

Zero-Knowledge for Homomorphic Key-Value Commitments with Applications to Privacy-Preserving Ledgers

Matteo Campanelli¹, Felix Engelmann², and Claudio Orlandi³

¹ Protocol Labs, matteo@protocol.ai

² ITU, Denmark, fe-research@nlogn.org

³ Aarhus University, Denmark, orlandi@cs.au.dk

Abstract. Commitments to key-value maps (or, authenticated dictionaries) are an important building block in cryptographic applications, including cryptocurrencies and distributed file systems.

In this work we study short commitments to key-value maps with two additional properties: double-hiding (both keys and values should be hidden) and homomorphism (we should be able to combine two commitments to obtain one that is the “sum” of their key-value openings). Furthermore, we require these commitments to be short and to support efficient transparent zero-knowledge arguments (i.e., without a trusted setup).

As our main contribution, we show how to construct commitments with the properties above as well as efficient zero-knowledge arguments over them. We additionally discuss a range of practical optimizations that can be carried out depending on the application domain.

Finally, we formally describe a specific application of commitments to key-value maps to scalable anonymous ledgers. We show how to extend QuisQuis (Fauzi et al. ASIACRYPT 2019). This results in an efficient, confidential multi-type system with a state whose size is independent of the number of transactions.

Keywords: Zero-knowledge · Key-Value map · Commitments.

1 Introduction

In this work we propose constructions for efficient commitments to key-value maps (with specific features) and for efficient zero-knowledge arguments that can prove properties on committed key-value maps.

KEY-VALUE MAPS. We can loosely consider a key-value map as the equivalent of a dictionary in some programming languages (e.g., Python): a way to map arbitrary keys—e.g., strings—to values—e.g., scalars. For example, the balance of a user in a wallet application could be represented by a key-value map as $kv = \{(USD, 100), (BTC, 10)\}$, where each of the different asset types (the keys) are associated to an amount (the values). In this paper we will generally assume that values are in an algebraic group endowed with an addition operation $+$.

Our Focus: Short, Homomorphic, Doubly-Hiding Commitments. A commitment to a key-value map is roughly similar to an ordinary commitment: it cannot be opened to two different key-value maps (binding) and it should not leak anything about neither the keys nor the values in it. In the case of key-value maps, however, we are interested in some additional functional and efficiency-related requirements:

- *Large key universe:* our commitments should support a large universe of keys, potentially superpolynomial in the security parameter⁴. This implies that the algorithms of the commitment scheme should have a runtime independent of (or logarithmic in) the size of the key universe.
- *Short commitments:* our commitments should have size independent not only of the size of the key universe, but also of the *density* of the key-value map. The density is the number of elements whose value is not zero (e.g., the density of kv in the example above was 2).
- *Homomorphic commitments:* we require our commitments to support an homomorphic operation \circ . For example assume each commitment encodes a wallet and that we have two wallets c, c' with $c = \text{Com}(\{(USD, 100), (BTC, 10)\})$ and $c' = \text{Com}(\{(USD, 20), (ETH, 1)\})$. Then we can compute the commitment $c^* \leftarrow c \circ c' = \text{Com}(\{(USD, 120), (BTC, 10), (ETH, 1)\})$ without knowing the opening of any of the commitments. Requiring homomorphism rules out Merkle Trees as a solution. Homomorphic properties of commitments to “structured objects” have wide applications in cryptography (see, e.g., [KZG10] for homomorphic polynomial commitments). The homomorphic property is a natural one and allows many useful applications: as an example we describe applications to privacy-preserving cryptocurrencies in Section 7 and an additional class of application scenarios in Section 1.1.
- *Efficient and transparent⁵ zero-knowledge proofs:* we should be able to prove (and verify) efficiently arbitrary properties over commitments of key-value maps. We are interested in zero-knowledge proofs—which allow to prove properties over a secret value without leaking it—and where both keys and values are part of the secret. For example, one can prove that two committed key-value maps hold the same value for some (hidden) key \tilde{k} . More formally, given as public input commitments c, c' and a public function f , one can prove knowledge of a key \tilde{k} such that c, c' are commitments to key-value maps kv, kv' respectively and $kv[\tilde{k}] = f(kv'[\tilde{k}])$.

While different subsets of these properties have been studied in literature, our contribution is to investigate constructions that require them *all*. Our goal is to provide concretely efficient tools useful in different application domains.

⁴ This is a way to describe our setting asymptotically. We stress, however, that is not necessary: an interesting setting for our constructions is just one where the universe of keys is *concretely* large.

⁵ In a *transparent* argument system the setup does not need to be produced by a trusted party. This property is interesting in the case of *non-interactive* argument systems, which are the focus of this work.

KEY-HIDING PROPERTIES. Here we clarify what we mean by key-hiding properties and discuss how existing solution fail to solve our problem. We have three *key sets* of interest: the set of all the keys in the universe (which we will assume to be $\{0, 1\}^*$ or a field \mathbb{F} from now on), the set K_{active} of *active keys*, defined as all the keys that are being used in the system, and the set K_{com} of *committed keys*, defined as the non-zero keys in any *given* commitment. For example, in a wallet setting, K_{active} consists of all the keys (asset types) encoded in *some* wallet, while K_{com} would consist of those encoded in a *specific given* wallet. Depending on whether we want to hide the active or the committed keys or both we get four different settings, which we discuss below (see also Fig. 1).

PUBLIC ACTIVE KEYS. In the case where both the active keys and the committed keys are public, Pedersen commitments are already a solution to our problem. The system parameters will contain group elements h, g_1, \dots, g_n where there is a known association between k_i and g_i for all active keys k_i . We commit by computing $c = h^r \prod_i g_i^{v_i}$, and proving properties of values is trivial to do using existing sigma protocols since the verifier is allowed to learn the keys. In the case in which the active keys are public but the committed keys are private, Pedersen commitment can still be used but the (proving) complexity of the ZK proof would be linear in the number of active keys⁶. One of our contributions is to show how to bring this down to the size of the committed set.

PRIVATE ACTIVE KEYS. It does not make sense to consider the case where the set of active keys are private but the committed keys are public. The most interesting case is the one in which both of these sets are private. In this setting, it would be possible to commit using a non-homomorphic version of Pedersen commitment. We thus have $2n + 1$ generators $(h, g_1, f_1, \dots, g_n, f_n)$ and we commit computing $c = h^r \prod_i g_i^{v_i} f_i^{k_i}$. Now it is possible to efficiently prove statements but the commitment is not homomorphic (and therefore not applicable in our settings of interest). Our main contribution is to provide a better solution for this case.

1.1 Applications of Our Work

APPLICATION: MULTI-TYPE QUISQUIS. The privacy-preserving transaction system QuisQuis [FMMO19] crucially relies commitments endowed with an homomorphic property. It builds upon accounts to which tokens are deposited in a transaction without interaction of the receiver. Compared to other privacy-preserving transaction systems like Zcash [BCG⁺14] or Monero [NM⁺16], the design of QuisQuis achieves a state size linear in the number of participants instead of monotonically growing over time. We extend this system through a notion of currency types so such that different currencies share a common anonymity set. This allows for a dynamic creation of confidential tokens by any participant without setting up a full separate system. For this application, we

⁶ This is true for the aforementioned approach with sigma-protocols as well as for other straightforward applications of NIZKs.

also present a secret key based key-value map commitment⁷. In combination with efficient NIZKs, to show that transactions conserve all value, we achieve small transaction sizes. We formally describe this application in Section 7.

APPLICATION: PUBLICLY VERIFIABLE EVOLVING DATABASE. Consider a database (representable as key-value store) which receives numerous updates and where we want the content of the database to remain private but we also are interested in the database publicly “evolving through time”.

As an example of the above, consider a register of tax-related information where users are identified by their SSN. The set of valid identities grows dynamically, which results in a high overhead if the public parameter changes every time. Users provide their SSN to their employer who uses it to report the salary. At the end of the month, each employer creates a hiding key-value map commitment with one key per employee and their earned amount as value, e.g. $\delta C_{\text{Corp X, May}} = \text{Com}(\{\text{SSN}_{\text{Alice}} : 3142, \text{SSN}_{\text{Bob}} : 2718\})$. The company may either prove that for the employee’s identity the correct amount was committed (without revealing the identities of co-workers), or reveal the full opening to their employees. Every company publishes these commitments to a persistent log. At the end of the year, the tax authority homomorphically adds all published commitments and can then generate proofs on a single commitment instead of all commitments from all companies. The required value opening is provided by the tax payers and the randomness by the companies. Employees with multiple sources of income get the amounts homomorphically added. Different categories of income may be separated by namespaces in the key.

1.2 Technical Overview

OUR CONSTRUCTION OF KEY-VALUE MAP COMMITMENTS. In order to commit to a key-value map $\{v_k\}_{k \in K}$ we assume a group \mathbb{G} where the discrete logarithm is hard and a hash function \mathcal{H} modeled as a random oracle mapping keys to group elements. We then compute a commitment as $c = \prod_{k \in K} \mathcal{H}(k)^{v_k} h^r$ where h is a random generator of the group and r is a random scalar. This can be seen as a (vector) Pedersen commitment with random key-dependent generators and it has short homomorphic commitments. In the next paragraphs, we show how we can construct efficient zero-knowledge proofs for circuits over such commitments.

MODULAR TRANSPARENT ZERO-KNOWLEDGE ARGUMENTS FOR COMMITTED KEY-VALUE MAPS. Fix a (large) field \mathbb{F} and consider a circuit C over key-value maps (we assume that \mathbb{F} is also both the key and the value space of the key-value map). We assume the syntax of C to be of the type $C(\text{kv}_1, \dots, \text{kv}_\ell, \omega)$, the kv_i -s as private key-value maps and ω as an additional private witness (ω is a vector of field elements). Given such a circuit we are interested in proving an augmented circuit that takes as public input commitments to the ℓ key-value

⁷ I.e., a commitment which can be opened using a secret key in place of the randomness.

maps and proves their opening in addition to the relation from circuit C . More specifically, we propose an argument system for:

$$\begin{aligned}
& C^*(c_1, \dots, c_\ell; (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega) \\
& := C(\text{kv}_1, \dots, \text{kv}_\ell, \omega) \wedge \bigwedge_{i \in [\ell]} c_i = \text{Com}(\text{kv}_i, \rho_i)
\end{aligned} \tag{1}$$

where the part after the semicolon is considered the private witness. A more concrete intuition on the circuit above is: given committed key-value pairs we can prove properties of their values, their keys and any relation between these and other private values (contained in ω).

OUR TEMPLATE FOR ZERO-KNOWLEDGE ARGUMENTS ON COMMITTED KEY-VALUE MAPS. We now describe how to prove properties on committed key-value maps. The following refers to the setting with private-active-keys/private-committed-keys (see lower-right quadrant in Figure 1). We denote the construction for the doubly-private case by $\text{Z}\bullet\text{-dp}$ (first part alternatively spelled “ZKeyWee”, as a pun on “ZK for KV”, and pronounced “zee-kee-wee”; “dp” stands for doubly-private).

	Public K_{active}	Private K_{active}
Public K_{com}	$c = h^r \prod_i g_i^{v_i} \ \& \ \Sigma\text{-protocols}$	<i>(Uninteresting case)</i>
Private K_{com}	c as in Fig. 2 & $\text{Z}\bullet\text{-set}$ (Sec. 1.2+ Appx. A)	c as in Fig. 2 & $\text{Z}\bullet\text{-dp}$ (Sec. 4)

Fig. 1: How to construct commitments to key-value maps & NIZKs over them within our settings and with our requirements of interests. (K_{active} : “active” keys set, i.e., all the keys committed somewhere; K_{com} committed keys set, i.e. those that open the commitments we are using in a proof *right now*). The related constructions are specified in the second lines with our two contributions $\text{Z}\bullet\text{-set}$ and $\text{Z}\bullet\text{-dp}$.

Our construction follows a basic blueprint, which we now exemplify through a concrete case. Consider a committed key-value map $\{v_k\}_{k \in K}$ and the problem of proving in zero-knowledge that all its values v_k are in some range $\{0, \dots, 2^\mu - 1\}$. We proceed in two steps: we first let the prover send the verifier what we call *key-tags*, these are masked version of the non-zero keys in the committed key-value map. The prover will also show that they are valid maskings of some set of keys. By providing key-tags, we can then break the rest of the relation in two parts: *a)* showing knowledge of values (and randomness) that combined with the key-tags produce the commitment c (part of the public input); *b)* showing that these values are in range. We now elaborate on each of these steps.

