Non-Malleable Vector Commitments via Local Equivocability

Lior Rotem[⋆,⋆⋆] and Gil Segev[⋆]

School of Computer Science and Engineering,
Hebrew University of Jerusalem, Jerusalem 91904, Israel.
{lior.rotem,segev}@cs.huji.ac.il

Abstract. Vector commitments (VCs), enabling to commit to a vector and locally reveal any of its entries, play a key role in a variety of both classic and recently-evolving applications. However, security notions for VCs have so far focused on passive attacks, and non-malleability notions considering active attacks have not been explored. Moreover, existing frameworks that may enable to capture the non-malleability of VCs seem either too weak (non-malleable non-interactive commitments that do not account for the security implications of local openings) or too strong (non-malleable zero-knowledge sets that support both membership and non-membership proofs).

We put forward a rigorous framework capturing the non-malleability of VCs, striking a careful balance between the existing weaker and stronger frameworks: We strengthen the framework of non-malleable non-interactive commitments by considering attackers that may be exposed to local openings, and we relax the framework of non-malleable zero-knowledge sets by focusing on membership proofs. In addition, we strengthen both frameworks by supporting (inherently-private) updates to entries of committed vectors, and discuss the benefits of non-malleable VCs in the context of both UTXO-based and account-based stateless blockchains, and in the context of simultaneous multi-round auctions (that have been adopted by the US Federal Communications Commission as the standard auction format for selling spectrum ranges).

Within our framework we present a direct approach for constructing non-malleable VCs whose efficiency essentially matches that of the existing standard VCs. Specifically, we show that any VC can be transformed into a non-malleable one, relying on a new primitive that we put forth. Our new primitive, *locally-equivocable commitments with all-but-one binding*, is evidently both conceptually and technically simpler compared to multi-trapdoor mercurial trapdoor commitments (the main building block underlying existing non-malleable zero-knowledge sets), and admits more efficient instantiations based on the same number-theoretic assumptions.

⋆ Supported by the European Union's Horizon 2020 Framework Program (H2020) via an ERC Grant (Grant No. 714253).

⋆⋆ Supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

1 Introduction

Vector commitments (VCs) [LY10, CF13] enable to non-interactively commit to a vector $(x_1, \ldots, x_q)$ while offering the useful property of *local opening*: The committer can reveal any individual entry x_i without the overhead of revealing the entire vector. At the same time, VCs are also required to be *position binding*: The committer should not be able to reveal any entry of an even maliciously-committed vector to more than a single value.

The main measure of efficiency for VCs, which makes them extremely useful for a variety of applications but highly non-trivial to construct, is their succinctness: Both the size of the commitment and the size of the local openings should be sublinear in the number q of elements in the committed vector. Whereas the classic notion of a Merkle tree [Mer87] can be seen as a VC in which the size of the commitment is independent of q and the size of local openings scales logarithmically with q, Libert and Yung [LY10] and Catalano and Fiore [CF13] presented constructions in which both sizes are independent of q.

Starting already with Merkle's early work, VCs consistently play a key role in a wide range of applications as a communication-efficient method for authenticating rather large amounts of data by allowing users to retrieve small parts of the data alongside short proofs of authenticity. Such applications include, for example, verifiable databases and authenticated data structures (e.g., [NN98, MND$^+$04, BGV11, SvDJ$^+$12, KSS$^+$16, CFG$^+$20]), zero-knowledge sets (e.g., [MRK03, LY10, CRF$^+$11, CHL$^+$13]), cryptographic accumulators [BdM93] (which have many applications on their own right – see for example [BP97, GR97, CL02, DKN$^+$04, Ngu05, ABC$^+$12, MGG$^+$13, FVY14] and the references therein), stateless blockchains (e.g., [STS99, Tod16, But17, BBF19, TAB$^+$20]), and succinct arguments (e.g., [Kil92, Mic94, BBF19, LM19, OWW$^+$20]).

Non-malleable commitments. Another long line of research regarding commitment schemes, initiated by the seminal work of Dolev, Dwork and Naor [DDN00], deals with the construction of *non-malleable* commitments. Roughly speaking, a commitment scheme is non-malleable if an adversary which receives a commitment to some value x, cannot produce a commitment to some "non-trivially related" value x'. Non-malleable commitments have established themselves as instrumental in a host of cryptographic tasks, especially those requiring to protect against man-in-the-middle attacks. Numerous constructions of non-malleable commitments have been suggested over the years, satisfying various flavors of security notions and achieving different efficiency tradeoffs, based on wide range of cryptographic assumptions (e.g., [CIO98, DDN00, FF00, CKO$^+$01, Bar01, CF01, PR05, PR08, PPV08, LPV08, LP09, PW10, Wee10, LP11, GLO$^+$12, GPR16, COS$^+$17, Khu17] and the many references therein).

This work: Non-malleable vector commitments. The fundamental importance of VCs and of non-malleable commitments motivates the study of non-malleable VCs with the premise of significantly strengthening the security

and improving the efficiency of the wide range of applications in which they play a key role. For example, non-malleable VCs would directly give rise to verifiable databases, authenticated data structures and cryptographic accumulators offering non-malleability guarantees. As additional, less direct examples, in Section 1.2 we discuss the benefits of using non-malleable VCs as building blocks in the contexts of stateless blockchains and simultaneous multi-round auctions.

However, the notion of non-malleable VCs has not yet been explored, and the existing framework and constructions of standard non-malleable commitments do not take into account the significant security implications of local openings. A closely-related notion, which has been thoroughly explored, is that of non-malleable zero-knowledge sets (ZKS), introduced in the beautiful work of Gennaro and Micali [GM06] (extending the notion of standard ZKS [MRK03]). Non-malleable ZKS can be seen as a substantial strengthening of non-malleable VCs, supporting *non-membership* proofs in addition to membership proofs. The work of Gennaro and Micali initiated an exciting line of research leading to constructions of non-malleable ZKS based on gradually weaker assumptions and with increasingly better parameters (see [LY10, CF13] and the references therein). However, these constructions rely on the useful yet intricate notion of multi-trapdoor mercurial trapdoor commitments [GM06], specifically tailored to support non-membership proofs (see also [CHL+13, CDV06] for basic background on mercurial commitments). As prominent applications of VCs generally do not require non-membership proofs (as we exemplify in Section 1.2), this raises the following question:

Can non-malleable VCs be constructed within a simplified framework
both **conceptually** *(e.g., simpler and more intuitive notions)*
and **technically** *(e.g., direct and more efficient constructions)?*

1.1 Our Contributions

Notion of non-malleability for VCs. We put forward a strong notion of non-malleability for vector commitment schemes. Our framework strikes a careful balance between the weaker notion of non-malleable non-interactive commitments [CIO98, CKO+01] and the considerably stronger notion of non-malleable zero-knowledge sets [GM06]. Concretely, we generalize the notion of non-malleable non-interactive commitments by incorporating the adversarial adaptivity and additional information resulting from local openings. That is, the key difference from the notion of non-malleable non-interactive commitments is that we aim at achieving non-malleability against adversaries which may have already been exposed to several local openings. Looking ahead, this key difference is the reason that simple attempts of combining VCs and non-malleable commitments do not seem to suffice for realizing our notion (as we demonstrate in Section 3.2).

Warm-up: Merkle trees are non-malleable in the random-oracle model. As a first step within our framework, we examine the non-malleability of existing vector commitments schemes and observe that they are easily malleable (some of

them by design in order to support public updates). Then, as a warm-up towards our main result, we show that a Merkle tree does satisfy our requirements when its underlying hash function is modeled as a random oracle [BR93] (and we show that this does not generally hold in the standard model):

Theorem 1.1 (informal). *Let* H *be a hash function and let* $\mathsf{tree}\mathcal{VC}$ *be the Merkle tree vector commitment scheme that obtained via* H*. Then,* $\mathsf{tree}\mathcal{VC}$ *is a non-malleable vector commitment scheme when* H *is modeled as a random oracle.*

Theorem 1.1 demonstrates the feasibility of realizing our notion of non-malleable vector commitments via a direct construction whose proof is not explicitly based on multi-trapdoor mercurial trapdoor commitments. However, the non-malleability of this construction heavily relies on the random-oracle model and, more importantly, the construction has local openings whose size scales logarithmically with the number q of elements in the committed vector.

Main result: Efficient non-malleable VCs via locally equivocability. We present a direct approach for constructing non-malleable VCs whose efficiency essentially matches that of the existing standard VCs. Inspired by constructions of non-malleable zero-knowledge sets [GM06, LY10, CF13] (and, more generally, of non-malleable cryptographic primitives [DDN00]), we show that any vector commitment scheme can be transformed into a non-malleable one, relying on a new primitive that we put forth. Our new primitive, *locally-equivocable commitments with all-but-one binding*, is evidently both conceptually and technically simpler when compared to multi-trapdoor mercurial trapdoor commitments, as we discuss below. We prove the following theorem:

Theorem 1.2 (informal). *Any vector commitment scheme can be transformed into a non-malleable one using: (1) a locally-equivocable commitment scheme with all-but-one binding, (2) a one-time strongly-unforgeable signature scheme, and (3) a universal one-way hash family.*

We note that our notions of non-malleability and our construction extend to accumulators [BdM93], which can be viewed as VCs for vectors whose length is not necessarily bounded ahead of time. Specifically, in our construction, the underlying VC can be replaced with an accumulator, and the underlying locally-equivocable commitment scheme can be replaced with one that supports an a-priori unbounded number of commitments (this is already the case with our number-theoretic constructions).

Intuitive, simple & efficient: Locally-equivocable commitments with all-but-one binding. Our new notion of commitments is obtained by augmenting the standard notion of tag-based commitments with the following two requirements:

- **Local equivocability:** A committer can generate several equivocal commitments with respect to a single common-reference string.

– **All-but-one binding:** Equivocal commitments generated with respect to a predetermined tag τ should be binding with respect to any other tag even when given the trapdoor associated with τ.

This new notion is evidently both conceptually and technically simpler than the notion of multi-trapdoor mercurial trapdoor commitments. From the conceptual perspective, it has a short and intuitive description. This is evident not only from the above informal description, but also from the fact that in addition to the standard setup, commitment and decommitment procedures, our notion consists of only 3 additional procedures, whereas the notion of a multi-trapdoor mercurial trapdoor commitment consists of 7 additional procedures (already in its non-vector variant) together with a non-trivial number of correctness and security requirements.

From the technical perspective, on the one hand we observe that our new notion strengthens Fischlin's notion of identity-based trapdoor commitments [Fis01, Ch. 2.6]; whereas on the other hand we nevertheless show that Fischlin's highly-efficient number-theoretic constructions satisfy our strengthened notion [Fis01, Ch. 3.3]. Specifically, this yields constructions based on the discrete logarithm assumption and on the RSA assumption, in which commitments consist of a *single* group element. This should be contrasted with the known constructions of multi-trapdoor mercurial trapdoor commitments based on the same assumptions in which commitments consist of *two* group elements. The difference between producing one or two group elements might not be significant on its own, but both in our construction and in those based on multi-trapdoor mercurial trapdoor commitments the underlying commitment scheme is used for producing q commitments (where q is the number of elements in the committed vector), and this translates into a more significant difference between producing q and $2q$ group elements.

In addition to these highly-efficient number-theoretic constructions, we also present a construction based on the existence of any standard commitment scheme (and thus based on the existence of any one-way function [Nao91, HIL$^+$99]). However, this construction is mainly of theoretical significance as it supports only an a-priori bounded of number q of equivocal commitments, and the length of its common-reference string is linear in this bound. Such guarantees still suffice for our non-malleable vector commitment, but lead to somewhat impractical efficiency guarantees.

Extension: Non-malleable dynamic VCs. Catalano and Fiore [CF13] constructed VCs in which individual entries of the committed vector can be updated *publicly* (i.e., without knowledge of the committer's private state). Such public updates, however, are inherently incompatible with the motivation underlying the notion of non-malleability, and indeed with our definition of non-malleable VCs. In light of this inherent limitation, we show that our framework and construction can nevertheless support updates in a *private* manner, requiring knowledge of the private state generated by the committer in order to update entries of the underlying vector.

