Linear-map Vector Commitments and their Practical Applications

Matteo Campanelli¹, Anca Nitulescu¹, Carla Ràfols², Alexandros Zacharakis², and Arantxa Zapico^{2*}.

 ¹ Protocol Labs {matteo, anca}@protocol.ai
 ² Universitat Pompeu Fabra {carla.rafols, alexandros.zacharakis, arantxa.zapico}@upf.edu

Abstract. Vector commitments (VC) are a cryptographic primitive that allows one to commit to a vector and then "open" some of its positions efficiently. Vector commitments are increasingly recognized as a central tool to scale highly decentralized networks of large size and whose content is dynamic. In this work, we examine the demands on the properties that a vector commitment should satisfy in the light of the emerging plethora of practical applications and propose new constructions that improve the state-of-the-art in several dimensions and offer new tradeoffs. We also propose a unifying framework that captures several constructions and we show how to generically achieve some properties from more basic ones. On the practical side, we focus on building efficient schemes that do not require a new trusted setup (we can reuse existing ceremonies for other pairing-based schemes, such as "powers of tau" run by real-world systems such as Zcash or Filecoin).

1 Introduction

Vector commitment schemes [18, 7] (or VC) allow a party to commit to a vector **v** through a short digest and then open some of its elements guaranteeing *position binding*³ (one should not be able to open a commitment at position *i* to two different values $v_i \neq v'_i$). For this primitive to be interesting the proof of opening—or just "opening"—should be of size sublinear in *m*, the size of the committed vector. A vector commitment with *subvector opening* also supports a short opening for arbitrary *subsets* of positions *I* (rather than individual ones only). More specifically this opening should be of size independent, not only of *m*, but of |I|. We denote commitment schemes with such property as SVC [16](also called VC with batch opening in [3]).

Functional Vector Commitments, first introduced by Libert, Ramanna and Yung in [17], capture the ability to compute commitments to vectors and later

^{*} Arantxa Zapico has been funded by a Protocol Labs PhD Fellowship PL-RGP1-2021-062. Alexandros Zacharakis has been partially funded by Protocol Labs Research Grant PL-RGP1-2021-048

³ For the applications considered in this work, hiding properties are not necessary. In particular, our commitments are deterministic.

perform openings of linear functions (inner-products) $f : \mathbb{F}^m \to \mathbb{F}^n$ of these vectors, for some field \mathbb{F} .

Both vector commitments with subvector openings and functional commitments for inner-products can be captured as vector commitments with openings for a more general class of function families, linear maps. Lai and Malavolta [16] were the first to introduce Linear Map Commitments (LMC). In such a scheme, the prover is able to open the commitment to some vector \mathbf{v} to the output of multiple linear functions or, equivalently, to the output of one linear-map $f: \mathbb{F}^m \to \mathbb{F}^n$, by producing a single short proof. In this work, we revisit Lai and Malavolta [16] LMC notion and augment it to a full-featured vector commitment generic definition that recovers all previously-defined schemes and more. We call our primitive Linear Map Vector Commitment and use LVC for short⁴.

1.1 Motivation for Better Vector Commitments

Vector commitments are very useful to scale highly decentralized networks of large size and whose content is dynamic [8, 3, 5, 13] (such dynamic content can be the state of a blockchain, amount stored on a wallet, the value of a file in a decentralized storage network, etc.). Beyond the basic requirement that openings should be efficient, in this work we also discuss how to achieve some additional properties of LVC. We discuss some of the most prominent applications of LVC to motivate and justify the importance of these properties in practice.

Verifiable Databases. One of the applications that can be significantly improved by Vector Commitments is Verifiable Databases (VDB). In this setting, a client outsources the storage of a database to a server while keeping the ability to access and change some of its records, i.e. query functions of the data and update some of the data and ensure the server does not tamper with the data. Solutions using (binding) commitment schemes provide security but not efficiency in such a setting. A popular instantiation that achieves both of them is a Merkle tree [19], but this is not expressive enough to allow for functional openings.

For a VC scheme to be the ideal solution for VDB application, we require it to additionally support efficient updates and expressive openings. For example, an LVC scheme that allows the client to update records of the database in sublinear time and to verify linear-map queries at almost the same cost as simple position openings is a great improvement over current solutions.

Stateless Cryptocurrency. A recent application that motivated more efficient constructions of VC schemes is *stateless cryptocurrency*, i.e. a payment system based on a distributed ledger where neither validators of transactions nor system users need to store the full ledger state. The ideal vector commitment scheme that provides the best trade-off between storage, bandwidth, and computation in this setting should have all of the following properties: it must have a small

2

 $^{^4}$ We prefer LVC rather than LMC to emphasize the Vector Commitment aspect of our notion.

commitment size, short proofs, efficient computation for openings and it should allow for proof updates and for aggregation to minimise communication in the transactions and *maintainability* for the proofs, that allows updating all prestored proofs in sublinear time.

Proof of Space. Proof of Space (PoS) is a protocol that allows miners (storage providers) to convince the network that they are dedicating physical storage over time in an efficient way. In a nutshell, a miner commits to a file (data) that uses a specified amount of disk space and then the miner proves that it continues to store the data by answering to recurring audits that consist of random spotchecks. A PoS construction based on vector commitments, as described in [11], requires short opening proofs for subvectors to be stored in a blockchain, cross-commitments aggregation techniques and the possibility to implement space-time tradeoffs to reduce the proving time for the miner (ideally sublinear in the size of the vector).

"Caching" Optimizations. In some applications, e.g. when performing HTTP queries, clients use the so-called $prefetching^5$ and receive from a server not only the values of interest but other related values that could potentially be queried in the near future (e.g., values in a neighboring range of the queried values). Vector commitments with efficient proofs for special ("caching") subset openings allow to add verifiability to such queries in a way that does not affect the speed of the server since the proving procedure for a bigger subset is close or the same as for individual positions.

1.2 Desired Properties and Limitations

At the very least a basic LVC should be *efficient* (small proof size and low opening/verifying computational needs). Obviously, the same design goals as with other cryptographic protocols apply, i.e. ideally one would like to prove security under as standard assumptions as possible.

Reusable setup refers to the common reference string that many pairing-based schemes use as public parameters. Ideally, one would like to have a transparent setup (consisting of uniformly distributed elements) that does not rely on any trusted parameter generation. It is common to sacrifice this goal for efficiency and settle for a trusted setup (producing a SRS, or structured reference string) that can be generated in a ceremony. But such ceremonies are complicated to implement⁶, so it is interesting to design LVC that do not have special SRS distributions and can reuse existing setups for other primitives.

Expressivity refers to the opening possibilities. One would like VC to be as expressive as possible, meaning that it should be possible to open to functions

 $^{^5\ {\}tt https://developer.mozilla.org/en-US/docs/Web/HTTP/Link_prefetching_FAQ}$

⁶ This remains true even if many setups are updatable [14] and they can be generated and updated non-interactively in a secure way as long as one party is honest. There might be issues if not enough parties participate in generating the SRS or updates are not properly validated.

of the vector as general as possible (subvector openings, linear or arbitrary functions).

Proof Aggregation captures the ability to "pack" two or more proofs together obtaining a new proof for their combined claims (e.g. $\mathbf{f}(\mathbf{v}) = y$ and $\mathbf{f}'(\mathbf{v}) = y'$). This should be done without knowledge of the opening of the vector and aggregation cost should be sublinear in the vector length. Importantly, the resulting proof should not significantly grow each time we perform an aggregation. One-hop aggregation allows only to aggregate fresh proofs. Ideally, one would also want to aggregate already aggregated proofs.

Updatability allows to efficiently update opening proofs: if C is a commitment to v and a position needs to be updated resulting in a new commitment C', an updatable VC must provide a method to update an opening π_f for a function f that is valid for C into a new opening for the same function that is valid for the new commitment C'. The new opening should be computed by only knowing the portion of the vector that is supposed to change and in time faster than recomputing the opening from scratch.

Maintainability aims at amortizing the proving costs in systems where committed values have a long life span and evolve over time. This is achieved by means of dedicated memory to reduce the computation time needed to open proofs. Concretely, the property requires that (1) one can efficiently store some values to reduce the cost of computing any individual openings (2) after updating a single position of the committed vector, it should be possible to update all proofs in time sublinear in the size of the vector (less than computing a single proof from scratch in some cases).

Homomorphic properties apply to commitments as well as to proofs. An LVC has homomorphic commitments if it is possible to meaningfully combine commitments without knowing their openings: that is, from commitments C_1 and C_2 to \mathbf{v}_1 and \mathbf{v}_2 , any party must be able to compute a commitment to $\alpha \mathbf{v}_1 + \beta \mathbf{v}_2$ for any $\alpha, \beta \in \mathbb{F}$. The scheme has homomorphic openings if it is possible to derive a proof that $f(\mathbf{v}_1 + \mathbf{v}_2) = \mathbf{y}_1 + \mathbf{y}_2$ from proofs for the claims $f(\mathbf{v}_1) = \mathbf{y}_1$ and $f(\mathbf{v}_2) = \mathbf{y}_2$. Finally, a vector commitment scheme has homomorphic proofs when it is possible to combine proofs of statements for different functions but same vector. As we will see, this property is interesting for its implications.

1.3 Our Contributions

4

Theoretical Advances. On the theoretical frontier, we unify previous definitions and augment them with additional properties. The basic notion we use is Linear Map Vector Commitments (LVC) and is inspired by the work of Lai and Malavolta [16]. We then define additional properties on top of this definition and explore their relations. Specifically, we augment this notion with *updatability* and aggregation properties, including a novel notion *-unbounded aggregation*capturing the ability to aggregate already aggregated proofs but relaxing incremental aggregation [5] in the sense that the verifier is allowed to do work linear in the number of aggregation hops (i.e. aggregation is "history" dependent), also, disaggregation is not possible. We show that having additional homomorphic properties is highly desirable, by arguing that any LVC that satisfies them: (1) can be augmented with unbounded aggregation as well as updatability; (2) can support general linear map openings (i.e. for any $f : \mathbb{F}^m \to \mathbb{F}^n$) as long as it supports inner product openings (i.e. for $f' : \mathbb{F}^m \to \mathbb{F}$). This allows us to focus on efficient constructions for inner products with homomorphic properties.

VC Constructions. First, we present two pairing-based LVC constructions for inner products based on the properties of monomial and Lagrange polynomial basis and prove that they satisfy all the relevant homomorphic properties to obtain unbounded aggregation and support general linear maps. In terms of expressivity, these constructions generalize previous work [22, 23] by supporting linear functions instead of only position or subvector openings. VC for this class of functions are core components of important primitives such as arguments of knowledge for Inner Product (IP) relations or aggregation arguments [9].

Second, we present two novel maintainable constructions by exploiting the tensor structure of multivariate and univariate polynomials. These constructions allow a stronger, more flexible form of maintainability: they support an arbitrary memory/time trade-off for openings, meaning that one can decide how much memory it wants to use to reduce the opening time.

The multivariate case is a generalization of Hyperproofs [22] in several dimensions. Roughly speaking maintanability is achieved in Hyperproofs by constructing a binary tree of proofs where at the leaves there are the values of individual positions. We present a single construction that can be instantiated in several ways (recovering Hyperproofs as a special case) with these features: (i) the tree can be of any arity, so proofs are shorter⁷; (ii) the leaves can be commitments for any LVC and not only individual openings, to achieve a fully flexible trade-off. As a result of (ii), the scheme is more expressive (as it can support openings to linear functions/subvector openins at leaf level if the underlying commitment supports it).

The univariate construction presents a similar generalization of previous work by [24] but it has the additional feature that the setup is independent of the trade-off, and can be decided by the prover on the fly.

Practical Improvements. As in some applications like Proof of Space, the subset of opened positions is not very meaningful and its distribution is expected to be known in advance, we study how to improve verification efficiency for certain special subsets I openings in our inner-product constructions. For some structured sets I, we achieve a verifier that performs half of the work it does for arbitrary sets J of the same size in the Lagrange construction, and only a constant number of group operations in the one that uses the monomial basis.

Second, we mitigate the challenges of deploying these constructions due to their need of a trusted setup. With the exception of the multivariate variant

⁷ If one uses the Inner Pairing Product argument of Bünz et al. [4] on top of PST commitments as suggested in Hyperproofs the difference in proof size is not so relevant, but IPP will be much cheaper to run.

VC Scheme	Setup	Aggregation	Updates	Assumption	Functional	Special Sets
					Opening	Opening (size n)
PoS aggSVC [5]	Trusted	Incremental Same-Com	hint	RSA	SVC	O(n)
Pointproofs [13]	Trusted	One-hop Cross-Com	key	pairings	×	O(n)
Stateless aggSVC [23]	Trusted	One-hop Same-Com	key	pairings	SVC	O(1)
Our Lagrange LVC	Reusable	UnboundedCross-Com	key	AGM	LVC	O(1)
Our Monomial LVC	Reusable	${\color{blue}Unbounded} Cross-Com$	keyless	AGM	LVC	O(1)

Table 1: Comparison of our LVCs with other aggregatable VC schemes (aggSVC) designed for Stateless Cryptocurrencies and Proof of Space applications. All schemes have O(1)-sized proofs that verify in O(1) time and can update commitments in O(1) time.

VC Scheme	Setup	Homomorphic .	Aggregation	$ \pi $	Prove	OpenAll	UpdateAll
Merkle Trees	Transparent	×	SNARK	$\log m$	O(k)	O(m)	$O(k + \log m')$
Hyperproofs [22]	Trusted	\checkmark	IPP	$\log m$	O(k)	$O(m\log m')$	$O(\log m')$
Our Multivariate LVC	Trusted	\checkmark	IPP	$\log_\ell m'$	O(k)	$O(m\log m')$	$O(\log m')$
Our Univariate LVC	Reusable	\checkmark	IPP	$\log m'$	O(k)	$O(m\log m')$	$O(\log m')$

Table 2: Comparison of our schemes with other maintainable VC. We consider vectors of dimension $m = k \cdot m'$ where m' is the amount of memory dedicated for storing proofs. All schemes are aggregatable using generic techniques, SNARKs or Inner Pairing Products [4]. All times/sizes omit the dependence on the security parameter λ . We omit constant additive terms from proof sizes. In the multivariate construction, ℓ refers to a constant parameter.

of the maintainable construction, all our constructions can reuse trusted setups such as "powers of tau" that were run for pairing-based SNARK schemes used in real-world applications.⁸, as opposed to for example [13], in which a certain middle power of τ needs to be missing in the SRS.

In the full version of this work we demonstrate the practical benefits of our special subset construction by providing an implementation and comparisons with current solutions.

