_{z_1} + \Delta_{z_3}, s_2 = r_1 \Delta_{z_1} + \Delta_{z_4}$ to the challenger, which satisfies $g^{s_1} h^{s_2} = a_3 d_2^e$. Thus, because of the knowledge extractor for proof of opening, we know we can rewind the adversary and compute $\delta_{s_1} = (s_1 - s'_1)/(e - e')$ as well as $\delta_{s_2} = (s_2 - s'_2)/(e - e')$ and β_3 . Because the adversary proved opening of d_2 , we have x_3, r_3, b_{d_2} such that $\delta_{s_1} = x_3, \delta_{s_2} = r_3, b_{d_2} = \beta_3$. We can then extract δ_{z_3} with the following equation: $\delta_{z_3} = x_3 - x_1 \delta_{z_1} = (z_3 - z'_3)/(e - e')$

This is because $\delta_{s_1} = x_1$ implies that:

$$\begin{aligned} x_3(e - e') &= s_1 - s'_1 = x_1 z_1 + z_3 - x_1 z'_1 - z'_3 \\ x_3(e - e') - x_1 z_1 + x_1 z'_1 &= z_3 - z'_3 \\ x_3(e - e') - x_1(z_1 - z'_1) &= z_3 - z'_3 \\ x_3(e - e') - x_1 \delta_{z_1} * (e - e') &= z_3 - z'_3 \\ x_3 - x_1 \delta_{z_1} &= (z_3 - z'_3)/(e - e') \end{aligned}$$

We then know that:

$$\delta_{s_2} = (s_2 - s'_2)/(e - e') = (r_1 z_1 + z_4 - r_1 z'_1 - z'_4)/(e - e')$$

And that $r_3 = \delta_{s_2}$ and thus:

$$\begin{aligned} r_3(e - e') &= (r_1 z_1 + z_4 - r_1 z'_1 - z'_4) \\ r_3(e - e') - r_1 z_1 + r_1 z'_1 &= (z_4 - z'_4) \\ r_3(e - e') - r_1(z_1 - z'_1) &= (z_4 - z'_4) \end{aligned}$$

And we know that $\delta_{z_1} = (z_1 - z'_1)/(e - e')$, so:

$$r_3(e - e') - \delta_{z_1} * (e - e') = (z_4 - z'_4)$$

$$\delta_{z_4} = r_3 - \delta_{z_1} = (z_4 - z'_4)/(e - e')$$

This gives us that $d_2 = g^{x_1 \delta_{z_1} + \delta_{z_3}} h^{r_1 \delta_{z_1} + \delta_{z_4}} \beta_3$. Which must agree with x_3, r_3, b_{d_2} . Because we know that $\delta_{z_1} = x_2$ from the opening of t , we know that $d_2 = g^{x_1 x_2 + \delta_{z_3}} h^{r_1 x_2 + \delta_{z_4}} b_{d_2}$.

We will now rewind the second equation, $c_1^{2z_1} g^{2z_3} = a_2 d_1^{2e}$ to extract values and prove them sound. We know that $g^{x_1} = c_1/M$ from the opening of c .

Since we know that Δ_{z_1} and Δ_{z_3} are divisible by Δ_e , we can proceed to extract the structure of d_1 .

$$\begin{aligned}
c_1^{z_1} g^{z_3} &= a_2 d_1^e \\
M^{z_1} g^{x_1 z_1} g^{z_3} &= a_2 d_1^e \\
M^{z_1 - z'_1} g^{x_1(z_1 - z'_1)} g^{z_3 - z'_3} &= d_1^{e - e'} \\
M^{(z_1 - z'_1)} g^{x_1(z_1 - z'_1)} g^{(z_3 - z'_3)} &= d_1^{(e - e')} \\
M^{(z_1 - z'_1)/(e - e')} g^{x_1(z_1 - z'_1)/(e - e')} g^{(z_3 - z'_3)/(e - e')} &= d_1 \\
bM^{\delta_{z_1}} g^{x_1 \delta_{z_1}} g^{\delta_{z_3}} &= d_1 \\
bM^{x_2} g^{x_1 x_2} g^{\delta_{z_3}} &= d_1 \\
bg^{x_1 x_2 + \delta_{z_3}} &= d_1 / M^{x_2}
\end{aligned}$$

We can see that $b \in \{-1, 1\}$ since $b^{e - e'} = 1$ and thus, d is a correct commitment to $|M^{x_2}|$.

Honest verifier zero knowledge. If the ranges are adjusted correctly, our construction achieves HVZK, similar to the proof in [DF02].

D Additional HEC definitions, constructions, and proofs

D.1 Security Properties of HEC Scheme

In this section, we provide formal definitions for the security properties of the HEC scheme which are unchanged from [KLN23].

$\text{SECX}_b^{\text{Adv}}(\lambda)$	$\text{DIRECTZ}_b^{\text{Adv}}(\lambda)$
1: $hecp\text{ar} \leftarrow \text{HECSETUP}(1^\lambda)$	1: $hecp\text{ar} \leftarrow \text{HECSETUP}(1^\lambda)$
2: $(f, x_0, x_1, \text{st}) \leftarrow \text{Adv}(1^\lambda, hecp\text{ar})$	2: $(f, x, y, r_X, \text{st}) \leftarrow \text{Adv}(1^\lambda, hecp\text{ar})$
3: if $f \in F, x_0, x_1 \in \text{domain}_{f,x}$	3: if $f \in F, x \in \text{domain}_{f,x}, y \in \text{domain}_{f,y}$
4: $X, _ \leftarrow \text{HECENC}(hecp\text{ar}, f, x_b)$	4: $X, _ = \text{HECENC}(hecp\text{ar}, f, x; r_X)$
5: return $\text{Adv}(hecp\text{ar}, X, \text{st})$	5: $Z_0 \leftarrow \text{HECEVAL}(hecp\text{ar}, f, X, y)$
6: return $\text{Adv}(\perp, \text{st})$	6: $Z_1 \leftarrow \text{HECDIRECT}(hecp\text{ar}, X, f(x, y))$
	7: return $\text{Adv}(hecp\text{ar}, Z_b, \text{st})$
	8: return $\text{Adv}(\perp, \text{st})$
$\text{SECXY}_b^{\text{Adv}}(\lambda)$	
1: $hecp\text{ar} \leftarrow \text{HECSETUP}(1^\lambda)$	
2: $(f, x_0, x_1, \text{st}) \leftarrow \text{Adv}(1^\lambda, hecp\text{ar})$	
3: if $f \in F, x_0, x_1 \in \text{domain}_{f,x}$	
4: $X, _ \leftarrow \text{HECENC}(hecp\text{ar}, f, x_b)$	
5: $(y_0, y_1, \text{st}) \leftarrow \text{Adv}(X, \text{st})$	
6: if $y_0, y_1 \in \text{domain}_{f,y}$	
7: $Z \leftarrow \text{HECEVAL}(hecp\text{ar}, f, X, y_b)$	
8: return $\text{Adv}(Z, \text{st})$	
9: return $\text{Adv}(\perp, \text{st})$	

Fig. D.1: HEC correctness, consistency and security games

Definition 14 (Security of x , security of x and y from third parties, and security of DIRECTZ). Consider Fig. D.1. HEC provides security for x if for any PPT Adv, $|p_{\text{Adv},0}^{\text{SEC}X}(\lambda) - p_{\text{Adv},1}^{\text{SEC}X}(\lambda)|$ is negligible. HEC provides security for x and y from third parties if or any PPT Adv, $|p_{\text{Adv},0}^{\text{SEC}XY}(\lambda) - p_{\text{Adv},1}^{\text{SEC}XY}(\lambda)|$ is negligible. HEC provides security of DIRECTZ if or any PPT Adv, $|p_{\text{Adv},0}^{\text{DIRECTZ}}(\lambda) - p_{\text{Adv},1}^{\text{DIRECTZ}}(\lambda)|$ is negligible.

Explanation for DirectZ. This is an algorithm we need in order to use a HEC in our construction of PPBs. Intuitively, recall that the security of PPBs requires that there be a simulator that can simulate the output of Escrow just given $z = f(x, y)$, without knowledge of x or y . DirectZ allows the simulator to compute the encryption of z directly. For example, if $z = f(x, y)$ where f is a one-way function of y for any fixed x , then access to just the Eval function is not sufficient to compute the encryption of z , since Eval requires y as input, and no such pre-image y cannot be computed from z because f is a One-Way Function.

D.2 Multi-attribute HEC Scheme

In this section, we provide a HEC scheme that satisfies Def. 6 and supports multiple attributes. Including multiple attributes increases the size of values that can be escrowed. In the case of ElGamal, this becomes $\text{poly}(\lambda)^\ell$ and in the case of Camenisch-Shoup, this becomes $(\mathbb{Z}_n)^\ell$. Notice in the case of ElGamal, this allows us to efficiently encrypt and decrypt public keys. This is still not as efficient as in the case of Camenisch-Shoup as the key has to be broken up into logarithmically sized chunks in the case of ElGamal. This makes proving properties of keys escrowed with the ElGamal scheme inefficient while with Camenisch-Shoup, the key can be encrypted while retaining more algebraic structure.

This allows for our Camenisch-Shoup scheme to potentially achieve more efficient proofs for extended properties such as retrospective blueprints.

Our function family for multi-attributes is $\{f_{n,k,\ell}\}_{n,k,\ell \in \mathbb{Z}}$, where n is the length of the auditor's list $\mathbf{x} = \{x_1, \dots, x_n\}$ and k is the bit length of each user attribute y_i^{attr} , where the user's input consists of the user's identifier y_{id} and ℓ attributes: $y = (y_{id}, y_1^{\text{attr}}, \dots, y_\ell^{\text{attr}})$. $f_{n,k,\ell}$ is defined as follows:

$$f_{n,k,\ell}(\mathbf{x}, y) = \begin{cases} y & y_{id} \in \mathbf{x} \\ \emptyset & \text{otherwise} \end{cases} \quad (1)$$

We construct a HEC scheme for this function in Fig. D.2. In our previous construction in Alg. 3 we have a commitment to E which is the commitment C . Remember, C is a commitment to the auditor's polynomial $p(\chi)$ evaluated at the user's identity y_{id} . Thus, E will be an encryption of zero if the user is on the watchlist ($y_{id} \in \mathbf{x}$). In Fig. D.2, we then scale C with the different randomization factors ($\{r_{E,i}\}_{i \in [\ell]}$) yielding the new commitments: $\{C_i\}_{i \in [\ell]}$ to these scaled encryptions. If the user is not on the watchlist, these ℓ commitments now encryptions of random values. We then homomorphically add each scaled encryption

C_i with the encryptions of attributes $\{Y_i\}_{i \in [\ell]}$ to ensure that they can only be decrypted if the user is on the watchlist. We need to use separate randomization scalars for each attribute because we will reveal each encryption. If the encryptions used the same random scalar, the adversary could homomorphically remove them by dividing one encryption by the other. Using independent randomness ensures that each of these commitments are scaled by a random factor and are independent of one another. We still need to include an encryption of E scaled by a random factor $Z_{nf} = r_{nf} \odot E$ to ensure non-framing. Because we only compute one commitment to E , when modifying the Ψ_2 proof from Sec. 3 to work for multiple attributes, we only need to perform the proof of correct encryption of E once. Then, we simply use our auxiliary proofs of commitments to ciphertexts to prove that the rest of the encryptions of attributes are correct, without needing to reprove the commitment to E . This makes our Ψ_2 scheme's communication size equal to $O(\log(x) + \ell)$ for multiple attributes.

Fig. D.2: Multi-attribute HEC functions

HECEVAL($hecp\text{ar}, f_{n,k,\ell}, X, y; r_Z$)	HECDEC($hecp\text{ar}, d, Z$)
1 : parse $X = (\text{pk}_{AH}, A_1, \dots, A_{n+1})$, $y = (y_{id}, y_1^{at}, \dots, y_\ell^{at})$, $r_Z = (\{r_{E,1}, \dots, r_{E,\ell}\}, \{r_1, \dots, r_\ell\}, r_{nf})$ 2 : $E \leftarrow \bigoplus_{i=1}^{n+1} (A_i \odot y_{id}^i)$ 3 : $\forall i \in [\ell] : Y_i \leftarrow \text{Enc}(\text{pk}_{AH}, y_i^{at}; r_i)$ 4 : $\forall i \in [\ell] : Z_i \leftarrow ((r_{E,i} \odot E) \oplus Y_i)$ 5 : $Z_{nf} = r_{nf} \odot E$ 6 : return $Z = (\{Z_1, \dots, Z_\ell\}, Z_{nf})$	1 : parse $d = (\text{sk}_E, f_{n,k,\ell}, x)$, $Z = (\{Z_1, \dots, Z_\ell\}, Z_{nf})$ 2 : $\forall i \in [\ell] : y_{g,i}^{at} \leftarrow \text{Dec}(\text{sk}_E, Z_i)$ 3 : $y_g \leftarrow \text{Dec}(\text{sk}_E, Z_{nf})$ 4 : if $y_g \neq g(0)$ 5 : return \emptyset 6 : for $y_{id} \in x$ 7 : if $g(y_{id}) = y_g$ 8 : return $(y_{id}, y_1^{at}, \dots, y_\ell^{at})$ where $\forall i \in [\ell] : y_i^{at} \in \text{domain}_{f,y}$ $\wedge g(y_i^{at}) = y_{g,i}^{at}$ 9 : return \emptyset

D.3 Constructions of HEC Schemes

KLN Construction of HEC from Fully Homomorphic Encryption (FHE)

In Def. 15, we review the definition of Circuit-private fully homomorphic encryption (CP-FHE) as from [AHH⁺23, OPCPC14]. We then review the construction of HEC from CP-FHE from [KLN23] and then prove that it is a consistent HEC scheme.