We extend our definition of non-malleable VCs to support dynamic VCs as well, essentially requiring that non-malleability is maintained even when the adversary receives a vector commitment which has undergone adversarially-chosen updates. We then revisit our construction from Theorem 1.2 and show that if the underlying VC supports private updates,[1] then so does our resulting non-malleable VC (which is indeed non-malleable with respect to our extended definition).

Theorem 1.3 (informal). *Any privately-updatable vector commitment scheme can be transformed into a non-malleable privately-updatable one using: (1) a locally-equivocable commitment scheme with all-but-one binding, (2) a strongly-unforgeable signature scheme, and (3) a universal one-way hash family.*

1.2 Applications

The notion of non-malleable commitments is over three decades old [DDN00], and has found a variety of applications. Since our notion of non-malleable VCs strengthens this notion in the non-interactive setting, it can be applied in any case in which non-interactive non-malleable commitments can be used, while offering significant efficiency improvement via local openings. Specifically, VCs play a key role in a wide range of applications both as an intermediate building block and as a direct communication-efficient method for authenticating large amounts of data (allowing users to retrieve small parts of the data alongside short proofs of their authenticity). Here, we focus our attention on discussing the benefits of non-malleable VCs in the contexts of stateless blockchains and simultaneous multi-round auctions.

Stateless blockchains. VCs are used as a direct communication-efficient method for authenticating large amounts of data in stateless blockchains both in the UTXO model (e.g., Bitcoin [Nak08]) and in the account model (e.g., Ethereum [But14]).[2] In both models, transactions and smart-contracts consist of local opening of VCs, where the VCs represent a compressed version of a current state, and are stored by validating parties. Their local openings are verified either as unspent transactions in the UTXO model, or as account balances and various other user-specific properties in the account model (see for example, [BBF19, GRW+20, BBB+18, TAB+20], for extensive discussions and additional related work – which is far beyond the context of our work).

In such scenarios, the basic security properties of VCs are generally insufficient in order to guarantee cross-transaction independence (also known as transaction non-malleability [BCG+14]). Specifically, in such highly interactive scenarios, attackers may indeed observe both VCs and local openings, then manipulate the VCs to represent a malleated state (e.g., either in an implicitly-malicious manner

[1] Note that a VC which supports public updates trivially supports private updates.

[2] In fact, in some cases, accumulators are used instead of vector commitments. As noted about, our notions of non-malleability and our construction apply also to accumulators.

by issuing honest yet tailored transactions that lead to specific state updates, or in an explicitly-malicious manner by potentially controlling to some extent some of the verifying parties), and then produce local openings with respect to the malleated VCs – as captured by our notion for non-malleable VCs. Thus, relying on non-malleable VCs in the context of stateless blockchains can significantly reduce both storage and communication while guaranteeing cross-transaction independence.

Simultaneous multi-round auctions. One of the most classic and direct applications of (non-malleable) commitments is that of sealed-bid auctions [DDN00], and in this context our notion of non-malleable VCs seems particularly suitable for Simultaneous Multi-Round Auctions (SMRA) [Bic17, Ch. 6]. Such auctions provide a widespread multi-round format for selling multiple items. SMRAs were designed for the US Federal Communications Commission in the early 1990s, and since then they have become the standard auction format for selling spectrum worldwide.

SMRAs proceed in rounds, where in each round some or all bidders bid for multiple items, and each item may either be sold or not sold in each round depending of the specific rules of the auction and the submitted bids. After each round is closed the auctioneer discloses which items were won, who wins each of these items, and at what price. Depending on the specific rules of the auction, there are differences in the level of information revealed about other bidders' bids. In some cases all bids are publicly revealed after each round, whereas in other cases only prices of the currently winning bids are publicly revealed.

From the perspective of using vector commitments, submitting each bidder's bids for all available items in each round using a VC, and then publicly revealing local openings for the required (e.g., winning) bids according to the rules of the auction, can lead to significant communication savings (at least in the case of spectrum ranges, the number of ranges may be rather large – although not as large as in the context of using VCs for stateless blockchains). However, this enables a malicious bidder to malleate vector commitments (i.e., bids) provided in earlier rounds or even in the same round after having seen some of their local openings, and to generate a vector commitment (i.e., a bid) to related values together with corresponding local openings at a later stage – as captured by our notion of non-malleable VCs. Thus, relying on non-malleable VCs in the context of SMRAs can significantly reduce communication while guaranteeing cross-round and cross-bid independence.

1.3 Overview of Our Approach

In this section we provide a high-level overview of our notion of non-malleability and of our main construction of a non-malleable vector commitment scheme (Theorem 1.2). For brevity, the main ideas underlying our additional results are described within the corresponding sections.

The starting point of our work is the notion of a vector commitment scheme $\mathcal{VC} = (\mathsf{VC.Setup}, \mathsf{VC.Commit}, \mathsf{VC.Open}, \mathsf{VC.Verify})$ [CF13] with the following syn-

tax: The algorithms $\mathsf{VC.Setup}$ and $\mathsf{VC.Commit}$ are invoked in order to produce a common-reference string crs, and in order to produce a commitment vcom for a vector $(x_1, \ldots, x_q)$, respectively. In turn, the algorithms $\mathsf{VC.Open}$ and $\mathsf{VC.Verify}$ are then invoked in order to produce a local opening π_i for each entry $i \in [q]$ of the committed vector, and in order to verify it, respectively. In terms of security, a vector commitment scheme should provide *position binding*, essentially asking that no efficient algorithm can generate a commitment vcom together with two valid openings for the same entry $i \in [q]$ corresponding to different values x_i and x'_i. The main measure of efficiency for vector commitments, which makes them non-trivial to construct, is their succinctness. This is captured by asking for upper bounds on the sizes of the resulting commitments and local openings (e.g., asking that both sizes are nearly independent of the length q of the committed vector). We refer the reader to Section 2.2 for the formal description of the position binding and succinctness requirements.

Our notion of non-malleability. Based on the standard notion of non-malleability for non-interactive commitment schemes [CIO98, CKO+01], any non-malleable vector commitment scheme should at least satisfy the following informal property: An efficient adversary which receives a commitment vcom to a vector $\vec{x} = (x_1, \ldots, x_q)$, should not be able to produce (and then open) a vector commitment $\widehat{\mathsf{vcom}}$ to some vector $\vec{\widehat{x}} = (\widehat{x_1}, \ldots, \widehat{x_q})$ which is "non-trivially related" to $\vec{x}$. However, this property does not capture the adversarial adaptivity and additional information resulting from local openings. Therefore, our notion of non-malleability for vector commitments asks that the above property holds even when the adversary can request local openings for some of the entries of $\vec{x}$ before deciding on $\widehat{\mathsf{vcom}}$, and then open only some of the entries $\vec{\widehat{x}}$ after obtaining local opening for all other entries of $\vec{x}$.

This is formalized by considering a "real" security experiment involving an adversary and an "ideal" security experiment involving a simulator. At a high level, in the real experiment, the adversary is provided with a commitment vcom to a vector $\vec{x} = (x_1, \ldots, x_q)$, and is allowed to request local openings $(\pi_i)_{i \in \mathcal{I}}$ for any subset $\mathcal{I} \subseteq [q]$ of the entries of $\vec{x}$ for producing a commitment $\widehat{\mathsf{vcom}}$. Then, the adversary is provided with local openings for all other entries of $\vec{x}$, and outputs local openings $(\widehat{\pi}_j)_{j \in \mathcal{J}}$ for a subset $\mathcal{J} \subseteq [q]$ of the entries of a malleated vector $(\widehat{x_1}, \ldots, \widehat{x_q})$ (although, note that $\widehat{\mathsf{vcom}}$ is not required to actually correspond to any such malleated vector). In the ideal experiment, the simulator is provided only with a description of the distribution $\mathcal{D}$ from which $\vec{x}$ is sampled (i.e., without the commitment vcom) and the values $(x_i)_{i \in \mathcal{I}}$ (i.e., without the local openings $(\pi_i)_{i \in \mathcal{I}}$), and outputs malleated values $(\widehat{x}_j)_{j \in \mathcal{J}}$.

The outputs of both experiments consist of the values $(x_i)_{i \in [q]}$ and $(\widehat{x}_j)_{j \in [q]}$, where in the real experiment we replace with $\perp$ each value $\widehat{x}_j$ for which either $j \notin \mathcal{J}$ or $\widehat{\pi}_j$ does not properly verify, and in the ideal experiment we replace with $\perp$ each value $\widehat{x}_j$ for which $j \notin \mathcal{J}$. Our notion of non-malleability then asks that for any efficient adversary there exists an efficient simulator such that the outputs of the two experiments are computationally indistinguishable. We refer

the reader to Section 3 for our formal definition, and for an in-depth discussion of its various technical aspects (including, the underlying distribution $\mathcal{D}$, the relation between the sets $\mathcal{I}$ and $\mathcal{J}$, and more).

Our main construction. Given any vector commitment scheme $\mathcal{VC}$ we transform it into a non-malleable one as follows. In order to commit to a vector $(x_1, \ldots, x_q)$ we first sample a signing key sk and a corresponding verification key vk for a one-time strongly-unforgeable signature scheme. Then, for each $i \in [q]$ we generate a commitment c_i to the value x_i using a locally-equivocable commitment scheme $\mathcal{LE}$ with all-but-one binding (our newly-introduced primitive augmenting the standard notion of tag-based commitments with two additional requirements). Each of these q commitments is generated with respect to the tag $\tau = h(vk)$ for a universal one-way hash function h. Then, we commit to the vector $(c_1, \ldots, c_q)$ using the underlying vector commitment scheme $\mathcal{VC}$, and output the resulting vector commitment vcom, the verification key vk and a signature σ on vcom using the signing key sk.

In turn, for every $i \in [q]$, a local opening of the value x_i consists of the commitment c_i and its corresponding decommitment d_i, and of a local opening π_i of the commitment c_i with respect to the vector commitment vcom. The verification algorithm first verifies the one-time signature σ, and then verifies the decommitment d_i and the local opening π_i. We refer the reader to Section 5 for a formal description of our construction.

Note that from a foundational perspective, the required building blocks can all be based on the existence of any vector commitment scheme. Specifically, any vector commitment scheme implies the existence of a one-way function, which in turns implies the existence of a locally-equivocable commitment scheme with all-but-one binding, a one-time strongly unforgeable signature scheme and a universal one-way hash family. In addition, from a more practical perspective, the above building blocks can all be realized based on a variety of number-theoretic assumptions leading to practical implementations (see the full version of the paper for practical number-theoretic constructions of locally-equivocable commitments with all-but-one binding, and for a construction based on one-way functions).

Focusing on the main measures of efficiency for vector commitments, namely the lengths of resulting commitments and local openings, and the verification time of the local openings, we observe the following:

- A commitment produced by our scheme consists of a commitment produced by the underlying vector commitment scheme, and of a verification key and a signature which can be instantiated with any practical strongly-unforgeable signature scheme[3]. Thus, the length of commitments produced by our scheme is essentially dominated by that of the underlying vector commitment scheme, which can be as short as a single group element.

[3] See, for example, [BSW06, BS07] and the many references therein for a variety of practical strongly-unforgeable signature schemes both in the random-oracle model and in the standard model.

- A local opening produced by our scheme consists of a local opening produced by the underlying vector commitment scheme together with a commitment and a decommitment produced by the underlying locally-equivocable commitment scheme with all-but-one binding. Relying on existing constructions of vector commitment schemes and on our number-theoretic constructions of locally-equivocable commitment schemes with all-but-one binding (which can be found in the full version), leads to local openings that are essentially as short as three group elements.
- The verification of a local opening produced by our scheme consists of a verification of a local opening produced by the underlying vector commitment scheme, a decommitment of the underlying locally-equivocable commitment scheme with all-but-one binding, and a signature verification. Once again, relying on our number-theoretic constructions of locally-equivocable commitment schemes with all-but-one binding and on practical signature schemes, this is dominated by the verification time of the underlying vector commitment scheme.

Proving the security of our main construction. Recall that for proving the security of our construction, we have to show that for any efficient adversary there exists an efficient simulator for which the outputs of the above-mentioned real and ideal experiments are computationally indistinguishable. Given the informal flavor of the current exposition, we refer the reader to the full version of the paper for an overview of the simulator's description and of the indistinguishably of the two experiments (in addition, of course, to the formal proof of security). For avoiding additional notation and various additional technical details, here we focus only on the adversary's behavior in the real experiment.