Given a key-value map $\{v_k\}_{k \in K}$ with density n (the number of non-zero keys⁸) we provide n key-tags by sending $b_k = \mathcal{H}(k)h^{r_k}$ for a random r_k . The prover should also prove that each of them is of the prescribed form. We stress that, in order to do this, we use the heuristic technique of proving a random oracle in zero-knowledge. Next, the prover would show knowledge of values v_1, \dots, v_n and an appropriate r' such that $c = b_{k_1}^{v_1} \dots b_{k_n}^{v_n} h^{r'}$ and $v_i \in \{0, \dots, 2^\mu - 1\}$ for all $i \in [n]$. The latter relation—comprising reconstructing the commitment from the key-tags and the range proof—can for example be performed through a system like the generalized Bulletproofs in LMR [LMR19] or compressed Σ -protocols [ACR20]. These provide interfaces to prove bilinear circuits, of which we only use the non-bilinear gates, with logarithmic proof sizes. We stress that our focus is on *transparent* solutions, i.e., without a trusted setup; all our constructions can be instantiated in a fully transparent manner. We discuss an experimental evaluation in Section 6. We estimate our system can open $n = 100$ values and prove they are in range in approximately one minute. A very loose estimate for the size of the corresponding proof is $< 6\text{KB}$ using [LMR19] in the Ristretto curve (see also Table 1 and Section 4.3).

We remark on two properties of the template of our construction above. *First*, we can easily reduce its amortized cost by splitting it into an offline stage (independent of the commitments on which we are carrying out proofs) and an online stage. We further discuss these improvements in Section 5. *Second*, we can adapt and optimize our construction to the scenario with *public* active keys and private committed keys (lower-left quadrant in Figure 1). We describe this next.

A 2ND CONSTRUCTION WITH REGISTRATION OF ACTIVE KEYS. In some settings, although the whole universe of keys can be extremely large, the set of active keys K_{active} at any given time can have a manageable size and be publicly known. Consider for example applications (e.g. multi-asset transaction system) where there is an exponentially large set of potential asset types (keys), but only a manageable subset of them are present in the system (active) at any given time. Moreover, before becoming active in the system they plausibly need to be registered (for example, through a first “genesis” transaction for that specific asset type). In such settings we can leverage the partial knowledge on existing keys to improve efficiency. We do this by introducing an operation that preprocesses the parameters of the system (or CRS) specializing them for a specific set of active keys. Our proposed construction for this setting—denoted by **Z \odot -set** (for “registered *set*” of keys) and discussed in more detail in Appendix A—assumes the specialized CRS for the set K_{active} to contain an accumulator⁹ to the set of (unmasked) key-tags corresponding to K_{active} , i.e., to the

⁸ Here we consider the case where leaking the density of the key-value map is not a problem. We will also adapt our construction where this leakage does not occur if an upper bound on this density is known.

⁹ An accumulator is a cryptographic data structure that allows to commit to a set in a binding manner and to prove membership of an element efficiently. NB: we can compute accumulators *deterministically* from a set, i.e., without a trusted authority.

set $B_{\text{active}} = \{\mathcal{H}(k) : k \in K_{\text{active}}\}$. Thus, in the online stage, we can produce a proof on a commitment $c = \text{Com}(\{v_k\}_{k \in K})$ by: 1) producing masked key-tags $B' = \{\mathcal{H}(k)h^{r_k} : k \in K\}$ produced with some fresh randomness r_k ; 2) showing that each $b'_k \in B'$ is of the form $b \cdot h^{r_k}$ for some b in the accumulator; 3) showing knowledge of v_k -s such that $c = h^r \prod_k b_k^{v_k}$.

The main advantage of the construction above, **Z \odot -set**, is that it does not require the hashing $\mathcal{H}(k)$ for the key-tags to be proven in zero-knowledge (by exploiting the fact that the keys are public and “pre-registered”). This can result in savings in verification time of one order of magnitude compared to **Z \odot -dp** (we elaborate in Appendix A.4); such savings also apply to the multi-type transaction systems Multi-type QuisQuis (Section 7) and can be extended to other transaction systems (our techniques in **Z \odot -set** are compatible with other frameworks besides QuisQuis and they could be straightforwardly applied, e.g, to obtain a multi-type version of Veksel [CHA21]).

MULTI-TYPE QUISQUIS. In contrast to the original QuisQuis where an account stored a scalar value representing an amount, we generalize accounts to store tokens of different types in a key-value map. Each key corresponds to a type, also known as currency, and the value specifies how many tokens of the specific type are held by the account. An account `acct` holds a balance `kv`, represented as a key-value map. To transfer tokens from one account to another, a transaction includes both accounts as active inputs, denoted by the set \mathcal{P} . The transaction subtracts tokens from one account with a secret index s and adds them to the other one. The output of the transaction are two updated accounts belonging to the same users as the inputs; the input accounts are discarded. To achieve anonymity, the input consists of \mathcal{P} together with a potentially large anonymity set of accounts \mathcal{A} which keep their balance unchanged but provide set anonymity for the active accounts. The sender uses **Z \odot -set** with a circuit to prove that they have enough funds (knowledge of secret key and positive balance) and that the updates are consistent.

1.3 Related Work

Authenticated Data Structures Besides the aforementioned straw-man schemes based on Pedersen, a common approach to succinct key-value commitments uses Merkle trees. They are not homomorphic and opening them in zero-knowledge requires proving a number of hashes logarithmic in the number of committed elements, which can be expensive.

A related primitive is that of vector commitments [CF13]. A limitation of using vector commitments as key-value map commitments in general is that values need to be stored at positions that have already been agreed upon. That is, since we need to know for each key k what is the index i_k it refers to in the vector. This type of common agreement may be achieved in the setting in the bottom left quadrant in Fig. 1 (public active keys) but not in the bottom right one (private active keys). Also, while some constructions of vector commitments are homomorphic (e.g., [CF13]) they lose this property when hiding is added to

them (which is usually achieved by storing hiding commitments to the values of interest). Other constructions do not have this limitation ([LY10, BG18]) but, like vector commitments in general, only support public active keys. We finally remark that, vector commitments focus on a different notion of *succinctness* than the one that is the focus in this paper. Our focus is on a proof size that is sublinear in the size of the circuit we apply on the opening, but not necessarily sublinear in the number of committed elements.

There is a large body of work on succinct commitments to key-value maps, e.g.¹⁰, [AR20, CFG⁺20, BBF19, TXN20]. Differently from our work, constructions in literature are not homomorphic and do not directly support hiding of keys/values. We observe, however, that if one could do without homomorphism the latter problem could be mitigated for some of these constructions by applying masking of keys/values and zero-knowledge. This is true for example for some of the works based on groups of unknown-order [CFG⁺20, BBF19] where we can use techniques to compose algebraic accumulators proofs and succinct zero-knowledge proof systems described in [BCFK19].

The work in [AK20] formalizes encryption on distributed key-value maps with consistency properties; it is not concerned with homomorphism or efficient zero-knowledge. Other works on efficient Zero-Knowledge and key-value maps include Spice [SAGL18]. The authors use data-structures that hide the key but that are not homomorphic. Their constructions use a trusted setup.

Confidential Transaction Systems and Multiple Token Types Here we describe works related to our application, a multi-type version of QuisQuis.

Works on confidential transaction systems include Zcash [BCG⁺14], Monero [NM⁺16], Omniring [LRR⁺19], and Veksel [CHA21]. We now compare these works against the QuisQuis framework, which we extend in this work. The most critical aspect is sender anonymity. Zcash obtains the largest anonymity set among these works (as large as the UTXO set), but it does not have plausible deniability¹¹ and requires a trusted setup. Monero does not have these limitations; it is unclear how the anonymity in Monero fares against that in QuisQuis (see Discussion in [FMMO19]). Omniring improves the transaction size from being linear in the size of the anonymity set (Monero) to a logarithmic size. Both Zcash and Monero style systems, however, have transaction outputs that can (essentially) never be removed from the UTXO set. The payment system in Veksel, like QuisQuis, does not have this drawback. Differently from QuisQuis, Veksel achieves $O(1)$ transaction sizes, but at the price of weaker anonymity guarantees. In all these systems amounts are confidential, yet they lack a notion of type/currency.

The work in [PBF⁺19] introduces confidential types by using homomorphic commitments whose construction is the “single key” version of ours. Their design has also been used in SwapCT [EMP⁺21] and integrated in MimbleWimble [YYD⁺19]. Another construction of confidential types is that of Cloaked Assets

¹⁰ We refer the reader to [TXN20] for a survey of this rapidly growing field.

¹¹ Plausible deniability: no one can tell if a user meant to be involved in a transaction.

developed by Stellar, which separates types and values in two different data structures, similar to our non-homomorphic example. Therefore a transactor requires the openings of all inputs to create a conservation proof, providing no sender anonymity.

2 Notation and Preliminaries

Preliminaries on (Sparse) Key-Value Maps We assume a universe of keys \mathcal{K} and a universe of values \mathcal{V} such that in a key-value map, keys are a subset of \mathcal{K} and values are any element in \mathcal{V} ; they may be of size superpolynomial in a security parameter λ . We assume \mathcal{V} to be an additive group endowed with some operation $+$. A key-value map is defined as a function $\text{kv} : \mathcal{K} \rightarrow \mathcal{V}$. We call its *density* the number of elements that are mapped to a non-zero value in \mathcal{V} . Our focus is on *sparse* key-value maps whose density grows asymptotically with $\text{poly}(\lambda)$ (and in practice may be concretely small). We can represent a sparse key-value map as a set of pairs $\{(k, v_k)\}_{k \in K}$ where $K \subseteq \mathcal{K}$: this maps each element $k \in K$ to v_k and any other element to $0 \in \mathcal{V}$. We often use the more succinct notation $\{v_k\}_{k \in K}$ for a key-value map $\{(k, v_k)\}_{k \in K}$ over the set K (we assume that the set K is implicitly part of the description of $\{v_k\}_{k \in K}$). Hence the empty set \emptyset represents the key-value map with all elements in the universe initialized to zero; we denote the latter empty key-value map \emptyset_{kv} to be explicit. We denote by $-\text{kv}$ the key-value map that associates to each key k the value $-\text{kv}(k)$; we define a sum of key-value maps as follows: $\{v_k\}_{k \in K} + \{v'_{k'}\}_{k' \in K'}$ is defined as $(k, v_k + v'_{k'})_{k \in K \cup K'}$. A partition of a key-value map $\{v_k\}_{k \in K}$ is a pair of key-value maps $(\{v'_{k'}\}_{k' \in K'}, \{v''_{k''}\}_{k'' \in K''})$ such that (K', K'') is a partition of K .

Cryptographic Assumptions For convenience we use a different, but equivalent, formulation of the discrete logarithm assumption. Below \mathcal{G} denotes a group generator.

Assumption 1 (Generalized DLOG [BBB⁺18]) \forall PPT $\mathcal{A}, m \geq 2$:

$$\Pr \left[\begin{array}{l} \mathbb{G} \leftarrow \mathcal{G}(1^\lambda); (g_1, \dots, g_m) \leftarrow_{\mathcal{S}} \mathbb{G} \\ (a_1, \dots, a_m) \leftarrow \mathcal{A}(\mathbb{G}, g_1, \dots, g_m) \end{array} : \begin{array}{l} \exists j^* \in [m] \ a_{j^*} \neq 0 \wedge \\ \prod_{j \in [m]} g_j^{a_j} = 1_{\mathbb{G}} \end{array} \right] \leq \text{negl}(\lambda)$$

NIZKs Here we describe the basic notion of non-interactive zero-knowledge. In Section 4 we provide explicit syntax for the specific setting of NIZKs over committed key-value maps.

Definition 1. A NIZK for a relation family $\mathcal{R} = \{R_\lambda\}_{\lambda \in \mathbb{N}}$ is a tuple of algorithms $\text{NIZK} = (\text{Setup}, \text{Prove}, \text{VerProof})$ with the following syntax:

- $\text{NIZK.Setup}(1^\lambda) \rightarrow \text{crs}$ outputs a common-reference string crs ; if the argument system is transparent this can consist of uniform random elements.
- $\text{NIZK.Prove}(\text{crs}, \mathbb{x}, \mathbb{w}) \rightarrow \pi$ takes as input a string crs , an input description \mathbb{x} (in which we embed the whole public input), a witness \mathbb{w} such that $R_\lambda(\mathbb{x}, \mathbb{w})$; it returns a proof π .
- $\text{NIZK.VerProof}(\text{crs}, \mathbb{x}, \pi) \rightarrow b \in \{0, 1\}$ takes as input a string crs , a public input \mathbb{x} , a proof π ; it accepts or rejects the proof.