1.4 Related Work

Vector commitments were fully formalized in [7] and two first constructions were proposed under standard, constant-size, assumptions: CDH in bilinear groups and RSA respectively. Many follow-up works built on these constructions to obtain better efficiency and more properties such as subvector openings, functional openings, aggregation, updates and variants of these. A number of constructions [5, 3] use the properties of hidden order groups to achieve constructions with attractive features such as constant size parameters or incremental aggregation but are concretely less efficient than pairing-based constructions.

Merkle trees are quite efficient and only need a transparent setup. They also offer natural time-memory tradeoffs due to their tree structure. Nevertheless, VC schemes based on bilinear groups are more expressive in terms of openings, have homomorphic properties, allow for efficient updates for the proofs and aggregation mechanisms, so they are becoming an interesting alternative.

6

⁸ E.g., the one used by ZCash. https://z.cash or and Filecoin [10]

Expressivity. VC were generalized by Libert et al. [17], who formalize the notion of functional commitments (FC). They construct vector commitments with openings to linear-forms of the vector based on the Diffie-Hellman exponent assumption over pairing groups. Later, Lai and Malavolta [16] introduce subvector openings and show applications to building succinct-arguments of knowledge (similar applications were shown by [3]) in the bilinear group setting. They also generalize the notion of SVCs to allow the prover to reveal arbitrary linear maps computed over the committed vector. Previously, only Functional VC for single-output linear functions were proposed which did not account for provers that want to reveal multiple locations or function outputs of the committed vector in a concise way.

Updatability. Vector commitments that allow for updates are useful in applications such as stateless cryptocurrencies. A weak variant of updatability requires the algorithms that update the commitment and the opening to take as input an opening for the position in which the vector update occurs called *hints*. Recent RSA-based constructions are hint-updatable [3, 5]. Compared to hint updates, key-updates only need fixed update keys corresponding to the updated positions. Schemes based on bilinear groups require such fixed keys, and no extra information about the change made in the vector in order to update.

Aggregation. Vector Commitments with an additional aggregation property are very appealing for blockchain applications for their even shorter proofs of opening. Campanelli et al. [5] showed two constructions of incrementally aggregatable SVCs, that have constant-size parameters and work over groups of unknown order. Unfortunately, the practical efficiency of these constructions is still not sufficient for their deployment in real-world systems.

Gorbunov et al. [13] show how to extend the VC scheme of [18] to allow for cross-commitment aggregation. Like our constructions, they assume the Algebraic Group Model (AGM) [12] in bilinear groups and a random oracle. Their final SVC requires public parameters whose size is linear in the size of the committed vector, while cross-commitment aggregation allow for splitting up a long vector into shorter ones and simply aggregate the proofs. However, this approach allows only for one-hop aggregation, meaning that already aggregated proofs cannot be reused in further aggregations by external nodes.

Tomescu et al.[23] showed how to realize an *updatable* SVC with one-hop aggregation from bilinear groups. Their scheme has linear-sized public parameters, and it supports commitment updates, proof updates from a static linear-sized update key tied only to the updated position, in contrast with the dynamic update *hints* required by related works.

Maintainability. Apart from Merkle tree based Vector Commitments which are known to be maintainable, Srinivasan et. al. [22] show that the multilinear PST polynomial commitment [20] can be turned to a maintainable VC construction. Pre-computing all (single-position) opening proofs is done in quasilinear time

(contrary to the trivial quadratic time) and updating all proofs after a (single position) vector update needs only logarithmic time. Contrary to Merkle tree based approaches, the scheme has *homomorphic* properties. Furthermore, due to its algebraic structure, it supports one-hop aggregation through generic means, namely, Inner Pairing Product Arguments [4], albeit with a concretely expensive proving computation. Tomescu et al. [24] add the same attribute to KZG polynomial commitment schemes, resulting in an univariate construction with the same properties.

2 Preliminaries

8

Bilinear Groups. A bilinear group is given by a description $\mathbf{gk} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e)$ with additive notation such that p is prime, so $\mathbb{F} = \mathbb{F}_p$ is a field. $\mathbb{G}_1, \mathbb{G}_2$ are cyclic (additive) groups of prime order p. We use the notation $[a]_1, [b]_2, [c]_t$ for elements in $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T respectively. $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ is a bilinear asymmetric map (pairing), which means that $\forall a, b \in \mathbb{Z}_p, e([a]_1, [b]_2) \coloneqq [ab]_t$. We implicitly have that $[1]_t \coloneqq e([1]_1, [1]_2)$ generates \mathbb{G}_T . We use $[a]_{1,2}$ to refer to 2 group elements $[a]_1 \in \mathbb{G}_1, [a]_2 \in \mathbb{G}_2$. In our constructions, we denote by $\mathcal{G}(p)$ the algorithm that, given as input the prime value p, outputs a description $\mathbf{gk} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e)$.

Algebraic Group Model (AGM). The algebraic group model [12] lies between the standard model and the stronger generic group model. In AGM, we consider only so-called algebraic adversaries. Such adversaries have direct access to group elements and, in particular, can use their bit representation, like in the standard model. However, these adversaries are assumed to output new group elements only by applying the group operation to received group elements (like in the generic group model). This requirement is formalized as follows: Suppose an adversary \mathcal{A} is given some group elements $[x_1]_1 \dots [x_m]_1 \in \mathbb{G}_1$. Then, for every new group element $[z]_1 \in \mathbb{G}_1$ that the adversary outputs, it must also output $z_1 \dots z_m \in \mathbb{F}$ such that $[z]_1 = \sum_{i=1}^m [z_i x_i]_1$.

3 Definitions: Linear-map Vector Commitments

In the following, we define what we call Linear-map Vector Commitments (LVC) schemes. Notably, this definition has been introduced by Lai and Malavolta in [16] (except that there the name is *Linear Map Commitments*) to capture further functionalities of vector commitments, whose definition before only account for proofs of *position openings* (Vector Commitments) or more generally *subvector openings* (Sub-vector commitments). We introduce the definition and security properties of LVC. Importantly, we do not consider the hiding property as for our applications all vectors are public.

Linear-map Vector Commitment A linear-map vector commitment scheme for function families $\mathcal{F} \subset \{f : \mathcal{M}^m \to \mathcal{M}^n\}$ is a tuple of PPT algorithms (LVC.KeyGen, LVC.Commit, LVC.Open, LVC, Vf) that work as follows:

- LVC.KeyGen $(1^{\lambda}, \mathcal{F}) \rightarrow (\text{prk}, \text{vrk})$: The setup algorithm takes the security parameter λ , a family of functions \mathcal{F} implicitly defining the message space \mathcal{M} , and the maximum vector length $m = \text{poly}(\lambda)$, and outputs a pair of keys (prk, vrk).
- LVC.Commit(prk, \mathbf{v}) \rightarrow (C, aux): On input the proving key prk, and a vector $\mathbf{v} = (v_1, v_2, \ldots, v_m) \in \mathcal{M}^m$, returns a commitment C and auxiliary information aux. This algorithm is *deterministic*.
- LVC.Open(prk, aux, f, \mathbf{y}) $\rightarrow \pi_f$: Takes as input prk, the auxiliary information aux, a function $f \in \mathcal{F}$, and a claimed result $\mathbf{y} \in \mathcal{M}^n$. It outputs a proof π_f that $f(\mathbf{v}) = \mathbf{y}$.
- LVC.Vf(vrk, C, f, \mathbf{y}, π_f) $\rightarrow 0/1$: Takes as input the verification key vrk, C, function $f, \mathbf{y} \in \mathcal{M}^n$, and proof π_f . It accepts or rejects.
- A LVC scheme must satisfy the following properties:

Definition 1 (LVC correctness). An LVC scheme is perfectly correct if for all $\lambda \in \mathbb{N}$, for any family of functions $\mathcal{F} \subset \{f : \mathcal{M}^m \to \mathcal{M}^n\}$ and any $\mathbf{v} \in \mathcal{M}^m$,

$$\Pr\left[\mathsf{LVC.Vf}(\mathsf{vrk},\mathsf{C},f,\mathbf{y},\pi_f) = 1 \left| \begin{array}{l} (\mathsf{prk},\mathsf{vrk}) \leftarrow \mathsf{LVC.KeyGen}(1^{\lambda},\mathcal{F}) \\ (C,\mathsf{aux}) \leftarrow \mathsf{LVC.Commit}(\mathsf{prk},\mathbf{v}) \\ \pi_f \leftarrow \mathsf{LVC.Open}(\mathsf{prk},\mathsf{aux},f,\mathbf{y}) \end{array} \right] = 1.$$

Definition 2 (LVC (strong) function binding.). A linear map commitment LVC satisfies strong function binding if, for any PPT adversary \mathcal{A} , for all $\lambda \in \mathbb{N}$, for all integers $K \in \text{poly}(\lambda)$, and for any family of functions \mathcal{F} , the following probability is negligible in λ :

$$\Pr \begin{bmatrix} \forall k \in [K] : \\ \mathsf{LVC.Vf}(\mathsf{vrk}, \mathsf{C}, f_k, \mathbf{y}_k, \pi_{f_k}) = 1 \\ \land \ \not\exists \ \mathbf{v} \in \mathcal{M}^m \ s. \ t. \\ \forall k \in [K] : \ f_k(\mathbf{v}) = \mathbf{y}_k \end{bmatrix} \begin{pmatrix} (\mathsf{prk}, \mathsf{vrk}) \leftarrow \mathsf{LVC.KeyGen}(1^\lambda, \mathcal{F}) \\ \left(\mathsf{C}, \{f_k, \mathbf{y}_k, \pi_{f_k}\}_{k \in [K]}\right) \leftarrow \mathcal{A}(\mathsf{prk}, \mathsf{vrk}) \end{bmatrix}$$

The definition above can be relaxed to hold only for *honestly-generated* commitments C, raising to the *weak function binding* notion. In the weak definition, the adversary \mathcal{A} returns a vector \mathbf{v} while the commitment C is computed via LVC.Commit. In this work, constructions are proven strong function binding.

3.1 Homomorphic Properties for LVC

Homomorphic Commitments. Linear-map vector commitment schemes that satisfy homomorphic commitments allow to combine commitments of two vectors into a single one of their sum (or any linear combination). Namely, for all λ , and (vrk, prk) \leftarrow LVC.KeyGen(1^{λ}, \mathcal{F}), if (C₁, aux₁) \leftarrow LVC.Commit(prk, v₁) and (C₂, aux₂) \leftarrow LVC.Commit(prk, v₂), then $\tilde{C} = (\alpha C_1 + \beta C_2)$ is a valid commitment to $\tilde{v} = (\alpha v_1 + \beta v_2)$ for any $\alpha, \beta \in \mathcal{M}$.

In this work, we are particularly interested in LVC that also have *homomorphic proofs* for different functions applied to a committed vector and *homomorphic openings* for the same function applied to different initial vectors.

Homomorphic Proofs. An LVC scheme has homomorphic proofs if it allows recombine two proofs π_1 , π_2 corresponding to linear maps f_1 , f_2 into a new proof $\tilde{\pi}$ that opens to a linear combination of f_1 and f_2 applied to the same committed vector. Namely, for all λ , $\mathcal{F} \subset \{f : \mathcal{M}^m \to \mathcal{M}^n\}$ and all vectors $\mathbf{v} \in \mathcal{M}^m$, and $(\mathsf{vrk}, \mathsf{prk}) \leftarrow \mathsf{LVC}.\mathsf{KeyGen}(1^\lambda, \mathcal{F}), (\mathsf{C}, \mathsf{aux}) \leftarrow \mathsf{LVC}.\mathsf{Commit}(\mathsf{prk}, \mathbf{v}),$ if $\pi_1 \leftarrow \mathsf{LVC}.\mathsf{Open}(\mathsf{prk}, \mathsf{aux}, f_1, \mathbf{y}_1)$ and $\pi_2 \leftarrow \mathsf{LVC}.\mathsf{Open}(\mathsf{prk}, \mathsf{aux}, f_2, \mathbf{y}_2)$, then for all $\alpha, \beta \in \mathcal{M}$:

 $\tilde{\pi} = (\alpha \pi_1 + \beta \pi_2)$ verifies LVC.Vf(vrk, C, $\tilde{f} = (\alpha f_1 + \beta f_2), \tilde{\mathbf{y}} = (\alpha \mathbf{y}_1 + \beta \mathbf{y}_2), \tilde{\pi} = 1.$

Homomorphic Openings. An LVC scheme has homomorphic openings if we can combine opening proofs for the same linear-map f applied to two different vectors \mathbf{v}_1 and \mathbf{v}_2 to obtain a new proof of opening $\tilde{\pi}$ that verifies with respect to the linear combination \tilde{C} of the two initial commitments C_1, C_2 and show the result of \mathbf{f} applied to the linear combination of the vectors \mathbf{v}_1 and \mathbf{v}_2 .

More formally, for all λ , $\mathcal{F} \subset \{f : \mathcal{M}^m \to \mathcal{M}^n\}$, vectors $\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{M}^m$, and (vrk, prk) \leftarrow LVC.KeyGen $(1^{\lambda}, \mathcal{F})$, if $\pi_1 \leftarrow$ LVC.Open $(\text{prk}, \text{aux}_1, f, \mathbf{y}_1)$ and $\pi_2 \leftarrow$ LVC.Open $(\text{prk}, \text{aux}_2, f, \mathbf{y}_2)$, where $(C_1, \text{aux}_2) \leftarrow$ LVC.Commit $(\text{prk}, \mathbf{v}_1)$ and $(C_2, \text{aux}_2) \leftarrow$ LVC.Commit $(\text{prk}, \mathbf{v}_2)$, then for all $\alpha, \beta \in \mathcal{M}$:

 $\tilde{\pi} = (\alpha \pi_1 + \beta \pi_2)$ verifies LVC.Vf(vrk, $\tilde{\mathsf{C}} = (\alpha \mathsf{C}_1 + \beta \mathsf{C}_2), f, \tilde{\mathbf{y}} = (\alpha \mathbf{y}_1 + \beta \mathbf{y}_2), \tilde{\pi} = 1.$

4 Generic Constructions from Homomorphic Proofs

Many natural schemes (such as [23, 13], PST commitments or our constructions in Section 5) have homomorphic proofs or openings. This motivates us to consider generic constructions that enhance any LVC scheme with homomorphic properties. We start by defining the notions of unbounded aggregation for same and cross-commitments and then we show how to add such properties to LVC schemes that have homomorphic proofs for the former and, additionally, homomorphic commitments for the latter.

4.1 New Notion: Unbounded Aggregation

The intuition for our definition is that, given t proofs, commitments or openings, we can aggregate them by performing a linear combination with random coefficients. Importantly, these coefficients have to be chosen after the claims are fixed and for that we rely on the RO model, as it is often the case for aggregation in the literature.