Definition 15 (Circuit-private fully homomorphic encryption (CP-FHE)).

Algorithms (FHEKeyGen, FHEEnc, FHEDec, FHEEval) form a secure fully homomorphic public-key encryption scheme [Gen09, BV11, BGV12, GSW13] if:

Input-output specification: FHEKeyGen($1^\lambda, \Lambda$) takes as input the security parameter and possibly system parameters Λ and outputs a secret key sk_{FHE} and a public key pk_{FHE} . FHEEnc(pk_{FHE}, b) takes as input the public key and a bit $b \in \{0, 1\}$ and outputs a ciphertext c . FHEDec(sk_{FHE}, c) takes as input a ciphertext c and outputs the decrypted bit $b \in \{0, 1\}$. FHEEval($\text{pk}_{FHE}, \mathcal{C}, c_1, \dots, c_n$) takes as input a public key, a Boolean circuit $\mathcal{C} : \{0, 1\}^n \mapsto \{0, 1\}$, and n ciphertexts and outputs a ciphertext c_C ; correctness (below) ensures that c_C is an encryption of $\mathcal{C}(b_1, \dots, b_n)$ when c_i encrypts b_i .

Correctness of evaluation: For any integer n (polynomial in λ) for any circuit \mathcal{C} with n inputs of size that is polynomial in λ , for all $x \in \{0, 1\}^n$, the event that $\text{FHEDec}(\text{sk}_{FHE}, C) \neq \mathcal{C}(x)$ where $(\text{sk}_{FHE}, \text{pk}_{FHE})$ are outputs of FHEKeyGen, ciphertexts c_i are outputs of FHEEnc(pk_{FHE}, x_i), and c_C is output of FHEEval($\text{pk}_{FHE}, \mathcal{C}, c_1, \dots, c_n$), has probability 0.

Security: FHE must satisfy the standard definition of semantic security.

Compactness: What makes fully homomorphic encryption non-trivial is the property that the ciphertext c_C should be of a fixed length that is independent of the size of the circuit \mathcal{C} and of n . More formally, there exists a polynomial $s(\lambda)$ such that for all circuits \mathcal{C} , for all $(\text{sk}_{FHE}, \text{pk}_{FHE})$ output by FHEKeyGen(λ) and for all input ciphertexts c_1, \dots, c_n generated by FHEEnc(pk_{FHE}, \cdot), c_C generated by FHEEval($\text{pk}_{FHE}, \mathcal{C}, c_1, \dots, c_n$) is at most $s(\lambda)$ bits long.

Circuit-privacy: As defined by [Gen09, OPP14, BdMW16, DD22] an FHE scheme is circuit private for a circuit family \mathcal{C} if for any PPT algorithm Adv that outputs $(R, \mathcal{C}_0, \mathcal{C}_1, (x_1, r_1), \dots, (x_n, r_n))$, the probability of distinguishing the homomorphic evaluation of \mathcal{C}_0 on $\{c_i = \text{FHEEnc}(\text{pk}_{FHE}, x_i; r_i)\}_{i \in [n]}$ where pk_{FHE} is computed as FHEKeyGen($1^\lambda; R$) cannot be distinguished from the corresponding evaluation of \mathcal{C}_1 on the same ciphertexts, as long as $\mathcal{C}_0(x_1, \dots, x_n) = \mathcal{C}_1(x_1, \dots, x_n)$.

Construction of HEC for any f from CP-FHE. For a Boolean function $g : \{0, 1\}^{\ell_x} \times \{0, 1\}^{\ell_y} \mapsto \{0, 1\}$, an ℓ_y -bit string y and a value $z \in \{0, 1\}^2$, let $\mathcal{C}_{y,z}^g(x)$ be the Boolean circuit that outputs $g(x, y)$ if $z_1 = 0$, and z_2 otherwise.

Recall that our goal is to construct a secure f -HEC scheme with a direct encryption algorithm; suppose that the length of the output of f is ℓ ; for $1 \leq j \leq \ell$, let $f_j(x, y)$ be the Boolean function that outputs the j^{th} bit of $f(x, y)$. Suppose we are given an FHE scheme that is circuit-private for the families of circuits $\{\mathcal{C}_j\}$ defined as follows: $\mathcal{C}_j = \{\mathcal{C}_{y,z}^{f_j}(x) : y \in \{0, 1\}^{\ell_y}, z \in \{0, 1\}^2\}$.

HECSETUP(1^λ) \rightarrow $hecp\text{ar}$: Generate the FHE parameters $hecp\text{ar}$, if needed.

HECENC($hecp\text{ar}, f, x$) \rightarrow (X, d) : Generate $(\text{sk}_{FHE}, \text{pk}_{FHE}) \leftarrow$ FHEKeyGen($1^\lambda, hecp\text{ar}$). Let $|x| = n$; set $c_i \leftarrow$ FHEEnc(pk_{FHE}, x_i). Output $X = (\text{pk}_{FHE}, c_1, \dots, c_n)$, and decryption key $d = \text{sk}_{FHE}$.

$\text{HECEVAL}(hecp\text{ar}, f, X, y) \rightarrow Z$: Parse $X = (\text{pk}_{FHE}, c_1, \dots, c_n)$. For $j = 1$ to ℓ ,
 compute $Z_j \leftarrow \text{FHEEval}(\text{pk}_{FHE}, \mathcal{C}_{y,00}^{f_j}, c_1, \dots, c_n)$. Output $Z = (Z_1, \dots, Z_\ell)$.
 $\text{HECDEC}(hecp\text{ar}, d, Z) \rightarrow z$: Output $(\text{FHEDec}(d, Z_1), \dots, \text{FHEDec}(d, Z_\ell))$.
 $\text{HECDIRECT}(hecp\text{ar}, X, z) \rightarrow Z$: Parse $X = (\text{pk}_{FHE}, c_1, \dots, c_n)$. For $j = 1$ to ℓ ,
 compute $Z_j \leftarrow \text{FHEEval}(\text{pk}_{FHE}, \mathcal{C}_{0^\ell, 1z_j}^{f_j}, c_1, \dots, c_n)$. Output $Z = (Z_1, \dots, Z_\ell)$.

Theorem 7. For a FHE scheme, $(\text{FHEKeyGen}, \text{FHEEnc}, \text{FHEDec}, \text{FHEEval})$ with the Correctness property, for a circuit family $\{\mathcal{C}_j^f : f \in F\}$ (as defined in [KLN23]), the construction in [KLN23] is a consistent HEC for the family F .

Proof. Let us assume the existence of an adversary \mathcal{A} that is able to produce a $(f, x, \text{st}, r, y, r_Z)$ such that $Z \leftarrow \text{HECEVAL}(hecp\text{ar}, f, X, y; r_Z)$ but $\text{HECDEC}(hecp\text{ar}, d, Z) \neq f(x, y)$. We can then construct an adversary \mathcal{A}' from adversary \mathcal{A} which outputs x, y and Φ_y^f where the output of the circuit $\Phi_y^f(x) = f(x, y)$.

This gives us a tuple (x, y, Φ_y^f) for which the keys $\text{sk}_{FHE}, \text{pk}_{FHE} \in \text{FHEKeyGen}(\lambda)$, $c \in \text{FHEEnc}(\text{FHEPK}, x)$ from the output of $\text{HECENC}(hecp\text{ar}, f, x; r)$ and $c_\Phi \in \text{FHEEval}(\text{FHEPK}, \Phi, c)$ from Z are as required, but $\text{FHEDec}(\text{FHESK}, c_\Phi) \neq \Phi(x)$. Since the correctness of FHE (as provided in Appx. D.3) is defined over all possible inputs x and y , all randomness tapes, and for all circuits Φ , the tuple (x, y, Φ_y^f) is clearly a violation of the correctness condition. This proves that the HEC construction is indeed consistent.

As shown by [KLN23] both Security of x , SECX and the security of x and y from third parties, SECXY is obtained by the semantic security of the FHE. The security of DIRECTZ follows from the circuit privacy.

Additional proofs for consistent HEC scheme

Lemma 10 (Security of DIRECTZ for Fig. D.3) *Our construction in Fig. D.3 achieves security of DIRECTZ defined in Def. 14.*

Proof of Lem. 10 We prove the theorem for the two separate cases of when the user is in the watchlist and when they are not.

For the former, since the user is on the watchlist, $f(x, y) \neq 0$. In HECEVAL, (Z_{id}, Z_{at}) is an encryption of $f(x, y)$ and in HECDIRECT, Z_{nf} is an encryption of 0. Considering the experiments $\text{DIRECTZ}_0^{\text{Adv}}$ and $\text{DIRECTZ}_1^{\text{Adv}}$, since the ciphertext of $f(x, y)$ is output in both cases, the indistinguishability of the experiments can be reduced to the IND-CPA security of the underlying encryption scheme.

In the case where the user is not on the watchlist, $f(x, y) = \emptyset$. Since separate randomness is used for each of r_1, r_2 , and r_3 in HECEVAL, therefore each ciphertext is the encryption of a random value, $Z_{id} = r_1 P(y_{id}) + y_{id}$, $Z_{at} = r_2 P(y_{id}) + at$ and $Z_{nf} = r_3 P(y_{id})$ because $P(y_{id}) \neq 0$. This makes the three ciphertext values indistinguishable from random in $\text{DIRECTZ}_0^{\text{Adv}}$. In experiment $\text{DIRECTZ}_1^{\text{Adv}}$, the HECDIRECT function simply encrypts random values when $f(x, y) = 0$. Therefore, the two experiments are indistinguishable and we achieve security of DIRECTZ.

$\text{HECDEC}(hecp\text{ar}, d, Z)$	$\text{HECEVAL}(hecp\text{ar}, f_{n,k,\ell}, X, y; r_{\hat{Z}})$
1 : parse $d = (\text{sk}_E, f_{n,k,\ell}, x)$, $Z = (Z_{id}, Z_{at}, Z_{nf})$ 2 : $y'_{id} \leftarrow \text{Dec}(\text{sk}_E, Z_{id})$ 3 : $y'_{at} \leftarrow \text{Dec}(\text{sk}_E, Z_{at})$ 4 : $y' \leftarrow \text{Dec}(\text{sk}_E, Z_{nf})$ 5 : if $y' \neq g(0)$ 6 : return \emptyset 7 : for $y_{id} \in x$ 8 : if $g(y_{id}) = y'_{id}$ 9 : return (y_{id}, y_{at}) where $y_{at} \in \text{domain}_{f,y,at}$ $\wedge g(y_{at}) = y'_{at}$ 10 : return \emptyset	1 : parse $X = (\text{pk}_{AH}, A_1, \dots, A_{n+1})$, $y = (y_{id}, y_{at})$, $r_{\hat{Z}} = (r_{id}, r_{at}, r_1, r_2, r_3)$ 2 : if $r_3 = 0$, return \perp 3 : $E \leftarrow \bigoplus_{i=1}^{n+1} (A_i \odot y_{id}^i)$ 4 : $Y_{id} \leftarrow \text{Enc}(\text{pk}_{AH}, y_{id}; r_{id})$ 5 : $Y_{at} \leftarrow \text{Enc}(\text{pk}_{AH}, y_{at}; r_{at})$ 6 : $Z_{id} \leftarrow (r_1 \odot E) \oplus Y_{id}$ 7 : $Z_{at} \leftarrow (r_2 \odot E) \oplus Y_{at}$ 8 : $Z_{nf} = r_3 \odot E$ 9 : return $Z = (Z_{id}, Z_{at}, Z_{nf})$
$\text{HECENC}(hecp\text{ar}, f_{n,k}, x)$	$\text{HECDIRECT}(hecp\text{ar}, X, z)$
1 : $(\text{pk}_{AH}, \text{sk}_E) \leftarrow \text{KeyGen}(1^\lambda)$ 2 : $s \leftarrow \mathcal{M}_{\text{pk}_{AH}}$ 3 : $P \leftarrow s \prod_{i=1}^n (\chi - x_i)$ 4 : for i in $\{1, \dots, n+1\}$ 5 : $A_i \leftarrow \text{Enc}(\text{pk}_{AH}, P_i)$ 6 : return $(X = (\text{pk}_{AH}, A_1, \dots, A_{n+1}),$ 7 : $d = (\text{sk}_E, f_k, x))$	1 : parse $X = (\text{pk}_{AH}, A_1, \dots, A_{n+1})$ 2 : $z = (z_1, z_2, z_3)$ 3 : if $z = \emptyset$ 4 : $\beta_1 \leftarrow \mathcal{M}_{\text{pk}_{AH}}$ 5 : $\beta_2 \leftarrow \mathcal{M}_{\text{pk}_{AH}}$ 6 : $\beta_3 \leftarrow \mathcal{M}_{\text{pk}_{AH}}$ 7 : return $(\text{Enc}(\text{pk}_{AH}, \beta_1),$ 8 : $\text{Enc}(\text{pk}_{AH}, \beta_2), \text{Enc}(\text{pk}_{AH}, \beta_3))$ 9 : return $(\text{Enc}(\text{pk}_{AH}, g(z_1)),$ 10 : $\text{Enc}(\text{pk}_{AH}, g(z_2)), \text{Enc}(\text{pk}_{AH}, g(z_3)))$