Whenever the relation family is obviously defined, we talk about a “NIZK for a relation R ”. We require a NIZK to be *complete*, that is, for any $\lambda \in \mathbb{N}$ and $(\mathbb{x}, \mathbb{w}) \in R_\lambda$ it holds with overwhelming probability that $\text{VerProof}(\text{crs}, \mathbb{x}, \pi)$ where $\text{crs} \leftarrow \text{Setup}(1^\lambda)$ and $\pi \leftarrow \text{Prove}(\text{crs}, \mathbb{x}, \mathbb{w})$. Other properties we require are: *knowledge-soundness* and *zero-knowledge*. Informally, the former states we can efficiently “extract” a valid witness from a proof that passes verification; the latter states that the proof leaks nothing about the witness (this is modeled through a simulator that can output a valid proof for an input in the language without knowing the witness). Notationally, we separate public and private inputs in relations and proving algorithm through a semicolon.

Knowledge-Soundness. For all $\lambda \in \mathbb{N}$ and for all (non-uniform) efficient adversaries \mathcal{A} , auxiliary input $z \in \{0, 1\}^{\text{poly}(\lambda)}$, there exists a (non-uniform) efficient extractor \mathcal{E} such that

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda); (\mathbb{x}, \pi) \leftarrow \mathcal{A}(z, \text{crs}) \\ \mathbb{w} \leftarrow \mathcal{E}(z, \text{crs}) \end{array} : \begin{array}{l} R_\lambda(\mathbb{x}, \mathbb{w}) \neq 1 \wedge \\ \text{Vfy}(\text{crs}, \mathbb{x}, \pi) = 1 \end{array} \right] \leq \text{negl}(\lambda)$$

Zero-Knowledge. There exists a PPT simulator \mathcal{S} such that for any $\lambda \in \mathbb{N}$, PPT \mathcal{A} , auxiliary input $z \in \{0, 1\}^{\text{poly}(\lambda)}$, and it holds $p_0 = p_1$ where

$$p_b := \Pr \left[\begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda); (\mathbb{x}, \mathbb{w}) \leftarrow \mathcal{A}(z, \text{crs}) \\ \pi \leftarrow X_b(\text{crs}, \mathbb{x}, \mathbb{w}) \text{ if } R_\lambda(\mathbb{x}, \mathbb{w}) \text{ o.w. } \perp \end{array} : \mathcal{A}(z, \text{crs}, \pi) = 1 \right]$$

$X_0(\text{crs}, \mathbb{x}, \mathbb{w}) := \mathcal{S}(z, \text{crs}, \mathbb{x})$ and $X_1(\text{crs}, \mathbb{x}, \mathbb{w}) := \text{Prove}(\text{crs}, \mathbb{x}, \mathbb{w})$.

On efficiency of NIZKs. The efficiency (proving/verification runtimes and proof size) of a NIZK often depends on the size of the description of a relation in *constraints* (these roughly correspond to the multiplication gates of its circuit representation). We will refer to this notion later in the text. See also [BBB⁺18].

3 Key-Value Commitments

Here we define homomorphic commitments to key-value maps where both keys and values are hidden. In the appendix (Appendix D.2) we also present an extended construction, which we use to build Multi-Type Quis-Quis.

Definition 2 (Commitment to Key-Value Maps (kvC)). *The following is a syntax for our key-value maps*

$\text{Setup}(1^\lambda) \rightarrow \text{pp}$ generates public parameters.

$\text{Com}(\text{pp}, \{v_k\}_{k \in K}; r) \rightarrow c$ commits to the key-value map with randomness r .

We keep the randomness implicit whenever it does not affect clarity and we assume it to be sampled from an additive group.¹²

Definition 3 (Hiding). A key-value map commitment is hiding if for all key-value maps $\{v'_{k'}\}_{k' \in K'}$, $\{v''_{k''}\}_{k'' \in K''}$ (even of different size), for $\text{pp} \leftarrow \text{Setup}(1^\lambda)$ the following two distributions are computationally indistinguishable:

$$\{\text{Com}(\text{pp}, \{v'_{k'}\}_{k' \in K'})\} \approx \{\text{Com}(\text{pp}, \{v''_{k''}\}_{k'' \in K''})\}$$

Definition 4 (Binding). A key-value map commitment is (computationally) binding if for any PPT adversary \mathcal{A} , it holds that

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda) \\ (c, \{v'_{k'}\}_{k' \in K'}, r') \\ \{v''_{k''}\}_{k'' \in K''}, r'' \leftarrow \mathcal{A}(\text{pp}) \end{array} : \begin{array}{l} c = \text{Com}(\text{pp}, \{v'_{k'}\}_{k' \in K'}, r') \wedge \\ c = \text{Com}(\text{pp}, \{v''_{k''}\}_{k'' \in K''}, r'') \wedge \\ \{v'_{k'}\}_{k' \in K'} \neq \{v''_{k''}\}_{k'' \in K''} \end{array} \right] \leq \text{negl}(\lambda)$$

Definition 5 (Homomorphism). We say a commitment to a key-value map is homomorphic if there exists an operation \circ such that Setup always produces pp such that for all maps $\{v_k\}_{k \in K}$, $\{v'_{k'}\}_{k' \in K'}$ and randomness r, r' it holds that

$$\text{Com}(\{v_k\}_{k \in K}; r) \circ \text{Com}(\{v'_{k'}\}_{k' \in K'}; r') = \text{Com}(\{v_{k^*}^*\}_{k^* \in K^*}; r + r')$$

$$\text{where } K^* = K \cup K' \text{ and } (k^*, v_{k^*}^*) = \begin{cases} (k^*, v_{k^*}) & \text{if } k^* \in K \setminus K' \\ (k^*, v'_{k^*}) & \text{if } k^* \in K' \setminus K \\ (k^*, v_{k^*} + v'_{k^*}) & \text{if } k^* \in K \cap K' \end{cases}$$

3.1 Construction

We recap some of the properties we are interested in obtaining in our construction: (i) support large key universe; (ii) small commitments and small parameters; (iii) homomorphic (Definition 5); (iv) support efficient non-interactive zero-knowledge proof of knowledge of opening (in particular they should run in time linear in the density of the key-value map).

In Figure 2 we propose a construction based on random-oracle with the properties above. Given a prime p we consider the universe of values $\mathcal{V} = \mathbb{Z}_p$, a group \mathbb{G} isomorphic to it and for which the GDLOG assumption holds, a hash function \mathcal{H} modeled as a random oracle and an arbitrary key universe \mathcal{K} such that $\mathcal{H} : \mathcal{K} \rightarrow \mathbb{G}$. We prove Theorem 1 in Appendix B.

Theorem 1. *If \mathcal{H} is a random oracle and under the GDLOG assumption the construction in Figure 2 is a kvC with value universe \mathbb{Z}_p .*

¹² We will use this when defining the homomorphic property of commitments.

$\text{Setup}(1^\lambda) \rightarrow \text{pp}$: samples group \mathbb{G} ; $g, h \leftarrow \mathbb{G}$; return $\text{pp} = (\mathbb{G}, g, h)$.
 $\text{Com}(\text{pp}, \{v_k\}_{k \in K}; r) \rightarrow c$: return $\prod_{k \in K} \mathcal{H}(k)^{v_k} h^r$.

Fig. 2: Our construction for kvC.

4 Arguments on Key-Value Commitments (Double Private Setting)

Here we formalize and construct zero-knowledge arguments over key-value map commitments for the setting in which there is no information on the keys available in the system and those we are using in our proof.

Circuits over Key-Value Maps To support arbitrary computation on committed key-value maps, we provide an interface which supports any arithmetic circuit of the following form. The keys and values kv as well as an additional witness ω are field elements in \mathbb{F} . The circuit consists of multiplication gates of the form $\mathbb{F} \times \mathbb{F} \rightarrow \mathbb{F}$. They have an unbounded outbound degree and any linear relations are directly expressed between outputs and inputs.

We write any circuit using multiplication gates in the domain $\mathcal{KV}^\ell \times \mathbb{F}^{n_\omega}$ as $C^\mathbb{F}(\text{kv}_1, \dots, \text{kv}_\ell, \omega)$. This circuit depends on the desired property the openings should have. Here ω is an additional private witness that may not depend on the opening to the key-value maps.

4.1 Arguments for Circuits over Committed Key-Value Maps

Here we present the overview of our argument system which works over committed key-value maps and takes an arbitrary inner circuit operating on the openings of the commitments. To be clear, we spell out its formalization explicitly but it is a special case of Definition 1.

Given an inner circuit $C^\mathbb{F}$ as described above, our high level interface for proofs has the following form:

$\text{kvNIZK.Setup}(1^\lambda) \rightarrow \text{crs}$ takes a security parameter λ and outputs a crs .

$\text{kvNIZK.Prove}(\text{crs}, C^\mathbb{F}, c_1, \dots, c_\ell, (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega) \rightarrow \pi$ takes the crs and a circuit $C^\mathbb{F}$ as well as ℓ commitments c_i and their openings (kv_i, ρ_i) and an auxiliary witness ω . It outputs a proof π

$\text{kvNIZK.VerProof}(\text{crs}, C^\mathbb{F}, c_1, \dots, c_\ell, \pi) \rightarrow b \in \{0, 1\}$ takes the crs , a circuit $C^\mathbb{F}$, and ℓ commitments c_i . It outputs a bit b depending on the validity of the proof π .

The relation we want to prove is defined by the circuit C^* in Equation (1). To clarify our notation we re-define correctness for arguments for committed key-value maps.

Definition 6 (kvNIZK Correctness). A kvNIZK is correct if, for any $\lambda \in \mathbb{N}$ with $\text{crs} \in \text{kvNIZK.Setup}(1^\lambda)$, circuit $C^{\mathbb{F}}$, key-value maps $\vec{\text{kv}} \in \mathcal{KV}^\ell$ and randomness $\vec{\rho} \in \mathbb{F}^\ell$ with $\forall i \in [\ell] : c_i = \text{Com}(\text{kv}_i, \rho_i)$ and any $\omega \in \mathbb{F}^{n_\omega}$ for which $C^*(c_1, \dots, c_\ell; (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega) = 1$ it holds that

$$\begin{aligned} \text{kvNIZK.VerProof}(\text{crs}, C^{\mathbb{F}}, c_1, \dots, c_\ell, \pi) &= 1 \text{ where} \\ \pi &\leftarrow \text{kvNIZK.Prove}(\text{crs}, C^{\mathbb{F}}, c_1, \dots, c_\ell, (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega). \end{aligned}$$

As for NIZKs, we require knowledge-soundness and zero-knowledge.

4.2 Construction with Intermediate Key-Tags

Our construction has two stages. First the prover creates some key-tags b_k and proves that they are well formed (i.e., they are obtained by hashing a key and masking with a random group element, both known to the prover). These key-tags are then used in a subsequent proof for the opening of the commitment and the actual relation. Intuitively, since the prover knows how the key-tags have been produced, the prover is able to compute openings of the input commitments under the new “base” (h, b_1, \dots, b_n) – as opposed to the original base $(g, \mathcal{H}(k_1), \dots, \mathcal{H}(k_n))$ – by appropriately computing the randomizers as a function of the values in the commitment and the randomness used for producing the key-tags. This allows to avoid proving properties of the hash function in the second part of the proof. The full construction **Z \odot -dp** is presented in Figure 3. For sake of presentation and w.l.o.g., in the construction we assume that all our key-value maps include the same keys k_1, \dots, k_n .

Theorem 2. *Under the GDLOG assumption, if $\text{NIZK}_{\text{tags}}$ and NIZK_C are secure (correct, zero-knowledge, knowledge-sound) NIZKs for their required relation families, then the construction **Z \odot -dp** is a secure kvNIZK for arbitrary circuits over the key-value map commitment in Fig. 2.*

Proof. Correctness: Correctness follows by inspection. In particular, note that when proving R_C the prover “opens” the commitments c_i under the base defined by the key-tags b_k ’s. Since the b_k ’s are generated with the same h as the original commitment and the prover knows the openings of the b_k ’s, the prover can find the right value to be used as exponent for h .

Knowledge-Soundness: To prove knowledge-soundness, assume the existence of extractors $\mathcal{E}_{\text{tags}}, \mathcal{E}_C$ for the two sub-relations. We build an extractor \mathcal{E}^* that on input a statement $(C^{\mathbb{F}}, c_1, \dots, c_\ell)$ and accepting proof $(b_1, \dots, b_n, c^*, \pi_{\text{tags}}, \pi_C)$, outputs $((\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega)$. Our \mathcal{E}^* works as follows. It first extracts through $(k'_1, \dots, k'_n, \vec{r}', s') \leftarrow \mathcal{E}_{\text{tags}}(b_1, \dots, b_n, c^*, \pi_{\text{tags}})$ and then the rest through $((\text{kv}'_1, \rho'_1), \dots, (\text{kv}'_\ell, \rho'_\ell), \vec{r}, s, \omega) \leftarrow \mathcal{E}_C(C^{\mathbb{F}}, c_1, \dots, c_\ell, b_1, \dots, b_n, c^*, \pi_C)$ such that the two relations hold for the extracted witnesses. We first argue it holds that $(k'_1, \dots, k'_\ell, s') = (k_1, \dots, k_\ell, s)$, since otherwise we can construct an adversary that breaks the binding property of the Pedersen commitment c^* . Then, we show how to extract valid openings for the input commitments c_i . Remember

The protocol is described in the random oracle model where all parties have access to $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{G}$.