In our work, we go a step further and show how this procedure can be done over already aggregated proofs. Actually, aggregating *already aggregated* proofs consists off just sampling new coefficients and using them for fresh linear combinations. Importantly, the verifier needs to have access to the aggregation history: it has to recompute the coefficient corresponding to each initial proof π , which is the product of all the coefficients used in the aggregations it was involved in. Note that this also adds a small overhead to the verifier: it makes a linear (in the number of aggregation "hops") number of hash computations.

11

Example for same-commitment aggregation: Consider vector \mathbf{v} committed in C , functions f_1, f_2 and f_3 ; let π_1, π_2 and π_3 be proofs that $f_1(\mathbf{v}) = \mathbf{y}_1, f_2(\mathbf{v}) = \mathbf{y}_2$ and $f_3(\mathbf{v}) = \mathbf{y}_3$. An aggregated proof for $f_2(\mathbf{v}) = \mathbf{y}_2, f_3(\mathbf{v}) = \mathbf{y}_3$, would be $\pi_1^* = \pi_2 + \gamma_1 \pi_3$, for $\gamma_1 = \mathsf{H}(\mathsf{C}, \{(f_2, \mathbf{y}_2), (f_3, \mathbf{y}_3)\})$. In a second step, we can aggregate a proof that $f_1(\mathbf{v}) = \mathbf{y}_1$, by performing $\pi_2^* = \pi_1 + \gamma_2 \pi_1^*$, for $\gamma_2 = \mathsf{H}(\mathsf{C}, (f_1, \mathbf{y}_1), \gamma_1)$. At the verification step, the verifier would reconstruct the coefficients of each initial proof in π_2^* . For instance, $\delta_1 = 1, \delta_2 = \gamma_1 \gamma_2, \delta_3 = \gamma_2$. Then, the verifier can run the LVC.Vf algorithm to check whether $\pi_2^* = \pi_1 + \gamma_2 \pi_1^* = \pi_1 + \gamma_1 \gamma_2 \pi_2 + \gamma_2 \pi_3$ is a valid proof that function $f = f_1 + \gamma_1 \gamma_2 f_2 + \gamma_2 f_3$ evaluated at the vector committed in C opens to $y = y_1 + \gamma_1 \gamma_2 y_2 + \gamma_2 y_3$. For this last step to work we need the homomorphic proof property and the verifier to have access to the aggregation "history".

To describe our history of claims we move to *trees* of statements $\{f_j, \mathbf{y}_j\}_{j=1}^t$. In these trees, leaves are pairs of function-output (f, \mathbf{y}) . As in the usual case internal nodes are defined as an ordered list of subtrees. An empty history/tree is referred to as null. We denote trees using the syntax $T_{f,\mathbf{y}}$ and the operation that "merges" two subtrees in order adding a new root as " \therefore ". The following definition formalizes the above and will be useful in our construction. We remark that we include the commitment in each of the leaves of the trees $T_{f,\mathbf{y}}$. This does not increase the input size for cross-commitment aggregation where this information is necessary (for same-commitment aggregation the commitment is not necessary). This also allows to model more closely the "claims" for the cross-commitment case where each proof is for a statement ($\mathsf{C}, f, \mathsf{y}$).

Definition 3. Given a tree T we associate to each of its internal nodes a hash label h defined so that $h(L \therefore R) := H(C, L, R)$. We then associate to each of the leaves in the tree a label

$$\delta(\mathsf{leaf}) := \prod_{i=1,\dots,t} h(x_i)^{r(x_i,\mathsf{leaf})}$$

where the x_i -s are the internal nodes along the path from leaf to the root (root included and starting from the bottom), the predicate r(x, leaf) is 1 if leaf is a right child of x and 0 otherwise.

Remark 1 (Unbounded vs One-hop vs Incremental). Previous works have defined other types of aggregation. In one-hop aggregation (or batching) [3] aggregated proofs cannot be aggregated further. Incremental aggregation [5] does not have this limitation. The difference between the latter and our notion is that incremental aggregation does not require to keep track of the order in which the aggregation has been applied (for verification or further aggregation). On the other hand, we do require to track order, but we argue that this is not an overhead in many settings. In particular, even incremental aggregated, albeit in no order. Adding a structure to the claims roughly adds a number of bits linear in the length of the opening for additional separators (see also examples on tree histories above). When we consider unbounded-aggregatable LVC, we assume KeyGen outputs additional parameters for aggregations in pp. The aggregation algorithm will follow this syntax⁹:

LVC.Agg(pp,
$$T_{f,\mathbf{y}}, \pi, T_{f',\mathbf{y}'}, \pi') \rightarrow \pi^*$$

We subsequently modify the syntax for the verification algorithm in an (unbounded) aggregatable LVC as follows:

LVC.Vf(vrk, C,
$$T_{f,\mathbf{y}}$$
 \therefore $T'_{f,\mathbf{y}}$, π^*) \rightarrow $b \in \{0,1\}$

with $T_{f,\mathbf{y}}$ replacing f, \mathbf{y} .

We require the following correctness property and that function binding still holds.

Definition 4 (Unbounded Aggregation Correctness). For any $T_{f,\mathbf{y}}, T_{f',\mathbf{y}'}$ and any π, π' :

$$\Pr \begin{bmatrix} (\mathsf{LVC}.\mathsf{Vf}(\mathsf{vrk},\mathsf{C},T_{f,\mathbf{y}},\pi) = 1 \land \\ \mathsf{LVC}.\mathsf{Vf}(\mathsf{vrk},\mathsf{C},T_{f,\mathbf{y}}',\pi') = 1 \end{pmatrix} \Rightarrow \\ \mathsf{LVC}.\mathsf{Vf}(\mathsf{vrk},\mathsf{C},T_{f,\mathbf{y}}',\pi') = 1 \end{pmatrix} \Rightarrow \\ \begin{bmatrix} (\mathsf{prk},\mathsf{vrk},\mathsf{pp}) \leftarrow \mathsf{LVC}.\mathsf{KeyGen}(1^{\lambda},\mathcal{F}) \\ (\mathsf{C},\mathsf{aux}) \leftarrow \mathsf{LVC}.\mathsf{Commit}(\mathsf{prk},\mathbf{v}) \\ \pi^* \leftarrow \mathsf{LVC}.\mathsf{Agg}(\mathsf{pp},T_{f,\mathbf{y}},\pi,T_{f',\mathbf{y}'},\pi') \end{bmatrix} = 1 \end{bmatrix}$$

Definition 5 (Unbounded Aggregation Function Binding). For any $T_{f,\mathbf{y}}$, $T_{f',\mathbf{y}'}$ the following probability is negligible in λ :

$$\Pr \begin{bmatrix} \mathsf{LVC.Vf}(\mathsf{vrk},\mathsf{C},T_{f,\mathbf{y}} \therefore T'_{f,\mathbf{y}},\pi^*) = 1 & (\mathsf{prk},\mathsf{vrk},\mathsf{pp}) \leftarrow \mathsf{LVC.KeyGen}(1^{\lambda},\mathcal{F}) \\ \wedge \nexists \mathbf{a} \text{ s.t. } f(\mathbf{a}) = \mathbf{y} \wedge f'(\mathbf{a}) = \mathbf{y}' & (\mathsf{C},\pi^*,T_{f,\mathbf{y}},T'_{f,\mathbf{y}}) \leftarrow \mathcal{A}(\mathsf{pp},\mathsf{prk},\mathsf{vrk}) \end{bmatrix}$$

Definition: Cross-Commitment Aggregation. Unbounded aggregation can be performed across different commitments as well. This property is called *Crosscommitment Aggregation* and makes sense when we have a set of commitments C'_1, \ldots, C'_t that we want to open at one or more maps f, as it allows to compute a succinct proof of opening for linear-maps from different vectors committed separately. Below we show our syntax which directly expands on our samecommitment aggregation described above. Function binding and correctness are also straightforward to expand. We let $T_{f,\mathbf{y}}$ include our commitments in the leaves (see also next section).

Cross-commitment aggregation: LVC.CrossAgg Cross-commitment verification: LVC.CrossVfy(

$$\begin{aligned} & \mathsf{LVC.CrossAgg}(\mathsf{pp}, T_{f,\mathbf{y}}, \pi, T_{f',\mathbf{y}'}, \pi') \to \pi^* \\ & \mathsf{LVC.CrossVfy}(\mathsf{vrk}, \left(\mathsf{C}_j'\right)_j, T_{f,\mathbf{y}}, \pi^*) \to 0/1 \end{aligned}$$

4.2 Unbounded Aggregation for LVC

We now describe unbounded aggregation algorithms for *any* LVC scheme that satisfies the homomorphic properties of Section 3.1.

⁹ The algorithms can be generalized for more proofs. Proof size remains the same, also for cross-commitment aggregation.

 $\mathsf{LVC}.\mathsf{KeyGen}(1^{\lambda},\mathcal{F}) \to (\mathsf{prk},\mathsf{vrk},\mathsf{pp},\{\mathsf{upk}_j\}_{j=1}^m): \text{ Additionally generate the description of a hash function } \mathsf{H}(\cdot) \text{ and set it as } \mathsf{pp}.$

 $\mathsf{LVC.Agg}(\mathsf{pp}, T_{f,\mathbf{y}}, \pi, T_{f',\mathbf{y}'}, \pi') \to \pi^*:$

Compute $\gamma = \mathsf{H}(\mathsf{C}, T_{f,\mathbf{y}}, T_{f',\mathbf{y}'})$ and output $\pi^* = \pi + \gamma \pi'$.

LVC.Vf(vrk, C, $T_{f,\mathbf{y}}$: $T_{f',\mathbf{y}'}$, π^*) $\rightarrow b$

Return $b \leftarrow \text{LVC.Vf}(\text{vrk}, \mathsf{C}, f^*, y^*, \pi^*)$ where:

- let $\{\operatorname{\mathsf{leaf}}_i = (\mathsf{C}, f_i, \mathbf{y}_i)\}_{i=1}^{\ell}$ be all the leaves in $T_{f, \mathbf{y}} \therefore T_{f', \mathbf{y}'}$.
- For each *i* let $\delta_i := \delta(\mathsf{leaf}_i)$ be the value defined as in Definition 3.

$$f^* := \sum_i \delta_i f_i \qquad y^* := \sum_i \delta_i \mathbf{y}_i$$

Theorem 1. When applied to a function binding LVC scheme with homomorphic proofs, (LVC.Agg, LVC.Vf) satisfies Unbounded Aggregation Correctness (as in Def. 4) and Function Binding (Def. 5) in the ROM.

Proof. Correctness follows by inspection, using the fact that the LVC satisfies homomorphic proof, so we omit it.

For function binding, let $(\mathsf{C}, \pi^*, T_{f,\mathbf{y}}, T_{f',\mathbf{y}'})$ be an output of \mathcal{A} such that LVC.Vf $(\mathsf{vrk}, \mathsf{C}, T_{f,\mathbf{y}} \therefore T_{f',\mathbf{y}'}, \pi^*)=1$. By construction this implies IP.Vf $(\mathsf{vrk}, \mathsf{C}, \sum_i \delta_i f_i, \sum_i \delta_i \mathbf{y}_i, \pi^*)=1$. Because IP is function binding, except with negligible probability, there exists a vector \mathbf{a} such that $f(\mathbf{a}) = \mathbf{y}$, for $\mathbf{y} = \sum_i \delta_i \mathbf{y}_i, f(\mathbf{X}) = \sum_i \delta_i f_i(\mathbf{X})$ then there exists \mathbf{a} such that $\sum_{i=1}^t \delta_i f_i(\mathbf{a}) = \sum_{i=1}^t \delta_i \mathbf{y}_i$.

Since H is a random oracle, the coefficients δ_i do not depend on \mathbf{y}_i, f_i . And by the Schwartz-Zippel lemma, except with probability r/\mathbb{F} , $f_i(\mathbf{a}) = \mathbf{y}_i$ for all i, which concludes the proof.

Cross-Commitment Aggregation for LVC. For the case of cross-commitment aggregation, we proceed similarly but we also need to homomorphically operate on the commitments (recall that hashing on trees implicitly hashes the commitments too since we include them there).

 $\mathsf{LVC.CrossAgg}(\mathsf{pp}, T_{f,\mathbf{y}}, \pi, T_{f',\mathbf{y}'}, \pi') \to \pi^*:$

Compute $\gamma = \mathsf{H}(T_{f,\mathbf{y}}, T_{f',\mathbf{y}'})$

Output $\pi^* = \pi + \gamma \pi$

LVC.CrossVfy(vrk, $(\mathsf{C},\mathsf{C}',T_{f,\mathbf{y}}:T_{f',\mathbf{y}'},\pi^*) \to b$

- let $\mathsf{leaf}_1, \ldots, \mathsf{leaf}_\ell$ be all the leaves in $T_{f,\mathbf{y}} \therefore T_{f',\mathbf{y}'}$. We add to each leaf leaf_i and additional subindex j that refers to which commitment the proof in leaf_{ij} corresponds to. Note that we still consider ℓ leaves.
- each leaf_{ij} is of the form (C_j, f_i, \mathbf{y}_i)
- For each *i* let $\delta_{ij} := \delta(\mathsf{leaf}_{ij})$ be the value defined as in Definition 3.
- Compute

$$f_j^* := \sum_i \delta_{ij} f_i \qquad y_j^* := \sum_i \delta_{ij} \mathbf{y}_i$$

- Return 1 iff $b_j = 1$ for all $b_j \leftarrow \text{LVC.Vf}(\text{vrk}, \mathsf{C}_j, f_j^*, y_j^*, \pi^*)$.

14M. Campanelli, A. Nitulescu, C. Ràfols, A. Zacharakis, A. Zapico

Efficiency. For our constructions, the verification equations for computing $b_i =$ IP.Vf(vrk, C^*, f^*, y^*, π^*) are two pairing equations where the elements in the right side can be aggregated, and thus the verifier performs only $\ell + 1$ pairings.

Security. The security of this augmented construction follows analogously to that for same-commitment aggregation, with the additional requirement for the LVC scheme to have homomorphic commitments and openings.

4.3From Inner-Products to Arbitrary Linear-Maps

In this section we show we can obtain LVC schemes for any family of functions $\mathcal{F} \subset \{f: \mathbb{F}^m \to \mathbb{F}^n\}$ starting from simpler constructions that have homomorphic proofs and openings.

Our starting point are LVC schemes for $\mathcal{F}_{\mathsf{IP}} = \{f : \mathbb{F}^m \to \mathbb{F}\}$, or innerproduct VC schemes, that we will denote as IP = (IP.KeyGen, IP.Commit, IP.Open, IP.Vf). All this algorithms work as the ones for LVC, except that instead of $f \in \mathcal{F}_{\mathsf{IP}_{m,n}}$, they use the vector $\mathbf{f} \in \mathbb{F}^m$ so that $f(\mathbf{v}) = \mathbf{f} \cdot \mathbf{v}$.