Fig. D.3: HEC algorithms

Lemma 11 (Security of x and y for Fig. D.3) *Our construction in Fig. D.3 achieves security of x and y from third parties.*

Proof of Lem. 11. Let us assume there exists an adversary for whom $|p_{\text{Adv},0}^{\text{SECXY}}(\lambda) - p_{\text{Adv},1}^{\text{SECXY}}(\lambda)|$ is non-negligible. This implies that either (i) the adversary can distinguish an encryption of x_0 from x_1 or (ii) the adversary can distinguish an encryption of y_0 from y_1 . From Lem. 12, the adversary distinguishing an encryption of x_0 from an encryption of x_1 can be reduced to the IND-CPA game of the underlying scheme. This holds similarly for y_0 and y_1 .

Lemma 12 (Security of x for Fig. D.3) *Our construction in Fig. D.3 achieves Security of y*

Proof of Lem. 12. Let us assume there exists an adversary Adv for whom $|p_{\text{Adv},0}^{\text{SECX}}(\lambda) - p_{\text{Adv},1}^{\text{SECX}}(\lambda)|$ is non-negligible. Let x_0 and x_1 be the input for which Adv wins the SECX game by correctly distinguishing the ciphertext of x_0 from the ciphertext of x_1 . In that case, we can construct an IND-CPA adversary Adv' that wins the IND-CPA game by using the same input x_0 and x_1 . This is possible since Adv does not possess the secret key for the HEC scheme. Thus, IND-CPA security of the underlying encryption scheme implies the SECX security of the HEC scheme.

E Additional Constructions and Proofs of Security for R_P , R_f , and Ψ_2

We present the omitted verification function for R_P (described in Section 3.2) in Alg. 5.

Algorithm 5 $V_P^*(C_y, c_0, \dots, c_{n-1}, C_P, \text{aux}, \pi, \tau) \rightarrow \{0, 1\}$

```

1: if  $n = 1$ ,
2:   parse  $\pi = (\pi_1)$ 
3:   Verify  $\pi_1$  using the eqrep protocol in Appx. C.3
4:   return 0 if  $\pi_1$  didn't verify, otherwise, return 1
5: parse  $\pi = ((C_1, C_2, C_3, C'_P, \pi_\alpha), \pi')$ 
   Query the random oracle with the current transcript ( $\tau$ ) of the proof
6:  $\alpha \leftarrow H(C_1, C_2, C_3, \tau)$ 
   Compute the encryptions of the new coefficients for a reduced degree polynomial
7:  $\forall i \in [n/2 - 1], c'_i = \boxed{x'_i} = \boxed{x_i} \oplus (\boxed{x_{i+n/2}} \odot \alpha)$ 
8: Verify  $\pi_\alpha$  using  $C_P, C_1, C_2, C_3, C'_P, c_0, \dots, c_{n-1}, c'_0, \dots, c'_{n/2-1}$ 
   and the verification function for commitments to AHE described in Section 3.1.
9: If  $\pi_\alpha$  failed to verify, return 0, otherwise, continue.
10: parse  $\tau = (\pi, \tau')$ 
11: return  $V_P^*(C_y, c'_0, \dots, c'_{n-1}, C'_P, \text{aux}, \pi', \tau') \rightarrow \{0, 1\}$ 

```

E.1 Proof of security for the R_P proof system

Proof of Thm. 1 (Completeness and ZK). Completeness is clear by inspection.

The zero knowledge property of Alg. 2 relies on the hiding and zero knowledge property of our underlying ciphertext and scalar commitment scheme and associated protocols described in Sec. 3.1 and constructed in Sec. 4. Since we have committed to all values and do all proofs with a NIZK scheme with a trapdoor that allows our simulator to produce proofs for relations not in the language, we can simply choose random elements as our commitments and simulate all proofs. We show the simulator for PoK_P and PoK_P^* in Algs. 7 and 6 for completeness. We can see that if we replace the real commitments and proofs one-by-one with

hybrids, an adversary that can distinguish these hybrids can defeat either the hiding of the commitment or the zero knowledge of the proof systems.

We quickly review the Schwartz-Zippel lemma [Sch80,Sho97] in Lemma 13. We will use this in our proof of black-box simulation extractability proof for Alg. 2 in Thm. 2

Lemma 13 (Schwartz-Zippel [Sch80,Sho97]) *For two distinct polynomials, $r(\chi)$, $r'(\chi)$, over a field, \mathbb{F} of size p , the probability that $r(\alpha) = r'(\alpha)$ when α is sampled randomly from \mathbb{F} is d/p where d is the larger degree out of either polynomial, $d = \max\{\deg r, \deg r'\}$. Where “distinct polynomials” means there exists some power where the coefficients for r and r' differ.*

We need one more form of the Schwartz-Zippel lemma in order to prove our construction sound for Camenisch-Shoup encryptions which we show in Lemma 14

Lemma 14 (Schwartz-Zippel for \mathbb{Z}_n) *For two distinct polynomials, $r(\chi)$, $r'(\chi)$, over a ring, \mathbb{Z}_n where $n = pq$ for p, q prime, the probability that $r(\alpha) = r'(\alpha)$ when α is sampled randomly from \mathbb{Z}_n is d/p where d is the larger degree out of either polynomial, $d = \max\{\deg r, \deg r'\}$ and WLOG $q \geq p$. Where “distinct polynomials” means there exists some power where the coefficients for r and r' differ.*

Proof of Lemma 14. Let us label the polynomial, $r(\chi) - r'(\chi)$, as $t(\chi)$. We can see that because $t(\alpha) = 0 \pmod n$, we have that $t(\alpha) = 0 \pmod p$ and $t(\alpha) = 0 \pmod q$ since $p|n$ and $q|n$. Let us define a map from $\mathbb{Z}_n[x]$ to $\mathbb{Z}_p[x]$, ϕ_p where for $t(\chi) = t_0 + t_1\chi + \dots + t_d\chi^d$ we have that $\phi_p(t(\chi)) = \sum s_i\chi^i$ where $s_i = t_i \pmod p$. Thus, if $t(\alpha) = u \pmod n$, then $s(\alpha) = u \pmod p$. We also know that the polynomial, $t(\chi)$ in $\mathbb{Z}_n[\chi]$ is not identically zero for one of the two polynomial $\phi_p(t)$ or $\phi_q(t)$. If this were not true, then the coefficients of $t(\chi)$ in \mathbb{Z}_n would be multiples of both p and q (since p, q prime and $pq = n$) and thus the coefficients would be multiples of n . This would mean the coefficients would be zero in \mathbb{Z}_n but we’ve assumed that $t(\chi) \in \mathbb{Z}_n[x]$ is not identically zero. WLOG we’ll assume $\phi_p(t)$ is a non-zero polynomial in $\mathbb{Z}_p[\chi]$. We thus know that we can map this polynomial onto a non-zero polynomial in $\mathbb{Z}_p[\chi]$. We’ll call this polynomial $s(\chi) \in \mathbb{Z}_p[\chi]$. Thus, we know that $s(\alpha) = 0 \pmod p$ since $t(\alpha) \in \mathbb{Z}[x]$ is some multiple of n and $p|n$. Because $s(\alpha) = 0 \pmod p$ and $s(\chi) \pmod p$ is not identically zero, we can use Lemma 13 for the field \mathbb{Z}_p to determine the probability of this evaluating to 0 (for a random evaluation point) is d/p . Because this must be true if $t(\alpha) = 0 \pmod n$, this must only occur with at most d/p probability. By choosing p to be the smaller prime factor of n , we’ve proven our bound in Lemma 14. \square

Proof of Thm. 2 (Simulation extractability of PoK_p^).* This property of Alg. 2 relies on the BB-extraction and binding of our underlying ciphertext and scalar

commitment scheme and associated protocols described in Sec. 3.1 and constructed in Sec. 4. We can use the simulator (SimPoK_P) in Alg. 7 for this reduction. Because our simulator is zero knowledge, the BB-simulation-extractability adversary gets no advantage when given these proofs.

To do this, we'll prove that C is correctly computed and that we can extract the witnesses for the relation. We can prove that we can extract recursively. As a base case, we see that when ProveRecursive is called with $n = 1$. We can see on line 1 that in this case, the correct computation of P is directly computed.

Thus, if we can prove that C is correctly computed, assuming that C'_P is correctly computed, we can use induction to conclude that the original commitment given to the recursion from $\Psi_2.P$ (on line 4 of Alg. 3) was correctly computed. From the proof, π_{rec} , we know that $P'(y) = e_1 + \alpha e_3$. We see that α is computed from a hash of the transcript, including C_1 and C_3 . Thus, the adversary cannot make e_1 or e_3 depend on α , since this would reduce to either distinguishing a random oracle or double opening C_1 or C_3 . We now rewrite these polynomials and fix

y to reform these as: $q(\chi) = e_1 + \chi e_3$ and $q'(\chi) = \sum_{i=0}^{n/2-1} y^i a_i + \sum_{i=0}^{n/2} \chi y^i a_{i+n/2}$. For

the proof to succeed, $q(\chi)$ must equal $q'(\chi)$ when evaluated at the random value, α . We know from the Schwartz-Zippel lemma (Lemma 13) that the probability of this occurring when $q(\chi)$ is distinct from $q'(\chi)$ is negligible in the size of the ring, \mathbb{Z}_r . Thus, with overwhelming probability, these must be equivalent polynomials. Because α is multiplied by the right term and not the left, and (with overwhelming probability) the polynomials are equivalent, this further proves

that $e_1 = \sum_{i=0}^{n/2-1} y^i a_i$ and $e_3 = \sum_{i=0}^{n/2} y^i a_{i+n/2}$. This is because e_1 is the 0-degree

coefficient in $q(\chi)$ and $\sum_{i=0}^{n/2-1} y^i a_i$ is the 0-degree coefficient in $q'(\chi)$ (with similar

reasoning for e_3 and $\sum_{i=0}^{n/2} y^i a_{i+n/2}$ for being the 1-st degree coefficient of $q(\chi)$ and $q'(\chi)$). We then see that π_C proves that $e_2 = e_3 \odot y^{n/2}$. Thus, $e_2 = e_3 \odot y^{n/2}$ and

since we proved e_3 correctly with π_C , we now know that $e_2 = \sum_{i=0}^{n/2} \chi y^{i+n/2} a_{i+n/2}$.

We then see that π_{rec} proves that $e = e_1 + e_2$, which proves that $e = \sum_{i=0}^n \chi y^i a_i$,

thus, proving C to be correctly formed. Thus, after extracting all witnesses from the underlying NIZKs, we know that these are correct witnesses for the relation.

E.2 Proofs for the Ψ_2 proof system

Proof of Thm. 8 (Completeness and ZK). Our scheme is correct by inspection. We see that because the proof only consists of commitments and zero-knowledge proofs, it is zero-knowledge as well.