Consider an adversary $\mathcal{A}$ that is provided in the real experiment with a commitment vcom to a vector $\vec{x} = (x_1, \ldots, x_q)$. Recall that, in our construction, the commitment vcom is of the form $\mathsf{vcom} = \mathsf{vcom}_0 \| vk \| \sigma$, where vcom_0 is a commitment produced using the underlying vector commitment scheme $\mathcal{VC}$ to the vector of commitments $(c_1, \ldots, c_q)$ produced using the locally-equivocable scheme $\mathcal{LE}$ to $(x_1, \ldots, x_q)$ with respect to the tag $h(vk)$ (for a universal one-way hash function h included in the common-reference string), vk is a verification key for a one-time strongly-unforgeable signature scheme, and σ is a signature on vcom_0 produced using the corresponding signing key. The adversary $\mathcal{A}$ requests local openings $(\pi_i)_{i \in \mathcal{I}}$ for some subset $\mathcal{I} \subseteq [q]$ of the entries of $\vec{x}$, and produces a commitment $\widehat{\mathsf{vcom}} = \widehat{\mathsf{vcom}_0} \| \widehat{vk} \| \widehat{\sigma}$. Then, the adversary is provided with local openings for all other entries of $\vec{x}$, and outputs local openings $(\widehat{\pi}_j)_{j \in \mathcal{J}}$ for a subset $\mathcal{J} \subseteq [q]$. Our proof considers the following three cases (the first and second cases are straightforward, and the third case is the main technical argument):

- **Case 1: $\widehat{vk} = vk$.** This case reduces to the one-time strong unforgeability of the signature scheme, unless $\widehat{\mathsf{vcom}_0} = \mathsf{vcom}_0$ or the signature σ does not verify properly (and in these cases our simulator guarantees that the outputs of the real and ideal experiments are identical).

- **Case 2: $\widehat{vk} \neq vk$ but $h(\widehat{vk}) = h(vk)$.** This case reduces to the universal one-wayness of h.
- **Case 3: $h(\widehat{vk}) \neq h(vk)$.** In this case we rely on the position binding of the underlying vector commitment scheme $\mathcal{VC}$, and on the equivocability[4] and all-but-one binding of the locally-equivocable scheme $\mathcal{LE}$. Our main observation is that essentially any advantage that may be obtained in the real experiment must follow from the adversary's ability to choose the values $(\widehat{x_j})_{j\in\mathcal{J}}$ to which it opens the commitment $\widehat{\mathsf{vcom}}$ *after* issuing $\widehat{\mathsf{vcom}}$. That is, any such advantage must follow from the adversary's ability to produce a commitment $\widehat{\mathsf{vcom}}$ and then to provide local openings to more than a single tuple of values $(\widehat{x_j})_{j\in\mathcal{J}}$. These local openings are obtained by relying on the fact that generating $c_1, \ldots, c_q$ using the equivocation algorithms of $\mathcal{LE}$ is indistinguishable from the real experiment and does not bind them to a single tuple of values with respect to the tag $\tau = h(vk)$. Thus, we can rewind the adversary to obtain corresponding local openings with respect to the tag $\widehat{\tau} = h(\widehat{vk})$. But, if $\mathcal{A}$ can open, say, the j-th location of $\widehat{\mathsf{vcom}}$ in two different ways, we show that this contradicts either the position binding of $\mathcal{VC}$ or the all-but-one binding of $\mathcal{LE}$.

1.4 Open Problems

Our framework and constructions lead to various open problems, and here we discuss two such problems focusing on further extending our approach both in the context of vector commitments and in the more general context of non-interactive non-malleable commitments.

Non-malleable subvector commitments. The recent works of Lai and Malavolta [LM19] and of Boneh, Bünz and Fisch [BBF19] introduced the notion of VCs with *subvector openings*. These are VCs which allow the committer to open k entries of the committed vector simultaneously, with a proof whose length is sublinear in k. Our construction, being quite modular, does not seem to support such concise openings, and an interesting open problem is to construct non-malleable VCs that do support subvector openings. A possible starting point may be the recent work Gorbunov et al. [GRW+20], presenting the notion of commitments with aggregatable proofs. Constructing commitments which satisfy both this notion and our notion of local equivocability with all-but-one binding would seem to enable the construction of non-malleable subvector commitments, using our underlying approach for constructing non-malleable VCs.

Implications to non-malleable commitments. Finally, note that any non-malleable vector commitment scheme is also a non-interactive non-malleable commitment scheme (when the vector is of length 1). In that respect, our work

[4] In our formal proof, we actually rely on the equivocability guarantee earlier in order to enable the simulator to invoke the adversary in the ideal experiment.

presents a general and unified framework for constructing non-interactive non-malleable commitments, capturing both the generic construction of Di Cresccenzo, Ishai and Ostrovsky from any one-way function [CIO98] and the efficient number-theoretic constructions of Di Cresccenzo, Katz, Ostrovsky and Smith [CKO$^+$01]. As such, it may enable to construct efficient non-interactive non-malleable commitments based on new assumptions (e.g., isogenies or lattice-based assumptions) by constructing equivocable tag-based commitments with all-but-one binding based on such assumptions.

1.5 Paper Organization

The remainder of this paper is organized as follows. First, in Section 2 we present the basic notation and standard cryptographic primitives that are used throughout the paper. In Section 3 we present our framework for non-malleable VCs, show that existing VCs do not satisfy our requirements, and demonstrate that simple attempts of combining VCs and non-malleable commitments do not suffice for realizing our notion. In Section 4 we introduce our notion of a locally-equivocable commitment scheme with all-but-one binding, and in Section 5 we present our construction of a non-malleable VC.

Due to space limitations some of our contributions appear in the full version of this paper. In particular, in the full version we give a formal proof of security for our construction of a non-malleable VC. We also show that a Merkle tree is a non-malleable VC in the random-oracle model, and present our constructions of locally-equivocable commitment schemes with all-but-one binding. Finally, we show that our framework and construction extend to the dynamic setting.

2 Preliminaries

In this section we present the basic notions and standard cryptographic tools that are used in this work. For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \ldots, n\}$. For a distribution X we denote by $x \leftarrow X$ the process of sampling a value x from the distribution X. Similarly, for a set $\mathcal{X}$ we denote by $x \leftarrow \mathcal{X}$ the process of sampling a value x from the uniform distribution over $\mathcal{X}$. A function $\nu : \mathbb{N} \to \mathbb{R}^+$ is *negligible* if for any polynomial $p(\cdot)$ there exists an integer N such that for all $n > N$ it holds that $\nu(n) \leq 1/p(n)$.

2.1 Equivocable Commitment Schemes

We rely on the standard notion of a (non-interactive) equivocable commitment scheme which can be realized based on the existence of any one-way function [Nao91, CIO98]. An equivocable commitment scheme over a domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ is a 5-tuple $\mathcal{EQ} = (\mathsf{EQ.Setup}, \mathsf{EQ.Commit}, \mathsf{EQ.Decommit}, \mathsf{EQ.Equiv}_1, \mathsf{EQ.Equiv}_2)$ of polynomial-time algorithms defined as follows:

- The algorithm $\mathsf{EQ.Setup}$ is a probabilistic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$ and outputs a common-reference string crs.

- The algorithm $\mathsf{EQ.Commit}$ is a probabilistic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$, a common-reference string crs, an element $x \in \mathcal{X}_\lambda$, and outputs a commitment c and a decommitment d.
- The algorithm $\mathsf{EQ.Decommit}$ is a deterministic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$, a common-reference string crs, a commitment c and a decommitment d, and outputs an element $x \in \mathcal{X}_\lambda$ or the rejection symbol $\perp$.
- The algorithm $\mathsf{EQ.Equiv}_1$ is a probabilistic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$, and outputs a common-reference string $\widehat{\mathsf{crs}}$, a commitment $\widehat{c}$ and a state st.
- The algorithm $\mathsf{EQ.Equiv}_2$ is a deterministic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$, a state st and an element $x \in \mathcal{X}_\lambda$, and outputs a decommitment $\widehat{d}$.

Correctness. We rely on the standard notion of correctness for commitment schemes. That is, for any security parameter $\lambda \in \mathbb{N}$ and for any $x \in \mathcal{X}_\lambda$ it should hold that

$$\Pr\left[\mathsf{EQ.Decommit}(1^\lambda, \mathsf{crs}, c, d) = x\right] = 1,$$

where $\mathsf{crs} \leftarrow \mathsf{EQ.Setup}(1^\lambda)$ and $(c, d) \leftarrow \mathsf{EQ.Commit}(1^\lambda, \mathsf{crs}, x)$, and the probability is taken over the internal randomness of all algorithms.

Equivocability. We rely on the following notion of equivocability [CIO98, CKO+01]:

Definition 2.1. *A commitment scheme* $\mathcal{EQ} = (\mathsf{EQ.Setup}, \mathsf{EQ.Commit}, \mathsf{EQ.Decommit}, \mathsf{EQ.Equiv}_1, \mathsf{EQ.Equiv}_2)$ *over a domain* $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ *is* equivocable *if the following requirements hold:*

- **Equivocation correctness:** *For any* $\lambda \in \mathbb{N}$ *and* $x \in \mathcal{X}_\lambda$ *it holds that*

$$\Pr\left[\mathsf{EQ.Decommit}(1^\lambda, \widehat{\mathsf{crs}}, \widehat{c}, \widehat{d}) = x\right] = 1,$$

  where $(\widehat{\mathsf{crs}}, \widehat{c}, \mathsf{st}) \leftarrow \mathsf{EQ.Equiv}_1(1^\lambda)$ *and* $\widehat{d} := \mathsf{EQ.Equiv}_2(1^\lambda, \mathsf{st}, x)$*, and the probability is taken over the internal randomness of all algorithms.*
- **Equivocation indistinguishability:** *For any probabilistic polynomial-time algorithm* $\mathcal{A}$ *there exists a negligible function* $\nu(\cdot)$ *such that*

$$\mathbf{Adv}^{\mathsf{Equiv}}_{\mathcal{EQ},\mathcal{A}}(\lambda) \stackrel{\mathsf{def}}{=} \left|\Pr\left[\mathsf{Equiv}^{(0)}_{\mathcal{EQ},\mathcal{A}}(\lambda)\right] - \Pr\left[\mathsf{Equiv}^{(1)}_{\mathcal{EQ},\mathcal{A}}(\lambda)\right]\right| \leq \nu(\lambda)$$

  for all sufficiently large $\lambda \in \mathbb{N}$*, where for each* $b \in \{0, 1\}$ *the experiment* $\mathsf{Equiv}^{(b)}_{\mathcal{EQ},\mathcal{A}}(\lambda)$ *is defined as follows:*
  1. $(x, \mathsf{st}_\mathcal{A}) \leftarrow \mathcal{A}(1^\lambda)$.
  2. $\mathsf{crs}_0 \leftarrow \mathsf{EQ.Setup}(1^\lambda)$.
  3. $(c_0, d_0) \leftarrow \mathsf{EQ.Commit}(1^\lambda, \mathsf{crs}, x)$.
  4. $(\mathsf{crs}_1, c_1, \mathsf{st}_1) \leftarrow \mathsf{EQ.Equiv}_1(1^\lambda)$.

5. $d_1 = \mathsf{EQ.Equiv}_2(1^\lambda, \mathsf{st}_1, x)$.
6. $b' \leftarrow \mathcal{A}(\mathsf{st}_\mathcal{A}, \mathsf{crs}_b, c_b, d_b)$.
7. *Output* b'.

Binding. We rely on the standard notion of computational binding for commitment schemes.