We can write the linear-map $f: \mathbb{F}^m \to \mathbb{F}^n$ as $f = (f_1, f_2, \dots, f_n)$, where each f_i is an inner product function. If the IP scheme has homomorphic proofs, and we set π_i to be the proof that $f_i(\mathbf{v}) = \mathbf{f}_i \cdot \mathbf{v} = y_i$, an aggregation of $\{\pi_i\}_{i=1}^n$ is a proof of the statement $f(\mathbf{v}) = \mathbf{y}$. Later, in the following section, we show two possible constructions of IP vector commitments schemes that can be used to instantiate the framework in this section.

An IP aggregation algorithm for one-hop aggregation¹⁰ of proofs works as follows:

IP.Agg(pp, $\{\mathbf{f}_i, y_i\}_{i=1}^n, \pi = (\pi_i)_{i=1}^n \to \pi'$:

Parse pp = H, where H is a hash function, compute $\gamma = H(C, \{f_i, y_i\}_{i=1}^n)$ Output $\pi' = \sum_{i=1}^{n} \gamma^{i-1} \pi_i$ IP.VfAgg(vrk, C, { \mathbf{f}_i, y_i }_{i=1}^n, π') $\rightarrow b$:

Compute
$$\gamma = \mathsf{H}(\mathsf{C}, \{\mathbf{f}_i, y_i\}_{i=1}^n), \quad \mathbf{f}' = \sum_{i=1}^n \gamma^{i-1} \mathbf{f}_i, \quad y' = \sum_{i=1}^n \gamma^{i-1} y_i$$

Output $b \leftarrow \mathsf{IP.Vf}(\mathsf{vrk}, \mathsf{C}, \mathbf{f}', y', \pi').$

Using IP.Agg, we present an alternative way of computing concise proofs of LVC for more general functions $f : \mathbb{F}^m \to \mathbb{F}^n$, based on aggregation.

LVC.KeyGen $(1^{\lambda}, \mathcal{F}) \rightarrow (prk, vrk, pp)$:

- Run (prk, vrk) \leftarrow IP.KeyGen $(1^{\lambda}, \mathcal{F}_{IP})$ and generate aggregation parameters pp = H (a hash function). Output (prk, vrk, pp).

LVC.Commit(prk, \mathbf{v}) \rightarrow (C, aux) :

 $- \operatorname{Run}(\mathsf{C},\mathsf{aux}) \leftarrow \mathsf{IP}.\mathsf{Commit}(\mathsf{prk},\mathbf{v}) \text{ and output }(\mathsf{C},\mathsf{aux}).$

LVC.Open(prk, pp, aux, $f, \mathbf{y}) \rightarrow \pi$:

- Parse $f = (f_1, f_2, \dots, f_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$. Consider \mathbf{f}_i as the vector representing inner-product function f_i . Run $\pi_i \leftarrow \mathsf{IP.Open}(\mathsf{prk},\mathsf{aux},\mathbf{f}_i,y_i)$ for $i \in [n]$ and output $\pi \leftarrow \mathsf{IP}.\mathsf{Agg}(\mathsf{pp}, \{\mathbf{f}_i, y_i\}_{i=1}^n, (\pi_i)_{i=1}^n)$.

LVC.VfAgg(vrk, pp, C,
$$f, \mathbf{y}, \pi) \rightarrow b$$
:

 $^{^{10}}$ Naturally, this can be seen as a particular case of unbounded aggregation.

- Parse $f = (f_1, f_2, \dots, f_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$. Consider \mathbf{f}_i as the vector representing function f_i . Output $b \leftarrow \mathsf{IP.VfAgg}(\mathsf{vrk}, \mathsf{C}, {\mathbf{f}_i, y_i}_{i=1}^n, \pi)$

4.4 Updability for LVC

We consider updatability as an extra property of the LVC scheme. The KeyGen algorithm additionally computes the update keys, while two extra algorithms are defined as follows:

LVC.UpdCom(upk, C, j, δ) \rightarrow C': takes as input C, a position $j \in [m]$, update key upk, and a constant $\delta \in \mathcal{M}$. It outputs C' as a commitment for $\mathbf{v}' = \mathbf{v} + \delta \mathbf{e}_j^{11}$.

LVC.UpdOpen(upk, $j, \delta, f, \mathbf{y}, \pi$) $\rightarrow \pi'$: Takes as input upk, j, δ , a function f, a valid opening pair (\mathbf{y}, π) for f and outputs a proof π' for the new opening $\mathbf{y}' = f(\mathbf{v} + \delta \mathbf{e}_i)$.

Update correctness. Let $(prk, vrk, upk) \leftarrow LVC.KeyGen(1^{\lambda}, \mathcal{F})$, and let (C, j, f, y, π) be a tuple such that LVC.Vf $(vrk, C, f, y, \pi) = 1$. Then LVC satisfies update correctness if for any $\delta \in \mathcal{M}$,

$$\Pr \begin{bmatrix} \mathsf{LVC.Vf}(\mathsf{vrk},\mathsf{C}',f,\mathbf{y}',\pi') = 1 \\ \land \mathbf{y}' = \mathbf{y} + \delta f(\mathbf{e}_j) \end{bmatrix} \begin{bmatrix} \mathsf{C}' \leftarrow \mathsf{LVC.UpdCom}(\mathsf{upk}_j,\mathsf{C},j,\delta) \\ \pi' \leftarrow \mathsf{LVC.UpdOpen}(\mathsf{upk}_j,j,\delta,f,\mathbf{y},\pi) \end{bmatrix} = 1.$$

Updates for IP. We present a generic construction of the updatability algorithms for inner-product schemes. We state that even though algorithms can be generalized to LVC for arbitrary functions, for ease of exposition we only present it for inner-product openings, rather than generic linear-maps.

It is easy to see that commitments can be updated when one value of the vector changes by simply applying the linear-homomorphic property of the underlying IP scheme. Given C such that $(C, aux) \leftarrow LVC.Commit(prk, v)$, when position t of the vector changes, i.e. $\mathbf{v}' = \mathbf{v} + \delta \mathbf{e}_t$ we can compute a commitment to the new vector \mathbf{v}' as $C' = (C + \hat{C})$ where $(\hat{C}, a\hat{u}x) \leftarrow LVC.Commit(prk, \mathbf{e}_t)$ is given as an update key.

Moreover, it is possible to update existing proofs using the homomorphic openings property of the IP scheme: when position t of the vector changes as above, to update a prior proof we simply add to π a proof $\hat{\pi}$ corresponding to the opening of $f(\delta \mathbf{e}_t)$. The resulting $\pi' = \pi + \hat{\pi}$ corresponds to the opening of the sum $f(\mathbf{v}') = f(\mathbf{v}) + \delta f(\mathbf{e}_t)$ with respect to the updated commitment $C' = C + \hat{C}$.

We extend IP arguments to satisfy updatability by asking the IP.KeyGen algorithm to additionally generate updatable keys and introduce IP.UpdCom and IP.UpdOpen that work the following way; $IP_{i}(x_{i}) = \frac{1}{2} \int_{-\infty}^{\infty} \frac{1}{2} \int_{-\infty}$

 $\mathsf{IP}.\mathsf{KeyGen}(1^{\lambda},\mathcal{F}_{\mathsf{IP}}) \to (\mathsf{prk},\mathsf{vrk},\{\mathsf{upk}_j\}_{j=1}^m):$

- Additionally generate public update keys upk: Set $\pi_{u_{ij}} \leftarrow \text{IP.Open}(\text{prk}, aux_j, \mathbf{e}_i, u_{ij} = \mathbf{e}_i \cdot \mathbf{e}_j), \forall i, j \in [m]$, Define $\mathsf{upk}_j = \{\pi_{u_{ij}}\}_{i=1}^m$ for all $j \in [m]$, and output $(\mathsf{prk}, \mathsf{vrk}, \{\mathsf{upk}_i\}_{i=1}^m)$.

 $[\]mathsf{IP}.\mathsf{UpdCom}(\mathsf{prk},\mathsf{C},t,\delta)\to\mathsf{C}':$

 $[\]overline{}^{11}$ This notion can be generalized to more than one position.

- Set $\hat{\mathsf{C}} \leftarrow \mathsf{IP}.\mathsf{Commit}(\mathsf{prk}, \mathbf{e}_t)$, and output $\mathsf{C}' = \mathsf{C} + \delta \hat{\mathsf{C}}$.

IP.UpdOpen(upk_t, t, δ , C, f, y, π) $\rightarrow \pi'$:

- Parse $\mathsf{upk}_t = \{\pi_{u_{it}}\}_{i=1}^m$ and compute $\hat{\pi} = \sum_{i=1}^m f_i \pi_{u_{it}}$. Set $\pi' = \pi + \delta \hat{\pi}$ as proof for $y' = y + \mathbf{f} \cdot \delta \mathbf{e}_t$ and output π' .

Theorem 2. If IP satisfies function binding and has homomorphic commitments and openings, the extension above satisfies update correctness.

Proof. The proof follows directly by the definition of homomorphic proof and IP.UpdCom, IP.UpdOpen.

Constructions for Inner-Pairing VC $\mathbf{5}$

In this section, we present two constructions of LVC for inner products, that is, for functions $f \subset \mathcal{F}_{\mathsf{IP}} = \{f : \mathbb{F}^m \to \mathbb{F}\}$. We denote as $\mathsf{IP} = (\mathsf{IP}.\mathsf{KeyGen}, \mathsf{IP})$ IP.Commit, IP.Open, IP.Vf) a vector commitment scheme with inner product openings. All the algorithms work as the ones for LVC, except that they take as inputs the vector of coefficients of the linear function $f \in \mathcal{F}_{\mathsf{IP}}$, $f(\mathbf{v}) = \mathbf{f} \cdot \mathbf{v}$, i.e. use the vector $\mathbf{f} \in \mathbb{F}_n^m$.

The first one is in the monomial basis and the other based on the univariate sumcheck of [1, 21] that considers vectors encoded as polynomials in the Lagrange basis. We prove they are indeed linear vector commitment arguments with homomorphic proofs and openings. Therefore, they can be used as a starting point to obtain further aggregation properties as shown in Section 4.1 and, in particular, lead to two different more generic linear-map vector commitment schemes.

Monomial Basis 5.1

For the first scheme, we consider vectors $\mathbf{a} \in \mathbb{F}^m$ encoded as a polynomial in the monomial basis, that is as $a(X) = \sum_{i=1}^{m} a_i X^{i-1}$.

 $\mathsf{IP}.\mathsf{KeyGen}(1^{\lambda},\mathcal{F}_{\mathsf{IP}}) \to (\mathsf{prk},\mathsf{vrk}):$

Generate group description $\mathsf{gk} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e) \leftarrow \mathcal{G}(p)$ Sample $\tau \leftarrow \mathbb{F}$

Output prk = $(\{[\tau^i]_{1,2}\}_{i=0}^{m-1})$, vrk = $([\tau^{m-1}]_1, \{[\tau^i]_2\}_{i=0}^m)$. IP.Commit(prk, **a**) \rightarrow (C_a, **aux**): Compute C_a = $\sum_{i=1}^m a_i[\tau^{i-1}]_1$ and output (C_a, **a**). IP.Open(prk, aux, $\mathbf{b}, y) \rightarrow \pi$:

Find R(X), H(X) such that $\deg(R) < m - 1$ and

$$\left(\sum_{i=1}^{m} a_i X^{i-1}\right) \left(\sum_{i=1}^{m} b_i X^{m-i}\right) - y X^{m-1} = R(X) + X^m H(X).$$

Define $\hat{R}(X) = XR(X)$ Output $\pi = ([R(\tau)]_1, [H(\tau)]_1, [\hat{R}(\tau)]_1).$ IP.Vf(vrk, C_a, \mathbf{b}, y, π) $\rightarrow 0/1$: Compute $C_b = \sum_{i=1}^m b_i [\tau^{m-i}]_1$, parse $\pi = ([R]_1, [H]_1, [\hat{R}]_1)$ and output 1 if and only if

$$e(\mathsf{C}_a,\mathsf{C}_b) - e(y[\tau^{m-1}]_1,[1]_2) = e([R]_1,[1]_2) + e([H]_1,[\tau^m]_2)$$
 and
 $e([R]_1,[\tau]_2) = e([\hat{R}]_1,[1]_2).$

Remark 2. The second verification equation is meant to ensure that $[R]_1$ is the evaluation at τ of a polynomial of degree at most m-2. Note that it is important for security that only m-1 powers of τ are available to the adversary in \mathbb{G}_1 , otherwise the second equation does not guarantee that the degree is at most m-2. Importantly, even though we present our construction using a srs with powers of τ up to m-1 in \mathbb{G}_1 and m in \mathbb{G}_2 , it can easily be adapted for a bigger srs and, in particular, existing trusted setups where the same powers of τ are available in both groups are enough. For the second check, if m+k powers of τ are given in \mathbb{G}_1 , $\hat{R}(X)$ should be defined as $X^{k+2}R(X)$.

We implement this construction for single positions and compare it with individual position openings in Merkle tree-based vector commitments in the full version of the paper. Below, we state the theorems for security guarantees, and refer the reader to the full version for the formal proofs.

Theorem 3. The construction above satisfies Completeness, Homomorphic Proofs and Homomorphic Openings.

Theorem 4. The construction above satisfies Strong Function Binding in the AGM under the (m-1,m)-BSDH Assumption([2]).

Intuition. Without loss of generalization, we consider an adversary that provides two proofs π_1 , π_2 with IP.Vf(vrk, $C_a, \mathbf{b}_k, y_k, \pi_{b_k}$) = 1 for k = 1, 2 and there is no $\mathbf{a} \in \mathbb{F}^m$ s. t. $\mathbf{a} \cdot \mathbf{b}_k = y_k$.

For commitment C_a and proofs $\pi_k = ([R_k]_1, [H_k]_1, [\hat{R}_k]_1)$ under the AGM we can extract polynomials $a(X), H_k(X), R_k(X), \hat{R}_k(X)$ of degree up to m-1such that the proof elements are their evaluations in \mathbb{G}_1 at secret point τ . On the other hand, the second verification equation in our IP scheme assures that $R_k(X)X = \hat{R}_k(X)$. Because $\deg(\hat{R}_k) \leq m-1 \deg(R_1), \deg(R_2) \leq m-2$. Consider $\mathbf{a} \in \mathbb{F}^m$ the vector of the coefficients of $a(X) = \sum_{i=1}^m a_i X^{i-1}$.