Proof of Thm. 9 (BB-PSL). We assume in this theorem that we can extract a witness for the relation R_f in a black-box way (Thm. 2) by instantiating the

Algorithm 6 $\text{SimPoK}_P(C_y, c_0, \dots, c_{n-1}, C_P) \rightarrow \pi$

- 1: $C_P \leftarrow \text{Com}_{AH}(*; 0)$ where $*$ is a random value
 - 2: For $i = 1$ to $\log n$, let $(C_{y^{2^i}}, r_i) = \text{Com}(*)$
 and let $\pi_{y^{2^i}} \leftarrow \text{Sim}[(z, r_{i-1}, r_i) : C_{y^{2^{i-1}}} = \text{Com}(z; r_{i-1}) \wedge C_{y^{2^i}} = \text{Com}(z^2; r_i)]$.
 - 3: Initialize $\text{aux} = (\{C_{y^{2^i}}\}, \{\pi_{y^{2^i}}\})$.
 - 4: **return** $\text{SimPoK}_P^*(C_y, c_0, \dots, c_{n-1}, C_P, \text{aux})$
-

Algorithm 7 $\text{SimPoK}_P^*(C_y, c_0, \dots, c_{n-1}, C_P, \text{aux}) \rightarrow \pi$

- 1: **if** $n = 1$, **return** (aux, π_1) where $\pi_1 \leftarrow \text{Sim}[r : \text{Com}_{AH}(\overline{x_0}, r) = C_P]$
 - 2: $(C_1, \rho_1) = \text{Com}_{AH}(*)$
 - 3: $(C_2, \rho_2) = \text{Com}_{AH}(*)$
 - 4: $(C_3, \rho_3) = \text{Com}_{AH}(*)$
 - 5: $\alpha \leftarrow H(\tau)$
 - 6: $\forall i \in [n/2 - 1]$, $c'_i = \overline{x'_i} = \overline{x_i} \oplus (\overline{x_{i+n/2}} \odot \alpha)$
 - 7: $(C'_P, r') = \text{Com}_{AH}(*)$
 - 8: $\pi_\alpha \leftarrow \text{Sim}[r, \rho_1, \rho_2, \rho_3, r', r_y, y, \overline{e}, \overline{e_1}, \overline{e_2}, \overline{e_3}, \overline{e'}]$
 - 9: $\text{Com}_{AH}(\overline{e}, r) = C_P \wedge \text{Com}_{AH}(\overline{e'}, r') = C'_P \wedge \forall 1 \leq i \leq 3 : \text{Com}_{AH}(\overline{e_i}, \rho_i) = C_i$
 - 10: $\wedge \overline{e} = \overline{e_1} \oplus \overline{e_2}$
 - 11: $\wedge \overline{e_2} = y^{n/2} \odot \overline{e_3}$ ▷ proven relative to $C_{y^{n/2}}$ in aux
 - 12: $\wedge \overline{e'} = \overline{e_1} \oplus (\alpha \odot \overline{e_3})$
 - 13: Append $(C_1, C_2, C_3, C'_P, \pi_\alpha)$ to aux
 - 14: **return** $(\text{SimPoK}_P^*(C_y, c'_0, \dots, c'_{n/2-1}, C'_P, \text{aux}))$
-

NIZKs in Alg. 3 with the proof function for R_f in Alg. 8 (Appendix E.4). Thus, we know that the ciphertext (C_{sky}) containing $g(y)$ is correct, and thus, our straight line extractor (defined in 2) can extract $g(y) = g^*(y, r_{\hat{z}})$ by decrypting this ciphertext.

E.3 Construction of NIZKs in Ψ_2 Proof Scheme

In Algs. 3 and 2, we perform the following four proofs:

$$\begin{aligned}
 \pi_{\hat{z}} &\leftarrow \text{NIZK}[O, O_Y, r_1, r_2, r_3, E, Y, O, O_{id}, y_{id}, y : \\
 &\wedge \text{Com}(y_{id}, O_{id}) = C_{id} \wedge \text{Com}(y, r_y) = C_y \wedge y = (y_{at}, y_{id}) \\
 &\wedge \text{Com}_{AH}(E, O) = C \wedge \text{Com}_{AH}(Y, O_Y) = C_Y \\
 &\wedge Y = \text{Enc}_{AH}(\text{pk}_{AH}, y; r_3) \\
 &\wedge ((E \odot r_1) \oplus Y, E \odot r_2) = \hat{Z}]
 \end{aligned}$$

$\pi_\alpha \leftarrow \text{NIZK}[O, O_1, O_2, O_3, O', r_y, y, O_z, z, r_z, \boxed{e}, \boxed{e_1}, \boxed{e_2}, \boxed{e_3}] :$

$$\text{Com}_{AH}(\boxed{e}, O) = C \wedge \text{Com}_{AH}(\boxed{e'}, O') = C'$$

$$\wedge \forall 1 \leq i \leq 3 : \text{Com}_{AH}(\boxed{e_i}, O_i) = C_i$$

$$\wedge \text{Com}_{AH}(\boxed{e'}, O') = C'; \text{Com}(z, O_z) = C_z;$$

$$\wedge \boxed{e} = \boxed{e_1} \oplus \boxed{e_3}$$

$$\wedge \boxed{e_3} = z \odot \boxed{e_2}$$

$$\wedge \boxed{e'} = \boxed{e_1} \oplus (\alpha \odot \boxed{e_2})$$

$\pi_2 \leftarrow \text{NIZK}[y, r_z, y, O, r_y : \text{Com}(y, r_y) = C_y$

$$\wedge \text{Com}_{AH}(\boxed{e}, O) = C \wedge \boxed{e} = \bigoplus_{i=0}^n \boxed{a_i} \odot y^i$$

$$\wedge \text{Com}(y; r_z) = C_z]$$

$\pi_z \leftarrow \text{NIZK}[z, z', r_z, r'_z : \text{Com}(z, r_z) = \text{Com}(z', r'_z) = C_z \wedge z = z' * z']$

For $\pi_{\hat{Z}}$, we see that we need to construct a proof for $\hat{Z} = ((E \odot r_1) \oplus Y, E \odot r_2)$. We can prove each element in \hat{Z} separately. Proving $\hat{Z}_1 = r_2 \odot E$ is straightforward. The prover creates a commitment to r_2 and then invokes our multiplication protocol $\text{Prove}_{AH}^{\text{mult}}$ for our ciphertext commitments and scalar commitments.

For proving $\hat{Z}_2 = r_1 \odot E \oplus Y$, the prover needs to create intermediate commitments to $r_1 \odot E$ and Y , C_E and C_Y . The proof for $r_1 \odot E$ will be formed similar to our proof for $r_2 \odot E$, i.e. by committing to r_1 and then invoking our $\text{Prove}_{AH}^{\text{mult}}$ protocol. The prover then proves that C is committed to the product of C_E and C_Y . This can be done using our $\text{Prove}^{\text{add}}$ function for commitments to ciphertexts.

For π_α , we have a number of addition and multiplication proofs. The prover is trying to show that $\boxed{e} = \boxed{e_1} \oplus \boxed{e_3}$ and $\boxed{e_3} = \boxed{ze_2}$. This is straightforward as the prover can apply $\text{Prove}^{\text{add}}$ and $\text{Prove}^{\text{mult}}$ directly to C_1, C_2, C_3 , and C_z . For the last proof $\boxed{e'} = \boxed{e_1} \oplus \boxed{\alpha e_2}$, the prover will need to create an intermediate commitment to $\boxed{\alpha e_2}$ and use $\text{Prove}^{\text{mult}}$ to prove that it was computed with C_2, C_1 and α (we can create a canonical commitment to α to reuse our multiplication protocol as-is). We then compute $\text{Prove}^{\text{add}}$ on this intermediate commitment and C_1 to prove the final product contained in C' .

For π_2 , the prover performs a number of intermediate commitments for each $i \in [n]$. The prover computes and proves intermediate commitments to each $H_i = y^i \odot \boxed{a_i}$ and then computes and proves intermediate commitments to larger products of the elements, $D_i = \bigoplus_{j=1}^i \boxed{a_j y^j}$ for $i \in [n]$, using the previous commitment to prove the next, i.e. $D_i = D_{i-1} \oplus H_i$.

Note that this could be done much more efficiently by having the verifier compute the homomorphic operation on the commitment themselves, but using intermediate commitments and proofs assumes less about our underlying commitments. While this is easy to do with Pedersen commitments, the size of

value committed to by Damgård-Fujisaki commitments grows as homomorphic operations are performed on them. Having the prover use $\text{Prove}^{\text{add}}$ and $\text{Prove}^{\text{mult}}$ ensures that our values stay low as discussed in Sec. 4.2 in Remark 1.

For π_z , the prover creates commitment C'_z to $y^{(n+1)/4}$ and proves via $\text{Prove}^{\text{mult}}$ that the value in this commitment multiplied by itself is equivalent to C_z .

E.4 Construction and proof of security for the R_f proof system

We construct our multivariable proof system in Alg. 8. In this proof function, we prove a special class of polynomials, which is simpler to present, though just as powerful. In this class of polynomials, we break the polynomial down in terms of monomials (polynomials with a single term) of powers of the different y_i variables. Specifically, each polynomial is defined by a vector of coefficients, (a_1, \dots, a_n) , and a vector of powers of y_i 's, for each a_i , $((d_{1,1}, \dots, d_{1,k}), \dots, (d_{n,1}, \dots, d_{n,k}))$ such that $d_{i,j}$ is the power of y_j in the monomial with coefficient a_i . The resulting form of the polynomial looks as: $f = \sum_{i=1}^n a_i x_i \prod_{j=1}^k y_j^{d_{i,j}}$. We then show that any polynomial (which is linear in the x_i 's) can be proven correct using this proof by possibly duplicating x_i 's and adding an extra encryption of 1 to the x_i 's to ensure the polynomial can have a degree-0 term in any x_i . As in Alg. 2, we assume that the prover also has already created a commitment to each $\{y_i, y_i^2, y_i^4, y_i^8, \dots, y_i^{d_i}\}$ where d_i is the largest power of y_i in the polynomial and proved that it was correct, and these commitments and proofs are included in the aux variable passed to the proof and they are implicit and used in line 17 in Alg. 8. We also prove the relation such that the verifier only has a commitment to c_f instead of the actual ciphertext, similar to PoK_P^* in Sect. 3.2. This allows us to recursively call PoK_f^* without revealing intermediate ciphertexts.

In this proof of knowledge, we reduce the degree of y_1 by half at each step. We assume that the maximum degree of each variable, y_i , is a power of 2¹⁷. After a logarithmic number of recursions, we'll have that y_1 only has degree 1 when calling the proof. This will be divided out in line 10 of the proof (in Alg. 8) and thus, we'll be left with f_4 (the polynomial we recurse on) being a degree 0 polynomial in y_1 . Thus, on the next recursive step, we'll trigger the conditional on line 4 and will remove y_1 from the witnesses (and polynomial). Thus, our proof will remove variables, y_i , one-by-one, until we have 0 left, in which we'll trigger the conditional on line 1, in which we're almost finished since at this point, f is a function of linear operations on the x_i values which the verifier can compute. The prover simply needs to prove that the C_f is a commitment to the c_f computed by the verifier. We note the steps that a verifier can also compute with a star (*).

Complexity analysis. We can see that at each step, we reduce the degree of one of the y_i variables by half. By the end, all of the y_i variables have been

¹⁷ If not, we can add a “dummy” monomial with the smallest power of 2 in each variable such that this degree is larger than any degree of that variable in the original polynomial. This dummy monomial can simply have a coefficient of 0 to ensure it doesn't affect the outcome.

removed from the polynomial and thus because our polynomial is linear in the x_i 's, the verifier can compute the encryption themselves, meaning our proof is independent of n . Thus, our complexity will be $O(k \log(d_{\max}))$ where d_{\max} is the maximum degree among all y_i variables in the polynomial.

Proof of theorems 3 and 4. For zero knowledge, it's easy to see that because we're committing to every encryption and variable, and using ZKPs to manipulate them, our proof is also ZK. On the last recursion, the verifier does see an encryption in the clear, which seems to contradict zero-knowledge, but we can see that this is simply a combination of the original coefficients (x_i) and random outputs from the random oracle. For BB extraction, we can prove this by induction. If $f_4(\dots)$ is correctly computed, and C_4^* is truly a commitment to $c_1^* \oplus \alpha c_3^*$. Then, we know that $f_3(\dots)$ and $f_1(\dots)$ must be correctly computed (due to similar logic as the proof for PoK_{f_y}). Thus, because we've also proven that $f_2(\dots) = y^{d_{\max}/2}$ and $f(\dots) = f_1(\dots) + f_2(\dots)$, we've proven correctness of $f(\dots)$. When $d_{\max} = 0$, we simply relabel our witnesses, removing one which isn't necessary to prove $f(\dots)$ anymore. As our base case, we have that if there are no y_i variables left, we can prove correctness of the encryption of c_f .

F Commitments to Ciphertexts

F.1 Additional Notes and Proofs for Simplified Camenisch-Shoup

Subgroup Generators g and h are both in the group $|QR_{n^2}| = \{x : x \in QR_{n^2} \wedge x < n^2/2\}$ We see that g is in $|QR_{n^2}|$ because it is equal to $|(g')^{2n}|$. Squaring $g' \in \mathbb{Z}_{n^2}$ ensures that the result is in QR_{n^2} and taking the absolute value of an element in QR_{n^2} ensures the result is in $|QR_{n^2}|$. We prove that $h \in QR_{n^2}$ using Lemma 2 and $h \in |QR_{n^2}|$ follows from the fact that $|1+n| = 1+n$.