Definition 2.2. *A commitment scheme* $\mathcal{EQ} = (\mathsf{EQ.Setup}, \mathsf{EQ.Commit}, \mathsf{EQ.Decommit}, \mathsf{EQ.Equiv}_1, \mathsf{EQ.Equiv}_2)$ *over a domain* $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ *is* binding *if for any probabilistic polynomial-time algorithm* $\mathcal{A}$ *there exists a negligible function* $\nu(\cdot)$ *such that*

$$\mathbf{Adv}^{\mathsf{Bind}}_{\mathcal{EQ},\mathcal{A}} \stackrel{\mathsf{def}}{=} \Pr\left[\mathsf{Bind}_{\mathcal{EQ},\mathcal{A}}(\lambda) = 1\right] \leq \nu(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$*, where the experiment* $\mathsf{PosBind}_{\mathcal{EQ},\mathcal{A}}(\lambda)$ *is defined as follows:*

1. $\mathsf{crs} \leftarrow \mathsf{EQ.Setup}(1^\lambda)$
2. $(c, (d, x), (d', x')) \leftarrow \mathcal{A}(1^\lambda, \mathsf{crs})$.
3. *Output* 1 *if the following conditions hold:*
   - $x \neq x'$ *and* $x, x' \in \mathcal{X}_\lambda$.
   - $\mathsf{EQ.Decommit}(1^\lambda, \mathsf{crs}, c, d) = x$.
   - $\mathsf{EQ.Decommit}(1^\lambda, \mathsf{crs}, c, d') = x'$.

   Otherwise, output 0.

2.2 Vector Commitment Schemes

We follow the notion of a vector commitment scheme as formalized by Libert and Yung [LY10] and Catalano and Fiore [CF13]. As discussed in Section 1.1, here we consider the static setting (i.e., vector commitment schemes without updates). In the full version of the paper, we extend our approach to the dynamic setting.

Definition 2.3. *A vector commitment scheme over a domain* $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ *is a quadruple* $\mathcal{VC} = (\mathsf{VC.Setup}, \mathsf{VC.Commit}, \mathsf{VC.Open}, \mathsf{VC.Verify})$ *of algorithms defined as follows:*

- *The algorithm* VC.Setup *is a probabilistic algorithm that receives as input the security parameter* $\lambda \in \mathbb{N}$ *and a polynomial* $q = q(\lambda)$ *and outputs common-reference string* crs.
- *The algorithm* VC.Commit *is a probabilistic algorithm that receives as input the security parameter* $\lambda \in \mathbb{N}$*, a common-reference string* crs *and a vector* $(x_1, \ldots, x_q) \in (\mathcal{X}_\lambda)^q$*, and outputs a commitment* vcom *and a state* st.
- *The algorithm* VC.Open *is a probabilistic algorithm that receives as input the security parameter* $\lambda \in \mathbb{N}$*, a common-reference string* crs*, a commitment* vcom*, a state* st *and an index* $i \in [q]$*, and outputs a proof* π.
- *The algorithm* VC.Verify *is a deterministic algorithm that receives as input the security parameter* $\lambda \in \mathbb{N}$*, a common-reference string* crs*, a commitment* vcom*, an index* $i \in [q]$*, an element* $x \in \mathcal{X}_\lambda$ *and a proof* π*, and outputs a bit* $b \in \{0, 1\}$.

Correctness. A vector commitment scheme $\mathcal{VC} = (\mathsf{VC.Setup}, \mathsf{VC.Commit}, \mathsf{VC.Open}, \mathsf{VC.Verify})$ over a domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ is correct if for any $\lambda \in \mathbb{N}$, for any polynomial $q = q(\lambda)$, for any vector $(x_1, \ldots, x_q) \in (\mathcal{X}_\lambda)^q$, and for any index $i \in [q]$, it holds that

$$\Pr\left[\mathsf{VC.Verify}\left(1^\lambda, \mathsf{crs}, \mathsf{vcom}, i, x_i, \pi\right) = 1\right] = 1,$$

where $\mathsf{crs} \leftarrow \mathsf{VC.Setup}(1^\lambda)$, $(\mathsf{vcom}, \mathsf{st}) \leftarrow \mathsf{VC.Commit}(1^\lambda, \mathsf{crs}, (x_1, \ldots, x_q))$ and $\pi \leftarrow \mathsf{VC.Open}(1^\lambda, \mathsf{crs}, \mathsf{vcom}, \mathsf{st}, i)$; and the probability is taken over the randomness of all algorithms.

Security. Catalano and Fiore introduced the following notion of position binding for capturing the security of vector commitment schemes.

Definition 2.4. *A vector commitment scheme* $\mathcal{VC} = (\mathsf{VC.Setup}, \mathsf{VC.Commit}, \mathsf{VC.Open}, \mathsf{VC.Verify})$ *over a domain* $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ *is* position binding *if for any polynomial* $q = q(\lambda)$ *and for any probabilistic polynomial-time algorithm* $\mathcal{A}$ *there exists a negligible function* $\nu(\cdot)$ *such that*

$$\mathbf{Adv}^{\mathsf{PosBind}}_{\mathcal{VC},q,\mathcal{A}} \stackrel{\mathsf{def}}{=} \Pr\left[\mathsf{PosBind}_{\mathcal{VC},q,\mathcal{A}}(\lambda) = 1\right] \leq \nu(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$*, where the experiment* $\mathsf{PosBind}_{\mathcal{VC},q,\mathcal{A}}(\lambda)$ *is defined as follows:*

1. $\mathsf{crs} \leftarrow \mathsf{VC.Setup}(1^\lambda, q)$
2. $(\mathsf{vcom}, i, x_i, x'_i, \pi, \pi') \leftarrow \mathcal{A}(1^\lambda, q, \mathsf{crs})$.
3. *Output* 1 *if the following conditions hold:*
   - $x_i \neq x'_i$.
   - $\mathsf{VC.Verify}\left(1^\lambda, \mathsf{crs}, \mathsf{vcom}, i, x_i, \pi\right) = 1$.
   - $\mathsf{VC.Verify}\left(1^\lambda, \mathsf{crs}, \mathsf{vcom}, i, x'_i, \pi'\right) = 1$.

   Otherwise, output 0.

Succinctness. The main measure of efficiency for vector commitments, which makes them non-trivial to construct, is their succinctness. This may be captured by asking for upper bounds $\ell_{\mathsf{Commit}}(\lambda, q)$ and $\ell_{\mathsf{Open}}(\lambda, q)$ on the size of the commitment and the size of the local openings, respectively, as follows.

Definition 2.5. *A vector commitment scheme* $\mathcal{VC} = (\mathsf{VC.Setup}, \mathsf{VC.Commit}, \mathsf{VC.Open}, \mathsf{VC.Verify})$ *over a domain* $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ *is* $(\ell_{\mathsf{Commit}}, \ell_{\mathsf{Open}})$*-succinct if for any* $\lambda \in \mathbb{N}$*, for any polynomial* $q = q(\lambda)$*, for any common-reference string* crs *produced by* $\mathsf{VC.Setup}(1^\lambda, q)$*, for any vector* $(x_1, \ldots, x_q) \in (\mathcal{X}_\lambda)^q$*, and for any commitment and state* $(\mathsf{vcom}, \mathsf{st})$ *produced by* $\mathsf{VC.Commit}(1^\lambda, \mathsf{crs}, (x_1, \ldots, x_q))$ *the following two requirements are satisfied:*

- *The bit-length of* vcom *is at most* $\ell_{\mathsf{Commit}}(\lambda, q)$.
- *For any index* $i \in [q]$ *and for any proof* π *produced by* $\mathsf{VC.Open}(1^\lambda, \mathsf{crs}, \mathsf{vcom}, \mathsf{st}, i)$*, the bit-length of* π *is at most* $\ell_{\mathsf{Open}}(\lambda, q)$.

2.3 One-Time Strongly-Unforgeable Signature Schemes

We rely on the standard notion of a one-time strongly-unforgeable signature scheme, which is known to exist based on the existence of any one-way function [Lam79, NY89, Rom90] (and thus, in particular, based on any of the number-theoretic assumptions that we consider in this paper). A signature scheme is a tuple $\mathcal{SIG} = (\mathsf{Sig.Gen}, \mathsf{Sig.Sign}, \mathsf{Sig.Verify})$ of algorithms defined as follows:

- The algorithm $\mathsf{Sig.Gen}$ is a probabilistic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$ and outputs a pair (sk, vk) of a signing key and a verification key.
- The algorithm $\mathsf{Sig.Sign}$ is a (possibly) probabilistic algorithm that receives as input a signing key sk and a message m and outputs a signature σ.
- The algorithm $\mathsf{Sig.Verify}$ is a deterministic algorithm that receives as input a verification key vk, a message m and a signature σ, and outputs a bit $b \in \{0,1\}$.

In terms of correctness, the standard requirement for signature schemes asks that

$$\Pr\left[\mathsf{Sig.Verify}_{vk}(m, \mathsf{Sig.Sign}_{sk}(m)) = 1\right] = 1$$

for every $\lambda \in \mathbb{N}$ and for every message m, where $(sk, vk) \leftarrow \mathsf{Sig.Gen}(1^\lambda)$, where the probability is taken over the internal randomness of all algorithms. In terms of security, we rely on the following standard notion of one-time strong unforgeability.

Definition 2.6. *A signature scheme* $\mathcal{SIG} = (\mathsf{Sig.Gen}, \mathsf{Sig.Sign}, \mathsf{Sig.Verify})$ *is* one-time strongly unforgeable *if for every probabilistic polynomial-time algorithm* $\mathcal{A}$ *there exists a negligible function* $\nu(\cdot)$ *such that*

$$\mathbf{Adv}^{\mathsf{Forge}}_{\mathcal{SIG},\mathcal{A}}(\lambda) \stackrel{\mathsf{def}}{=} \Pr\left[\mathsf{Forge}_{\mathcal{SIG},\mathcal{A}}(\lambda) = 1\right] \leq \nu(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$*, where the experiment* $\mathsf{Forge}_{\mathcal{SIG},\mathcal{A}}(\lambda)$ *is defined as follows:*

1. $(sk, vk) \leftarrow \mathsf{Sig.Gen}(1^\lambda)$.
2. $(m, \mathsf{st}_{\mathcal{A}}) \leftarrow \mathcal{A}(1^\lambda, vk)$.
3. $(m^*, \sigma^*) \leftarrow \mathcal{A}(\mathsf{st}_{\mathcal{A}}, \sigma)$*, where* $\sigma \leftarrow \mathsf{Sig.Sign}_{sk}(m)$.
4. *If* $\mathsf{Sig.Verify}_{vk}(m^*, \sigma^*)$ *and* $(m^*, \sigma^*) \neq (m, \sigma)$ *then output* 1 *and otherwise output* 0.

2.4 Universal One-Way Hash Functions

We rely on the standard notion of universal one-way hash functions, which is known to exist based on the existence of any one-way function [NY89, Rom90] (and thus, in particular, based on any of the number-theoretic assumptions that we consider in this paper). A hash family from domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ to range $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda\in\mathbb{N}}$ is a collection $\mathcal{H} = \{\mathcal{H}\}_{\lambda\in\mathbb{N}}$ where each $\mathcal{H}_\lambda$ consists of functions $h : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda$. For simplifying our notation we let $h \leftarrow \mathcal{H}_\lambda$ denote the process of

sampling a function h from $\mathcal{H}_\lambda$ without explicitly describing a sampling algorithm, where h denotes both the description of the sampled function and its evaluation algorithm.

Definition 2.7. *A hash family $\mathcal{H}$ from domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ to range $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda\in\mathbb{N}}$ is a* universal one-way hash family *if for every probabilistic polynomial-time algorithm $\mathcal{A}$ there exists a negligible function $\nu(\cdot)$ such that*

$$\mathbf{Adv}^{\mathsf{UOWHF}}_{\mathcal{H},\mathcal{A}}(\lambda) \stackrel{\mathsf{def}}{=} \Pr\left[\mathsf{UOWHF}_{\mathcal{H},\mathcal{A}}(\lambda) = 1\right] \leq \nu(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$, where the experiment $\mathsf{UOWHF}_{\mathcal{H},\mathcal{A}}(\lambda)$ is defined as follows:

1. $(x, \mathsf{st}) \leftarrow \mathcal{A}(1^\lambda)$.
2. $h \leftarrow \mathcal{H}_\lambda$.
3. $x' \leftarrow \mathcal{A}(\mathsf{st}, h)$.
4. *If $x \neq x'$ and $h(x) = h(x')$ then output* 1*, and otherwise output* 0*.*

3 Non-Malleable Vector Commitments

In this section we begin by presenting our notion of non-malleability for vector commitment schemes. Then, in Section 3.1 we show that existing vector commitment schemes do not satisfy it (some of them by design in order to support public updates). As mentioned in Section 1.1, the key difference from the standard notion of non-malleable non-interactive commitments is that we aim at achieving non-malleability even with respect to adversaries which have already been exposed to several local openings. This key difference is the reason that simple attempts of combining VCs and non-malleable commitments, that we discuss in Section 3.2, do not suffice for realizing our new notion.