Consider $\mathbf{a} \in \mathbb{F}^m$ the vector of the coefficients of $a(X) = \sum_{i=1}^m a_i X^{i-1}$. Then, from the first verification equation in both of the proofs we have that $p_1(X), p_2(X)$ have a common root in τ , where for k = 1, 2:

$$P_k(X) = \left(\sum_{i=1}^m a_i X^{i-1}\right) \left(\sum_{i=1}^m b_{ki} X^{m-i}\right) - y_k X^{m-1} - R_k(X) + X^m H_k(X).$$

If for some k, $p_k(X)$ is not the zero polynoial, since τ is one of its roots, we can solve the discrete logarithm problem and extract τ from $[P_k(\tau)]_1$. Thus, it must be the case that $p_k(X) \equiv 0$ for k = 1, 2. Because $\deg(R_b), \deg(R_c) < m-1$ and $\deg(X^m H_b(X)), \deg(X^m H_c(X)) > m-1$, we have that the coefficient for X^{m-1} in polynomial $P_k(X)$ is $\sum_{i=1}^m a_i b_{ki} - y_k = 0$. Indeed, there exists a vector **a** such that $\mathbf{a} \cdot \mathbf{b}_k = y_k$, contradicting the initial assumption that the adversary \mathcal{A} breaks the strong functional binding.

17

Updates Without Keys. We remark that we do not need any additional update keys added to the setup. Indeed, the update key is made by proofs of inner products between cannonic vectors $\mathbf{e}_i \cdot \mathbf{e}_i = 1$ or $\mathbf{e}_i \cdot \mathbf{e}_j = 0$. In our construction for encodings in the monomial basis, a proof that $\mathbf{e}_i \cdot \mathbf{e}_i = 1$ consists on R(X) = H(X) = 0. On the other hand, to prove that $\mathbf{e}_i \cdot \mathbf{e}_j = 0$ for $i \neq j$ the proof is (the evaluation in the group of) either $R(X) = X^{m+i-j}$ if j > i, or $H(X) = X^{i-j}$ if i > j. As such powers of τ are already included in prk, upk = \emptyset .

5.2 Lagrange Basis

In this second scheme, for a Lagrange basis $\{\lambda_i(X)\}_{i=1}^m$ over a multiplicative group $\mathbb{H} = \{\mathbf{h}_1, \ldots, \mathbf{h}_m\}$ of size m in \mathbb{F} we encode a vector $\mathbf{a} \in \mathbb{F}^m$ as a polynomial $a(X) = \sum_{i=1}^m a_i \lambda_i(X)$. Recall that when \mathbb{H} is a multiplicative subgroup, $\lambda_i(0) = m^{-1}$ for all $i \in [m]$. Moreover, if we set $t(X) = \prod_{i=1}^m (X - \mathbf{h}_i)$ we have that $\lambda_i(X)\lambda_j(X) \equiv 0 \mod t(X)$, and $\lambda_i(X)^2 \equiv \lambda_i(X) \mod t(X)$. The construction below, presented in [21], exploits these properties in the proof of openings for inner-products:

IP.KeyGen $(1^{\lambda}, \mathcal{F}_{\mathsf{IP}_m}) \to (\mathsf{prk}, \mathsf{vrk})$:

Generate group description $\mathsf{gk} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e) \leftarrow \mathcal{G}(p)$, define multiplicative group $\mathbb{H} = \{\mathsf{h}_1, \ldots, \mathsf{h}_m\}$ in \mathbb{F} , and compute Lagrange polynomials $\{\lambda_j(X)\}_{j=1}^m$ over \mathbb{H} .

Sample $\tau \leftarrow \mathbb{F}$ and output $\mathsf{prk} = (\{[\tau^i]_{1,2}\}_{i=1}^m, \{[\lambda_i(\tau)]_1\}_{i=1}m-1, [\tau^m]_2)$ and $\mathsf{vrk} = ([1]_{1,2}, \{[\tau^i]_2, [\lambda_i(\tau)]_2\}_{i=1}^m).$

IP.Commit(prk, **a**) \rightarrow (C_a, **aux**): Compute C_a = $\sum_{i=1}^{m} a_i [\lambda_i(\tau)]_1$ and output (C_a, **a**). IP.Open(prk, aux, **b**, y) $\rightarrow \pi$:

Find R(X), H(X) such that $\deg(R) < m - 1$ and

$$\left(\sum_{i=1}^{m} a_i \lambda_i(X)\right) \left(\sum_{i=1}^{m} b_i \lambda_i(X)\right) - m^{-1}y = XR(X) + t(X)H(X)$$

Define $\hat{R}(X) = XR(X)$ and output $\pi = ([R(\tau)]_1, [H(\tau)]_1, [\hat{R}(\tau)]_1)$. IP.Vf(vrk, $C_a, \mathbf{b}, y, \pi) \to 0/1$: Calculate $C_b = \sum_{i=1}^m b_i [\lambda_i(\tau)]_2$

Parse $\pi = ([R]_1, [H]_1, [\hat{R}]_1)$ and output 1 if and only if

$$e(\mathsf{C}_a,\mathsf{C}_b) - e(m^{-1}y[1]_1,[1]_2) = e([R]_1,[1]_2) + e([H]_1,[t(\tau)]_2), \text{ and}$$

 $e([R]_1,[\tau]_2) = e([\hat{R}]_1,[1]_2).$

The proof of completeness can be found in [21]. Below, we state the theorems for Strong Function Binding and homomorphic proofs and openings, and refer the reader to [6] for the formal proofs.

Theorem 5. The construction above has Homomorphic Proofs and Openings.

Theorem 6. The construction above satisfies Strong Function Binding in the AGM under the (m-1,m)-BSDH Assumption ([2])

Intuition. The proof goes as the one for Theorem 4 except that

$$p_k(X) = \left(\sum_{i=1}^m a_i \lambda_i(X)\right) \left(\sum_{i=1}^m b_{ki} \lambda_i(X)\right) - m^{-1} y_k - X R_k(X) + t(X) H_k(X).$$

Once more, if one of the polynomials $p_k(X)$ is not the zero polynomial, since τ is one of its roots, we can solve the discrete logarithm problem and extract τ from $[P_k(\tau)]_1$.

Then, $P_k(X) \equiv 0$, for both k = 1, 2. Because deg $(R_k) < m - 1$ and

$$\left(\sum_{i=1}^{m} a_i \lambda_i(X)\right) \left(\sum_{i=1}^{m} b_{ki} \lambda_i(X)\right) \equiv \sum_{i=1}^{m} a_i b_{ki} \lambda_i(X) \mod t(X),$$

 $\sum_{i=1}^{m} a_i b_{ki} \lambda_i(X) - m^{-1} y_k = X R_k(X), \text{ which implies } \sum_{i=1}^{m} a_i b_{ki} \lambda_i(0) - m^{-1} y_k = 0.$ As \mathbb{H} is a multiplicative subgroup, $\lambda_i(0) = m^{-1}$ for all $i \in [m]$ and thus $\sum_{i=1}^{m} a_i b_{ki} = y_k$. Then, there exists **a** such that $\mathbf{a} \cdot \mathbf{b}_k = y_k$ for k = 1, 2, contradicting the initial claim that adversary \mathcal{A} is successful.

Updatability with Short Keys In this construction, a proof that $\mathbf{e}_i \cdot \mathbf{e}_i = 1$ is the encoding in a group of the polynomial $R_i(\tau)$, for $R_i(X) = (\lambda_i(X) - 1)/X$. On the other hand, the proof that $\mathbf{e}_i \cdot \mathbf{e}_j = 0$ for $i \neq j$ is $[H(\tau)]_1$, for $H(X) = ((\lambda_i(X)\lambda_j(X))/t(X))$. Including the evaluation of all these polynomials in upk would require a srs of quadratic size. Still, as noted in [23], these keys can be computed in constant time from a linear-size update key.

6 Subvector Openings

In this section, we present schemes for VC with Subvector Openings (SVC), starting from the constructions of Section 5. We will consider SVC as a special case of LVC. The class of functions that open a set of positions $I = \{i_1, \ldots, i_n\}$ of a committed vector $\mathbf{v} \in \mathbb{F}^m$ is given by the linear-map f_I with

$$f_I: \mathbb{F}^m \to \mathbb{F}^n, \qquad f_I(\mathbf{v}) = (\mathbf{e}_{i_1} \cdot \mathbf{v}, \dots \mathbf{e}_{i_n} \cdot \mathbf{v})$$

where for each $k \in [n]$, e_{i_k} is the i_k th vector of the canonical basis \mathbb{F}^m .

Naturally, for a vector $\mathbf{v} \in \mathbb{F}^m$, we can construct proofs of openings of subvectors $\mathbf{v}_I = (v_i)_{i \in I}$ by aggregating different inner product proofs for vectors \mathbf{e}_{i_k} for $i_k \in I$ using the techniques in Section 4.1. We refer to these aggregated proofs as *non-native* subvector openings, given that they require a random oracle and in particular, are no longer algebraic and homomorphic. As opposed to them, we call *native* subvector opening, a scheme that is algebraic and homomorphic.

In what follows, we improve on Subvector Openings in some special scenarios, achieving native aggregation for new schemes and reducing the verifier complexity in existing ones.

19

6.1 Native SV Openings for the Monomial Basis

For the construction of Section 5.1, we introduce native subvector openings for subsets with consecutive position $I = \{i, i + 1, ..., i + k\}$. That is, for $\tilde{\mathbf{c}} = (c_i)_{i \in I}$ such that there exist $\mathbf{u}_1, \mathbf{u}_2$ with $\mathbf{c} = (\mathbf{u}_1, \tilde{\mathbf{c}}, \mathbf{u}_2)$. To prove an opening of $\tilde{\mathbf{c}}$, we only need commitments to $R(X) = \sum_{s=1}^{i-1} c_i X^{m-i+s-1}$ and $H(X) = \sum_{i=i+k+1}^{m} c_{m-i+s+1} X^{s-1}$, which are shifted-encodings of $\mathbf{u}_1, \mathbf{u}_2$. The verifier checks that $\deg(R) < m-1$, computes $\tilde{C}(X) = \sum_{s=i}^{i+k} \tilde{c}_s X^{s-i}$ and $\tilde{C} = [\tilde{C}(\tau)]_1$ and checks whether $e(\mathbf{C} - \tilde{\mathbf{C}}, [\tau^{m-i}]_1) = e([R]_1, [1]_2) + e([H]_1, [\tau^{m+k}]_2)$.

Note that, given individual proofs of openings as in Section 5.1, that is, $[R_s(\tau)]_1, [H_s(\tau)]_1$ such that $C(X)X^{m-s} - c_sX^{m-1} = R_s(X) + X^mH_s(X)$ and $\deg(R_s) < m-1$, for the commitments defined above we have $[R]_1 = [R_i(\tau)]_1$ and $[H]_1 = [H_{i+k}(\tau)]_1$, that is, proofs can be aggregated at no cost for the prover.

6.2 Non-native SV Openings for the Monomial Basis

For the LVC scheme of Section 5.1, the techniques of Section 4.1 allow us to redefine the **Open** and Vf algorithms to work for an arbitrary subset of positions $I \subset [m]$. More specifically, the prover will simply run IP.**Open**(prk, aux, $\mathbf{e}_{i_k}, \mathbf{v}$) for $k = 1, \ldots, n$ to obtain (v_{i_k}, π_{i_k}) and π_{i_k} a proof of correct computation of v_{i_k} . Then, use the random oracle to sample a randomness $\gamma \in \mathbb{F}$ and output $\pi_I = \sum_{k=1}^n \gamma^{k-1} \pi_{i_k}$.

The verifier will receive $\pi_I = ([R]_1, [H]_1, [\hat{R}]_1)$, compute $y = \sum_{k=1}^n \gamma^{k-1} v_{i_k}$, and check as before $e([R]_1, [\tau]_2) = e([\hat{R}]_1, [1]_2)$ and

$$e\left(\mathsf{C},\sum_{k=1}^{n}\gamma^{k-1}[\tau^{m-i_{k}}]_{2}\right)-e\left(y[\tau^{m-1}]_{1},[1]_{2}\right)=e\left([R]_{1},[1]_{2}\right)+e\left([H]_{1},[\tau^{m}]_{2}\right).$$

Note that verifier's work is dominated by the computation of $\sum_{k=1}^{n} \gamma^{k-1} [\tau^{m-i_k}]_2$, so we analyze for which sets $I \subset [m]$ this computation can be cheaper than $|I| \mathbb{G}_2$ -exponentiations. Without loss of generality, we can re-assign $\gamma^{k-1} \to \gamma^{m-i_k}$, and thus our verifier now needs to compute $\sum_{k=1}^{n} [(\gamma X)^{m-i_k}]_2 = \sum_{i \in I} [(\gamma X)^{m-i}]_2$.

Now, note that if $I_{k,s,n} \subset [m]$ is an arithmetic progression, i.e. it is such that for a given ratio s, a starting power k and a number n of desired elements, $I_{k,s,n} = \{k, s+k, \ldots, (n-1)s+k\}$, then

$$\sum_{i \in I_{k,s,n}} (\gamma X)^{m-i} = (\gamma X)^k \frac{1 - (\gamma X)^n}{1 - (\gamma X)^s}.$$

This reduces the work of the verifier to compute $\sum_{i \in I_{k,s,n}} (\gamma X)^{m-i}$ to constant. Note that the verifier cannot compute $(1 - (\gamma X)^s)^{-1}$, so we multiply all the terms of the equation by $1 - (\gamma X)^s$. I.e., the verifier computes y =

 $\sum_{i \in I_{k,s,n}} \gamma^{m-i} y_i$ and checks whether

$$e([C]_1, \gamma^k[\tau^k]_2 - \gamma^{k+n}[\tau^{k+n}]_2) - e([\tau^{m-1}]_1y - [\tau^{n+s-1}]_1\gamma^s y, [1]_2)$$

= $e([R]_1, 1 - \gamma^s[\tau^s]_2)) + e([H]_1, [\tau^n]_2 - \gamma^s[\tau^{n+s}]_2).$

6.3 Lagrange Basis

Native. In the Lagrange Basis, one can use the native subset openings of [23]. There, the verifier needs to compute computation the vanishing polynomial $t_I(X) = \prod_{i \in I} (X - h_i)$. To reduce verifer's work we focus on those subsets $I \subset [m]$ such that $t_I(X)$ can be calculated in less than |I| computations. One answer to this question comes from cosets. That is, given $\mathbb{H} = \{1, \omega, \omega^2, \ldots, \omega^{m-1}\}$ group of roots of unity where $m = 2^n$, let \mathbb{H}_k be the subgroup of order 2^k of \mathbb{H} , where k goes from 0 to n. Then, for each $0 \leq s < 2^{n/k}$ we can construct the coset $I = \omega^s \mathbb{H}_k$, whose vanishing polynomial is $t_I(X) = X^{2^k} - (\omega^s)^{2^k}$. Verifier accepts if and only if $e(\mathsf{C} - \tilde{\mathsf{C}}, [1]_2) = e([H]_1, [x^{2^k}]_2 - \omega^{s2^k})$.