Additional notes about the simplified Camenisch-Shoup scheme. Another modification we've made to the Camenisch-Shoup cryptosystem is that we remove the third element from ciphertexts. Camenisch and Shoup [CS03] construct their scheme with a third element to prove CCA security. We've removed the third element from these ciphertexts as we do not need CCA security for our scheme. Since we don't need the third element to correctly decrypt honest ciphertexts, we can simply drop the element and attain CPA security. CPA security is sufficient for our purposes since we provide proofs of correct encryption.

Correctness and CPA security of simplified Camenisch-Shoup in Fig. 4.1a. Since the third element is only used in [CS03] for CCA security, our decryption algorithm works for honest encryptions. This is because $h^m = (1+n)^m = \sum_{i=0}^m \binom{m}{i} 1^{m-i} n^i = 1 + mn + (m-1)n^2 + \dots = 1 + mn \pmod{n^2}$ and y^r can be cancelled out with u^x . We can see that taking the absolute value of ciphertexts does not affect this correctness because part of the decryption squares the ciphertexts. Because $c^2 = (|c|)^2$, after squaring the ciphertexts our decryption algorithm works correctly. Security holds via a straightforward reduction from the CCA security of the original Camenisch-Shoup scheme.

Algorithm 8 $\text{PoK}_f^*(params, f, X, W)$

-
- parse** $f = \sum_{i=1}^n a_i x_i \prod_{j=1}^k y_j^{d_{i,j}}$; in other words, f consists of n monomials (m_1, \dots, m_n) and for $1 \leq i \leq n$, the i^{th} monomial involves is linear in x_i ; it is a product of x_i and the monomials of y -variables, $m_i(y_1, \dots, y_k) = \prod_{j=1}^k y_j^{d_{i,j}}$ where $d_{i,j}$ is the degree of variable y_j in the i^{th} monomial.
- parse** $X = (\text{pk}_{AH}, \overline{x_1}, \dots, \overline{x_n}, C_1, \dots, C_k, C_f)$
and $W = (y_1, \dots, y_k, c_f, r_1, \dots, r_y, r_f)$.
 $W = (y_1, \dots, y_k, r_1, \dots, r_k, c_f = [f(x_1, \dots, x_n, y_1, \dots, y_k)], r_f)$
- 1: **if** $k = 0$,
 - 2: **return** Prove that C_f is the commitment to $c_f = \sum_{i=1}^n a_i x_{j_i}$ (the verifier can compute c_f autonomously).
 - 3: Let d_{\max} be the maximum degree of y_1 in any monomial.
 - 4: **if** $d_{\max} = 0$ (i.e. y_1 does not appear in f),
 - 5: **return** $\text{PoK}_f^*(params, f', X', W')$ where $f' = f$, $X' = (\text{pk}_{AH}, \overline{x_1}, \dots, \overline{x_n}, C_2, \dots, C_k, C_f)$, $W' = (y_2, \dots, y_k, r_1, \dots, r_k, c_f = [f(x_1, \dots, x_n, y_1, \dots, y_k)], r_f)$.
 - 6: Recursive step:
 - 7: * Let (e'_1, \dots, e'_t) be the indices such that y_1 in the monomials $(m_{e'_1}, \dots, m_{e'_t})$ has degree $\geq d_{\max}/2$. Let (e^*_1, \dots, e^*_s) be the indices of the remaining monomials $(m_{e^*_1}, \dots, m_{e^*_s})$ with degree $< d_{\max}/2$ over y_1 . Note that $s + t = n$.
 - 8: * Let $f_1(x_1, \dots, x_n, y_1, \dots, y_k) = \sum_{i=1}^t a_{e'_i} x_{e'_i} \prod_{j=1}^k y_j^{d_{e'_i,j}}$
 - 9: * Let $f_2(x_1, \dots, x_n, y_1, \dots, y_k) = \sum_{i=1}^s a_{e^*_i} x_{e^*_i} \prod_{j=1}^k y_j^{d_{e^*_i,j}}$
 - 10: * Let $f_3(x_1, \dots, x_n, y_1, \dots, y_k) = \sum_{i=1}^s a_{e^*_i} x_{e^*_i} (\prod_{j=1}^k y_j^{d_{e^*_i,j}}) / y_1^{d_{\max,1}/2}$
 - 11: Compute $\forall i \in [3], c_i^* = [(f_i(x_1, \dots, x_n, y_1, \dots, y_k))]$ computed homomorphically from the input to the prover, and let $\forall i \in [3], (C_i^*, \kappa_i) = \text{Com}(c_i^*)$.
 - 12: Let $\alpha = H(\tau)$ where τ is a transcript of the proof so far (along with the statement and parameters) that includes C_1^* , C_2^* and C_3^* .
 - 13: * Let x'_1, \dots, x'_t be a reordering of x_1, \dots, x_n such that x'_1, \dots, x'_t correspond to the monomials in which y_1 was of degree $< d_{\max}/2$, and x'_{t+1}, \dots, x'_n correspond to those where the degree was $\geq d_{\max}/2$.
 - 14: * Let $(x_1^*, \dots, x_n^*) = (x'_1, \dots, x'_t, \alpha x'_{t+1}, \dots, \alpha x'_n)$. Compute $[\overline{x_1^*}], \dots, [\overline{x_n^*}]$, and let X^* be the same as X except that $[\overline{x_1}], \dots, [\overline{x_n}]$ are replaced by $[\overline{x_1^*}], \dots, [\overline{x_n^*}]$, so the order in which the encrypted x variables appear in X^* corresponds to the order in which they appear in the monomials of f_4 .
 - 15: * Let $f_4(x_1, \dots, x_n, y_1, \dots, y_k) = f_1(x_1, \dots, x_n, y_1, \dots, y_k) + \alpha f_3(x_1, \dots, x_n, y_1, \dots, y_k)$.
 - 16: Compute $c_4^* = \text{Enc}(f_4(x_1, \dots, x_n, y_1, \dots, y_k))$ homomorphically using X^* , and $(C_4^*, r_4^*) = \text{Com}(c_4^*)$.
 - 17: Prove that $c_2^* = c_3^* \odot y_1^{d_{\max}/2}$ using the commitments, C_i^* and openings, κ_i , using Prove_{AH}^{mult} , yielding π_α .
 - 18: Prove that $c_f = c_1^* \oplus c_2^*$ using the commitments, C_f , C_i^* and openings, r_f, κ_i , using Prove_{AH}^{add} , yielding π_f .
 - 19: Prove that $c_4^* = c_1^* \oplus \alpha c_3^*$ using the commitments, C_4 , C_i^* and openings, r_4, κ_i , using Prove_{AH}^{add} , yielding π_4 .
 - 20: **return** $(\pi_f, \pi_\alpha, \pi_4, \text{PoK}_f^*(params, f_4, X^*, W))$
-

Remark 1 (Reducing the size of scalars). Our protocols for commitments must have a maximum size of the witnesses (the committed values). We label this as T . This bound ensures that our protocols remain zero knowledge. For our Camenisch-Shoup scheme, this will need to be $T = \mathbb{Z}_n$ since \mathbb{Z}_n is our message space for these ciphertexts. We run into a problem with $|QR_{n^2}|$ commitments that we didn't have with \mathbb{G}_p commitments here because the scalar commitments we use (Damgård-Fujisaki commitments) do not directly commit to the message space of Camenisch-Shoup commitments. Thus, in order to keep exponents small after an exponentiation proof, we'll also include a proof of modular arithmetic over n in our exponentiation proof. This ensures that the values needed in the proofs never grow large enough to violate our zero knowledge property. This proof of modular arithmetic works by computing a commitment to n and then proving that a remainder of n in a commitment is equal to the original commitment summed with a multiple of n . This ensures that honest provers can reduce the size of the commitments while still proving equivalence modulo n . As an example, let a prover have two $|QR_{n^2}|$ commitments and one scalar commitment, $C_M = (|Mg^{s_M}a_M|, (g')^{s_M}(h')^{r_M})$, $C_N = (|Ng^{s_N}a_N|, (g')^{s_N}(h')^{r_N})$, $C_y = (g')^y(h')^r$. To prove that $|N| = |M^{y \bmod n}|$, the prover will construct $|QR_{n^2}|$ commitment $C_P = (|Pg^{s_P}a_P|, (g')^{s_P}(h')^{r_P})$ where $|P| = |M^y|$ and $C_Q = (|Qg^{s_Q}a_Q|, (g')^{s_Q}(h')^{r_Q})$ where $|Q| = |M^n|$. They will then prove that $|N| = |M^{y \bmod n} * (M^n)^k|$ where $k = y - (y \bmod n)$. This can be done generically using *eqrep-n** described in Sec. 2.1. Notice that a prover could select an incorrect k value in this proof. This is not a problem because larger scalars only affects zero knowledge and not soundness. Thus any honestly created commitments and proofs will remain zero knowledge and any malicious proofs will remain sound.

F.2 Security Proof of Damgård-Fujisaki Commitments for $\mathcal{G} = \mathbb{Z}_{n^2}$

We present our modified version of Damgård-Fujisaki commitments which lie in \mathbb{Z}_{n^2} in Fig. F.1. In this construction, 2^B is roughly the order of $\phi(n^2)$ (where ϕ is Euler's totient function) though 2^B is computable without knowing $\phi(n^2)$ (as defined in [DF02]).

Damgård and Fujisaki [DF02] list four properties sufficient for an Abelian group to create an integer commitment scheme. They then prove that the group \mathbb{Z}_n satisfies these properties. We will prove these properties for the group \mathbb{Z}_{n^2} .

The assumptions Damgård and Fujisaki required to prove their integer commitment scheme secure are shown below. They [DF02] provide a construction and prove that if a group meets all four requirements, their construction is secure. We will modify these requirements slightly and prove that \mathbb{Z}_{n^2} satisfies them. In these assumptions, C is some number which is super polynomial in the security parameter, but smaller than the primes, p, q, p', q' .

Damgård-Fujisaki commitment properties:

Fig. F.1: Simplified Damgård-Fujisaki commitments in \mathbb{Z}_{n^2}

<p>Setup(1^λ) \rightarrow <i>params</i> :</p> <hr/> <p>1: Sample $O(\lambda)$-bit SG primes p', q' and compute $p = 2p' + 1, q = 2q' + 1, n = pq$. 2: Sample random $g, h \in \mathbb{Z}_{n^2}$. 3: return <i>params</i> = (g, h)</p> <p>Commit(<i>params</i>, m) \rightarrow (C, O) :</p> <hr/> <p>1: To commit to integer, m, compute: $C = g^m h^r$ where $r \leftarrow_{\\$} [2^{B+\lambda}]$ 2: Let the opening be $O = r$ 3: return (C, O)</p>
--

1. **Strong root property** - Let Adv be any PPT algorithm. After generating the group with security parameter, λ , then, with a description of the group, \mathcal{G} , (without the trapdoor) and a random $h \in \mathcal{G}$, Adv is tasked with outputting $y \in \mathcal{G}$ and a number, $t > 1$, such that $y^t = h$. The probability of this occurring is negligible.
2. **Small order property** - Let Adv be an PPT algorithm. With a description of the group, \mathcal{G} , Adv is tasked with outputting $b \in \mathcal{G}, \sigma \in \mathbb{Z}$ such that $b \neq 1, b^2 \neq 1, 0 < \sigma < C$, and $b^\sigma = 1$. The probability of this occurring is negligible.
3. **No large even powers in orders** - Any element in \mathcal{G} of the form a^{2^t} has odd order.
4. **Many elements with only large prime factors in orders** - If h is chosen randomly in \mathcal{G} , then there is an overwhelming $(1 - O(2^{-\lambda}))$ probability that the order of h has no prime factors less than C .

Damgård and Fujisaki [DF02] prove that \mathbb{Z}_n satisfies these properties where $n = pq$ and $p \equiv q \equiv 3 \pmod{4}$ and p, q are safe primes. The primes, p and q , are not given to the adversary in these assumptions.

We now prove that these properties hold for \mathbb{Z}_{n^2} with n formed the same way as in Damgård-Fujisaki [DF02]. We review the strong RSA assumption (Assumption 1 of [DF02]), and prove a useful lemma (Lemma 15).

Assumption 1 (Strong RSA assumption[DF02]) *Given $n = pq$ (where $|n| = O(2^\lambda)$), and a number, $t \in \mathbb{Z}_n$, no PPT algorithm can find a pair, v, e such that $v^e = t$ and $e > 1$ with non-negligible probability in λ .*

Lemma 15 *If $a = b \pmod{n^2}$, then $a = b \pmod{n}$.*

Proof of Lemma 15 Take values $a, b \neq 0 \in \mathbb{Z}_{n^2}$ such that $a = b \pmod{n^2}$. This implies that $a = mn^2 + d, b = on^2 + d$ for some $m, o \in \mathbb{Z}$ where $0 < d < n^2$.