- Setup(1^λ):**
1. compute $\text{crs}_{\text{tags}} \leftarrow \text{NIZK.Setup}_{\text{tags}}(1^\lambda)$, $\text{crs}_C \leftarrow \text{NIZK.Setup}(1^\lambda)$.
 2. generate the commitment key $h \xleftarrow{\$} \mathbb{G}$ (setup for kv commitments).
 3. sample random generators $\vec{g} \xleftarrow{\$} \mathbb{G}^n$ where n is the maximum number of committed keys (\leq number of active keys) in a commitment.
 4. Output $\text{crs} = (\text{crs}_{\text{tags}}, \text{crs}_C, h, \vec{g})$

Prove($\text{crs}, C^{\mathbb{F}}, c_1, \dots, c_\ell, (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega$):

1. From all key-value maps kv_i , extract the set of non-zero keys K_i and define $(k_1, \dots, k_n) := \bigcup_{i \in [\ell]} \{k \in K_i\}$. Parse $\text{kv}_i = \{(k, v_{i,k}) : k = k_1, \dots, k_n\}$ for each $i = 1, \dots, \ell$
2. Sample random values $\vec{r} \xleftarrow{\$} \mathbb{F}^n$ and create the key-tags $b_{k_i} = \mathcal{H}(k_i)h^{r_i}$. Additionally, create a vector Pedersen commitment to all keys (called pre-image commitment) $c^* = h^s \prod_{i=1}^n g_i^{k_i}$ with randomness $s \xleftarrow{\$} \mathbb{F}$.
3. Generate a proof that the key-tags are well formed by proving knowledge of the pre-images k_i i.e., generate a proof

$$\pi_{\text{tags}} \leftarrow \text{NIZK.Prove}_{\text{tags}}(\text{crs}, b_{k_1}, \dots, b_{k_n}, c^*; k_1, \dots, k_n, \vec{r}, s)$$

for the relation

$$R_{\text{tags}}(b_{k_1}, \dots, b_{k_n}, c^*; k_1, \dots, k_n, \vec{r}, s) := \\ (c^* = h^s \prod_{i=1}^n g_i^{k_i} \wedge \forall i \in [n] : b_{k_i} = \mathcal{H}(k_i)h^{r_i})$$

4. Generate a proof

$$\pi_C \leftarrow \text{NIZK.Prove}_C(\text{crs}, C^{\mathbb{F}}, c_1, \dots, c_\ell, b_{k_1}, \dots, b_{k_n}, c^*; \\ (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \vec{r}, s, \omega)$$

for the relation

$$R_C(C^{\mathbb{F}}, c_1, \dots, c_\ell, b_{k_1}, \dots, b_{k_n}, c^*; (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \vec{r}, s, \omega) := \\ c^* = h^s \prod_{i=1}^n g_i^{k_i} \wedge \forall i \in [\ell] : c_i = h^{\rho_i - \sum_{j=1}^n r_j v_{i,k_j}} \prod_{j=1}^n b_{k_j}^{v_{i,k_j}} \quad (2) \\ \wedge C^{\mathbb{F}}((\text{kv}_i)_{i \in [\ell]}, \omega) = 1$$

5. Return both proofs including the key-tags and the pre-image commitment $\pi := (b_{k_1}, \dots, b_{k_n}, c^*, \pi_{\text{tags}}, \pi_C)$

VerProof($\text{crs}, C^{\mathbb{F}}, c_1, \dots, c_\ell, \pi$):

1. Parse π as $(b_{k_1}, \dots, b_{k_n}, c^*, \pi_{\text{tags}}, \pi_C)$.
2. Verify $b_0 \leftarrow \text{NIZK.VerProof}_{\text{tags}}(\text{crs}, b_{k_1}, \dots, b_{k_n}, c^*, \pi_{\text{tags}})$
3. Verify $b_1 \leftarrow \text{NIZK.VerProof}_C(\text{crs}, C^{\mathbb{F}}, c_1, \dots, c_\ell, b_{k_1}, \dots, b_{k_n}, c^*, \pi_C)$
4. return $b_0 \wedge b_1$

Fig. 3: Z^{dp}, our construction for kvNIZK

that thanks to the knowledge-soundness of the second proof system (for relation R_C) we know that for all commitments c_i it holds $c = h^{\rho' - \sum_{j=1}^n r_j v_{k_j}} \prod_{j=1}^n b_{k_j}^{v_{k_j}}$ (we remove the index i here to improve readability).

Thanks to the knowledge soundness of the first proof system (for relation R_{tags}) we know that $b_{k_j} = \mathcal{H}(k_j) h^{r'_j}$, thus we can rewrite c as

$$c = h^{\rho' - \sum_{j=1}^n r_j v_{k_j}} \prod_{j=1}^n \mathcal{H}(k_j)^{v_{k_j}} h^{r'_j \cdot v_{k_j}} = h^{\rho' - \sum_{j=1}^n (r_j - r'_j) v_{k_j}} \prod_{j=1}^n \mathcal{H}(k_j)^{v_{k_j}}$$

Thus our extractor can output $((\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega)$ with $\rho_i = \rho'_i - \sum_{j=1}^n (r_j - r'_j) v_{k_{i,j}}$ as the witness for the overall relation. Note that the proof does not guarantee that $\vec{r} = \vec{r}'$. However this is not a problem since we are still guaranteed that the extracted witness for the overall relation is a valid one. *Zero-Knowledge*: follows from the ZK property of the underlying arguments and the hiding property of the commitments. Details can be found in Appendix B.

4.3 How to Instantiate the Subprotocols in $\mathbf{Z\textcircled{d}p}$

To instantiate the well formedness of the tags, i.e. the relation R_{tags} , we propose to use a cryptographic hash function such as MiMC [AGR⁺16], Poseidon [GKR⁺21], GMiMC [AGP⁺19], or Marvellous [AAB⁺19] which are optimized for zero-knowledge proofs. They provide hashing to a field element ($\mathcal{H}^{\mathbb{F}} : \{0, 1\}^* \rightarrow \mathbb{F}$). A subsequent circuit then proves that the field element is equivalent to the group element in the key-tags. For the example of elliptic curves, this is the relation $[b_{k_i}] = [(\mathcal{H}^{\mathbb{F}}(k_i), 0)] + r_i \cdot [h]$ where brackets enclose group elements and $(\mathcal{H}^{\mathbb{F}}(k_i), 0)$ is an explicit description of the curve point from a field element $\mathcal{H}^{\mathbb{F}}(k_i)$. The combined constraints are then proven by e.g. Bulletproofs [BBB⁺18].

The circuit for the relation R_C can be implemented through a generalization of Bulletproofs [LMR19] or compressed Σ -protocols [ACR20]. They provide an interface supporting bilinear circuits with five gates to enable arbitrary computation of which we use a subset of gates for non-bilinear circuits only. Given a circuit C constructed from the available gates, they then provide an efficient protocol with communication complexity $6\lceil \log_2(|C|) \rceil + 28$ group elements.

5 Improvements in Practice: Offline/Online Stages

The proving algorithms of our constructions (both $\mathbf{Z\textcircled{d}p}$ and $\mathbf{Z\textcircled{d}set}$) follow a two-step template. In step (a) the prover provides key-tags b_k -s and proves they are valid. This can be done independently of the commitment. In a following step (b) we compute a proof about properties of the commitment opening (and that actually depends on the commitments). Crucially, in the latter step, the prover uses the key-tags as (rerandomized) “anchors” to the keys. We observe that we can exploit the fact that (a) does not depend on the commitments to

the key-value maps (but only on the keys that they will contain) to preprocess this step.

Consider, for instance, the running example from the introduction where the commitments contain multi-currency wallets. Assume that the prover knows that *in the future* they will want to prove some properties about some of their wallets (which are expected to keep changing between now and the proving time). Moreover, while a large set K_{active} of asset types might be circulating in the system, the prover knows that they will only hold a very specific and relatively small subset K_{pre} of these keys (e.g., maybe only ETH, USD and EUR). If that is the case they can preemptively perform an “offline proving stage” that would be valid for all of the *online* proofs they will have to carry out later. Specifically, in **Z \odot -dp** the prover performs step (a) above offline as follows. On input set K_{pre} , the prover provides a set of $|K_{\text{pre}}|$ key-tags $B = \{b_k : k \in K_{\text{pre}}\}$, defined as usual as $b_k = \mathcal{H}(k)h^{r_k}$ together with a proof π_{tags} that they were constructed honestly. The output of this step is therefore $\pi_{\text{off}} = (B, \pi_{\text{tags}})$. At a later time, when input commitments c_1, \dots, c_ℓ are available, the prover uses the pre-computed set of key-tags B to produce a step (a) (the production of key-tags) for each of the commitments c_i . In order to preserve zero-knowledge, step (b) is modified to rerandomize the related b_k -s first. The rerandomization hides the mapping between online proofs, however the verifier learns that all commitments of the online phase contain the same set of keys. We can similarly adapt **Z \odot -set** by performing a proof of membership and masking (step 1 in Fig. 4) before hand.

The advantage of this approach is to use a single offline stage for many proofs. The efficiency savings of this stage (both for proving and verification time) can be significant since it involves proving/verifying hashing in zero-knowledge. For example, we conservatively estimate approximately 5k constraints using Poseidon hashing on a Ristretto curve [GKR⁺21]. Each of these hashes can be proved in the order of hundreds of milliseconds (see, e.g., Table 5 in [GKR⁺21]). For $n \approx 100$ this involves for example saving half a million constraints¹³ amounting to around half a minute of proving time. Savings for verification time are comparable.

Naturally the offline stage preprocesses an *upper bound* U on the total number of active keys, since each commitment may have openings to key-value maps of different density. This may incur a high overhead cost if U is far from the actual densities (because we still need to process U key-tags as input to the circuit). The gains from an offline stage can differ accordingly and should be weighed depending on the application.

Efficiency summary of our Constructions We summarize the (asymptotic) efficiency of our constructions in Table 1. We present it for the case with offline processing, but summing the offline and online columns corresponds to the setting without an offline stage.

¹³ For some applications this can be huge—for comparison, the ZCash circuit [HBHW21] has approximately 100k constraints.

Table 1: Efficiency of our constructions and comparison with non-homomorphic solutions (when they are applicable). Above we describe proof sizes, proving time and verification times during offline and online stage. We also describe the additional costs for proving the homomorphism in zero-knowledge with a non-homomorphic solution ($\text{kv}_{\text{add}}(n)$). All values are implicitly in big O notation and denote operations in a prime-order group unless underlined. Rows marked with \star refer to this work. The construction of $\text{Z}\otimes\text{-dp}$ is in Fig. 3. “ $\text{Z}\otimes\text{-set (Acc)}$ ” refers to the instantiation of $\text{Z}\otimes\text{-set}$ with NIZKs over accumulators in unknown-order groups we describe in Appendix C (the more general construction is in Fig. 4 in Appendix A). “Pedersen (Non-Hom.)” refers to the non-homomorphic solution based on Pedersen described in the introduction. Typical values for our parameters could be $M \approx 1000$ and $n \approx 100$.

K_{active}	K_{com}	Construction	$ \pi_{\text{off}} $	$ \pi_{\text{onl}} $	$ \pi(\text{kv}_{\text{add}}(n)) $
priv	priv	$\text{Z}\otimes\text{-dp} \star$	$\log(n(\mathcal{H} + \mathbb{G}_{\text{add}}))$	$\log(C)$	—
priv	priv	Pedersen (Non-Hom.)	—	$\log(C)$	$\log(n)$
publ	priv	$\text{Z}\otimes\text{-set (Acc)} \star$	$n + \log(n \mathbb{G}_{\text{add}})$	$\log(C)$	—

K_{active}	K_{com}	Construction	\mathcal{V}_{off}	\mathcal{V}_{onl}	$\mathcal{V}(\text{kv}_{\text{add}}(n))$
priv	priv	$\text{Z}\otimes\text{-dp} \star$	$n(\mathcal{H} + \mathbb{G}_{\text{add}})$	$ C $	—
priv	priv	Pedersen (Non-Hom.)	—	$ C $	n
publ	priv	$\text{Z}\otimes\text{-set (Acc)} \star$	$n \mathbb{G}_? + n \mathbb{G}_{\text{add}} $	$ C $	—

K_{active}	K_{com}	Construction	\mathcal{P}_{off}	\mathcal{P}_{onl}	$\mathcal{P}(\text{kv}_{\text{add}}(n))$
priv	priv	$\text{Z}\otimes\text{-dp} \star$	$n(\mathcal{H} + \mathbb{G}_{\text{add}})$	$ C $	—
priv	priv	Pedersen (Non-Hom.)	—	$ C $	n
publ	priv	$\text{Z}\otimes\text{-set (Acc)} \star$	$\frac{(M - n + n \log n) \mathbb{G}_? }{+n \mathbb{G}_{\text{add}} }$	$ C $	—

N :	Size of key universe \mathcal{K}
M :	Number of active / registered keys (K_{active})
n :	Number of keys in the opening of a key-value map commitment
	$N \gg M \geq n$
$ \mathcal{H} $:	Number of constraints for hashing to a group element
$ \mathbb{G}_{\text{add}} $:	Number of constraints for summing two group elements
$ \mathbb{G}_? $:	Cost of exponentiation in unknown-order group
$\text{kv}_{\text{add}}(n)$:	Sum operation among key-value maps of size n
C :	Circuit computed on key-value map.