Loosely speaking, a vector commitment scheme is non-malleable if an efficient adversary which receives a vector commitment vcom to a vector $\vec{x} = (x_1, \ldots, x_q)$, cannot produce (and open) a vector commitment $\widehat{\mathsf{vcom}}$ to some vector $\vec{\widehat{x}} = (\widehat{x_1}, \ldots, \widehat{x_q})$ which is "non-trivially related" to $\vec{x}$. This property should hold even when the adversary can request local openings for some of the entries of $\vec{x}$ before deciding on $\widehat{\mathsf{vcom}}$; and open only some of the entries $\vec{\widehat{x}}$. Definition 3.1 below uses the term "valid distribution" which is formally clarified following the definition. As discussed in Section 1.1, we start by considering the static setting of vector commitments without updates. In the full version, we extend our approach to the dynamic setting.

Definition 3.1. *A vector commitment $\mathcal{VC} = (\mathsf{VC.Setup}, \mathsf{VC.Commit}, \mathsf{VC.Open}, \mathsf{VC.Verify})$ over a domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ is* non-malleable *if for any polynomially-bounded integer $q = q(\lambda)$ and for any probabilistic polynomial-time algorithm $\mathcal{A}$ there exist a probabilistic polynomial-time algorithm $\mathcal{S}$ such that the following holds:*

For any probabilistic polynomial-time algorithm $\mathcal{R}$ and for any valid distribution $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda\in\mathbb{N}}$ over $\{(\mathcal{X}_\lambda)^q\}_{\lambda\in\mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that

$$\begin{aligned}&\mathbf{Adv}^{\mathsf{NM}}_{\mathcal{VC},q,\mathcal{A},\mathcal{S},R,\mathcal{D}}(\lambda)\\&\quad\stackrel{\mathsf{def}}{=}\left|\Pr\left[\mathcal{R}\left(\mathsf{Real}_{\mathcal{VC},q,\mathcal{A},\mathcal{D}}(\lambda)\right)=1\right]-\Pr\left[\mathcal{R}\left(\mathsf{Ideal}_{\mathcal{VC},q,\mathcal{S},\mathcal{D}}(\lambda)\right)=1\right]\right|\leq\nu(\lambda)\end{aligned}$$

for all sufficiently large $\lambda\in\mathbb{N}$, where the experiments $\mathsf{Real}_{\mathcal{VC},q,\mathcal{A},\mathcal{D}}(\lambda)$ and $\mathsf{Ideal}_{\mathcal{VC},q,\mathcal{S},\mathcal{D}}(\lambda)$ are defined as follows:

The Experiment $\mathsf{Real}_{\mathcal{VC},q,\mathcal{A},\mathcal{D}}(\lambda)$:

1. $\mathsf{crs}\leftarrow\mathsf{VC.Setup}(1^\lambda,q)$.
2. $(x_1,\ldots,x_q)\leftarrow\mathcal{D}_\lambda$.
3. $(\mathsf{vcom},\mathsf{st})\leftarrow\mathsf{VC.Commit}(1^\lambda,\mathsf{crs},(x_1,\ldots,x_q))$.
4. $(\mathcal{I},\mathsf{st}_\mathcal{A})\leftarrow\mathcal{A}(1^\lambda,\mathsf{crs},\mathsf{vcom})$ *where* $\mathcal{I}\subseteq[q]$.
5. $\pi_i\leftarrow\mathsf{VC.Open}(1^\lambda,\mathsf{crs},\mathsf{vcom},\mathsf{st},i)$ *for each* $i\in[q]$.
6. $\left(\widehat{\mathsf{vcom}},\mathcal{J},\mathsf{st}_\mathcal{A}\right)\leftarrow\mathcal{A}\left(\mathsf{st}_\mathcal{A},(x_i)_{i\in\mathcal{I}},(\pi_i)_{i\in\mathcal{I}}\right)$, *where* $\mathcal{J}\subseteq[q]$.
7. $((\widehat{x_j})_{j\in\mathcal{J}},(\widehat{\pi_j})_{j\in\mathcal{J}})\leftarrow\mathcal{A}\left(\mathsf{st}_\mathcal{A},(x_i)_{i\in\overline{\mathcal{I}}},(\pi_i)_{i\in\overline{\mathcal{I}}}\right)$, *where* $\overline{\mathcal{I}}=[q]\setminus\mathcal{I}$.
8. *If* $\widehat{\mathsf{vcom}}=\mathsf{vcom}$ *or if* $\mathsf{VC.Verify}\left(1^\lambda,\mathsf{crs},\widehat{\mathsf{vcom}},j,\widehat{x_j},\widehat{\pi_j}\right)=0$ *for some* $j\in\mathcal{J}$, *then output* $((x_1,\ldots,x_q),(\bot)^q,\mathcal{I})$.
   Otherwise, output $((x_1,\ldots,x_q),(\widehat{x_1},\ldots,\widehat{x_q}),\mathcal{I})$, *where* $\widehat{x_j}=\bot$ *for each* $j\in[q]\setminus\mathcal{J}$.

The Experiment $\mathsf{Ideal}_{\mathcal{VC},q,\mathcal{S},\mathcal{D}}(\lambda)$:

1. $(x_1,\ldots,x_q)\leftarrow\mathcal{D}_\lambda$.
2. $(\mathcal{I},\mathsf{st}_\mathcal{S})\leftarrow\mathcal{S}(1^\lambda,\mathcal{D})$.
3. $(\mathcal{J},(\widehat{x_j})_{j\in\mathcal{J}})\leftarrow\mathcal{S}(\mathsf{st}_\mathcal{S},(x_i)_{i\in\mathcal{I}})$.
4. *Output* $((x_1,\ldots,x_q),(\widehat{x_1},\ldots,\widehat{x_q}),\mathcal{I})$ *where* $\widehat{x_i}=\bot$ *for every* $i\in[q]\setminus\mathcal{J}$.

Succinctness. Recall that the main measure of efficiency for vector commitments, which makes them non-trivial to construct, is their succinctness: Both the size of the commitment and the size of the local openings should be sublinear in the number q of elements in the committed vector. That is, the standard notion of vector commitments does not require any hiding guarantees [CF13], and thus can be trivially satisfied if succinctness is not required (in this case a vector commitment scheme can simply output the vector itself). When additionally requiring a vector commitment scheme to hide all entries of the committed vector for which local openings were not provided, the task becomes non-trivial even when succinctness is not required (since this introduces a selective decommitment problem whenever an attacker can request local openings after having seen the commitment).

Our notion of non-malleability implies, in particular, such a hiding guarantee, and is therefore non-trivial to realize even when succinctness is not required. Nevertheless, as discussed in Section 1.1, the non-malleable vector commitments

resulting from our transformation are essentially as succinct as the existing standard vector commitments that do not require any hiding guarantees.

Valid distributions. Definition 3.1 considers *valid* distributions, and here we formally define this notion. On the face of it, one can hope to consider all distributions that are samplable in polynomial time. However, exactly as in case of non-malleable zero-knowledge sets [GM06], our notion of non-malleable vector commitments faces a "selective decommitment" problem (since it considers attackers which may be exposed to several adaptively-chosen local openings). One approach to overcome this difficulty, which is the approach that we follow in this work, is to restrict our attention to considering the natural subclass of all efficiently samplable distributions that was considered by Gennaro and Micali [GM06]). This subclass consists of all distributions that are not only efficiently samplable, but also all of their marginal distributions are efficiently samplable.

That is, we say that a distribution $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda\in\mathbb{N}}$ over $\{(\mathcal{X}_\lambda)^{q}\}_{\lambda\in\mathbb{N}}$ is valid if the following holds: For every $\lambda \in \mathbb{N}$, for every $\vec{x} = (x_1, \ldots, x_{q(\lambda)})$ in the support of $\mathcal{D}_\lambda$, and for every subset $\mathcal{I} = (i_1, \ldots, i_{|\mathcal{I}|}) \subseteq [q(\lambda)]$, it is possible to efficiently sample a vector $\vec{y}$ from the conditional distribution $\mathcal{D}_\lambda|(\forall i \in \mathcal{I} : y_i = x_i)$. We denote the process of sampling the entries of $\vec{y}$ in $\overline{\mathcal{I}} = [q] \setminus \mathcal{I}$ by $(y_j)_{j\in\overline{\mathcal{I}}} \leftarrow \mathcal{D}|(\mathcal{I}, (x_i)_{i\in\mathcal{I}})$. Note that this requirement is fairly reasonable, and in particular, it is satisfied by any product distribution $\mathcal{D}$ over $(\mathcal{X}_\lambda)^q$.

An alternative approach, as pointed out by Gennaro and Micali, is to rely on an underlying commitment scheme that provides a certain form of security against selective decommitment attacks. In their context, it seems that the underlying commitment scheme would have to be at least both mercurial and provide security against selective decommitment attacks (realizing this alternative approach for non-malleable zero-knowledge sets still remains an interesting open problem). Similarly, in our context it would have to be at least locally equivocable with all-but-one binding (as we define in Section 4) and provide security against selective decommitment attacks. We leave the exploration of this alternative approach as an avenue for further research.

$\mathcal{J}$ cannot be chosen later. Note that we allow the adversary $\mathcal{A}$ in the experiment $\mathsf{Real}_{\mathcal{VC},q,\mathcal{A},R,\mathcal{D}}(\lambda)$ to choose the subset $\mathcal{J}$ at the latest stage possible. This is true because had we let A choose $\mathcal{J}$ in Step 7 of the experiment, then $\mathcal{A}$ could have encoded information about $(x_j)_{j\in\overline{\mathcal{I}}}$ within their choice of $\mathcal{J}$. For example, assume that we let the adversary choose $\mathcal{J}$ in Step 7 of $\mathsf{Real}_{\mathcal{VC},q,\mathcal{A},R,\mathcal{D}}(\lambda)$ (after observing $(x_j)_{j\in\overline{\mathcal{I}}}$), and consider an adversary which chooses $\mathcal{J}$ to be of size 1 if the parity of the bit-description of $x_{j_1}\|\cdots\|x_{j_{|\overline{\mathcal{I}}|}}$ is 1, and chooses $\mathcal{J}$ to be of size 0 if this parity is 0, where $\overline{\mathcal{I}} = \{j_1, \ldots, j_{|\overline{\mathcal{I}}|}\}$. Of course, this cannot be simulated, since the simulator never gets access to $x_{j_1}, \ldots, x_{j_{|\overline{\mathcal{I}}|}}$.

Invalid openings. Whenever the adversary $\mathcal{A}$ provides an invalid opening for *any* index in $\mathcal{J}$, then the output of the real experiment is set to be of the form $((x_1, \ldots, x_q), (\perp)^q, \mathcal{I})$. We argue that this choice is indeed a necessary one. To see

why that is the case, consider the following alternative (and faulty) approach: For all $j \in \mathcal{J}$ for which $\mathcal{A}$ provides invalid openings set $\widehat{x_j} = \perp$, but for all indices for which $\mathcal{A}$ provides valid openings, keep the $\widehat{x_j}$'s in the output of the experiment as is (that is, as outputted by $\mathcal{A}$ in Step 7). The problem with this approach is that it effectively gives $\mathcal{A}$ the power to choose $\mathcal{J}$ in Step 7 of the experiment, for example by outputting $\mathcal{J} = [q]$ in Step 6 and then providing valid openings for a different set $\mathcal{J}' \subsetneq [q]$ in Step 7. As explained above, such a definition cannot be satisfied, as it allows $\mathcal{A}$ to encode information about $(x_j)_{j \in \overline{\mathcal{I}}}$ via the set of validly-opened positions.