This implies that $a = m'n + d, o'n + d$ where $m' = mn, o' = on$. If we take the remainder of $d \bmod n$, as $d = ln + \rho$ for some $l \in \mathbb{Z}$ where $0 < \rho < n$, we find that the following equation holds: $a = (m' + l)n + \rho, (o' + l)n + \rho$. Since division with remainder is unique for $0 \leq \rho < n$, we've shown that a and b are equal mod n .

Proof of DF Property 1 for \mathbb{Z}_{n^2} . Assume we have a PPT algorithm that given $t \in \mathbb{Z}_{n^2}$ can produce a $g \in \mathbb{Z}_{n^2}, y$ such that $g^y = t \bmod \mathbb{Z}_{n^2}$. We are then tasked with creating a reduction to strong RSA in \mathbb{Z}_n . Let our reduction take t in \mathbb{Z}_n and give $t + bn \bmod n^2$ to this adversary where b is a random number drawn from 0 to $n-1$. The adversary then provides g, y such that $g^y = (t + bn) \bmod n^2$. Since this equality holds in \mathbb{Z}_{n^2} , it holds in \mathbb{Z}_n as well due to Lemma 15. We can see that $t + bn = t \bmod n$. Thus $g^y = t \bmod n$. Lastly, we have to prove that $(t + bn)$ is distributed indistinguishably from a uniform drawing from \mathbb{Z}_{n^2} . We can see that $t + bn$ can "reach" almost every element of \mathbb{Z}_{n^2} since if $t = n-1$ and $b = n-1$, then $t + bn = n-1 + (n-1)n = n-1 + n^2 - n = n^2 - 1$ and if $t = 1, b = 0$, we get 1. Then, we see that there are no duplicates of $t + bn$ across this range since no $t, b, t', b' \in \{0, \dots, n-1\}$ exist such that $t + bn = t' + b'n$. There are $(n-1)n$ possible combinations of t and b from our ranges. Thus, each value mapped to by $t + bn$ uniformly maps to a random element of \mathbb{Z}_{n^2} except for values of \mathbb{Z}_{n^2} where n is a factor. There are only n samples of \mathbb{Z}_{n^2} that are divisible by n out of a total of n^2 instances and thus the probability of drawing one of these samples is negligible and our assumed strong RSA adversary in \mathbb{Z}_{n^2} must be able to solve problems when the challenge is not a multiple of n with non-negligible probability.

Proof of DF Property 2 for \mathbb{Z}_{n^2} . The only possible orders of elements in \mathbb{Z}_{n^2} are $2, 4, p, q, p', q'$ or some product of these. If the adversary outputs a b with $\sigma = 2$, we see that it must be that $b^2 = 1$ and thus this is not a valid solution. If σ is a multiple of p, q, p' , or q' , then $\sigma > C$ and thus this solution doesn't work for this property. Thus, the only possible values for σ is 4. We can see that, in this case, if b^2 is a non-trivial root of 1 (i.e. $b^2 \neq -1$) we can factor by rewriting $(b-1)(b+1) = 0 \bmod n^2$ thus ensuring that taking the gcd of $b-1$ or $b+1$ with p, q, p' , or q' yields a factorization. We see that if $b^4 = 1$ and $b^2 = -1$, this must be true in \mathbb{Z}_p and \mathbb{Z}_q due to the Chinese remainder theorem. We can see that because $p \equiv 3 \pmod{4}$, it must be that $p = 4k + 3$ and thus $(p-1)/2$ is odd and so $(-1)^{(p-1)/2} = -1$ implying that (-1) is not a quadratic residue mod p . Thus, if $b^4 = 1$ but $b^2 = -1$, this would be a contradiction and thus b^2 must be a non-trivial square root allowing us to factor.

Proof of DF Property 3 for \mathbb{Z}_{n^2} . We see that the order of $\phi(n^2)$ is $2pqp'q'$ and thus, if a^{2t} has even order, then a has order $4k$ but $4 \nmid 2pqp'q'$ and thus does not divide the order of the group and thus we have a contradiction and a^{2t} cannot have even order.

Proof of DF Property 4 for \mathbb{Z}_{n^2} . If we find a non-trivial square root of 1, we factor and we showed in the proof of DF Property 2 that if we find a 4-th root

of 1, it must be that when we square the value, we can factor. Thus, these must be hard to sample, otherwise, it would be trivial to factor. Thus, the only orders of sampleable elements (by a PPT algorithm) must be some product of p, q, p' and q' . We can simply set $C < p, q, p', q'$ and $p, q, p', q' \approx O(2^\lambda)$ to satisfy this.

F.3 Commitments to \mathbb{G}_p Elements and ElGamal Ciphertexts

In this section, we introduce commitments to group elements (in \mathbb{G}_p) and then construct a commitment scheme to ElGamal ciphertext in Fig. F.3 which relies on those commitments to group elements. Note that the generators g and h used in this section are distinct from those used in the encryption schemes in Sec. 4.1. In this section, g and h refer to commitment bases for a Pedersen commitment.

Commitments to \mathbb{G}_p group elements. In Alg. F.2 we present a commitment scheme for committing to group elements. Our parameters for the scheme are the same as a Pedersen commitment, yielding g and h . We then commit to a group element by computing $C_1 = Mg^s$ and $C_2 = g^s h^r$. We can see that C_2 is a Pedersen commitment and that s is hidden by C_2 . Thus, for any $M, C_1, C_2 \in \mathbb{G}_p$, there exists an s, r that forms a valid opening. We can see that using the opening information, the group element can be retrieved by computing $M = C_1/g^s$.

Proof of opening of an committed group element. We can create a ZK proof of knowledge of an opening of the commitment $C = (C_1, C_2) = \text{Com}_{\mathbb{G}_p}(M)$ by proving knowledge of an opening for C_2 as a Pedersen commitment, i.e. it is the proof of knowledge of representation of C_2 in bases g and h .

Proof of equality of committed group elements. Proving that two group commitments $C = (C_1, C_2) = (Mg^s, g^s h^r)$ and $C' = (C'_1, C'_2) = (M'g^{s'}, g^{s'} h^{r'})$ are committed to the same value ($M = M'$) reduces to a proof of knowledge of equality of representations: $\text{NIZK}[M, M', s, r, s', r' : C_1/C'_1 = g^{s-s'} \wedge C_2/C'_2 = g^{s-s'} h^{r-r'}]$. We can see that this proof works because $C_1/C'_1 = Mg^s/(M'g^{s'}) = g^{s-s'}$ and $C_2/C'_2 = g^s h^r/(g^{s'} h^{r'}) = g^{s-s'} h^{r-r'}$. If the second commitment were committed to a distinct value, then C_1/C'_1 would equal $Mg^s/(M'g^{s'}) = (M/M')g^{s-s'}$ which the adversary could not prove was equivalent to $g^{s-s'}$.

Proof of multiplication of committed group elements. We can also prove that a commitment $C_c = (C_{c,1}, C_{c,2}) = (cg^{s_c}, g^{s_c} h^{r_c})$ opens to the product c of two group elements a, b committed to by two other group element commitments, $C_a = (C_{a,1}, C_{a,2}) = (ag^{s_a}, g^{s_a} h^{r_a})$ and $C_b = (C_{b,1}, C_{b,2}) = (bg^{s_b}, g^{s_b} h^{r_b})$ using $\text{eqrep-}\mathbb{G}_p$. This can be done by having the verifier and prover compute $D_1 = C_{c,1}/(C_{a,1}C_{b,1}) = cg^{s_c}/(bg^{s_b}ag^{s_a})$ and $D_2 = C_{c,2}/(C_{a,2}C_{b,2}) = g^{s_c}h^{r_c}/(g^{s_a}h^{r_a}g^{s_b}h^{r_b})$. We can see that if the relation is true, c will be cancelled out by ab in D_1 , leading to D_1 being simply the result of an exponentiation of g (we'll label this exponent $\beta_1 = s_c - s_a - s_b$). Further, we see that if the relation is true, D_2 is a Pedersen commitment to β_1 . The prover then proves the relation: $\text{PoK}_{\text{eqrep-}\mathbb{G}_p}[s_a, s_b, s_c, r_a, r_b, r_c, \beta_1, \beta_2 : D_1 = g^{\beta_1} \wedge D_2 = g^{\beta_1} h^{\beta_2}]$ where $\beta_1 = s_c - s_a - s_b$ and $\beta_2 = r_c - r_a - r_b$. We

can see that if D_1 can be represented as g^{β_1} and D_2 can be represented as a Pedersen commitment to β_1 , we know that C_c is a commitment to ab .

Proof of exponentiation of committed group elements. We can also prove the exponentiation of a \mathbb{G}_p commitment using a scalar in a Pedersen commitment. This can be done by using the $eqrep\text{-}\mathbb{G}_p$ relation described in Sec. 2.1. An exponentiation proof takes group element commitments C_a to \mathbb{G}_p element, a , and C_b to element b . It also takes in a Pedersen commitment C_y to y . The goal of this proof is to prove that $a = b^y$. To do this, we prove that $\text{PoK}_{eqrep\text{-}\mathbb{G}_p}[y, r_y, \beta_1, \beta_2 : C_y = g^y h^{r_y} \wedge C_{a,1} = C_{b,1}^y g^{\beta_1} \wedge C_{a,2} = C_{b,2}^y g^{\beta_1} h^{\beta_2}]$ where $\beta_1 = s_a - y s_b$ and $\beta_2 = r_a - y r_b$ and where $C_y = g^y h^{r_y}$, $C_{a,1} = a g^{s_a}$, $C_{b,1} = b g^{s_b}$, $C_{b,2} = g^{s_b} h^{r_b}$, and $C_{a,2} = g^{s_a} h^{r_a}$.

Another notable feature of this commitment scheme is that the commitments are homomorphic, i.e. if $C = \text{Com}_{\mathbb{G}_p}(M; (s, r))$ and $C' = \text{Com}_{\mathbb{G}_p}(M'; (s', r'))$, then $C \cdot C' = \text{Com}_{\mathbb{G}_p}(MM'; (s + s', r + r'))$.

Fig. F.2: Commitments to \mathbb{G}_p elements

<p>$\text{Setup}_{\mathbb{G}_p}(1^\lambda) \rightarrow \text{params}$</p> <hr/> <ol style="list-style-type: none"> 1: Generate a group of prime order p, $\mathbb{G}_p = \langle g \rangle$. (or using an existing group e.g. from a bilinear pairing) 2: Generate a random element $h \in \mathbb{G}_p$ as the base for opening. 3: return $\text{params} = (\mathbb{G}_p, g, h)$ <p>$\text{Commit}_{\mathbb{G}_p}(\text{params}, M \in \mathbb{G}_p) \rightarrow C, O$</p> <hr/> <ol style="list-style-type: none"> 4: $s \leftarrow \mathbb{Z}_p; r \leftarrow \mathbb{Z}_p$ 5: $C \leftarrow (C_1, C_2) = (Mg^s, g^s h^r)$ 6: return $C, O = (s, r)$

Theorem 11 *Our construction in Fig. F.2 is binding.*

Proof of Thm. 11 If a PPT adversary can produce (C, M, M', s, s', r, r') such that $C_1 = Mg^s = M'g^{s'}$ and $C_2 = g^s h^r = g^{s'} h^{r'}$ where $M \neq M'$, we can double open C_2 as a Pedersen commitment. We see that if $M \neq M'$, then $s \neq s'$ because otherwise $M = C_1/g^s = C_1/g^{s'} = M'$. Thus, $s \neq s'$ and s, r, s', r' is a valid double opening for C_2 as a Pedersen commitment. The binding property of Pedersen commitments relies on the computational Diffie-Hellman assumption and so our \mathbb{G}_p commitments are computationally binding.

Theorem 12 *Our construction in Fig. F.2 is hiding.*

Proof of Thm. 12 For any $M, C_1, C_2 \in \mathbb{G}_p$, we see that $\exists s, r$ such that $C_1 = Mg^s, C_2 = g^s h^r$. This is because g is a generator for \mathbb{G}_p and thus $\exists s$ such that $g^s = C_1/M$. Because C_2 is a Pedersen commitment which is perfectly hiding, there exists an r such that $C_2 = g^s h^r$ for our picked s . Finally, because s is chosen randomly from \mathbb{Z}_p , we see that any M is equally likely given C and thus this commitment scheme is perfectly hiding.