Non-native. Given that the native subvector opening procedure above works for arbitrary subsets $I \subset [m]$, we don't consider aggregation of individual positions. The latter makes sense only when applying a linear function to the new subset. That is, when the verifier is given $C_{f,I}$, claimed to be a commitment to $\mathbf{f} \cdot \mathbf{c}_I$, for some linear function \mathbf{f} applied to the vector $\mathbf{c}_I = (c_i)_{i \in I}$.

7 Maintainable Vector Commitment Schemes

7.1 Multivariate Case

One of the key points of vector commitment schemes that allow to speedup subvector openings is the ability to pre-compute and store individual openings and later aggregate them to create subvector openings without incurring linear amount of computations each time. This is the case for the construction in [23], presented in Section ??, and also for the maintainable scheme of [24].

In constructions such as the ones presented in Section 5, the proof of opening of one position involves *all* other elements in the vector. That is, the polynomials committed to create the proof have coefficients that involve all the values of the committed vector $\mathbf{v} \in \mathbb{F}^m$. As a consequence, prover work is linear in the size of \mathbf{v} (as it has to evaluate polynomials of degree m). To alleviate this, Shrinivasan et. al. [22] utilize a tree-like structure for computing/communicating proofs which allows pre-computation in quasi-linear (instead of quadratic) time and efficient updates at the cost of a proof of size log m.

In this section, we extend the techniques of [22] to achieve trade-offs and efficiency improvements. Roughly speaking, we present a way to "compose" the tree-based commitments of [22] with constant size ones. We achieve this by considering trees that themselves have commitments for leaves instead of openings.

21

The intuition is the following: we divide the vector \mathbf{v} in small chunks $\{\mathbf{v}_j\} \in \mathbb{F}^k$. We then arrange these chunks in a tree as follows: each chunk corresponds to a leaf of the tree and each node is a succinct representation of its children. The root of the tree is the commitment to the vector. An opening proof only involves the elements in the path of the root to the leaf containing the position to be opened. That is, if we want to open value a in position i of $\mathbf{v} \in \mathbb{F}^{k \cdot m'}$, we prove that (1) c_j is the leaf that contains the commitment to the j chunk containing i and (2) c_j opens to a in the position corresponding to i. The former part can be pre-computed and efficiently maintained while the latter is computed on the fly.

This results in a construction with the following memory/time trade-off: for any $k, m' \in \mathbb{N}$ with $m = k \cdot m'$, any opening can be computed in time *independent of* m' after pre-computing and storing $O_{\lambda}(m')$ values (independent of k). Furthermore, a relaxed *maintainability* notion is satisfied: all stored values can be pre-computed efficiently in $O_{\lambda}(m \cdot \log m')$ time and updated in $O(\log m')$ time.

Additionally, we show how to use a higher arity tree (any constant ℓ contrary to the binary ones used in [22]) to further reduce the proof size by a constant factor, namely $O_{\lambda}(\log_{\ell} m')$ (assuming a constant size commitment for the leaf part), at the expense of a slightly worse prover time. We note that –apart from the evident advantage of shorter proofs– this results in smaller aggregation time for the prover and verifier when using inner pairing products.

Our starting point is the PST polynomial commitment [20]. The PST polynomial commitment is a natural generalization of the KGZ polynomial commitment [15] for multivariate polynomials, that is, it allows to commit to ν -variate polynomials of individual degrees less than ℓ . The core idea of the construction lies in the fact that for every $p(\mathbf{X}) \in \mathbb{F}[X_{\nu}, \ldots, X_1]$ and $\mathbf{x} = (x_{\nu}, \ldots, x_1) \in \mathbb{F}^{\nu}$, $p(\mathbf{x}) = y$ if and only if there exist polynomials $H_{\nu}(\mathbf{X}), \ldots, H_1(\mathbf{X})$ such that

$$p(\mathbf{X}) - y = \sum_{j=1}^{\nu} H_j(\mathbf{X}) \cdot (X_j - x_j)$$

where the proof polynomials $H_i(\mathbf{X})$ are efficiently computable.

Using standard techniques to encode monomials in a cryptographically secure bilinear group (encode setting $\mathbf{X} = \boldsymbol{\tau}$ and publishing all the monomials $[\tau_{\nu}^{d_{\nu}}\cdots\tau_{1}^{d_{1}}]_{1}$ and $[\boldsymbol{\tau}]_{2}$) results in a polynomial commitment with proof of size roughly ν group elements.

Tree structure. To achieve the flexible memory/time trade-off, instead of having the vector values in the leaves of the tree, we replace them with elements $[\mathbf{r}]_1 \cdot \mathbf{v}_j$ where $[\mathbf{r}] \in \mathbb{G}_1^k$ is the commitment key of an arbitrary algebraic vector commitment scheme LVC. To open a position of \mathbf{v} , we use the PST approach to reach corresponding leaf j, and then the opening algorithm of LVC on \mathbf{v}_j .

One subtlety of replacing leaves with commitments is that a standalone PST proof is no longer binding, that is, the prover can undetectably claim arbitrary values that supposedly correspond to a leaf. We overcome this by using a low degree test to ensure that the claimed value for the leaf is uniquely defined.

Note that the root of the tree depends on the elements τ , **r**. Viewing both $\boldsymbol{\tau} = (\tau_{\nu}, \ldots, \tau_1)$ and $\mathbf{r} = (r_k, \ldots, r_1)$ as formal variables \mathbf{X}, \mathbf{R} , we can treat the root node (the commitment) as an evaluation of a polynomial. Now, note that this polynomial corresponds to the interpolation of the elements of the leaves in Σ^{ν} . Thus, the aforementioned polynomial is

$$p(\mathbf{X}, \mathbf{R}) = \boldsymbol{\lambda}(\mathbf{X}) \cdot (\mathbf{R} \cdot \mathbf{v}_1, \dots, \mathbf{R} \cdot \mathbf{v}_{\ell^{\nu}}) = (\boldsymbol{\lambda}(\mathbf{X}) \otimes \mathbf{R}) \cdot \mathbf{v}$$

The prover can still evaluate one by one the variables X_{ν}, \ldots, X_1 at $\sigma_{\nu}, \ldots, \sigma_1$ -as it would do in the simple PST case- and end up with a polynomial $q(\mathbf{R}) =$ $p(\boldsymbol{\sigma}, \mathbf{R}) = \mathbf{R} \cdot \mathbf{v}_i$. To ensure that q does not contain any X_i variable, we also include a low degree test in the proof. The evaluation of the latter polynomial at $[\mathbf{r}]_1$ corresponds to the leaf commitment at position σ and can be opened by employing the Open algorithm of the leaf commitment scheme with key $[\mathbf{r}]_1$.

Construction. First, we introduce some notation. Let $\Sigma \subseteq \mathbb{F}$ denote an interpolating set of size ℓ . Given $\boldsymbol{\sigma} = (\sigma_{\nu}, \ldots, \sigma_1) \in \Sigma^{\nu}$, we denote $\boldsymbol{\sigma}_{|i} = (\sigma_{\nu}, \ldots, \sigma_i) \in$ $\Sigma^{\nu-i+1}$. For $\mathbf{v} = (\mathbf{v}_{\sigma})_{\sigma \in \Sigma^{\nu}}$ with $\mathbf{v}_{\sigma} \in \mathbb{F}^k$ and $\sigma_1 \in \Sigma^i$ we denote with \mathbf{v}_{k,σ_1} the vector $(\mathbf{v}_{\sigma_1,\sigma_2})_{\sigma_2 \in \Sigma^{\nu-i}}$, that is, the concatenation of vectors \mathbf{v}_j whose ℓ -ary representation of the index j is prefixed with σ_1 . Finally, we denote with $\tau_{\nu,\ell}$ the ν -variate monomial basis of individual degree less than ℓ evaluated at $\tau_{\nu}, \ldots, \tau_1$. In all cases, we omit the subscript when it is clear from the context.

We present the construction next. While our aim is individual position openings, the construction supports a bigger family of functions: linear forms¹² applied to one of the k-sized chunks of the vector. Concretely, let $\mathcal{F}_{p,k} \subseteq \{f : \mathbb{F}^k \to \mathbb{F}\}$ be the family of linear forms supported by the leaf commitment scheme. We define the ℓ , ν -extended family as

$$\mathsf{Ext}_{\ell,\nu} - \mathcal{F}_{p,k} = \{ f : \mathbb{F}^{k \cdot \ell^{\nu}} \to \mathbb{F} \mid \exists f' \in \mathcal{F}_{p,k}, i \in \{1, \dots, \ell^{\nu}\} \text{ s.t.} \\ \forall \mathbf{v}_1, \dots, \mathbf{v}_{\ell^{\nu}} \in \mathbb{F}^k : f(\mathbf{v}_1, \dots, \mathbf{v}_{\ell^{\nu}}) = f'(\mathbf{v}_i) \}$$

Our construction is a linear vector commitment MVTree for the family $\mathsf{Ext}_{\ell^{\nu}}$ - $\mathcal{F}_{p,k}$, that uses as a black box an algebraic linear vector commitment scheme LVC' for the family $\mathcal{F}_{p,k}$.

- $\begin{array}{l} \mathsf{MVTree.KeyGen}(1^{\lambda},\mathsf{Ext}_{\ell^{\nu}}\text{-}\mathcal{F}_{p,k}) \to (\mathsf{prk},\mathsf{vrk}): \\ & (\mathsf{prk}' = [\mathbf{r}]_1,\mathsf{vrk}') \leftarrow \mathsf{LVC}'.\mathsf{KeyGen}(1^{\lambda},\mathcal{F}_{p,k}) \\ & \operatorname{Let} \boldsymbol{\lambda}(X) \text{ be the vector of Lagrange polynomials associated to } \boldsymbol{\Sigma}. \end{array}$
 - Sample $\tau_{\nu}, \ldots, \tau_1 \leftarrow \mathbb{F}$ and output $\mathsf{prk} = (\mathsf{prk}', [\lambda]_1 = [\lambda(\tau_{\nu}) \otimes \cdots \otimes$ $\mathsf{vrk} = (\mathsf{vrk}', [\tau_{\nu}]_2, \dots, [\tau_1]_2, [\tau_{\nu}^{\ell-1} \cdots \tau_1^{\ell-1}]_2),$ $\lambda(\tau_1)\otimes \mathbf{r}]_1, [\boldsymbol{\tau}\otimes \mathbf{r}]_1),$ $\mathsf{upk} = (\{[\lambda(\tau_j) \otimes \cdots \otimes \lambda(\tau_1) \otimes \mathbf{r}]_1\}_{j=\nu-1}^1),$

MVTree.Commit(prk, \mathbf{v}) \rightarrow (C, aux):

- For all
$$\sigma \in \Sigma^{\nu}$$
: compute $(C_{\sigma}, \mathsf{aux}_{\sigma}) \leftarrow \mathsf{LVC'}$. Commit(prk', \mathbf{v}_{σ}). Compute $\mathsf{C} = [p(\tau, \mathbf{r})]_1 = [\lambda]_1 \cdot \mathbf{v}$ and output $\mathsf{C}, \mathsf{aux} = (\{\mathsf{aux}_{\sigma}\}_{\sigma \in \Sigma^{\nu}}, \mathbf{v})$

 $^{^{12}}$ We use linear forms for simplicity, one could also consider general linear functions.

MVTree.Open(prk, aux, f, \mathbf{y}) $\rightarrow \pi$:

- Let $f(\mathbf{v}_1, \ldots, \mathbf{v}_{\ell^{\nu}}) = f'(\mathbf{v}_i)$ for $f' \in \mathcal{F}_{p,k}$ and $i = (\boldsymbol{\sigma})_{\ell}$ in ℓ -ary.
- Consider $\boldsymbol{\tau}, \mathbf{r}$ as formal variables $\mathbf{X} = (X_{\nu}, \dots, X_1), \mathbf{R} = (R_k, \dots, R_1).$
- Denote $p_{\nu+1}(\mathbf{X}, \mathbf{R}) = p(\mathbf{X}, \mathbf{R}) = (\boldsymbol{\lambda}(\mathbf{X}) \otimes \mathbf{R}) \cdot \mathbf{v}$
- For all $\nu \ge j \ge 1$:

Compute $p_j(X_{j-1},\ldots,X_1,\mathbf{R}) = \lambda(X_{j-1},\ldots,X_1,\mathbf{R}) \cdot \mathbf{v}_{\sigma_{|j|}}$ Compute $H_i(X_i,\ldots,X_1,\mathbf{R})$ as

$$H_j(X_j, \dots, X_1, \mathbf{R}) = \frac{p_{j+1}(X_j, \dots, X_1, \mathbf{R}) - p_j(X_{j-1}, \dots, X_1, \mathbf{R})}{(X_j - \sigma_j)}$$

and group element $[H_j]_1 = [H_j(\tau_j, \dots, \tau_1, \mathbf{r})]_1$ - Compute $\hat{\mathsf{C}}_{\boldsymbol{\sigma}} = [\tau_{\nu}^{\ell-1} \cdots \tau_1^{\ell-1} \cdot \mathbf{r}]_1 \cdot \mathbf{v}_{\boldsymbol{\sigma}}, \ \mathbf{\pi}' \leftarrow \mathsf{LVC'}.\mathsf{Open}(\mathsf{prk}', \mathsf{aux}_{\boldsymbol{\sigma}}, f', \mathbf{y})$ and output $\pi = ([H_{\nu}]_1, \ldots, [H_1]_1, \mathsf{C}_{\boldsymbol{\sigma}}, \hat{\mathsf{C}}_{\boldsymbol{\sigma}}, \pi').$

MVTree.Vf(vrk, C, $f, \mathbf{y}, \pi) \rightarrow 0/1$:

- Let $f(\mathbf{v}_1, \dots, \mathbf{v}_{\ell^{\nu}}) = f'(\mathbf{v}_i)$ for $f' \in \mathcal{F}_{p,k}$ and $i = (\boldsymbol{\sigma})_{\ell}$ in ℓ -ary. $b_{\mathsf{Path}} \leftarrow e(\mathsf{C} \mathsf{C}_{\boldsymbol{\sigma}}, [1]_2) = \sum_{j=1}^{\nu} e([H_j]_1, [\tau_j \sigma_j]_2)$ $b_{\mathsf{LD-Test}} \leftarrow e(\mathsf{C}_{\boldsymbol{\sigma}}, [\tau_{\nu}^{\ell-1} \cdots \tau_1^{\ell-1}]_2) = e(\hat{\mathsf{C}}_{\boldsymbol{\sigma}}, [1]_2)$ $b_{\mathsf{Leaf}} \leftarrow \mathsf{LVC'}.\mathsf{Vf}(\mathsf{vrk'}, \mathsf{C}_{\boldsymbol{\sigma}}, f', \mathbf{y}, \pi')$

- Output $b_{\mathsf{Path}} \wedge b_{\mathsf{LD-Test}} \wedge b_{\mathsf{Leaf}}$

We omit explicitly describing the update algorithm. Instead, we demonstrate in Thm. 8 how to efficiently update all proofs after modifying a position in the committed vector.