6 Experimental Evaluation

Here we show the practical feasibility of our construction. Our focus is on Z_{dp} ; we compare its efficiency to that of Z_{set} in Appendix A.4.

Recall that our construction-template uses two separate steps (see also beginning of Section 5 and Fig. 3): (a) validity of key-tags and (b) actual property on opening of commitment. We evaluate our construction on a representative application setting for cryptocurrencies, that is a 64-bit range proof as a circuit proven in step (b).

Let n be the size of the opening of the commitment (also equal to the number of elements we are showing are in range). We estimate the following runtimes. For step (a): $\approx n \cdot 700\text{ms}$ for proving and $n \cdot 100\text{ms}$ for verification; for step (b) $\approx n \cdot 235\text{ms}$ for proving and $n \cdot 89\text{ms}$ verification. We stress that proving times for step (a) are fully parallelizable (as we generate n independent proofs for key-tags).

These timings refer to those for a common laptop (i7-6820HQ CPU at 2.70 GHz) and aim at estimating an efficient instantiation through the zero-knowledge scheme LMR [LMR19] as $\text{NIZK}_{\text{tags}}$ and NIZK_{C} using Ristretto Curve as an underlying group.

How we derive timings. A similar derivation for Bulletproof timings was previously used in [LRR⁺19]. For each timing we use the formula $T(n) \approx n \cdot \text{num_of_constraints_circuit}_{\text{LMR}} \cdot \text{cost_per_constraint}$. **Deriving step (a):** for proving, $\text{cost_per_constraint}$ is measured to be $\approx 8.97/64$ ms/constraint (our experimental finding) for the implementation in [dal]. For verification $1.22/64$ ms/constraint. We estimate $\text{num_of_constraints_circuit}_{\text{LMR}}(\text{tag}) \approx 5k$: for a circuit for Poseidon hash [GKR⁺21] and fixed base exponentiation (for rerandomization) of curve points¹⁴ requires $L = 2806$ multiplicative constraints. To use this in LMR [LMR19] we need to encode this as a witness vector (a very conservative upper bound is $2L$, which we approximate to $5k$). **Deriving step (b):** for proving, $\text{cost_per_constraint}$ is measured to be $\approx 232/64$ ms/constraint. For verification $88/64$ ms/constraint. We derive these estimations from BL12-381¹⁵; we know this to be a fair estimate for Ristretto [dal]. $\text{num_of_constraints_circuit}_{\text{LMR}}(\text{range}_{64}) \approx 65$ constraints.

7 Application: Multi-Type QuisQuis

QuisQuis [FMMO19] is a privacy-preserving transaction system which allows for pruning old transactions, keeping the state of each participant linear in the

¹⁴ Since there is no public circuit implementation for Ristretto operations for this, we use arkworks [ark] BL12-381 implementation for this estimate. We expect this to be an upper bound on Ristretto points given their smaller field size— $\frac{255}{381}$ x smaller, more precisely. We measure this number using the implementation in [ark]

¹⁵ We use a different implementation [zen] on BL12-381 points as the implementation in [dal] is not compatible with BL12-381

number of users. This is a major advantage over other privacy-preserving transaction systems, which require a state size linear in the number of transactions. A QuisQuis transaction is a redistribution of tokens among a set of accounts. An account belongs to an owner and stores their tokens. Instead of consuming accounts and creating new ones, QuisQuis updates the accounts. These update operations need to change the balance of a peer without knowing their total balance. This is achieved with homomorphic commitments.

In contrast to the original QuisQuis where an account stored a scalar value representing an amount, we generalize accounts to store tokens of different types in a key-value map. Each key corresponds to a type, also known as currency, and the value specifies how many tokens of the specific type are held by the account. An account acct then belongs to a secret key sk and holds a balance kv , represented as a key-value map. To transfer tokens from one account to another, a transaction includes both accounts as active inputs, denoted by the set \mathcal{P} . The transaction subtracts tokens from one account with a secret index s and adds them to the other one. The output of the transaction are two updated accounts belonging to the same users as the inputs; the input accounts are discarded. To achieve anonymity, the input consists of \mathcal{P} together with a potentially large anonymity set of accounts \mathcal{A} which keep their balance unchanged but provide set anonymity for the active accounts.

As the central building block of our multi-type QuisQuis system, we present an updatable account based on our key-value commitments.

7.1 Multi-Type QuisQuis: Syntax

The original QuisQuis transaction protocol consists of the three algorithms (Setup, Trans, Verify). Their multi-type equivalent is as follows:

Setup($1^\lambda, \vec{\text{kv}}$) \rightarrow **state**: takes the security parameter λ and a vector of key-value balances $\vec{\text{kv}}$ and outputs an initial state **state**. One part of the state is a set of unspent accounts where each key-value balance has an account.

Trans($\text{sk}, \mathcal{P}, \mathcal{A}, \delta\vec{\text{kv}}$) \rightarrow **tx**: takes a secret key sk which corresponds to one account in the set of active accounts \mathcal{P} and an anonymity set \mathcal{A} with a vector of update key-value maps $\delta\vec{\text{kv}}$ to redistribute wealth. **Trans** outputs a transaction **tx**.

Verify(**state**, **tx**) \rightarrow \perp /**state'**: takes a **state** and a transaction **tx** and outputs a new **state'** or fails with \perp .

To support dynamic registration of new types, we require an additional algorithm **Register**, which is defined as:

Register(acct, k, v_k) \rightarrow **tx** takes an account acct and a new type k with amount v_k and outputs a transaction **tx**.

A registration transaction is accepted by **Verify** if the type k has not been registered before. We define the correctness of a transaction system more formally. Let for all $\lambda \in \mathbb{N}$ and $\vec{\text{kv}}$ with $R_{\text{rng}}^{\text{kv}}(\text{kv}_i) = 1$ be $\text{state} \leftarrow \text{Setup}(1^\lambda, \vec{\text{kv}})$. For all

accounts in \mathcal{P}, \mathcal{A} with index sets $\mathcal{P}^* := \{i \in [\text{Sort}(\mathcal{P} \cup \mathcal{A})] : \text{acct}_i \in \mathcal{P}\}$ and \mathcal{A}^* accordingly in a canonically ordered form with Sort . All accounts in $\mathcal{P} \cup \mathcal{A}$ are part of the UTXO set in state , all $\delta\vec{k}\mathbf{v}$ with $R_{\text{rng}}^{\text{kv}}(-\text{kv}_s) = 1$ and $R_{\text{rng}}^{\text{kv}}(\text{kv}_i) = 1$ for $i \in \mathcal{P}^*\{s\}$ and $\text{kv}_i = \emptyset_{kv}$ for $i \in \mathcal{A}^*$ and sk corresponding to an account $\text{acct}_s \in \mathcal{P}$ with enough tokens such that after the transaction there is no negative type $R_{\text{rng}}^{\text{kv}}(\text{kv}_s + \delta\text{kv}_s) = 1$, it holds that $\text{Verify}(\text{state}, \text{Trans}(\text{sk}, \mathcal{P}, \mathcal{A}, \delta\vec{k}\mathbf{v})) = \text{state}'$ where $\text{state}' \neq \perp$ and contains an updated UTXO set with all inputs $\mathcal{P} \cup \mathcal{A}$ removed and the transaction outputs added.

Multi-Type QuisQuis: Security The security of a QuisQuis-like transaction system consists of two main properties. The first property we need to achieve is *anonymity*. A transaction system is anonymous if an adversary cannot successfully distinguish two transactions. The transactions are created according to malicious instructions after the adversary has interacted with an oracle signing transactions on behalf of uncompromised participants.

The second property is *theft prevention*. This entails that (i) the adversary cannot steal tokens from uncompromised accounts; (ii) the adversary cannot create tokens out of thin air. Slightly more formally, we model these properties as follows. While interacting with the aforementioned signing oracle the balance of honest accounts (not controlled by the adversary) must not decrease. Additionally, the total amount of tokens must not increase from transaction to transaction. Notice, however, the number of tokens may increase as the result of mining or a token registration—the latter counts as a “genesis” transaction.

Our variant of QuisQuis with multiple token types shares many of the same properties as the non-type aware system. We refer the reader to the original QuisQuis paper [FMMO19] for details.

7.2 Construction

We construct the multi token QuisQuis scheme following the original QuisQuis but with two main adaptations: each account holds tokens in multiple types and making sure a transaction guarantees that the amounts of tokens are balanced for each of the token types.

The details of updatable accounts for key-value maps are presented in Appendix D. In a nutshell they have the same algorithms as the original QuisQuis construction but allow for multiple kind of tokens.

Setup The setup algorithm generates a list of updatable accounts, one for each initial balance key-value map.

Trans Our transaction structure follows that in QuisQuis where a “transaction” denotes a redistribution of wealth among all accounts involved ($\mathcal{P} \cup \mathcal{A}$). The transaction takes a vector of key-value maps, one for each account. The account is then updated according to the key-value map. Key-value maps that contain only valid positive values ($R_{\text{rng}}^{\text{kv}}(\delta\text{kv}_i) = 1$) are used to deposit tokens to receiving accounts. In order to ensure that the total number of tokens is preserved, we require that one key-value map holds negative values. This is to satisfy that the

sum of all key-value maps is zero, or $\sum_{i=1}^{|\delta\vec{kv}|} \delta kv_i = \emptyset_{kv}$. For the account with the negative key-value map (indexed by s in the canonically ordered set $\mathcal{P} \cup \mathcal{A}$), the transaction signature ensures that the owner of the account acct_s authorizes the spending by proving knowledge of the matching secret key sk . The algorithm $\text{Trans}((s, \text{sk}_s, \text{kv}_s), \mathcal{P}, \mathcal{A}, \vec{kv})$ performs the following steps:

1. Parse all input accounts $\mathcal{P} \cup \mathcal{A} = \{\text{acct}_1, \dots, \text{acct}_n\}$ and check the spending account is valid by $\text{VerifyAcct}(\text{acct}_s, (\text{sk}_s, \text{kv}_s)) = 1$. The transaction needs to be balanced: $\sum_{i=1}^{|\delta\vec{kv}|} \delta kv_i = \emptyset_{kv}$ and all key-value maps other than the spending account must be non-negative $\forall i \neq s : R_{\text{rng}}^{\text{kv}}(\delta kv_i) = 1$. To support large anonymity sets \mathcal{A} , we choose to disclose the upper bound on active accounts by showing that $\delta kv_i = \emptyset_{kv}$ for $i \in \mathcal{A}^*$ instead of a range proof. The spending account must be negative $R_{\text{rng}}^{\text{kv}}(-\delta kv_s) = 1$ and the resulting account must be valid $R_{\text{rng}}^{\text{kv}}(\text{kv}_s + \delta kv_s) = 1$, to prevent overspending.
2. Let $\text{outputs} = (\text{acct}_1^T, \dots, \text{acct}_n^T)$ be a canonical order of the accounts generated by $\text{UpdateAcct}(\mathcal{P} \cup \mathcal{A}, \delta\vec{kv}; r)$.
3. Let $\psi : [n] \rightarrow [n]$ be the permutation that maps the canonically ordered inputs to the canonically ordered outputs, i.e. input i has the same secret key as output $\psi(i)$.
4. Create a zero knowledge proof π showing that the transaction is well formed, i.e. that it satisfies the following relation:

$$R_{\text{tx-wf}} : \begin{cases} \mathbb{x} = (\text{inputs}, \text{outputs}), \mathbb{w} = (\text{sk}, \text{kv}_s, \delta\vec{kv}, r = (r_1, r_2), \psi) \text{ s.t.} \\ \text{VerifyUpdateAcct}(\text{acct}_i^S, \text{acct}_{\psi(i)}^T, \delta kv_i) = 1 \forall i \in \mathcal{P}^* \\ \wedge R_{\text{rng}}^{\text{kv}}(\delta kv_i) = 1 \forall i \in \mathcal{P}^* / \{s\} \wedge R_{\text{rng}}^{\text{kv}}(-\delta kv_s) = 1 \\ \wedge \delta kv_i = \emptyset_{kv} \forall i \in \mathcal{A}^* \wedge \sum_{i=1}^n \delta kv_i = \emptyset_{kv} \\ \wedge \text{VerifyAcct}(\text{acct}_{\psi(s)}^T, \text{sk}, \text{kv}_s + \delta kv_s) = 1. \end{cases}$$

The relation ensures that the permuted output account is correctly updated by the transferred balance δkv_i for all active accounts. It then ensures that the updated key-value maps are valid, i.e. there is one spending account at index s and no value is taken from other accounts. The balances of the accounts in the anonymity set must not change. To ensure that the spender has enough tokens, the proof checks that the updated spender account has no negative balance. The transaction consists of the inputs, outputs and the proof π .