Letting $\mathcal{J}$ intersect $\mathcal{I}$. At first glance, it might seem uncanny that we let the adversary choose the set $\mathcal{J}$ such that it includes locations for which the adversary has seen openings before producing $\widehat{\mathsf{vcom}}$ (i.e., it intersects $\mathcal{I}$). On the face of it, this allows for trivial attacks, since the adversary can trivially commit, via $\widehat{\mathsf{vcom}}$, to values that are related to $(x_i)_{i \in \mathcal{I}}$. However, Definition 3.1 "discounts" such trivial attacks from the adversary's advantage, by allowing the simulator to access values $(x_i)_{i \in \mathcal{I}}$ as well.

Choosing $\mathcal{I}$ adaptively. We note that Definition 3.1 can be strengthened, by allowing the adversary in $\mathsf{Real}_{\mathcal{VC},q,\mathcal{A},R,\mathcal{D}}(\lambda)$ to choose the set $\mathcal{I}$ in an adaptive manner. That is, to choose the indices included in $\mathcal{I}$ one by one, each index being chosen after $\mathcal{A}$ has observed the values x_i (and the associated proof π_i) for each previous chosen index i. Our construction in Section 5 remains secure under this strengthened definition, and its proof of security readily extends to it.

Reusability. One might consider a strengthening of Definition 3.1, by providing the adversary with many vector commitments $\mathsf{vcom}_1, \ldots, \mathsf{vcom}_k$ (and to local openings of their choice) to vectors $\vec{v_1}, \ldots, \vec{v_k}$, and requiring that they cannot produce (and later open) a vector commitment $\widehat{\mathsf{vcom}}$ to a vector $\vec{v}$ which is non-trivially related to $\vec{v_1}, \ldots, \vec{v_k}$. Such a strengthening is in line with the notion of a *reusable* non-malleable non-interactive commitment scheme [DG03] and more generally, with the notion of concurrent non-malleable commitments [DDN00]. We believe that our framework and constructions can be generalized to support such a definition, and we leave this task to future work.

3.1 Existing Schemes Do Not Satisfy Our Notion

Merkle trees in the standard model. Consider the Merkle tree construction of vector commitments with respect to a hash function $h : \{0,1\}^{2\lambda} \to \{0,1\}^{\lambda}$. That is, a commitment vcom to a vector $\vec{x} \in \{0,1\}^{\lambda \times q}$ is the root of the binary hash tree whose left leaves (i.e., leaves which are left children) are the values of $\vec{x}$; the right leaves are assigned some predetermined arbitrary values; and the value of each node is obtained by applying h to the concatenation of its

children.[5] In the full version we present a formal description of this construction, and show that if h is modeled as a random oracle, then this construction is indeed non-malleable per Definition 3.1. Alas, if h is instantiated via a standard-model collision resistant hash function, this is not necessarily the case. Loosely speaking, this is because the function h itself may be malleable.

As a concrete and simple example, consider the case in which $h(z) = z_1 \| h'(z_2 \| \cdots \| z_{2\lambda})$, where $z = z_1 \| \cdots \| z_{2\lambda} \in \{0,1\}^{2\lambda}$ and $h' : \{0,1\}^{2\lambda-1} \to \{0,1\}^{\lambda-1}$ is a collision-resistant hash function. It is not hard to verify that h is also collision resistant; but still, the vector commitment it induces is malleable. In fact, this vector commitment is not even completely hiding: Consider the following attacker which first request to see an opening of the first entry x_1 of $\vec{x}$ (by outputting $\mathcal{I} = \{1\}$ in Step 4 of the real experiment of Definition 3.1). This opening includes the value assigned to the sibling of the parent of x_1 (which is the parent of x_2); denote this value by $y = y_1 \| \cdots \| y_\lambda \in \{0,1\}^\lambda$. Then y_1 is equal to the first bit of x_2. This means that the adversary can commit from scratch to some vector $(\widehat{x_1}, \ldots, \widehat{x_q})$ such that the first bit of $\widehat{x_2}$ is also y_1 (and the other entries are chosen arbitrarily), satisfying a non-trivial relation with $\vec{x}$. This is just one simple example, and many more examples exist for the malleability of standard-model instantiation of Merkle trees.

Algebraic constructions. More recent algebraic constructions of vector commitments turn out to be malleable as well. To start, consider the constructions of Catalano and Fiore [CF13], based on either the discrete logarithm assumption or the RSA assumption. In both of these construction, a user commits to a vector $\vec{x}$ of integers, by computing $\mathsf{vcom} = \prod_{i \in [q]} g_i^{x_i}$, where $g_1, \ldots, g_q$ are publicly-known group elements. It is not hard to see, that an attacker receiving vcom can produce a commitment $\widehat{\mathsf{vcom}}$ to any affinely-related vector $a \cdot \vec{x} + \vec{z}$, by computing $\mathsf{vcom}^a \cdot \prod_{i \in [q]} g_i^{z_i}$.

Lai and Malavolta [LM19] recently generalized the constructions of Catalano and Fiore to Euclidean rings (they also presented an additional construction in bilinear groups, which falls into the same template as the constructions of Catalano and Fiore, and hence the same attack applies to it). Concretely, they consider a module over a ring R, consisting of an Abelian group $(\mathbb{G}, \times)$ and a binary operation $\circ : R \times \mathbb{G} \to \mathbb{G}$. A vector commitment to a vector $\vec{x} \in \mathcal{X}^q$ is then computed by the inner product $\left\langle \vec{x}, \vec{S} \right\rangle = (x_1 \circ S_1) \times \cdots \times (x_q \circ S_q)$, where $\mathcal{X} \subseteq R$ is a subset satisfying some natural property and $\vec{S}$ is a vector of publicly-known group elements. Unsurprisingly, the afore-described attack easily generalizes to this construction as well. For any $a \in R$ and $z \in R^q$, an attacker which receies a commitment vcom to a vector $\vec{x} \in \mathcal{X}^q$ can compute a commitment to any affinely-related vector $a \cdot \vec{x} + \vec{z}$, where $(+, \cdot)$ are the two ring operations, by computing $(a \circ \mathsf{vcom}) \times \left\langle \vec{z}, \vec{S} \right\rangle$. Note that this attack works as long as $a \cdot \vec{x} + \vec{z}$ lies in $\mathcal{X}$.

[5] We embed the entries of $\vec{x}$ only as left leaves as to avoid trivial attacks. Doing so, the opening of say, the i-th entry does not trivially reveal any other entries.

3.2 Simple Attempts That Fail

For obtaining an initial understanding of the challenges in constructing non-malleable vector commitments, consider the following two constructions which are based on rather simple and direct combinations of vector commitments and non-malleable commitments, and fail to satisfy Definition 3.1. In what follows, $\mathsf{nm}\mathcal{COM}$ is a standard non-malleable commitments scheme and $\mathcal{VC}$ is a (potentially malleable) vector commitment scheme.

Applying $\mathsf{nm}\mathcal{COM}$ and then $\mathcal{VC}$. As a first attempt, consider what happens when in order to commit to some vector $\vec{x}$, one first applies $\mathsf{nm}\mathcal{COM}$ locally to each entry of $\vec{x}$ to obtain q commitments $c_1, \ldots, c_q$; and then uses $\mathcal{VC}$ to commit to these commitments. The problem with this approach is that $\mathcal{VC}$ might be malleable. For example, if $\mathcal{VC}$ appends a random bit to the end of each commitment, then an adversary which receives a commitment vcom to a vector $\vec{x}$ produced using the approach described above, can easily produce a different commitment $\widehat{\mathsf{vcom}}$ to the same $\vec{x}$ by flipping the last bit of vcom. It might be also the case that $\mathcal{VC}$ is malleable in the following sense: Given a commitment vcom to a vector $\vec{x}$ produced using $\mathcal{VC}$, it is easy to "replace" some of the entries of the vector underlying vcom, resulting in a commitment to a related vector $\vec{x'}$ which identifies with $\vec{x}$ on some of its locations. If this is the case, then such an attack is also possible for the combined vector commitment scheme which first applies $\mathsf{nm}\mathcal{COM}$ locally.[6]

Applying $\mathcal{VC}$ and then $\mathsf{nm}\mathcal{COM}$. Consider a construction which, in order to commit to a vector $\vec{x}$, first applies $\mathcal{VC}$ to produce a commitment vcom_0 and commits to vcom_0 using $\mathsf{nm}\mathcal{COM}$ to produce a commitment vcom. Alas, this approach also does not meet Definition 3.1. The main issue is unique to the setting of non-malleable vector commitments: Per Definition 3.1, an adversary can request to see openings of individual entries of $\vec{x}$ before outputting their own commitment $\widehat{\mathsf{vcom}}$. These openings must include in particular the intermediate commitment vcom_0. Hence, if $\mathcal{VC}$ is malleable, then the adversary, having observed vcom_0 can come up with a different commitment $\widehat{\mathsf{vcom}_0}$ with respect to $\mathcal{VC}$ for some related vector $\vec{x'}$. Then, the adversary can simply commit to $\widehat{\mathsf{vcom}_0}$ using $\mathsf{nm}\mathcal{COM}$ to produce the desired commitment $\widehat{\mathsf{vcom}}$.

[6] Another issue which may arise, is that $\mathsf{nm}\mathcal{COM}$ might not be *concurrent* non-malleable (see, for example, [DDN00, PR05, PR08, LPV08] and the references therein). In this case, an adversary which observes some of the local commitments and openings produced via $\mathsf{nm}\mathcal{COM}$ may be able to come up with $\mathsf{nm}\mathcal{COM}$ commitments to related values. This issue, however, can be relatively easily resolved by using a commitment scheme which offers non-malleability even against adversaries which observe at most q commitments and openings.

4 Locally-Equivocable Commitments with All-But-One Binding

In this section we introduce the notion of a locally-equivocable commitment scheme with all-but-one binding, which serves as one of the main building-blocks underlying our construction of a non-malleable vector commitment scheme. Our notion is obtained by augmenting the standard notion of a non-interactive tag-based commitment scheme with two additional requirements, namely local equivocability and all-but-one binding.

In addition, we present both a somewhat theoretical realization of the our new notion based on the existence of any one-way function, and two efficient number-theoretic realizations: A construction based on the discrete logarithm assumption, and a construction based on the RSA assumption. In both cases, the common-reference string consists of 2-3 group elements (in addition to the description of the group), and a commitment consists of a single group element. As discussed in Section 1.1, these number-theoretic constructions were described by Fischlin in his Ph.D. thesis [Fis01] (and also used by Crescenzo, Katz, Ostrovsky and Smith [CKO$^+$01] in their number-theoretic constructions of non-malleable non-interactive commitment schemes[7]). Although our notion of a locally-equivocable commitment scheme with all-but-one binding strengthens Fischlin's notion of identity-based trapdoor commitments (as we discuss below), we nevertheless show that these constructions satisfy our notion. Our constructions are provided in the full version.

Formally, a locally-equivocable commitment scheme with all-but-one binding over a domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ and a tag space $\mathcal{T} = \{\mathcal{T}_\lambda\}_{\lambda\in\mathbb{N}}$ is a 6-tuple $\mathcal{LE} = (\mathsf{LE.Setup}, \mathsf{LE.Commit}, \mathsf{LE.Decommit}, \mathsf{LE.AltSetup}, \mathsf{LE.Equiv}_1, \mathsf{LE.Equiv}_2)$ of polynomial-time algorithms defined as follows:

- The algorithm $\mathsf{LE.Setup}$ is a probabilistic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$ and a polynomially-bounded integer $q = q(\lambda)$, and outputs a common-reference string crs.
- The algorithm $\mathsf{LE.Commit}$ is a probabilistic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$, a common-reference string crs, an element $x \in \mathcal{X}_\lambda$, an index $i \in [q]$ and a tag $\tau \in \mathcal{T}_\lambda$, and outputs a commitment c and a decommitment d.[8]
- The algorithm $\mathsf{LE.Decommit}$ is a deterministic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$, a common-reference string crs, a commitment c, a decommitment d, an index $i \in [q]$ and a tag $\tau \in \mathcal{T}_\lambda$, and outputs an element $x \in \mathcal{X}_\lambda$ or the rejection symbol $\perp$.

[7] Although Crescenzo et al. did not explicitly frame their construction as relying on an underlying equivocable commitment scheme, we follow a somewhat more fine-grained abstraction via our local equivocability and all-but-one binding properties.

[8] We note that the commitment and decommitment algorithms $\mathsf{LE.Commit}$ and $\mathsf{LE.Decommit}$ receive the index $i \in [q]$ as input for technical reasons that come up in our generic construction based on one-way functions.