We summarize the properties of the construction in the following theorems. Due to space limitations, we omit their proofs and refer the interested reader to the full version of this paper.

Theorem 7. Let LVC' be an algebraic vector commitment scheme that satisfies completeness, homomorphic openings and weak function binding for a function family $\mathcal{F}_{p,k}$. Then, MVTree satisfies (1) completeness, (2) Homomorphic Openings and (3) strong function binding for $\operatorname{Ext}_{\ell^{\nu}}$ - $\mathcal{F}_{p,k}$ in the AGM under the $(\ell - 1) \cdot \nu$ -BSDH assumption([2])

Theorem 8. Consider construction MVTree and let $\pi^{\sigma} = ([H_{\nu}^{\sigma}]_1, \ldots, [H_1^{\sigma}]_1, \ldots, [H_1^{\sigma}]_1)$ $C_{\sigma}, C_{\sigma}, \pi'_{\sigma})$ be some proof of opening for a leaf commitment in position σ written in ℓ -ary.

Then, computing all partial proofs $\left\{ ([H_{\nu}^{\sigma}]_2, \ldots, [H_1^{\sigma}]_1, \mathsf{C}_{\sigma}, \hat{\mathsf{C}}_{\sigma}) \right\}_{\sigma \in \Sigma^{\nu}}$ can be done in $O_{\lambda}(k \cdot \nu \cdot \ell^{\nu}) = O_{\lambda}(\nu \cdot m)$ time and storing them needs $O_{\lambda}(\ell^{\nu}) = O_{\lambda}(m/k)$ space. Furthermore, if we update C by adding δ in some position i^* , we can update all partial proofs in time $O_{\lambda}(\nu)$.

Efficiency of the Multivariate Construction. We only consider the case where $\ell = O(1)$. First, let's focus on the time needed to compute $[H_i]_1$. One can simply write the polynomial $p_j - p_{j-1}$ as a polynomial in $1, X_j, \ldots, X_j^{\ell-1}$ with polynomial coefficients in the other variables. Then, we can use standard

(univariate) polynomial division to divide each term with $X_j - \sigma_j$ in constant time. To encode it in the group, it is enough to note that the total degree of each term is $k \cdot \ell^{j-1}$, so we need to perform ℓ multi-exponentiations of this size totaling in $O(k \cdot \ell^j)$ operations.

That said, we demonstrate the efficiency of the construction. The commitment key consists of linear in m group elements. Opening needs $O(k \cdot \ell^j)$ operations for each iteration, totaling in $O(k \cdot \ell^{\nu})$ time. By inspection of the construction, proofs size is $\log_{\ell}(m/k) + 2 + |\pi'|$, where π' is the size of an opening of the leaf commitment. Finally, verification consists of (1) a $\log_{\ell}(m/k)$ -size pairing product equation, (2) a low degree test involving constant operations and (3) a verification of an opening of a leaf commitment.

Remark 3 (On aggregation). The first two verification tests are pairing product equations. Assuming the leaf commitment verification is also a pairing product equation, one can use inner pairing products [4] to aggregate many such equations as done in [22] and, thus, achieve one-hop cross commitment aggregation. While the aggregated proof size decreases exponentially, this comes at the cost of a significant overhead for the prover due to the need to work in the target group. Reducing the proof size from $\log_2 m$ to roughly $\log_{\ell}(m/k)$ (assuming constant size/verification for leaf commitment opening) can make aggregation significantly cheaper for the prover.

7.2 Univariate Maintainable Vector Commitments

In this section, we give an optimized construction that achieves the same memorytime tradeoffs for the prover that the scheme in Section 7.1, but for univariate polynomials. For that, we rely on the q-BSDH assumption for q = m - 1 ([2]), while we only needed $q = \log m$ plus the assumption that the leaf commitment is sound in the multivariate case.

Our work generalizes a previous univariate construction of [24] in a similar way as it generalizes hyperproofs. Namely, our construction truncates the tree at some level so that leaves are commitments and not individual positions.

For vectors of size m, we offer the following trade-off: for any ν, κ , such that $m = 2^{\nu+\kappa+1}$, one can derive openings of size $\nu + 5$ group elements. The prover can pre-compute and store $2^{\nu} - 1$ proofs, and then compute proofs by performing $O(\kappa 2^{\kappa})$ group operations. We show also how to compute all proofs with $O(\nu m)$ group operations (plus $O(m(\nu + \kappa))$) field operations). The proofs are maintainable, as an update in a position requires recomputing $O(\nu)$ proofs. One interesting feature is that the trusted setup depends only on m (the powers of τ) and not on ν, κ , so the right tradeoff can be decided on the fly.

Overview. Our construction builds a tree of commitments to a vector $\mathbf{v} \in \mathbb{F}^m$ build as follows. The root of the tree is a commitment $C = [\boldsymbol{\lambda}]_1 \mathbf{v}$, where $\boldsymbol{\lambda} = ([\lambda_1(\tau)]_1, \ldots, [\lambda_m(\tau)]_1)$, for $\{\lambda_j(X)\}$ the Lagrange interpolation polynomials for \mathbb{H} . The two children will be $C_0 = [\boldsymbol{\lambda}_0]_1 \mathbf{v}_0$ and $C_1 = [\boldsymbol{\lambda}_1]_1 \mathbf{v}_1$, which are commitments to \mathbf{v}_0 and \mathbf{v}_1 with keys $\boldsymbol{\lambda}_0$ and $\boldsymbol{\lambda}_1$ of half the size to be specified next. The two children of C_0 will be $C_{00} = [\lambda_{10}]_1 \mathbf{v}_{10}$, $C_{10} = [\lambda_{10}]_1 \mathbf{v}_{10}$ and so on. The leaves are commitments $C_{\mathbf{b}}$, $\mathbf{b} = (b_{\nu}, \ldots, b_0) \in \{0, 1\}^{\nu+1}$ to vectors of size 2^{κ} . For any leaf index $\mathbf{b} = (b_{\nu}, \ldots, b_0)$, we denote $\mathbf{b}_{|j} = (b_j \ldots b_0)$ the suffix¹³ of size j. Note that $C_{\mathbf{b}_{|j}}$ for $j = 0, \ldots, \nu - 1$ denotes all the commitments from the root to the leaf $C_{\mathbf{b}}$.

The division into vectors of half the size is done in bit reverse order according to the least significant bit of the binary representation of the index, b_0 . At the first level, there will be two vectors $\mathbf{v}_0, \mathbf{v}_1$ of size m/2 containing all positions of \mathbf{v} with suffix b_0 . At the next level, there will be four vectors $\mathbf{v}_{00}, \mathbf{v}_{01}, \mathbf{v}_{10}, \mathbf{v}_{11}$ of size m/4, and $\mathbf{v}_{b_1b_0}$ indicates all the positions of \mathbf{v} (in the natural order) that have as suffix b_1b_0 and so on.

The division into commitment keys of half the size will follow a similar pattern. At level 1, the group of roots of unity \mathbb{H} will be split into \mathbb{H}^0 and \mathbb{H}^1 , according to the least significant bit of the binary representation of the index of the root, i.e. \mathbb{H}^0 consists of all even and \mathbb{H}^1 all odd powers of ω . In particular, \mathbb{H}^0 are the roots of unity of size m/2, and $\mathbb{H}^1 = \omega \mathbb{H}^0$ is a coset. At level 2, the commitment keys will be associated to \mathbb{H}^{00} , \mathbb{H}^{01} , \mathbb{H}^{10} , \mathbb{H}^{11} and by the same reasoning, \mathbb{H}^{00} are the roots of unity of size m/4, $\mathbb{H}^{10} = \omega^2 \mathbb{H}^{00}$, $\mathbb{H}^{01} = \omega \mathbb{H}^{00}$ and $\mathbb{H}^{11} = \omega^3 \mathbb{H}^{00}$. More generally, we note that for any $0 \le j \le \nu$ and any string $(b_j, \ldots, b_0) \in \{0, 1\}^{j+1}$, $\mathbb{H}^{\mathbf{b}_{|j}} = \omega^s \mathbb{H}_r$, for $s = \sum_{i=0}^j b_i 2^i$ and $r = \frac{m}{2^{j+1}}$. The vanishing polynomial associated to $\mathbb{H}^{\mathbf{b}_{|j}}$ will be denoted $t_{\mathbf{b}_{|j}}(X) = X^r - (\omega^s)^r = X^{\frac{m}{2^{j+1}}} - \omega^{\frac{m \sum_{i=0}^j b_i 2^i}{2^{j+1}}}$. The Lagrange polynomials associated to the interpolation set $\mathbb{H}^{\mathbf{b}_{|j}}$ with the natural order will be written as $\lambda^{\mathbf{b}_{|j}}(X) = (\lambda_1^{\mathbf{b}_{|j}}(X), \ldots, \lambda_r^{\mathbf{b}_{|j}}(X))$ and the commitment key for node $\mathbf{b}_{|j}$ is $\lambda^{\mathbf{b}_{|j}} = [\lambda^{\mathbf{b}_{|j}}(\tau)]_1$.

As in the multivariate case, the idea to open the commitment to some function f that is a linear function of some chunk \mathbf{v}_i is to (1) open the root commitment to the leaf and (2) open the commitment to the leaf using the IP argument for the Lagrange basis of Section 5 or the construction of Tomescu et al. [23]. For (2), since at the leaf level the commitment is w.r.t to the key $\lambda^{\mathbf{b}}$ for some $\mathbf{b} = (b_{\nu}, \ldots, b_0)$, we prove the following lemma, that shows that the construction for inner products of Section 5 works for any coset of roots of unity.

Theorem 9. Let $\mathbb{H} \subset \mathbb{F}$ be a subset of roots of unity of size $m = 2^{\nu+\kappa+1}$, for some $\kappa, \nu \geq 0$. Given some $\mathbf{b} \in \{0,1\}^{\nu+1}$, define $s = \sum_{i=0}^{\nu} b_i 2^i$, $r = 2^{\kappa}$, $\mathbb{H}_r \subset \mathbb{H}$ the subgroup of roots of unity of size r, and $\mathbb{H}^{\mathbf{b}} = \omega^s \mathbb{H}_r$. Let $t_{\mathbf{b}}(X)$ be the vanishing polynomial at $\mathbb{H}^{\mathbf{b}}$ and $\lambda^{\mathbf{b}}(X)$ the associated Lagrange basis polynomials. Then, if $A(X) = \lambda^{\mathbf{b}}(X) \cdot \mathbf{a}$ and $B(X) = \lambda^{\mathbf{b}}(X) \cdot \mathbf{b}$, it holds that $\mathbf{a} \cdot \mathbf{b} = y$ if and only if there exist polynomials H(X), R(X) with $\deg(R) < r - 2$ such that

$$A(X)B(X) - r^{-1}y = XR(X) + t_{\mathbf{b}}(X)H(X).$$

¹³ Note that this notation is different than the one we used in the multivariate case. In the latter case, this notation denoted prefixes while here it denotes suffixes. We do this because in each case the corresponding notation makes presentation easier.

27

Therefore, at any leaf **b** we can open the commitment to any linear relation and verify with the same equation. To open C to a certain leaf commitment C_i , the idea is to implicitly show from root to leaf that $C_{\mathbf{b}_{|j|}}$, $C_{\mathbf{b}_{|j+1}}$ agree in $\mathbb{H}^{\mathbf{b}_{|j+1}}$. This is proven by showing that their difference is divisible by $t_{(1-b_{j+1})\mathbf{b}_{|j|}}(X)$. More specifically, we prove the following lemma, that shows how the parent and the children nodes ate each level relate through a simple equation:

Lemma 1. Consider two cosets $\mathbb{H}_{0\mathbf{b}_{|j}}$ and $\mathbb{H}_{1\mathbf{b}_{|j}}$. Let $C_{\mathbf{b}_{|j}}(X)$ be an encoding of vector $\mathbf{v}^{\mathbf{b}_{|j}}$, and $C_{0\mathbf{b}_{|j}}(X)$, $C_{1\mathbf{b}_{|j}}(X)$ those of vectors $\mathbf{v}^{0\mathbf{b}_{|j}}$ and $\mathbf{v}^{1\mathbf{b}_{|j}}$. For all $j = 0, \ldots, \nu$ it is true that

$$C_{\mathbf{b}_{|j}}(X) = t_{1\mathbf{b}_{|j}}(X) \frac{C_{0\mathbf{b}_{|j}}(X) - C_{1\mathbf{b}_{|j}}(X)}{2\omega^{sj}} + C_{1\mathbf{b}_{|j}}(X)$$
$$C_{\mathbf{b}_{|j}}(X) = t_{0\mathbf{b}_{|j}}(X) \frac{C_{0\mathbf{b}_{|j}}(X) - C_{1\mathbf{b}_{|j}}(X)}{-2\omega^{sj}} + C_{0\mathbf{b}_{|j}}(X)$$

Scheme Description. Formally, we present an LVC commitment scheme that works for the function family:

$$\mathsf{Ext}_{\nu} - \mathcal{F}_{p, 2^{\kappa}} = \{ f : \mathbb{F}^m \to \mathbb{F}, m = 2^{\kappa + \nu + 1} \mid \exists \mathbf{f} \in \mathbb{F}^{2^{\kappa}}, i \in 2^{\nu} \text{ s.t.} \\ \forall \mathbf{v}_1, \dots, \mathbf{v}_{2^{\nu}} \in \mathbb{F}^{2^{\kappa}} : f(\mathbf{v}_1, \dots, \mathbf{v}_{2^{\nu}}) = \mathbf{v}_i \cdot \mathbf{f} \}.$$

Algorithms LVC.KeyGen and LVC.Commit are the same as the Lagrange basis construction of Section 5 and are omitted. The commitment to v is $C = [\lambda^{\top}]_1 v$ together with the auxiliary input information aux. Note that step 4. of the open algorithm is IP.Open from Section 5.2.