Verify A transaction is valid in respect to a **state** if all accounts in **inputs** have not been used in another transaction and the proof π is valid.

Security Analysis Our key-value commitments provide the same hiding and binding properties as the commitments to single scalars used in QuisQuis. The construction is a parallel version of the single type case and thereby the theft security holds also for all keys in parallel. Regarding anonymity, we achieve the same properties as QuisQuis, if we define an upper bound of the number of types involved in a transaction. For transactions with few different types, we achieve this through padding. With a constant size transaction proof, our

new transactions have the same indistinguishability as the original QuisQuis transactions.

Acknowledgements. Research supported by: the Concordium Blockchain Research Center, Aarhus University, Denmark; the Carlsberg Foundation under the Semper Ardens Research Project CF18-112 (BCM); the European Research Council (ERC) under the European Unions’s Horizon 2020 research and innovation programme under grant agreement No 803096 (SPEC). This work was partly produced while Matteo Campanelli and Felix Engelmann were affiliated with Aarhus University.

References

- AAB⁺19. Abdelrahman Aly, Tomer Ashur, Eli Ben-Sasson, Siemen Dhooghe, and Alan Szepieniec. Design of symmetric-key primitives for advanced cryptographic protocols. Cryptology ePrint Archive, Report 2019/426, 2019. <https://eprint.iacr.org/2019/426>.
- ACR20. Thomas Attema, Ronald Cramer, and Matthieu Rabaud. Compressed sigma-protocols for bilinear circuits and applications to logarithmic-sized transparent threshold signature schemes. Cryptology ePrint Archive, Report 2020/1447, 2020. <https://eprint.iacr.org/2020/1447>.
- AGP⁺19. Martin R. Albrecht, Lorenzo Grassi, Léo Perrin, Sebastian Ramacher, Christian Rechberger, Dragos Rotaru, Arnab Roy, and Markus Schofnegger. Feistel structures for MPC, and more. In Kazue Sako, Steve Schneider, and Peter Y. A. Ryan, editors, *ESORICS 2019, Part II*, volume 11736 of *LNCS*, pages 151–171. Springer, Heidelberg, September 2019.
- AGR⁺16. Martin R. Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. MiMC: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part I*, volume 10031 of *LNCS*, pages 191–219. Springer, Heidelberg, December 2016.
- AK20. Archita Agarwal and Seny Kamara. Encrypted key-value stores. In Karthikeyan Bhargavan, Elisabeth Oswald, and Manoj Prabhakaran, editors, *INDOCRYPT 2020*, volume 12578 of *LNCS*, pages 62–85. Springer, Heidelberg, December 2020.
- AR20. Shashank Agrawal and Srinivasan Raghuraman. KVc: Key-Value Commitments for blockchains and beyond. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 839–869. Springer, Heidelberg, December 2020.
- ark. ark-works. <http://arkworks.rs>.
- 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.
- BBF19. Dan Boneh, Benedikt Bünz, and Ben Fisch. Batching techniques for accumulators with applications to IOPs and stateless blockchains. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 561–586. Springer, Heidelberg, August 2019.

- BCFK19. Daniel Benarroch, Matteo Campanelli, Dario Fiore, and Dimitris Kolonelos. Zero-knowledge proofs for set membership: Efficient, succinct, modular. Cryptology ePrint Archive, Report 2019/1255, 2019. <https://eprint.iacr.org/2019/1255>.
- BCG⁺14. Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014.
- BG18. Jonathan Bootle and Jens Groth. Efficient batch zero-knowledge arguments for low degree polynomials. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part II*, volume 10770 of *LNCS*, pages 561–588. Springer, Heidelberg, March 2018.
- BH11. Johannes Buchmann and Safuat Hamdy. A survey on iq cryptography. In *Public-Key Cryptography and Computational Number Theory*, pages 1–15, 2011.
- CF13. Dario Catalano and Dario Fiore. Vector commitments and their applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 55–72. Springer, Heidelberg, February / March 2013.
- CFG⁺20. Matteo Campanelli, Dario Fiore, Nicola Greco, Dimitris Kolonelos, and Luca Nizzardo. Incrementally aggregatable vector commitments and applications to verifiable decentralized storage. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 3–35. Springer, Heidelberg, December 2020.
- CHA21. Matteo Campanelli and Mathias Hall-Andersen. Veksel: Simple, efficient, anonymous payments with large anonymity sets from well-studied assumptions. *IACR Cryptol. ePrint Arch.*, 2021:327, 2021.
- dal. Dalek bulletproofs implementation. <https://github.com/zkcrypto/bulletproofs.git>.
- EMP⁺21. Felix Engelmann, Lukas Müller, Andreas Peter, Frank Kargl, and Christoph Bösch. SwapCT: Swap confidential transactions for privacy-preserving multi-token exchanges. *PoPETs*, 2021(4):270–290, October 2021.
- FMMO19. Prastudy Fauzi, Sarah Meiklejohn, Rebekah Mercer, and Claudio Orlandi. Quisquis: A new design for anonymous cryptocurrencies. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 649–678. Springer, Heidelberg, December 2019.
- GKR⁺21. Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. Poseidon: A new hash function for zero-knowledge proof systems. In *30th USENIX Security Symposium*, 2021.
- HBHW21. Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash protocol specification, version 2021.1.15, 2021.
- KZG10. Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Heidelberg, December 2010.
- 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.

- LRR⁺19. Russell W. F. Lai, Viktoria Ronge, Tim Ruffing, Dominique Schröder, Sri Aravinda Krishnan Thyagarajan, and Jiafan Wang. Omniring: Scaling private payments without trusted setup. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 31–48. ACM Press, November 2019.
- LY10. Benoît Libert and Moti Yung. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 499–517. Springer, Heidelberg, February 2010.
- NM⁺16. Shen Noether, Adam Mackenzie, et al. Ring confidential transactions. *Ledger*, 1:1–18, 2016.
- PBF⁺19. Andrew Poelstra, Adam Back, Mark Friedenbach, Gregory Maxwell, and Pieter Wuille. Confidential assets. In Aviv Zohar, Ittay Eyal, Vanessa Teague, Jeremy Clark, Andrea Bracciali, Federico Pintore, and Massimiliano Sala, editors, *FC 2018 Workshops*, volume 10958 of *LNCS*, pages 43–63. Springer, Heidelberg, March 2019.
- SAGL18. Srinath Setty, Sebastian Angel, Trinabh Gupta, and Jonathan Lee. Proving the correct execution of concurrent services in zero-knowledge. In *13th USENIX Symposium on Operating Systems Design and Implementation OSDI 18*, pages 339–356, 2018.
- TXN20. Alin Tomescu, Yu Xia, and Zachary Newman. Authenticated dictionaries with cross-incremental proof (dis)aggregation. Cryptology ePrint Archive, Report 2020/1239, 2020. <https://eprint.iacr.org/2020/1239>.
- YYD⁺19. Zheng Yi, Howard Ye, Patrick Dai, Sun Tongcheng, and Vladislav Gelfer. Confidential assets on MimbleWimble. Cryptology ePrint Archive, Report 2019/1435, 2019. <https://eprint.iacr.org/2019/1435>.
- zen. Zengo-x bulletproofs implementation. <https://github.com/ZenGo-X/bulletproofs>.

Appendix

A Arguments on Key-Value Commitments with Public Active Keys

In this section we consider the case where the set of active keys is known by both the prover and the verifier. We can then add a feature by which in a setup phase it is possible to register a set of keys which are going to be used for later proofs. This can improve the amortized efficiency when several proofs share the same set of active keys. We dub our resulting construction **Z_o-set**.

A.1 NIZKs with Key-Registration

We define NIZK for key-value maps which allow a key-registration phase. To this goal an algorithm `RegisterActiveSet` to the NIZK interface *deterministically* specializes a CRS to a set of active keys. Similar ideas have been used in previous work e.g., [BCFK19, CFG⁺20, CHA21] where a set-dependent specialized CRS is used.

Definition 7. *A NIZK for key-value maps with registration stage $\text{kvNIZK}^{\text{reg}}$ consists of the following algorithms:*

$\text{kvNIZK}^{\text{reg}}.\text{Setup}(1^\lambda) \rightarrow \text{crs}$ on input λ it outputs a crs crs .

$\text{kvNIZK}^{\text{reg}}.\text{RegisterActiveSet}(\text{crs}, K_{\text{active}}) \rightarrow \text{crs}^{\text{reg}}$ is a deterministic algorithm that takes as input a crs, an active set K_{active} and outputs a specialized CRS crs^{reg} ¹⁶.

$\text{kvNIZK}^{\text{reg}}.\text{Prove}(\text{crs}^{\text{reg}}, K_{\text{active}}, C^{\mathbb{F}}, c_1, \dots, c_\ell, (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega) \rightarrow \pi$ takes crs^{reg} , a circuit $C^{\mathbb{F}}$, ℓ commitments c_i , their openings (kv_i, ρ_i) and auxiliary witness ω . It outputs a proof π . Each of the kv_i should use keys that are subset of K_{active} .

$\text{kvNIZK}^{\text{reg}}.\text{VerProof}(\text{crs}^{\text{reg}}, C^{\mathbb{F}}, c_1, \dots, c_\ell, \pi) \rightarrow b \in \{0, 1\}$ takes the crs^{reg} , a circuit $C^{\mathbb{F}}$, and ℓ commitments c_i .

We require properties as in Definition 1: correctness (completeness), knowledge soundness, zero-knowledge. Definition 7 is *almost* a special case of the notion for kvNIZK s, but not quite. We could express it as a NIZK including K_{active} as input to verification and require whatever `RegisterActiveSet` does as part of the relation. However, as it's standard in NIZKs, that would require the verifier to run linearly in $|K_{\text{active}}|$. By taking it as input only a (potentially

¹⁶ It is w.l.o.g. that we require the registration algorithm to register in batch the whole set of active keys. It would be possible to define a version where we incrementally register a new key to a previously preprocessed CRS. Several constructions, e.g. algebraic accumulators, would efficiently support this syntax.

smaller) processing of it, crs^{reg} , the verifier can even run sublinearly in the set of active keys (which is the case in our construction). This motivates the special interface we introduce in Definition 7. Otherwise, the required properties are straightforward variants of those for NIZKs.

The relation we want to prove is again that defined by the circuit in Equation (1). For sake of clarity, we spell out correctness.

Definition 8 (Correctness). *We say that a NIZK for key-value maps with registration stage $\text{kvNIZK}^{\text{reg}}$ over key-value map commitment scheme C_{kv} is correct if for any $\lambda \in \mathbb{N}$ with $\text{crs} \in \text{kvNIZK.Setup}(1^\lambda)$, any $\text{crs} \in \text{C}_{\text{kv}}.\text{Setup}(1^\lambda)$, any set of keys K_{active} , any circuit $C^{\mathbb{F}}$, any key-value maps $\vec{\text{kv}} \in \mathcal{KV}^\ell$ with keys all in K_{active} and randomness $\vec{\rho} \in \mathbb{F}^\ell$ with $\forall i \in [\ell] : c_i = \text{C}_{\text{kv}}.\text{Com}(\text{crs}, \text{kv}_i, \rho_i)$ and any $\omega \in \mathbb{F}^{n_\omega}$ for which*

$$C(c_1, \dots, c_\ell; (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega) = 1$$

it holds that $\text{kvNIZK.VerProof}(\text{crs}^{\text{reg}}, C^{\mathbb{F}}, c_1, \dots, c_\ell, \pi) = 1$ where we compute $\text{crs}^{\text{reg}} \leftarrow \text{RegisterActiveSet}(\text{crs}, K_{\text{active}})$, $\pi \leftarrow \text{kvNIZK.Prove}(\text{crs}^{\text{reg}}, C^{\mathbb{F}}, c_1, \dots, c_\ell, (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), \omega)$.

A.2 Accumulator-based Construction

We require the two following building blocks (described below): *accumulators* and *NIZKs for membership over accumulators and rereandomization*.