- The algorithm $\mathsf{LE.AltSetup}$ is a probabilistic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$ and a polynomially-bounded integer $q = q(\lambda)$, and outputs a state st_0.
- The algorithm $\mathsf{LE.Equiv}_1$ is a probabilistic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$ a state st_0, a polynomially-bounded integer $q = q(\lambda)$ and a tag $\tau \in \mathcal{T}_\lambda$, and outputs a common-reference string $\widehat{\mathsf{crs}}$, commitments $\widehat{c_1}, \ldots, \widehat{c_q}$ and a state st_1.
- The algorithm $\mathsf{LE.Equiv}_2$ is a deterministic algorithm that receives as input the security parameter $\lambda \in \mathbb{N}$, an element $x \in \mathcal{X}_\lambda$, an index $i \in [q]$, a state st_1 and a tag $\tau \in \mathcal{T}_\lambda$, and outputs a decommitment $\widehat{d}$.

A commitment scheme as described above should satisfy the standard correctness requirement of commitment schemes. That is, for any security parameter $\lambda \in \mathbb{N}$, for any tag $\tau \in \mathcal{T}_\lambda$, for any polynomially-bounded $q = q(\lambda)$, for any $i \in [q]$ and for any $x \in \mathcal{X}_\lambda$ it holds that

$$\Pr\left[\mathsf{LE.Decommit}(1^\lambda, \mathsf{crs}, c, d, i, \tau) = x\right] = 1,$$

where $\mathsf{crs} \leftarrow \mathsf{LE.Setup}(1^\lambda, q)$ and $(c, d) \leftarrow \mathsf{LE.Commit}(1^\lambda, \mathsf{crs}, x, i, \tau)$, and the probability is taken over the internal randomness of all algorithms.

The following two definitions formally capture our local equivocability and all-but-one binding requirements.

Definition 4.1 (Local equivocability). *A commitment scheme* $\mathcal{LE} = (\mathsf{LE.Setup}, \mathsf{LE.Commit}, \mathsf{LE.Decommit}, \mathsf{LE.AltSetup}, \mathsf{LE.Equiv}_1, \mathsf{LE.Equiv}_2)$ *over a domain* $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ *and a tag space* $\mathcal{T} = \{\mathcal{T}_\lambda\}_{\lambda\in\mathbb{N}}$ *is* locally equivocable *if the following requirements hold:*

- **Equivocation correctness:** *For any* $\lambda \in \mathbb{N}$*, for any* $\tau \in \mathcal{T}_\lambda$*, for any polynomially-bounded* $q = q(\lambda)$*, for any* $i \in [q]$ *and for any* $x \in \mathcal{X}_\lambda$ *it holds that*

$$\Pr\left[\mathsf{LE.Decommit}(1^\lambda, \widehat{\mathsf{crs}}, \widehat{c_i}, \widehat{d}, i, \tau) = x\right] = 1,$$

  where $(\widehat{\mathsf{crs}}, \widehat{c_1}, \ldots, \widehat{c_q}, \mathsf{st}_1) \leftarrow \mathsf{LE.Equiv}_1(1^\lambda, \mathsf{LE.AltSetup}(1^\lambda), q, \tau)$ *and* $\widehat{d} = \mathsf{LE.Equiv}_2(1^\lambda, x, i, \mathsf{st}_1)$*, and the probability is taken over the internal randomness of all algorithms.*
- **Equivocation indistinguishability:** *For any probabilistic polynomial-time algorithm* $\mathcal{A}$*, there exists a negligible function* $\nu(\cdot)$ *such that for any polynomially bounded* $q = q(\lambda)$ *it holds that*

$$\begin{aligned}&\mathbf{Adv}^{\mathsf{LocalEquiv}}_{\mathcal{LE},q,\mathcal{A}}(\lambda)\\ &\quad\stackrel{\mathsf{def}}{=} \left|\Pr\left[\mathsf{IndParam}_{\mathcal{LE},q,\mathcal{A},0}(\lambda)\right] - \Pr\left[\mathsf{IndParam}_{\mathcal{LE},q,\mathcal{A},1}(\lambda)\right]\right| \leq \nu(\lambda)\end{aligned}$$

  for all sufficiently large $\lambda \in \mathbb{N}$*, where for any bit* $b \in \{0,1\}$ *the experiment* $\mathsf{IndParam}_{\mathcal{LE},q,\mathcal{A},b}(\lambda)$ *is defined as follows:*
  1. $(\tau, x_1, \ldots, x_q, \mathsf{st}_\mathcal{A}) \leftarrow \mathcal{A}(1^\lambda)$.
  2. $\mathsf{crs}_0 \leftarrow \mathsf{LE.Setup}(1^\lambda, q)$.

3. $(c_{0,i}, d_{0,i}) \leftarrow \mathsf{LE.Commit}(1^\lambda, \mathsf{crs}, x_i, i, \tau)$ *for each* $i \in [q]$.
4. $\mathsf{st}_0 \leftarrow \mathsf{LE.AltSetup}(1^\lambda, q)$.
5. $(\mathsf{crs}_1, c_{1,1}, \ldots, c_{1,q}, \mathsf{st}_1) \leftarrow \mathsf{LE.Equiv}_1(1^\lambda, \mathsf{st}_0, q, \tau)$.
6. $d_{1,i} = \mathsf{LE.Equiv}_2(1^\lambda, x_i, i, \mathsf{st}_1)$ *for each* $i \in [q]$.
7. $b' \leftarrow \mathcal{A}(\mathsf{st}_\mathcal{A}, \mathsf{crs}_b, (c_{b,i})_{i\in[q]}, (d_{b,i})_{i\in[q]})$.
8. *Output* b'.

Intuitively, the all-but-one binding property requires that an adversary which generates equivocable public parameters (via the $\mathsf{LE.Equiv}_1$ algorithm) using a tag τ of their choice, cannot break the binding property with respect to these parameters and a different tag $\tau' \neq \tau$.

Definition 4.2 (All-but-one binding). *A commitment scheme* $\mathcal{LE} = (\mathsf{LE.Setup}, \mathsf{LE.Commit}, \mathsf{LE.Decommit}, \mathsf{LE.AltSetup}, \mathsf{LE.Equiv}_1, \mathsf{LE.Equiv}_2)$ *over a domain* $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ *and a tag space* $\mathcal{T} = \{\mathcal{T}_\lambda\}_{\lambda\in\mathbb{N}}$ *is* all-but-one binding *if for any probabilistic polynomial-time algorithm* $\mathcal{A}$ *there exists a negligible function* $\nu(\cdot)$ *such that for polynomially-bounded* $q = q(\lambda)$ *it holds that*

$$\mathbf{Adv}^{\mathsf{ABOBind}}_{\mathcal{LE},q,\mathcal{A}}(\lambda) \stackrel{\mathsf{def}}{=} \Pr\left[\mathsf{ABOBind}^{\mathcal{LE}}_{q,\mathcal{A}}(\lambda) = 1\right] \leq \nu(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$, *where the experiment* $\mathsf{ABOBind}^{\mathcal{LE}}_{q,\mathcal{A}}(\lambda)$ *is defined as follows:*

1. $(\tau, \mathsf{st}_\mathcal{A}) \leftarrow \mathcal{A}(1^\lambda)$, *where* $\tau \in \mathcal{T}_\lambda$.
2. $\mathsf{st}_0 \leftarrow \mathsf{LE.AltSetup}(1^\lambda, q)$.
3. $\rho \leftarrow \{0,1\}^r$, *where* $r = r(\lambda)$ *is the number of random coins used by* $\mathsf{LE.Equiv}_1$ *on security parameter* $\lambda \in \mathbb{N}$.
4. $(\widehat{\mathsf{crs}}, \widehat{c_1}, \ldots, \widehat{c_q}, \mathsf{st}_1) = \mathsf{LE.Equiv}_1(1^\lambda, \mathsf{st}_0, q, \tau; \rho)$.
5. $(c, d, d', i, \tau') \leftarrow \mathcal{A}(\mathsf{st}_\mathcal{A}, \mathsf{st}_0, \rho)$.
6. $x = \mathsf{LE.Decommit}(1^\lambda, \widehat{\mathsf{crs}}, c, d, i, \tau')$ *and* $x' = \mathsf{LE.Decommit}(1^\lambda, \widehat{\mathsf{crs}}, c, d', i, \tau')$.
7. *Output* 1 *if* $\tau' \neq \tau$, $x \neq \perp$, $x' \neq \perp$ *and* $x \neq x'$. *Otherwise, output* 0.

Comparing our notion to identity-based and simulation-sound trapdoor commitments. Having formally defined our notion of a locally-equivocable commitment scheme with all-but-one binding, we can now compare it to Fischlin's notion of an identity-based trapdoor commitment scheme [Fis01, Ch. 2.6]. Both notions are obtained by augmenting the standard notion of a non-interactive tag-based commitment scheme with equivocability and all-but-one binding requirements. Our requirements, however, are more strict compared to those of Fischlin, both in terms of equivocability and in terms of all-but-one binding.

First, in terms of equivocability, Fischlin asks for an equivocation algorithm that produces an equivocable common-reference string and a *single* equivocable commitment which should be indistinguishable from an honestly-generated common-reference string and an honestly-generated commitment. However, for our construction of a non-malleable vector commitment scheme, producing a

single equivocable commitment seems insufficient. Thus, we ask for an equivocation algorithm that produces an equivocable common-reference string and q equivocable commitments (where $q = q(\lambda)$ is any predetermined polynomial) which should be indistinguishable from an honestly-generated common-reference string and an honestly-generated vector of q independent commitments. We note that such a requirement does not necessarily follow from the case $q = 1$ due to potential dependencies between the equivocable common-reference string and the single equivocable commitment that may be efficiently identifiable when producing more than a single equivocable commitment (this is evident in our generic construction based any non-interactive equivocable commitment scheme, where the common-reference string grows with q).

Second, in terms of all-but-one binding, Fischlin asks that when generating an equivocable common-reference string with respect to a predetermined tag τ, commitments with respect to all other tags should still be binding even when given the trapdoor associated with τ. For our construction we strengthen this requirements, and ask that commitments with respect to all other tags should still be binding even when given the trapdoor associated with τ and the internal randomness of the equivocation algorithm.

An additional related notion is that of a simulation-sound trapdoor commitment scheme, put forth by Garay, MacKenzie, and Yang [GMY03], which can be seen as augmenting standard trapdoor commitments [Rey01, Ch. A.5] with tags. Garay et al. also considered an enhanced binding property, requiring that binding with respect to a tag τ should be preserved, even if the attacker can obtain a single "fake" opening (using the trapdoor) for any commitment with respect to τ, as well as an unbounded number of openings for any commitment with respect to any other tag $\tau' \neq \tau$. This notion seems to be incomparable to our notion of locally-equivocable commitments with all-but-one binding. First, the trapdoor in simulation-sound trapdoor commitments is a global trapdoor generated by the honest parameters generation algorithm. There are no alternative procedures to generate equivocable parameters and commitments, and the trapdoor is not tied to any particular tag. This means that knowledge of the trapdoor allows one to open any (honestly generated) commitment to any value they desires. Second, whereas in our enhanced binding property the attacker receives the trapdoor associated with a tag τ of their choice, the attacker in the notion of Garay et al. does not receive the trapdoor, but only openings computed using it (this is unavoidable, since knowledge of the trapdoor in their notion allows the attacker to break binding with respect to all tags).