UVTree.Open(pk, b, aux, f, y) $\rightarrow \pi$: 1. Let $f(\mathbf{v}_{0...0}, \ldots, \mathbf{v}_{1...1}) = \mathbf{v}_{\mathbf{b}} \cdot \mathbf{f}$ for $\mathbf{f} \in$ $\mathbb{F}^{2^{\kappa}}$ and some $\mathbf{b} = (b_{\nu}, \ldots, b_0).$

- 2. For any $0 \le j \le \nu$, compute $C_{\mathbf{b}_{|j}} = [\lambda^{\mathbf{b}_{|j}}]_1 \mathbf{v}_{\mathbf{b}_{|j}}$. 3. Compute $[H]_1 = (C_0 C_1)/2$, and for any $0 \le j \le \nu 1$, compute $K_{b_{|j}} = (-1)^j (2\omega^{s_j})^{-1}$. Then define $[H_{\mathbf{b}_{|j}}]_1 = K_{b_{|j}} (C_{0b_j...b_0} C_{1b_j...b_0})$.
- 4. Find $R(X), H_{\mathbf{b}}(X)$ such that if

$$\left(\boldsymbol{\lambda}^{\mathbf{b}}(X) \ \mathbf{v}_{\mathbf{b}}\right) \left(\sum_{i=1}^{2^{\kappa}} f_i \lambda_i^{\mathbf{b}}(X)\right) - y 2^{-\kappa} = X R(X) + H_{\mathbf{b}}(X) t_{\mathbf{b}}(X).$$

Define $\hat{R}(X) = X^{m+1-2^{\kappa}} R(X)$.¹⁴ Define $\hat{\mathsf{C}}_{\mathbf{b}} = \tau^{m-2^{\kappa}} \mathsf{C}_{\mathbf{b}}$.

5. Output $\pi = ([H_{b_0}]_1, \dots, [H_{\mathbf{b}|_{\nu-1}}]_1, [H_{\mathbf{b}}(\tau)]_1, [R(\tau)]_1, [\hat{R}(\tau)]_1, \mathsf{C}_{\mathbf{b}}, \hat{\mathsf{C}}_{\mathbf{b}}).$ UVTree.Vf(vk, C, $f, \mathbf{y}, \pi) \rightarrow 0/1$:

- 1. Use the vector representation **f** of f and compute $C_f = \sum_{i=1}^{2^{\kappa}} f_i[\lambda_i^{\mathbf{b}}(\tau)]_2$.
- 2. Check that

$$e(\mathsf{C} - \mathsf{C}_{\mathbf{b}}, 1) = e([H]_1, [t_{b_0}(\tau)]_2) + \sum_{j=0}^{\nu-1} e([H_{\mathbf{b}_{|j}}]_1, [t_{\mathbf{b}_{|j+1}}(\tau)]_2) \quad (1)$$

¹⁴ We assume as in Section 5 that at most m-1 powers of τ are in the SRS in group \mathbb{G}_1 . The degree check is meant to ensure that R(X) is of degree at most $2^{\kappa} - 2$.

M. Campanelli, A. Nitulescu, C. Ràfols, A. Zacharakis, A. Zapico

28

$$e(\mathsf{C}_{\mathbf{b}},\mathsf{C}_{f}) - e(m^{-1}y[1]_{1},[1]_{2}) = e([R]_{1},[1]_{2}) + e([H_{\mathbf{b}}]_{1},[t_{\mathbf{b}}(\tau)]_{2})$$
(2)

$$e([R]_1, [\tau^{m+1-2^{\kappa}}]_2) = e([\hat{R}]_1, [1]_2)$$
(3)

$$e(\mathsf{C}_{\mathbf{b}}, [\tau^{m-2^{\kappa}}]_2) = e(\widehat{\mathsf{C}}_{\mathbf{b}}, [1]_2) \tag{4}$$

Maintainability. The cost of computing all proofs is $O(\nu m)$. For each piece \mathbf{v}_i with $O(\kappa 2^{\kappa})$ operations one can compute the coefficients in the monomial basis. Following Lemma 1, the parent node can be computed in cost dominated by $2^{\kappa} = \frac{m}{2^{\nu+1}}$ exponentiations from the expression of children nodes, and since there are 2^{ν} parent nodes the cost is dominated by $\frac{m}{2}$ exponentiations. Going one level up, the vector size doubles but the number of nodes is halved. We conclude that to compute all proofs one needs $O(\kappa 2^{\kappa} + \nu \frac{m}{2})$. The number of proofs to store (including leaf commitments) is $2^{\nu+1} - 1$.

Theorem 10. When instantiated with a function binding argument for inner product relations IP, the scheme above is a function binding LVC argument under the AGM if the (m - 1, m)-DLOG assumption holds.

Proof. Let \mathcal{A} be an adversary against the function binding game as in Definition 3. We will see, through game reductions, that the advantage of \mathcal{A} in strong function binding is negligible even for k = 2, that is, for two non-compatible functions f_1, f_2 . Note that for two functions to be non compatible they must be defined on the same block **b**.

 \mathcal{A} plays Game_0 , the strong function binding game as in Definition 3, and outputs $(\mathsf{C}, \{f_k, y_k, \pi_k\}_{k=1,2})$, where $\pi_1 = (\{[H_{\mathbf{b}|\mathbf{j}}]_1, \}_{j=0}^{\nu-1}, [H_{\mathbf{b}}]_1, [R]_1, [R]_1, [\hat{R}]_1, \mathsf{C}_{\mathbf{b}}, \hat{\mathsf{C}}_{\mathbf{b}}), \pi_2 = (\{[H'_{\mathbf{b}|\mathbf{j}}]_1\}_{j=0}^{\nu-1}, [H'_{\mathbf{b}}]_1, [R']_1, [R']_1, [\hat{R}']_1, \mathsf{C}'_{\mathbf{b}}, \hat{\mathsf{C}}'_{\mathbf{b}}), \text{ s.t. LVC.Verify}(\mathsf{vk}, \mathsf{C}, f_1, y_1, \pi_1) = 1, \mathsf{LVC.Verify}(\mathsf{vk}, \mathsf{C}, f_2, y_2, \pi_2) = 1, and wins if there exists no <math>\mathbf{v} \in \mathbb{F}^m$ such that $f_1(\mathbf{v}) = y_1$ and $f_2(\mathbf{v}) = y_2$.

Recall \mathcal{A} is algebraic and thus we assume one can extract polynomials $C_{\mathbf{b}}(X)$, $C'_{\mathbf{b}}(X)$, $\hat{C}_{\mathbf{b}}(X)$, $\hat{H}_{\mathbf{b}|_{j}}(X)$, $H'_{\mathbf{b}|_{j}}(X)$, H'_{\mathbf

Let Game_1 be exactly as Game_0 but the game aborts if $C_{\mathbf{b}}(X)$ or $C'_{\mathbf{b}}(X)$ are polynomials of degree more than $2^{\kappa} - 1$. If this is not the case, it is easy to find τ by observing that in this case either $C_{\mathbf{b}}(X)X^{m-2^{\kappa}} - \hat{C}(X)$ or $C'_{\mathbf{b}}(X)X^{m-2^{\kappa}} - \hat{C}'(X)$ is a non-zero polynomial with a root in τ so the difference between both games is bounded by the advantage of any adversary against the (m-1,m)-DLOG problem.

Let Game_2 be exactly as Game_1 but upon receiving π_1, π_2 , it checks if $\mathsf{C}_{\mathbf{b}}$ and $\mathsf{C}'_{\mathbf{b}}$ are equal and aborts otherwise. We next bound the probability of abort.

Define the polynomial $p(X) = C_{\mathbf{b}}(X) - C'_{\mathbf{b}}(X) - (H(X) - H'(X))t_{b_0}(X) + \sum_{j=0}^{\nu-1} (H_{\mathbf{b}_{|j|}}(X) - H_{\mathbf{b}_{|j+1}}(X))t_{\mathbf{b}_{|j|}}(X)$, which is the difference of verification equation (1) for each commitment. If $p(X) \neq 0$, the output of the adversary can be used to construct an adversary against the (m-1,m)-DLOG assumption, since τ is a root of p(X). On the other hand, if p(X) = 0, $C_{\mathbf{b}}(X) - C'_{\mathbf{b}}(X)$ can

be written as a sum of terms that are multiples of $t_{\mathbf{b}_{|j}}(X)$ for $j = 0, \ldots, \nu$. But all of these vanishing polynomials evaluate to 0 in $\mathbf{h} \in \mathbb{H}^{\mathbf{b}}$, since $t_{\mathbf{b}}(X)|t_{\mathbf{b}_{|j}}(X)$ for $j = 0, \ldots, \nu$. Therefore, $C_{\mathbf{b}}(X) - C'_{\mathbf{b}}(X)$ is also 0 when evaluated at the coset. But since this polynomial is of degree at most 2^k , $C_{\mathbf{b}}(X) = C'_{\mathbf{b}}(X)$ which implies that necessarily $C_{\mathbf{b}} = C'_{\mathbf{b}}$.

Therefore, in $Game_2$, except with negligible probability the leaf commitment is the same and the probability that the adversary wins is the same as in the strong function binding game of the inner product commitment.

References

- E. Ben-Sasson, A. Chiesa, M. Riabzev, N. Spooner, M. Virza, and N. P. Ward. Aurora: Transparent succinct arguments for R1CS. In Y. Ishai and V. Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 103–128. Springer, Heidelberg, May 2019. 16
- 2. D. Boneh and X. Boyen. Efficient selective identity-based encryption without random oracles. *Journal of Cryptology*, 24(4):659–693, Oct. 2011. 17, 18, 24, 25
- D. Boneh, B. Bünz, and B. Fisch. Batching techniques for accumulators with applications to IOPs and stateless blockchains. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 561–586. Springer, Heidelberg, Aug. 2019. 1, 2, 6, 7, 11
- 4. B. Bünz, M. Maller, P. Mishra, N. Tyagi, and P. Vesely. Proofs for inner pairing products and applications. In M. Tibouchi and H. Wang, editors, Advances in Cryptology - ASIACRYPT 2021 - 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6-10, 2021, Proceedings, Part III, volume 13092 of Lecture Notes in Computer Science, pages 65–97. Springer, 2021. 5, 6, 8, 25
- M. Campanelli, D. Fiore, N. Greco, D. Kolonelos, and L. Nizzardo. Incrementally aggregatable vector commitments and applications to verifiable decentralized storage. In S. Moriai and H. Wang, editors, ASIACRYPT 2020, Part II, volume 12492 of LNCS, pages 3–35. Springer, Heidelberg, Dec. 2020. 2, 4, 6, 7, 11
- M. Campanelli, A. Nitulescu, C. Ràfols, A. Zacharakis, and A. Zapico. Linear-map vector commitments and their practical applications. Cryptology ePrint Archive, Paper 2022/705, 2022. https://eprint.iacr.org/2022/705. 18
- D. Catalano and D. Fiore. Vector commitments and their applications. In K. Kurosawa and G. Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 55–72. Springer, Heidelberg, Feb. / Mar. 2013. 1, 6
- A. Chepurnoy, C. Papamanthou, and Y. Zhang. Edrax: A cryptocurrency with stateless transaction validation. Cryptology ePrint Archive, Report 2018/968, 2018. https://eprint.iacr.org/2018/968. 2
- V. Daza, C. Ràfols, and A. Zacharakis. Updateable inner product argument with logarithmic verifier and applications. In A. Kiayias, M. Kohlweiss, P. Wallden, and V. Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 527–557. Springer, Heidelberg, May 2020. 5
- 10. Filecoin. Filecoin powers of tau ceremony attestations, 2020. https://github.com/arielgabizon/perpetualpowersoftau. 6
- B. Fisch. PoReps: Proofs of space on useful data. Cryptology ePrint Archive, Report 2018/678, 2018. https://eprint.iacr.org/2018/678. 3

- 30 M. Campanelli, A. Nitulescu, C. Ràfols, A. Zacharakis, A. Zapico
- G. Fuchsbauer, E. Kiltz, and J. Loss. The algebraic group model and its applications. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Heidelberg, Aug. 2018. 7, 8
- S. Gorbunov, L. Reyzin, H. Wee, and Z. Zhang. Pointproofs: Aggregating proofs for multiple vector commitments. In J. Ligatti, X. Ou, J. Katz, and G. Vigna, editors, ACM CCS 2020, pages 2007–2023. ACM Press, Nov. 2020. 2, 6, 7, 10
- J. Groth, M. Kohlweiss, M. Maller, S. Meiklejohn, and I. Miers. Updatable and universal common reference strings with applications to zk-SNARKs. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 698–728. Springer, Heidelberg, Aug. 2018. 3
- A. Kate, G. M. Zaverucha, and I. Goldberg. Constant-size commitments to polynomials and their applications. In M. Abe, editor, ASIACRYPT 2010, volume 6477 of LNCS, pages 177–194. Springer, Heidelberg, Dec. 2010. 22
- R. W. F. Lai and G. Malavolta. Subvector commitments with application to succinct arguments. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019*, *Part I*, volume 11692 of *LNCS*, pages 530–560. Springer, Heidelberg, Aug. 2019. 1, 2, 4, 7, 8
- B. Libert, S. C. Ramanna, and M. Yung. Functional commitment schemes: From polynomial commitments to pairing-based accumulators from simple assumptions. In I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, and D. Sangiorgi, editors, *ICALP 2016*, volume 55 of *LIPIcs*, pages 30:1–30:14. Schloss Dagstuhl, July 2016. 1, 7
- B. Libert and M. Yung. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In D. Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 499–517. Springer, Heidelberg, Feb. 2010. 1, 7
- R. C. Merkle. A digital signature based on a conventional encryption function. In C. Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, pages 369–378. Springer, Heidelberg, Aug. 1988. 2
- C. Papamanthou, E. Shi, and R. Tamassia. Signatures of correct computation. In A. Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 222–242. Springer, Heidelberg, Mar. 2013. 7, 22
- 21. C. Ràfols and A. Zapico. An algebraic framework for universal and updatable SNARKs. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 774–804, Virtual Event, Aug. 2021. Springer, Heidelberg. 16, 18
- 22. S. Srinivasan, A. Chepurnoy, C. Papamanthou, A. Tomescu, and Y. Zhang. Hyperproofs: Aggregating and maintaining proofs in vector commitments. In 31st USENIX Security Symposium (USENIX Security 22), Boston, MA, Aug. 2022. USENIX Association. 5, 6, 7, 21, 22, 25
- A. Tomescu, I. Abraham, V. Buterin, J. Drake, D. Feist, and D. Khovratovich. Aggregatable subvector commitments for stateless cryptocurrencies. In C. Galdi and V. Kolesnikov, editors, *SCN 20*, volume 12238 of *LNCS*, pages 45–64. Springer, Heidelberg, Sept. 2020. 5, 6, 7, 10, 19, 21, 26
- A. Tomescu, R. Chen, Y. Zheng, I. Abraham, B. Pinkas, G. Golan-Gueta, and S. Devadas. Towards scalable threshold cryptosystems. In 2020 IEEE Symposium on Security and Privacy, pages 877–893. IEEE Computer Society Press, May 2020. 5, 8, 21, 25