The intuition behind our construction is to register the *already hashed* keys in an accumulator (without masking). This process can be performed publicly since producing the accumulator is deterministic. In order to output a proof, we produce a set of key-tags (masked version of the keys committed in the accumulator) and that these are appropriate rerandomized versions of the group elements in the accumulator.

A.2.1 Building Block: Accumulators

Definition 9 (Accumulator scheme). *An accumulator scheme AccScm over universe $\mathcal{U}_\lambda(\text{AccScm})$ (where λ is a security parameter) consists of a quadruple of PPT algorithms $\text{AccScm} = (\text{Setup}, \text{Accum}, \text{PrvMem}, \text{VfyMem})$ with the following syntax:*

$\text{Setup}(1^\lambda) \rightarrow \text{pp}_{\text{acc}}$ *generates public parameters pp_{acc} .*

$\text{Accum}(\text{pp}_{\text{acc}}, S) \rightarrow \text{acc}$ *deterministically computes accumulator acc for set $S \subseteq \mathcal{U}_\lambda(\text{AccScm})$.*

$\text{PrvMem}(\text{pp}_{\text{acc}}, S, S') \rightarrow \pi_{\text{acc}}$ *computes witness π_{acc} that shows $S' \subseteq S$.*

$\text{VfyMem}(\text{pp}_{\text{acc}}, \text{acc}, S', \pi_{\text{acc}}) \rightarrow b \in \{0, 1\}$ *verifies through witness whether subset S' is in the set accumulated in acc . We do not require parameter S to be in $\mathcal{U}_\lambda(\text{AccScm})$ from the syntax.*

An accumulator scheme should satisfy correctness—the accumulator works as expected—and soundness—no efficient adversary can pick set S and find a witness that checks on $\text{AccScm.Accum}(\text{pp}_{\text{acc}}, S)$ and $S' \not\subseteq S$ ¹⁷.

We assume the universe of the accumulator to be a group of prime order.

A.2.2 Building Block: NIZKs for membership and rerandomization

We will use a NIZK for the following relation, which intuitively states that a set of group elements (fresh key-tags) are a rerandomization of elements accumulated in the set.

$$\begin{aligned} R_{\text{mem-rand}}((\text{pp}_{\text{acc}}, h \in \mathbb{G}, \text{acc}); (B = (b_1, \dots, b_n) \in \mathbb{G}^n, \\ S = \{s_1, \dots, s_n\} \in \mathbb{G}^n, \pi_{\text{acc}}, (r_1, \dots, r_n) \in \mathbb{Z}_p^n)) \\ \text{s.t. } b_i = s_i \cdot h^{r_i} \text{ for } i \in [n] \text{ and } \text{VfyMem}(\text{pp}_{\text{acc}}, \text{acc}, S) = 1 \end{aligned}$$

We refer to the NIZK for the relation above as $\text{NIZK}_{\text{mem-rand}}$.

A.2.3 Construction We describe our construction in Fig. 4. In our description we assume the building blocks above and a general-purpose zero-knowledge scheme kvNIZK like the one assumed in Section 4.1.

The security of our construction can be argued analogously to the security proof of our construction in Fig. 3 given the security of $\text{NIZK}_{\text{mem-rand}}$. The latter NIZK proves membership of the keys in the accumulator without revealing anything about them and simulation is straightforward from this observation. By the binding property of the accumulator and the knowledge-soundness of $\text{NIZK}_{\text{mem-rand}}$ we are guaranteed that keys at proving time belong to the accumulated set.

A.3 Instantiation from Algebraic Accumulators in Groups of Unknown Order

It is possible to instantiate our construction above from accumulators in groups of unknown order [BBF19]. While we describe our instantiation in detail later in the appendix (see Appendix C), here we provide a high-level view of the proof system $\text{NIZK}_{\text{mem-rand}}$. Our techniques are adaptations of those in [BCFK19] and their concrete improvements in [CHA21]. The proof system consists of two parts. One part proves in zero-knowledge the bulk of accumulator membership. This consists of a low communication-complexity protocol where the proof partly consists of elements in $\mathbb{G}_?$ (by which we denote the unknown-order group). The other component is a Bulletproof where we both perform some sanity check on the values in the first proof (for example, some hidden values claimed by the prover should be in a certain range) and we perform a rerandomization

¹⁷ These definitions are standard; see [BBF19] for a formal treatment.

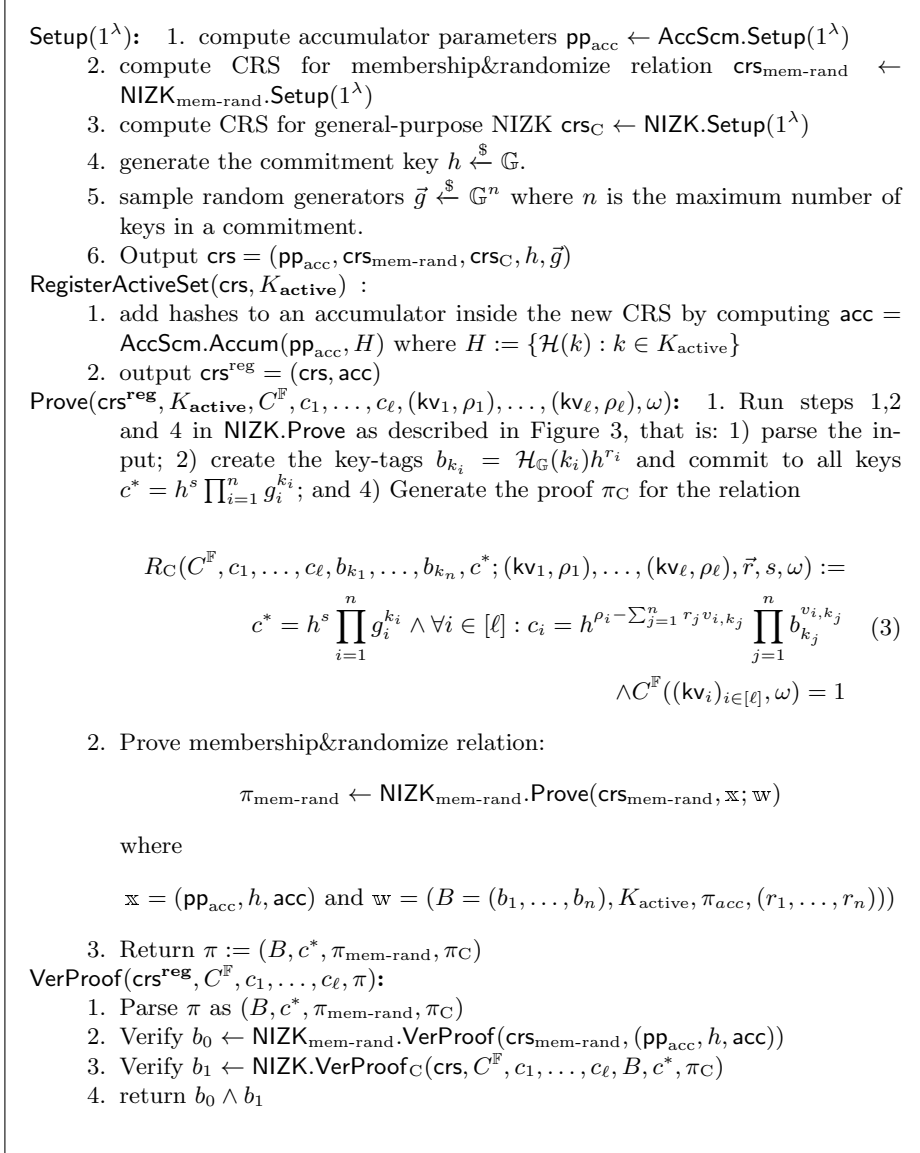


Fig. 4: **Z-set**, our accumulator-based construction for $\text{kvNIZK}^{\text{reg}}$ (i.e., with public previously registered active keys)

of the group elements. This construction can be instantiated in a completely transparent manner using class groups[BH11]¹⁸.

A breakdown of the communication complexity of the proof is as follows. Let n be the maximum number of keys in any of the key-value commitments input to the proof algorithm. The first component provides roughly δn elements in $\mathbb{G}_?$ and in the prime order group, for a small constant δ . The second component consists of roughly $\log(n|\mathbb{G}_{\text{add}}||C|)$ where $|C|$ is the size of the circuit representing the general relation on key-value maps and $|\mathbb{G}_{\text{add}}|$ is the cost of representing an addition in the prime-order group. The latter can be in the order of 1K constraints using some of the instantiations in [CHA21]. Concretely, for $n \approx 30$ we estimate the size of the proof to be lower than 180 KB plus the cost of a Bulletproof on the circuit C (of logarithmic size).

A.4 Discussion on Practical Benefits of \mathbf{Z}_{set}

Intuitively, using an accumulator with a public set of active keys is beneficial whenever the “cost” of a proof of membership & randomization (step 1 in Fig. 4) is lower than that of a proof for hashing & randomization (steps 1–3 in \mathbf{Z}_{dp} , Fig. 3). Below we discuss when this can be the case; see also Table 1. *The bottomline is: \mathbf{Z}_{set} can improve verification time trading proof size and proving time.*

For proving time Fig. 4 is a worse choice because proving membership introduces a linear dependency in $|K_{\text{active}}|$, not present in our construction from Fig. 3. This is not just a matter of asymptotics: by our estimates, this would also be the case for common concrete parameters for accumulators in unknown-order groups. Similar observations hold for proof size.

Where we may see substantial gains from applying \mathbf{Z}_{set} is in verification time. Ignoring the cost of the circuit C and only considering the additive overhead we achieve the following verification times. Without registration (\mathbf{Z}_{dp}), we would have n times: one hashing subcircuit—each of approximately 5k constraints—plus rerandomization (additional 1k constraints). This gives a total of $6nk$ constraints. With registration (\mathbf{Z}_{set}) we have: n group exponentiations in a group of unknown-order plus (proved with a special-purpose “accumulator-membership” NIZK; see also Appendix C), plus n additional rerandomizations and additional sanity check constraints (for a total around 1.5k constraints). To assess which solution is best one needs to compare the concrete costs of the n exponentiations in groups of unknown order vs the $n \times 4.5k$ constraints of the hash function. We know the first to be cheaper and in particular to be an order of magnitude so (each exponentiation requiring tens of milliseconds, vs hundreds of milliseconds for the case of verifying a single hash function). This time difference in verification can be substantial. Our time estimates refer to RSA-groups and are derived from [GKR⁺21, CHA21, CFG⁺20].

¹⁸ The more efficient instantiation based on RSA groups would require the existence of a trusted RSA modulus. Nonetheless, all the other components of the construction remain transparent.

B Additional Proofs

Proof of Theorem 1 As for standard Pedersen, the construction is perfectly hiding because of the masking factor h^r . To show why the construction is binding we argue as follows. Assume an adversary is able to provide two distinct openings for the same commitment. Each of these openings consists of both keys and values. The case where the keys are the same for both openings can be argued as in standard (vector) Pedersen. Consider the case where the two openings differ in the set of keys. We can write $c = \prod_{k \in K} \mathcal{H}(k)^{v_k} h^r = \prod_{k' \in K'} \mathcal{H}(k')^{v'_{k'}} h^{r'}$ which reduces to

$$\prod_{k \in K} \mathcal{H}(k)^{v_k} \prod_{k' \in K'} \mathcal{H}(k')^{-v'_{k'}} h^{r-r'} = 1_{\mathbb{G}} .$$

Since the random oracle evaluations are distributed as random group generators we can use this adversary to break Assumption 1.

Zero-Knowledge Proof in Theorem 2 For zero-knowledge, we build our simulator \mathcal{S}^* using the simulators for the two subprotocols, namely $\mathcal{S}_{\text{tags}}, \mathcal{S}_{\text{R}}$. On input a valid statement $(\text{crs}, C^{\mathbb{F}}, c_1, \dots, c_\ell)$, \mathcal{S}^* outputs $\pi = (b_1, \dots, b_n, c^*, \pi_{\text{tags}}, \pi_{\text{C}})$ where: b_1, \dots, b_n, c^* are random group elements, and we build simulated proofs $\pi_{\text{tags}} \leftarrow \mathcal{S}_{\text{tags}}(\text{crs}_{\text{tags}}, b_1, \dots, b_n, c^*)$ and $\pi_{\text{C}} \leftarrow \mathcal{S}_{\text{C}}(\text{crs}_{\text{C}}, C^{\mathbb{F}}, c_1, \dots, c_\ell, b_1, \dots, b_n, c^*)$. We can prove indistinguishability of the simulated and the real transcript through the following hybrids:

- *Hybrid 0.* This is the view in the real protocol.
- *Hybrid 1.* This is the same as the previous hybrid, but we replace π_{tags} with a simulated proof. Indistinguishability follows from the ZK property of the subprotocol for R_{tags} .
- *Hybrid 2.* This is the same as the previous hybrid, but we replace π_{C} with a simulated proof. Indistinguishability follows from the ZK property of the subprotocol for R_{C} .
- *Hybrid 3.* This is the same as the previous hybrid, but we replace b_1, \dots, b_n, c^* with random group elements. Indistinguishability follows since these group elements are distributed identically in both hybrid 2 and 3 due to the choice of the randomizers s, \vec{r} in the real protocol/Hybrid 2. Note moreover that when replacing these commitments with random group elements we do not change the validity of the statements to be proven by the two simulators $\mathcal{S}_{\text{tags}}$ and \mathcal{S}_{C} . That is, (b_1, \dots, b_n, c^*) and $(C^{\mathbb{F}}, c_1, \dots, c_\ell, b_1, \dots, b_n, c^*)$ are true statements for the two sub-relations also when the group elements are randomly sampled. This is due to the fact that both relations are trivial: since Pedersen commitments are perfectly hiding, for any $\text{kv}_1, \dots, \text{kv}_\ell$ there exist s, \vec{r} such that the relation holds.