So far, we've constructed commitments to elements of \mathbb{G}_p and discussed their associated proof protocols for opening and multiplication. Next we'll use these commitments and the intuition about their protocols to build commitments to ElGamal ciphertexts. We build these commitments to ElGamal ciphertexts in Fig. F.3. Verifying these proofs is a direct application of the *egrep*- \mathbb{G}_p verification protocol. We put square brackets $[\cdot]$ around secret values for proof functions. We can see in this ElGamal commitment scheme that we set it up by generating Pedersen commitment bases, g, h , while labeling the parameters for the ElGamal encryption scheme as g' and h' . To commit, we form a \mathbb{G}_p commitment to each the two elements of an ElGamal ciphertext, $c = (c_1, c_2)$, yielding C_1, C_2 as a commitment to c_1 and C_3, C_4 as a commitment to c_2 . Because our \mathbb{G}_p commitments are perfect hiding and computationally binding to elements of \mathbb{G}_p , our ElGamal commitments are perfectly hiding and computationally binding as well.

Proofs over commitments to ciphertexts. Inspecting our construction, we see that many of our proofs ($\text{Prove}_{ELG}^{\text{Com}}, \text{Prove}_{ELG}^{\text{add}}, \text{Prove}_{ELG}^{\text{mult}}$) consists of simply performing the proof on both group elements. For example, to prove knowledge of an opening of an ElGamal commitment, we open the Pedersen commitments of each \mathbb{G}_p commitment, C_2 and C_4 . This allows an extractor to recover s_1, s_2, r_1, r_2 allowing the extractor to compute $c_1 = C_1/g^{s_1}$ and $c_2 = C_3/g^{s_2}$. This is how we described opening those \mathbb{G}_p commitments earlier in this section. As another example, we see in $\text{Prove}_{ELG}^{\text{add}}$ that we want to prove that C_c is committed to ciphertext c where $c = ab$ and C_b is committed to ciphertext b and C_a is committed to ciphertext a . We label this add “addition” because multiplying two ciphertexts results in the addition of their encrypted messages. Intuitively, $\text{Prove}_{ELG}^{\text{mult}}$ requires the verifier to use the homomorphic properties of the commitment scheme to multiply two group elements and then requires the prover to prove that the resulting commitment is equivalent to C_a . We can see in this algorithm that $D_1 = C_{c,1}/(C_{a,1}C_{b,1})$ will be a power of g if (and only if) $c = ab$ because $D_1 = cg^{s_c}/(ag^{s_a}bg^{s_b}) = cg^{s_c - s_a - s_b}/(ab)$. The same is true for D_3 and D_4 .

Proving a ciphertext is an encryption of a Pedersen committed message. Proving that a committed ciphertext is an encryption of a Pedersen committed message somewhat breaks our ciphertext commitment scheme's paradigm of simply performing proofs on either element in the ciphertext. In this proof, $\text{Prove}_{ELG}^{\text{enc}}$, the prover must prove that the commitment is correctly formed for the message y (whereas in the other proofs, we assume the ciphertexts are correctly formed and proofs can be created without knowledge of the randomness of ciphertexts). Thus, we prove that $c_1 = (g')^{\rho c}$ and $c_2 = k^{\rho c} (h')^y$ where g' and h' are the generators for the encryption scheme (in the case of ElGamal, $g' = h'$ but in Sec. 4.2

we'll see that these may differ). We can see that verifying π ensures that the prover knows c (along with its randomness and message) such that is correct ElGamal encryption of y with randomness ρ_c and C_y is a scalar commitment to y .

Fig. F.3: Commitments to ElGamal ciphertexts

<p>Setup_{ElG}($1^\lambda, \text{params}_{ElG}$) \rightarrow params</p> <hr/> <p>parse $\text{params}_{ElG} = (\mathbb{G}_p, g', h')$</p> <ol style="list-style-type: none"> 1: $(g, h) \leftarrow \mathbb{G}_p$ 2: $\text{params} = (g, h, \text{params}_{ElG})$ 3: return params <p>Commit_{ElG}($\text{params}, c = (c_1, c_2)$) \rightarrow C, O</p> <hr/> <ol style="list-style-type: none"> 1: $s_1, s_2 \leftarrow \mathbb{Z}_p; r_1, r_2 \leftarrow \mathbb{Z}_p$ 2: $C \leftarrow (C_1, C_2, C_3, C_4)$ $= (c_1 g^{s_1}, g^{s_1} h^{r_1}, c_2 g^{s_2}, g^{s_2} h^{r_2})$ 3: return $(C, O = (s_1, s_2, r_1, r_2))$ <p>Prove_{ElG}^{Com}(params, C, M, O) \rightarrow π</p> <hr/> <p>parse $C = (C_1, C_2, C_3, C_4)$, $O = (s_1, s_2, r_1, r_2)$</p> <ol style="list-style-type: none"> 1: $\pi = \text{NIZK}_{\text{eqrep}}[s_1, s_2, r_1, r_2 : C_2 = g^{s_1} h^{r_1}, C_4 = g^{s_2} h^{r_2}]$ 2: return π <p>Prove_{ElG}^{enc}($\text{params}, \text{pk} = k, C_c, C_y, [c, \rho_c, y, O_c, O_y]$) \rightarrow π</p> <hr/> <p>parse $\text{params} = (g, h, \text{params}_{ElG})$ $\text{params}_{ElG} = (\mathbb{G}_p, g', h')$ $O_c = (s_{c,1}, s_{c,2}, r_{c,1}, r_{c,2})$ $c = ((g')^{\rho_c}, k^{\rho_c} (h')^y)$, $O_y = (r_y)$</p> <ol style="list-style-type: none"> 1: $\pi = \text{NIZK}[s_{c,1}, s_{c,2}, s_y, \rho_c, r_{c,1}, r_{c,2}, r_y, y :$ 2: $C_y = g^y h^{r_y}$ 3: $\wedge C_{c,1} = (g')^{\rho_c} g^{s_{c,1}}$ 4: $\wedge C_{c,2} = g^{s_{c,1}} h^{r_{c,1}}$ 5: $\wedge C_{c,3} = k^{\rho_c} (h')^y g^{s_{c,2}}$ 6: $\wedge C_{c,4} = g^{s_{c,2}} h^{r_{c,2}}$ 7: return π 	<p>Prove_{ElG}^{mult}($\text{params}, C_a, C_b, C_y, [c, a, b, y, O_a, O_b, O_y]$) \rightarrow π</p> <hr/> <p>parse $O_a = (s_{a,1}, s_{a,2}, r_{a,1}, r_{a,2})$ $O_b = (s_{b,1}, s_{b,2}, r_{b,1}, r_{b,2})$ $O_y = (r_y)$</p> <ol style="list-style-type: none"> 1: $\beta_1 = s_{a,1} - y s_{b,1}$ 2: $\beta_2 = r_{a,1} - y r_{b,1}$ 3: $\beta_3 = s_{a,2} - y s_{b,2}$ 4: $\beta_4 = r_{a,2} - y r_{b,2}$ 5: $\pi = \text{NIZK}[y, r_y, \beta_1, \beta_2, \beta_3, \beta_4 :$ 6: $C_y = g^y g^{r_y}$ 7: $\wedge C_{a,1} = (C_{b,1})^y g^{\beta_1}$ 8: $\wedge C_{a,2} = (C_{b,2})^y g^{\beta_1} h^{\beta_2}$ 9: $\wedge C_{a,3} = (C_{b,3})^y g^{\beta_3}$ 10: $\wedge C_{a,4} = (C_{b,4})^y g^{\beta_3} h^{\beta_4}]$ 11: return π <p>Prove_{ElG}^{add}($\text{params}, C_a, C_b, C_c, [a, b, c, O_a, O_b, O_c]$) \rightarrow π</p> <hr/> <p>parse $O_a = (s_{a,1}, s_{a,2}, r_{a,1}, r_{a,2})$ $O_b = (s_{b,1}, s_{b,2}, r_{b,1}, r_{b,2})$ $O_c = (s_{c,1}, s_{c,2}, r_{c,1}, r_{c,2})$</p> <ol style="list-style-type: none"> 1: $D_1 \leftarrow C_{c,1} / (C_{a,1} * C_{b,1})$ 2: $D_2 \leftarrow C_{c,2} / (C_{a,2} * C_{b,2})$ 3: $D_3 \leftarrow C_{c,3} / (C_{a,3} * C_{b,3})$ 4: $D_4 \leftarrow C_{c,4} / (C_{a,4} * C_{b,4})$ 5: $\beta_1 = s_{c,1} - s_{a,1} - s_{b,1}$ 6: $\beta_2 = r_{c,1} - r_{a,1} - r_{b,1}$ 7: $\beta_3 = s_{c,2} - s_{a,2} - s_{b,2}$ 8: $\beta_4 = r_{c,2} - r_{a,2} - r_{b,2}$ 9: $\pi = \text{NIZK}[\beta_1, \beta_2, \beta_3, \beta_4 :$ 10: $D_1 = g^{\beta_1}$ 11: $\wedge D_2 = g^{\beta_1} h^{\beta_2}$ 12: $\wedge D_3 = g^{\beta_3}$ 13: $\wedge D_4 = g^{\beta_3} h^{\beta_4}]$ 14: return π
---	---

Theorem 13 (Hiding of the commitments in Fig. F.3) *Our commitments to ElGamal ciphertexts in Fig. F.3 are statistically hiding.*

Proof (Proof of Thm. 13). We can see that (C_1, C_2) is identical to a \mathbb{G}_p commitment to c_1 and (C_3, C_4) is identical to a \mathbb{G}_p commitment to c_2 , we can see that they statistically hide c_1 and c_2 .

Theorem 14 (Binding of the commitments in Fig. F.3) *Our commitments to ElGamal ciphertexts in Fig. F.3 are computationally binding.*

Proof (Proof of Thm. 14). We can see that (C_1, C_2) is identical to a \mathbb{G}_p commitment to c_1 and (C_3, C_4) is identical to a \mathbb{G}_p commitment to c_2 , thus, if a PPT adversary can produce a double opening such that one of these commitments opens to some c'_1 or c'_2 in \mathbb{G}_p , we obtain a double opening for our \mathbb{G}_p commitments.

Theorem 15 (Zero-knowledge of Fig. F.3) *Our protocols in Fig. F.3 ($\text{Prove}_{\text{ElG}}^{\text{Com}}$, $\text{Prove}_{\text{ElG}}^{\text{enc}}$, $\text{Prove}_{\text{ElG}}^{\text{mult}}$, and $\text{Prove}_{\text{ElG}}^{\text{add}}$) are zero-knowledge against any PPT adversary.*

Proof (Proof of Thm. 15). We can see that in each of these NIZKs, we simply return a proof computed from the $\text{eqrep}-p^*$ protocol. Thus, we can use the simulator for this protocol to produce proofs in the zero knowledge games. Thus, if a PPT adversary can distinguish these simulated proofs from real proofs, we can break the zero knowledge of the $\text{eqrep}-p^*$ protocol.

Theorem 16 (Black box knowledge extraction of Fig. F.3) *Given a PPT adversary that can produce a proof that verifies for our protocols in Fig. F.3 ($\text{Prove}_{\text{ElG}}^{\text{Com}}$, $\text{Prove}_{\text{ElG}}^{\text{enc}}$, $\text{Prove}_{\text{ElG}}^{\text{mult}}$, and $\text{Prove}_{\text{ElG}}^{\text{add}}$) there exists an extractor with black-box access to the adversary that can extract a witness that proves the relations true.*

Proof (Proof of Thm. 16). Similar to our proof of zero-knowledge for these protocols, because these protocols simply return $\text{eqrep}-\mathbb{G}_p$ proofs, we can use the black-box extractor for these proofs to extract the witnesses. This extractor is described in Sec. 2.1.

F.4 Proofs of Hiding and Binding for $|QR_{n^2}|$ -commitments in Fig. 4.3

We provide number theory background in Appx. C.2.

Hiding proof for Fig. 4.3. To prove that our commitments are hiding, we show that, for any group element M , the commitment algorithm (which samples a commitment $C = (|Mg^s|, g^s h^r)$) provides a distribution that is statistically close to the distribution $(R_1, R_2) \in (|QR_{n^2}| \times \mathbb{Z}_{n^2})$ drawn uniformly at random.

We can see that since $n = pq$ where p, q are safe primes, then g with overwhelming probability generates QR_{n^2} due to Lemma 4. If s is large, g^s is indistinguishable from a random element of QR_{n^2} since s is much larger than $\text{ord}(g)$ (Lemma 5). Let \star be the “multiply-and-absolute-value” operation in that it takes two elements, multiplies them and then takes the absolute value. We see that $|QR_{n^2}|$ is a group under this operator as $x \star y = |x * y| = |x| * |y|$, $1 = |1|$, $(|QR_{n^2}|, \star)$ is closed since QR_{n^2} is closed and $|\cdot|$ maps QR_{n^2} to $|QR_{n^2}|$, and the inverse of any $x \in (|QR_{n^2}|, \star)$ is $|x^{-1}|$ where $x^{-1} \in QR_{n^2}$ ($|x * x^{-1}| = |1|$). Let $|\cdot| : QR_{n^2} \rightarrow |QR_{n^2}|$ be the map defined by the absolute value function. We see that $|\cdot|$ is a homomorphism as $|x * y| = |x| * |y|$ and $|1| = |1|$. We can see that $|x|$ is bijective as the only values of \mathbb{Z}_{n^2} that map to the same value have the form $-x$ and x , but if $x \in QR_{n^2}$, then $-x \notin QR_{n^2}$ since (-1) is not a quadratic residue (Lemma 3). We also defined $|QR_{n^2}|$ as the image of this function and thus because it is also injective, it is bijective. Thus, $(QR_{n^2}, *) \cong (|QR_{n^2}|, \star)$. This means that $|QR_{n^2}|$ is cyclic and any randomly sampled element of $|QR_{n^2}|$ is likely a generator due to Lemma 4. Thus, $M \star g^s$ is indistinguishable from a random element of $|QR_{n^2}|$.

We note that C_2 is simply a Damgård-Fujisaki integer commitment and thus is indistinguishable from a random element in \mathbb{Z}_{n^2} .

Binding proof for $|QR_{n^2}|$ -commitments in Fig. 4.3. If a PPT adversary can open a commitment $C = (C_1, C_2)$ to two values $M, M' \in \{|x| : x \in |QR_{n^2}|\}$ (providing openings, s, s', r, r') such that $M \neq M'$, we see that it must be that $C_1/g^s \neq C_1/g^{s'}$. If $s \neq s'$, we see that $C_2, (s, r), (s', r')$ is a double opening for the Damgård-Fujisaki integer commitment scheme. Because we proved that these Damgård-Fujisaki commitments are binding for \mathbb{Z}_{n^2} (In Appendix F.2), this double opening violation still holds even if C_1 and C_2 are created maliciously (i.e., they are not in QR_{n^2} , but instead some arbitrary element of \mathbb{Z}_{n^2}). Thus, $s = s'$ and it must be that $|C_1/g^s| = |C_1/g^{s'}|$. This tells us that $|M| = |M'| \in \mathbb{Z}_{n^2}$ and since $|M| = M \forall M \in |QR_{n^2}|$, we see that $M = M' \in |QR_{n^2}|$. Thus it is impossible (based on the strong RSA assumption) for a PPT adversary to double open our $|QR_{n^2}|$ commitments without double opening a Damgård-Fujisaki commitment.

F.5 Construction of Commitments to Camenisch-Shoup Ciphertexts

In Fig. F.4, we give a formal construction of our Camenisch-Shoup commitments.

F.6 Proofs for Commitments to Camenisch-Shoup Ciphertexts

We split Thm. 5 into the following theorems:

Theorem 17 (Zero-knowledge of proofs in Fig. F.4) *Our protocols in Fig. F.4 (Prove $_{CS}^{\text{Com}}$, Prove $_{CS}^{\text{enc}}$, Prove $_{CS}^{\text{mult}}$, and Prove $_{CS}^{\text{add}}$) are zero-knowledge against any PPT adversary.*

Fig. F.4: Commitments to Camenisch-Shoup ciphertexts

<p>Setup_{CS}($1^\lambda, \text{params}_{CS}, \text{params}_{DF}$) \rightarrow params</p> <hr/> <ol style="list-style-type: none"> 1: parse $\text{params}_{CS} = (\mathbb{Z}_{n^2}, g^*, h^*)$ 2: parse $\text{params}_{DF} = (\mathbb{Z}_{n^2}, g', h')$ 3: $g \leftarrow_{\\$} QR_{n^2}$ 4: $\text{params} = (\mathcal{G}, g, g^*, h^*, g', h')$ 5: return params <p>Commit_{CS}(params, c) $\rightarrow C, O$</p> <hr/> <ol style="list-style-type: none"> 1: parse $c = (c_1, c_2)$ 2: $s_1, s_2 \leftarrow_{\\$} [2^{B+\lambda}]; r_1, r_2 \leftarrow_{\\$} [2^{B+\lambda}]$ 3: $a_1, a_2, b_1, b_2 \leftarrow_{\\$} \{-1, 1\}$ 4: $C_1 \leftarrow a_1 c_1 g^{s_1}; C_2 \leftarrow b_1 (g')^{s_1} (h')^{r_1}$ 5: $C_3 \leftarrow a_2 c_2 g^{s_2}; C_4 \leftarrow b_2 (g')^{s_2} (h')^{r_2}$ 6: $C \leftarrow (C_1, C_2, C_3, C_4)$ 7: $O \leftarrow (a_1, a_2, s_1, s_2, r_1, r_2, b_1, b_2)$ 8: return (C, O) <p>Prove_{CS}^{add}($\text{params}, C_a, C_b, C_c, [a, b, c, O_a, O_b, O_c]$) $\rightarrow \pi$</p> <hr/> <ol style="list-style-type: none"> 1: parse $C_a = (C_{a,i})_{i \in [4]}$ 2: $C_b = (C_{b,i})_{i \in [4]}$ 3: $C_c = (C_{c,i})_{i \in [4]}$ 4: $O_a = (a_{a,i}, s_{a,i}, r_{a,i}, b_{a,i})_{i \in [2]}$ 5: $O_b = (b_{b,i}, s_{b,i}, r_{b,i}, b_{b,i})_{i \in [2]}$ 6: $O_c = (b_{c,i}, s_{c,i}, r_{c,i}, b_{c,i})_{i \in [2]}$ 7: $\forall i \in [4], D_i \leftarrow C_{c,i} / (C_{a,i} * C_{b,i})$ 8: $\gamma_1 \leftarrow s_{c,1} - s_{a,1} - s_{b,1}$ 9: $\gamma_2 \leftarrow r_{c,1} - r_{a,1} - r_{b,1}$ 10: $\gamma_3 \leftarrow s_{c,2} - s_{a,2} - s_{b,2}$ 11: $\gamma_4 \leftarrow r_{c,2} - r_{a,2} - r_{b,2}$ 12: $\beta_1 \leftarrow a_{c,1} / (a_{a,1} * a_{b,1})$ 13: $\beta_2 \leftarrow b_{c,1} / (b_{a,1} * b_{b,1})$ 14: $\beta_3 \leftarrow a_{c,2} / (a_{a,2} * a_{b,2})$ 15: $\beta_4 \leftarrow b_{c,2} / (b_{a,2} * b_{b,2})$ 16: $\pi = \text{NIZK}[\{\gamma_i, \beta_i\}_{i \in [4]}]$ 17: $D_1 = \beta_1 g^{\gamma_1}$ 18: $\wedge D_2 = \beta_2 (g')^{\gamma_1} (h')^{\gamma_2}$ 19: $\wedge D_3 = \beta_3 g^{\gamma_3}$ 20: $\wedge D_4 = \beta_4 (g')^{\gamma_3} (h')^{\gamma_4}$ 21: $\wedge \{\beta_i\}_{i \in [4]} \in \{-1, 1\}$ 22: return π 	<p>Prove_{CS}^{Com}($\text{params}, C, [M, O]$) $\rightarrow \pi$</p> <hr/> <ol style="list-style-type: none"> 1: parse $C = (C_1, C_2, C_3, C_4)$, 2: $O = (a_1, a_2, s_1, s_2, r_1, r_2, b_1, b_2)$ 3: $\pi = \text{NIZK}[O :$ $C_2 = b_1 (g')^{s_1} (h')^{r_1} \wedge C_4 =$ $b_2 (g')^{s_2} (h')^{r_2}$ $\wedge b_1 \in \{-1, 1\} \wedge b_2 \in \{-1, 1\}]$ 4: return π <p>Prove_{CS}^{mult}($\text{params}, C_a, C_b, C_y, [a, b, y, O_a, O_b, O_y, \{b_i\}_{i \in [4]}]$) $\rightarrow \pi$</p> <hr/> <ol style="list-style-type: none"> 1: parse $C_a = (C_{a,i})_{i \in [4]}$ 2: $C_b = (C_{b,i})_{i \in [4]}$ 3: $O_a = (a_{a,i}, s_{a,i}, r_{a,i}, b_{a,i})_{i \in [2]}$ 4: $O_b = (b_{b,i}, s_{b,i}, r_{b,i}, b_{b,i})_{i \in [2]}$ 5: $\gamma_1 \leftarrow s_{a,1} - y s_{b,1}; \gamma_2 \leftarrow r_{a,1} - y r_{b,1}$ 6: $\gamma_3 \leftarrow s_{a,2} - y s_{b,2}; \gamma_4 \leftarrow r_{a,2} - y r_{b,2}$ 7: $\beta_1 \leftarrow a_{a,1} / a_{b,1}; \beta_2 \leftarrow b_{a,1} / b_{b,1}$ 8: $\beta_3 \leftarrow a_{a,2} / a_{b,2}; \beta_4 \leftarrow b_{a,2} / b_{b,2}$ 9: $\pi = \text{NIZK}[\{\gamma_i, \beta_i\}_{i \in [4]}]$ 10: $C_y = b_y (g')^y (g')^{r_y}$ 11: $\wedge C_{a,1} = b_1 (C_{b,1})^y (g')^{\gamma_1}$ 12: $\wedge C_{a,2} = b_2 (C_{b,2})^y (g')^{\gamma_1} (h')^{\gamma_2}$ 13: $\wedge C_{a,3} = b_3 (C_{b,3})^y (g')^{\gamma_3}$ 14: $\wedge C_{a,4} = b_4 (C_{b,4})^y (g')^{\gamma_3} (h')^{\gamma_4}$ 15: $\wedge \beta_y, \beta_1, \beta_2, \beta_3, \beta_4 \in \{-1, 1\}$ 16: return π <p>Prove_{CS}^{enc}($\text{params}, \text{pk}_{AH} = k, C_a, C_y, [a, r_a, y, O_a, O_y, b_y, \{b_i\}_{i \in [4]}]$) $\rightarrow \pi$</p> <hr/> <ol style="list-style-type: none"> 1: parse $C_a = (C_{a,i})_{i \in [4]}$ 2: $O_a = (a_{a,i}, s_{a,i}, r_{a,i}, b_{a,i})_{i \in [2]}$ 3: $\pi = \text{NIZK}[O_a, s_y, r_a, r_y, y :$ $C_y = b_y (g')^y (h')^{r_y}$ $\wedge C_{a,1} = b_1 (g^*)^{r_a} (g')^{s_{a,1}}$ $\wedge C_{a,2} = b_2 (g')^{s_{a,1}} (h')^{r_{a,1}}$ $\wedge C_{a,3} = b_3 k^{r_a} (g^*)^y (g')^{s_{a,2}}$ $\wedge C_{a,4} = b_4 (g')^{s_{a,2}} (h')^{r_{a,2}}$ $\wedge b_y, b_1, b_2, b_3, b_4 \in \{-1, 1\}]$ 4: return π
---	---

Proof (Proof of Thm. 17). We can see that in each of these NIZKs, we simply return a proof computed from the *eqrep* protocol. Thus, we can use the simulator for this protocol to produce proofs in the zero knowledge games. Thus, if a PPT adversary can distinguish these simulated proofs from real proofs, we can break the zero knowledge of the *eqrep* protocol.

Theorem 18 (Black box knowledge extraction of proofs in Fig. F.4) *Given a PPT adversary that can produce a proof that verifies for our protocols in Fig. F.4 ($\text{Prove}_{CS}^{\text{Com}}$, $\text{Prove}_{CS}^{\text{enc}}$, $\text{Prove}_{CS}^{\text{mult}}$, and $\text{Prove}_{CS}^{\text{add}}$) there exists an extractor with black-box access to the adversary that can extract a witness that proves the relations true.*

Proof (Proof of Thm. 18). Similar to our proof of zero-knowledge for these protocols, because these protocols simply return *eqrep* proofs, we can use the black-box extractor for these proofs to extract the witnesses. This extractor is described in Sec. 2.1.

Theorem 19 (Hiding of the commitments in Fig. F.4) *Our commitments to Camenisch-Shoup ciphertexts in Fig. F.4 are statistically hiding.*

Proof (Proof of Thm. 19). We can see that (C_1, C_2) is identical to a $|QR_{n^2}|$ commitment to c_1 and (C_3, C_4) is identical to a $|QR_{n^2}|$ commitment to c_2 , we can see that they statistically hide c_1 and c_2 .

Theorem 20 (Binding of the commitments in Fig. F.4) *Our commitments to Camenisch-Shoup ciphertexts in Fig. F.4 are computationally binding.*

Proof (Proof of Thm. 20). We can see that (C_1, C_2) is identical to a $|QR_{n^2}|$ commitment to c_1 and (C_3, C_4) is identical to a $|QR_{n^2}|$ commitment to c_2 , thus, if a PPT adversary can produce a double opening such that one of these commitments opens to some c'_1 or c'_2 in $|QR_{n^2}|$, we obtain a double opening for our $|QR_{n^2}|$ commitments.