Since the distribution in Hybrid 3 is identical to a simulated execution, this concludes the proof of zero-knowledge.

C Zero-Knowledge over Algebraic Accumulators

The numbers in our Table 1 for the registration construction refer to instantiations described in this section.

Construction of accumulators for group elements Below we describe an instantiation of an accumulator for group elements pt on an elliptic curve. The original construction is that of accumulators in unknown-order groups from [BBF19], while the explicit description for elliptic curve points is from [CHA21]. The adaptations we describe in this section are important to obtain concretely efficient instantiations.

We define a permissible set \mathcal{P} of commitments which allows to avoid collisions. It is parametrized by an integer μ and consists of elliptic curve points where the x -coordinate is a μ -bit prime and the y -coordinate is the “canonically chosen” square root so that the point can be described by its x -coordinate alone: $\mathcal{P} := \{(x, y) \in (\mathbb{F}, \mathbb{F}) \mid x \in [2^{\mu-1}, 2^\mu) \wedge y \equiv 0 \pmod{2}\}$

One can use $\mu = 251$ bits for concrete instantiations in the Ristretto curve (which we can use with [LMR19]).

Below we restrict elements to prime numbers. This is standard and necessary for the soundness of the accumulator scheme. We observe that, when the accumulator is computed publicly (which is the case in our applications), it is not necessary to prove in zero-knowledge that its elements are primes. This can be done at registration time ensuring that, when adding a key k , its hash $\mathcal{H}(k)$ is prime. We can naturally add a padding to the key before registration time and, through rejection sampling, ensure we obtain a prime.

Setup (1^λ) \rightarrow pp	VfyMem (pp , acc , pt , π_{acc})
$(\mathbb{G}_?, g_?) \leftarrow \mathcal{G}_?(1^\lambda)$ return pp = $(\mathbb{G}_?, g_?)$	Parse pt as pt := (x, y) Accept iff $\pi_{acc}^x = \text{acc}$
Add (pp , pt , acc) \rightarrow acc'	PrvMem (pp , S , pt) \rightarrow π_{acc}
Parse pt as pt := (x, y) if pt $\notin \mathcal{P} \vee x$ not a prime then return \perp ; else return acc ^{x}	$S' := \{x' : (x', y') \in S \setminus \{\text{pt}\}\}$ $\text{prd} \leftarrow \prod_{x' \in S'} x'$ return $g_?$ ^{prd}

Fig. 5: Accumulator Instantiation for AccScm.

For efficiency, we observe that **PrvMem** can be executed in time growing with $(M - n) + n \log n$ for a subset of size n and accumulated set of size M (the quasilinear factor derives from **RootFactor** algorithm in [BBF19]).

Instantiating $\text{NIZK}_{\text{mem-rand}}$ In Appendix A, we need to efficiently prove in zero-knowledge that a key-tag (roughly, the hash of a key) is inside an accumulator ($\text{NIZK}_{\text{mem-rand}}$). For our instantiations in unknown-order groups, we can adopt a variant of the solution in [CHA21] (see their Section 5.2 and instantiations in 5.4.1 and 5.4.2). Our relation $R_{\text{mem-rand}}$ has only one difference from their “one-out-of-many” relation in Section 5.2: we need to prove the same relation for multiple key-tags (which correspond to what they call an “outer commitment”). In other words, we need a “many-out-of-many” relation.

The subsequent change to their instantiation is almost straightforward. Their one-out-of-many relation is composed of two parts: one described in 5.4.1 (roughly consisting of a Bulletproof execution on a gadget of approximately 1.5K constraints) the other in 5.4.2 (a NIZK whose proof consists of a constant number of unknown-order and prime-order group elements). To turn them into a many-out-of-many variant we do the following. The constraint system for 5.4.1 can simply be adapted to take as input n inputs, instead of just one. The batching properties of Bulletproof allow the proof size to be affected only minimally. For the second component, the one from 5.4.2, we can simply provide n runs of it. This part of the argument system grows linearly in n .

D Updatable Accounts

D.1 Preliminaries: Updatable Public Keys [FMMO19]

We need updatable public keys. Intuitively, such a key can be publicly updated and the update is indistinguishable from a newly created key. An updatable public key scheme consists of the following algorithms:

$\text{Setup}(1^\lambda) \rightarrow \text{pp}$ outputs the public parameters pp .
 $\text{Gen}(\text{pp}) \rightarrow (\text{sk}, \text{pk})$ generates a new key pair
 $\text{Update}(\text{pk}) \rightarrow (\text{pk}')$ updates the key pk to pk'
 $\text{VerifyKP}(\text{pk}, \text{sk}) \rightarrow b \in \{0, 1\}$ verifies that a key pair is consistent.
 $\text{VerifyUpdate}(\text{pk}, \text{pk}', r) \rightarrow b \in \{0, 1\}$ verifies that the public key pk was correctly updated to pk' with some randomness r .

The correctness follows directly from QuisQuis such that given a key pair, Update generates again a valid key pair. The update process is verifiable by VerifyUpdate and the updated key must keep the original secret key. The security requires (a) that an updated key is indistinguishable from a new, randomly generated one and (b) that the update process cannot change the secret key if the update is verified.

D.2 A secret-key version with updatable keys

We now describe an extended commitment to key-value map, which is the one we will use in our constructions for multi-type QuisQuis. In this extension we

allow a party with a secret key to open the commitment even without knowing the randomness used to commit to it. This will be useful to allow to verify that an account has a certain balance. In order to do this we let the setup return a pair of public–secret keys $\text{Setup}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$ and we use the notation $\text{VerifyCommSk}(\text{pk}, c, \text{sk}, \{v_k\}_{k \in K})$. We also require the commitment to be binding also with respect to this function, that is no efficient adversary can output $(\text{pk}, c, \text{sk}, \{v_k\}_{k \in K}, \text{sk}', \{v'_{k'}\}_{k' \in K'})$ with $\{v_k\}_{k \in K} \neq \{v'_{k'}\}_{k' \in K'}$ but such that $\text{VerifyCommSk}(\text{pk}, c, \text{sk}, \{v_k\}_{k \in K}) = \text{VerifyCommSk}(\text{pk}, c, \text{sk}', \{v'_{k'}\}_{k' \in K'}) = 1$. We provide a construction reminiscent of ElGamal commitments in Figure 6.

$\text{Setup}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$: samples a group \mathbb{G} and a generator g , $z \leftarrow_{\$} [|\mathbb{G}|]$ and return $(\text{pk} = (\mathbb{G}, g, h := g^z), \text{sk} = z)$.
 $\text{Com}(\text{pk}, \{v_k\}_{k \in K}; r) \rightarrow c$: return $(g^r, \prod_{k \in K} \mathcal{H}(k)^{v_k} h^r)$.
 $\text{VerifyCommSk}(\text{pk}, c = (u, v), \text{sk} = z, \{v_k\}_{k \in K})$: return 1 iff $\prod_{k \in K} \mathcal{H}(k)^{v_k} = v/u^z$.

Fig. 6: Secret-key variant of construction for kvC.

D.3 Formalization

With a key-value commitment scheme and updatable keys, we build updatable accounts which belong to a secret key and hold a key-value map of tokens. The important feature of an updatable account is, that anyone can rerandomize the account and create a proof of correct rerandomization. An updated account is not linkable to the previous version and indistinguishable from a newly created account. In combination with a shuffling proof, a set of accounts is permuted without changing ownership and without knowledge of openings. We provide formal details in the full version.

Definition 10. *Updatable Accounts consist of the following algorithms:*

$\text{GenAcct}(\text{pp}, \text{kv}) \rightarrow (\text{acct}, \text{sk})$ takes an amount map $\{v_k\}_{k \in K}$ and outputs an account acct and a secret key sk
 $\text{VerifyAcct}(\text{acct}, (\text{sk}, \text{kv})) \rightarrow b \in \{0, 1\}^*$ checks the account acct for consistency with the private key sk and the key-value map $\{v_k\}_{k \in K}$ such that $R_{\text{rng}}^{\text{kv}}(\text{kv}) = 1$.

These accounts express the same updatability as the original QuisQuis accounts, such that they remain valid for updates to the public key and homomorphic operations on the commitment. This leads to the following additional operations:

UpdateAcct($\{(\text{acct}_i, \delta\text{kv}_i)\}_{i=1}^n; r_1, r_2$) \rightarrow $\{\text{acct}'_i\}_{i=1}^n$ takes as input set of accounts $\text{acct}_i = (\text{pk}_i, c_i)$ and key-value maps δkv_i , representing the change of amount in different types, such that $R_{\text{rng}}^{\text{kv}}(\delta\text{kv}_i) = 1$ or $R_{\text{rng}}^{\text{kv}}(-\delta\text{kv}_i) = 1$; it outputs a new set of accounts $\{\text{acct}'_i\}_{i=1}^n$ with $\text{acct}'_i = (\text{Update}(\text{pk}_i; r_1)), c \circ \text{Com}(\delta\text{kv}_i, \text{pk}_i; r_2)$

VerifyUpdateAcct($\{(\text{acct}_i, \text{acct}'_i, \delta\text{kv}_i)\}_{i=1}^n; r_1, r_2$) \rightarrow $\{0, 1\}$ outputs 1 if $\{\text{acct}'_i\}_{i=1}^n = \text{UpdateAcct}(\{(\text{acct}_i, \delta\text{kv}_i)\}_{i=1}^n, \text{kv}; r_1, r_2)$ and $R_{\text{rng}}^{\text{kv}}(\delta\text{kv}_i) = 1$ or $R_{\text{rng}}^{\text{kv}}(-\delta\text{kv}_i) = 1$, 0 otherwise.

E Multi-Type QuisQuis Transaction Relation

With $(c_1, \dots, c_\ell) := (\delta\text{acct}_2, \dots, \delta\text{acct}_m, \text{acct}'_1^T, \delta\text{acct}_1)$, we construct a circuit which checks that the first update account δacct_1 is all negative and the remaining ones are all positive. Last, it checks that for each type, the update values sum to zero.

$$C^{\mathbb{F}}((k_1, (v_{1,k_1}, \dots, v_{\ell,k_1})), \dots, (k_n, (v_{1,k_n}, \dots, v_{\ell,k_n})), \omega) := \forall_{k \in K} \begin{cases} R_{\text{rng}}(-v_{\ell,k}) = 1 \\ \wedge \forall_{i \in [\ell-1]} R_{\text{rng}}(v_{i,k}) = 1 \\ \wedge \sum_{i \in [\ell-2] \cup \{\ell\}} v_{i,k} = 0 \end{cases}$$

Our kvNIZK construction builds upon Pedersen style key-value commitments, but for the multi-type QuisQuis setting we need public key based commitments as shown in Appendix D.2. Therefore we change the relation R_C in Equation (3) to prove that the input to the circuit is still equivalent to the commitment openings.

$$R_C(C^{\mathbb{F}}, ((g_1, h_1), c_1), \dots, ((g_\ell, h_\ell), c_\ell), b_{k_1}, \dots, b_{k_n}, c^*; (\text{kv}_1, \rho_1), \dots, (\text{kv}_\ell, \rho_\ell), r, \omega) := c^* = h^r \prod_{i=1}^n g_i^{k_i} \wedge \forall i \in [\ell] : c_i = (g_i^{\rho_i}, h_i^{\rho_i - \sum_{j=1}^n r_j v_{i,k_j}} \prod_{j=1}^n b_{k_j}^{v_{i,k_j}}) \wedge C^{\mathbb{F}}((k_1, (v_{1,k_1}, \dots, v_{\ell,k_1})), \dots, (k_n, (v_{1,k_n}, \dots, v_{\ell,k_n})), \omega) = 1$$