5 Our Construction of a Non-Malleable Vector Commitment Scheme

Our construction relies on the following building blocks:

- A vector commitment scheme $\mathcal{VC} = (\mathsf{VC.Setup}, \mathsf{VC.Commit}, \mathsf{VC.Open}, \mathsf{VC.Verify})$ over a domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ (see Section 2.2).[9]
- A locally-equivocable commitment scheme with all-but-one binding $\mathcal{LE} = (\mathsf{LE.Setup}, \mathsf{LE.Commit}, \mathsf{LE.Decommit}, \mathsf{LE.AltSetup}, \mathsf{LE.Equiv}_1, \mathsf{LE.Equiv}_2)$ over the domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda\in\mathbb{N}}$ and a tag space $\mathcal{T} = \{\mathcal{T}_\lambda\}_{\lambda\in\mathbb{N}}$ (see Section 4) with tags of length $t = t(\lambda)$ bits.
- A one-time strongly-unforgeable signature scheme $\mathcal{SIG} = (\mathsf{Sig.Gen}, \mathsf{Sig.Sign}, \mathsf{Sig.Verify})$ (see Section 2.3). Let $v = v(\lambda)$ denote the bit-length of the verification keys that are produced by $\mathsf{Sig.Gen}(1^\lambda)$.
- A universal one-way hash family $\mathcal{H} = \{\mathcal{H}_\lambda\}_{\lambda\in\mathbb{N}}$ (see Section 2.4), where each $\mathcal{H}_\lambda$ consists of functions mapping $v(\lambda)$-bit strings to $t(\lambda)$-bit strings for every security parameter $\lambda \in \mathbb{N}$.

As discussed in Section 1.3, from a foundational perspective, the above building blocks can all be based on the existence of any vector commitment scheme. Additional, from a more practical perspective, the above building blocks can all be realized based on a variety of number-theoretic assumptions leading to practical implementations.

Given the above building blocks, our construction of a non-malleable vector commitment scheme, denoted $\mathsf{nm}\mathcal{VC} = (\mathsf{nmVC.Setup}, \mathsf{nmVC.Commit}, \mathsf{nmVC.Open}, \mathsf{nmVC.Verify})$, is defined as follows:

A non-malleable vector commitment scheme $\mathsf{nm}\mathcal{VC}$

$\mathsf{nmVC.Setup}(1^\lambda, q)$:

1. Sample $\mathsf{crs}_{\mathsf{LE}} \leftarrow \mathsf{LE.Setup}(1^\lambda, q)$, $\mathsf{crs}_{\mathsf{VC}} \leftarrow \mathsf{VC.Setup}(1^\lambda, q)$ and $h \leftarrow \mathcal{H}_\lambda$.
2. Output $\mathsf{crs} = \mathsf{crs}_{\mathsf{LE}} \| \mathsf{crs}_{\mathsf{VC}} \| h$.

$\mathsf{nmVC.Commit}(1^\lambda, \mathsf{crs}, (x_1, \ldots, x_q))$:

1. Parse crs as $\mathsf{crs}_{\mathsf{LE}} \| \mathsf{crs}_{\mathsf{VC}} \| h$.
2. Sample $(sk, vk) \leftarrow \mathsf{Sig.Gen}(1^\lambda)$ and compute $\tau = h(vk)$.
3. For each $i \in [q]$ compute $(c_i, d_i) \leftarrow \mathsf{LE.Commit}(1^\lambda, \mathsf{crs}_{\mathsf{LE}}, x_i, i, \tau)$.
4. Compute $(\mathsf{vcom}_0, \mathsf{st}_0) \leftarrow \mathsf{VC.Commit}(1^\lambda, \mathsf{crs}_{\mathsf{VC}}, (c_1, \ldots, c_q))$ and $\sigma \leftarrow \mathsf{Sig.Sign}_{sk}(\mathsf{vcom}_0)$.
5. Output $(\mathsf{vcom}, \mathsf{st})$, where $\mathsf{vcom} = \mathsf{vcom}_0 \| vk \| \sigma$ and $\mathsf{st} = \mathsf{st}_0 \| c_1 \| \cdots \| c_q \| d_1 \| \cdots \| d_q$.

$\mathsf{nmVC.Open}(1^\lambda, \mathsf{crs}, \mathsf{vcom}, \mathsf{st}, i)$:

1. Parse crs as $\mathsf{crs}_{\mathsf{LE}} \| \mathsf{crs}_{\mathsf{VC}} \| h$, vcom as $\mathsf{vcom}_0 \| vk \| \sigma$ and st as $\mathsf{st}_0 \| c_1 \| \cdots \| c_q \| d_1 \| \cdots \| d_q$.
2. Compute $\pi_0 \leftarrow \mathsf{VC.Open}(1^\lambda, \mathsf{crs}_{\mathsf{VC}}, \mathsf{vcom}_0, \mathsf{st}_0, i)$.
3. Output $\pi = c_i \| d_i \| \pi_0$.

[9] We emphasize that the security of our construction does not rely on $\mathcal{VC}$ providing any flavor of hiding or succinctness, and this is discussed below in the overview of our proof.

$\mathsf{nmVC.Verify}(1^\lambda, \mathsf{crs}, \mathsf{vcom}, i, x, \pi)$:

1. Parse crs as $\mathsf{crs}_{\mathsf{LE}} \| \mathsf{crs}_{\mathsf{VC}} \| h$, vcom as $\mathsf{vcom}_0 \| vk \| \sigma$ and st as π as $c_i \| d_i \| \pi_0$.
2. Compute $\tau := h(vk)$.
3. Output 1 if all of the following conditions hold:
   - $\mathsf{Sig.Verify}_{vk}(\mathsf{vcom}_0, \sigma) = 1$.
   - $\mathsf{VC.Verify}(1^\lambda, \mathsf{crs}_{\mathsf{VC}}, \mathsf{vcom}_0, i, c_i, \pi_0) = 1$.
   - $\mathsf{LE.Decommit}(1^\lambda, \mathsf{crs}_{\mathsf{LE}}, c_i, d_i, i, \tau) = x$.

   Otherwise, output 0.

Finally, we note that for simplifying our construction and its proof, the length of the secret state $\mathsf{st} = \mathsf{st}_0 \| c_1 \| \cdots \| c_q \| d_1 \| \cdots \| d_q$ produced by the commitment algorithm $\mathsf{nmVC.Commit}$ in the above description depends linearly on q but this can be easily avoided whenever the committed vector $(x_1, \ldots, x_q)$ is additionally provided. Specifically, given $x_1, \ldots, x_q$, the entire sequence of values $c_1, \ldots, c_q, d_1, \ldots, d_q$ can be replaced with a single key K for a pseudorandom function PRF that will allow the algorithm $\mathsf{nmVC.Open}$ to recompute any of these values when needed. Specifically, instead of computing $(c_i, d_i) \leftarrow \mathsf{LE.Commit}(1^\lambda, \mathsf{crs}_{\mathsf{LE}}, x_i, i, \tau)$ by feeding the algorithm $\mathsf{LE.Commit}$ with a fresh random string $r_i \leftarrow \{0,1\}^*$, we can instead feed it with a pseudorandom string $r_i = \mathsf{PRF}_K(\mathsf{crs}_{\mathsf{LE}}, x_i, i, \tau)$ which is reproducible via knowledge of K and x_i.

Security. The following theorem captures the security of our construction, showing that it satisfies our notion of non-malleability for vector commitment schemes (recall Definition 3.1) based on the security of its underlying building blocks: (1) a vector commitment scheme $\mathcal{VC}$, (2) a locally-equivocable commitment scheme with all-but-one binding $\mathcal{LE}$, (3) a one-time strongly-unforgeable signature scheme $\mathcal{SIG}$, and (4) a universal one-way hash family $\mathcal{H}$.

Theorem 5.1. *For every probabilistic polynomial-time algorithm $\mathcal{A}$ and polynomial $q = q(\lambda)$, there exists a probabilistic polynomial-time algorithm $\mathcal{S}_{\mathcal{A}}$ such that the following holds: For any probabilistic polynomial-time algorithm $\mathcal{R}$, there are probabilistic polynomial-time algorithms $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4$ and $\mathcal{B}_5$ such that*

$$\begin{aligned}\mathbf{Adv}^{\mathsf{NM}}_{\mathsf{nm}\mathcal{VC},q,\mathcal{A},\mathcal{S}_{\mathcal{A}},\mathcal{R},\mathcal{D}}(\lambda) \leq\ & \mathbf{Adv}^{\mathsf{LocalEq}}_{\mathcal{LE},q,\mathcal{B}_1}(\lambda) + \mathbf{Adv}^{\mathsf{Forge}}_{\mathcal{SIG},\mathcal{B}_2}(\lambda) + \mathbf{Adv}^{\mathsf{UOWHF}}_{\mathcal{H},\mathcal{B}_3}(\lambda) \\ & + 2 \cdot \left(\mathbf{Adv}^{\mathsf{ABOBind}}_{\mathcal{LE},q,\mathcal{B}_4}(\lambda) + \mathbf{Adv}^{\mathsf{PosBind}}_{\mathcal{VC},q,\mathcal{B}_5}(\lambda) \right)\end{aligned}$$

for every $\lambda \in \mathbb{N}$.

Due to space limitations, the formal proof of Theorem 5.1 is provided in the full version.

References

[ABC+12] J. H. Ahn, D. Boneh, J. Camenisch, S. Hohenberger, abhi shelat, and B. Waters. Computing on authenticated data. In *Proceedings of the 9th Theory of Cryptography Conference*, pages 169–191, 2012.

[Bar01] B. Barak. How to go beyond the black-box simulation barrier. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 106–115, 2001.

[BBB+18] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 315–334, 2018.

[BBF19] D. Boneh, B. Bünz, and B. Fisch. Batching techniques for accumulators with applications to IOPs and stateless blockchains. In *Advances in Cryptology – CRYPTO '19*, pages 561–586, 2019.

[BCG+14] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 459–474, 2014.

[BdM93] J. Benaloh and M. de Mare. One-way accumulators: A decentralized alternative to digital signatures. In *Advances in Cryptology – EUROCRYPT '93*, pages 274–285, 1993.

[BGV11] S. Benabbas, R. Gennaro, and Y. Vahlis. Verifiable delegation of computation over large datasets. In *Advances in Cryptology – CRYPTO '11*, pages 111–131, 2011.

[Bic17] M. Bichler. Market Design: A Linear Programming Approach to Auctions and Matching. Cambridge University Press, 2017.

[BP97] N. Barić and B. Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In *Advances in Cryptology – EUROCRYPT '97*, pages 480–494, 1997.

[BR93] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 62–73, 1993.

[BS07] M. Bellare and S. Shoup. Two-tier signatures, strongly unforgeable signatures, and fiat-shamir without random oracles. In *Proceedings of the 10th International Conference on Theory and Practice of Public-Key Cryptography*, pages 201–216, 2007.

[BSW06] D. Boneh, E. Shen, and B. Waters. Strongly unforgeable signatures based on computational Diffie-Hellman. In *Proceedings of the 9th International Conference on Theory and Practice of Public-Key Cryptography*, pages 229–240, 2006.

[But14] V. Buterin. Ethereum: A next-generation smart contract and decentralized application platform. Available at `https://ethereum.org/en/whitepaper/`, 2014.

[But17] V. Buterin. The stateless client concept, 2017. Available at `https://ethresear.ch/t/the-stateless-client-concept/172`.

[CDV06] D. Catalano, Y. Dodis, and I. Visconti. Mercurial commitments: Minimal assumptions and efficient constructions. In *Proceedings of the 3rd Theory of Cryptography Conference*, pages 120–144, 2006.

[CF01] R. Canetti and M. Fischlin. Universally composable commitments. In *Advances in Cryptology – CRPYTO '01*, pages 19–40, 2001.

[CF13] D. Catalano and D. Fiore. Vector commitments and their applications. In *Proceedings of the 16th International Conference on Practice and Theory in Public-Key Cryptography*, pages 55–72, 2013.

[CFG+20] M. Campanelli, D. Fiore, N. Greco, D. Kolonelos, and L. Nizzardo. Incrementally aggregatable vector commitments and applications to verifiable decentralized storage. In *Advances in Cryptology – ASIACRYPT '20*, pages 3–35, 2020.

[CHL+13] M. Chase, A. Healy, A. Lysyanskaya, T. Malkin, and L. Reyzin. Mercurial commitments with applications to zero-knowledge sets. *Journal of Cryptology*, 26(2):251–279, 2013.

[CIO98] G. D. Crescenzo, Y. Ishai, and R. Ostrovsky. Non-interactive and non-malleable commitment. In *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing*, pages 141–150, 1998.

[CKO+01] G. D. Crescenzo, J. Katz, R. Ostrovsky, and A. D. Smith. Efficient and non-interactive non-malleable commitment. In *Advances in Cryptology – EUROCRYPT '01*, pages 40–59, 2001.

[CL02] J. Camenisch and A. Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In *Advances in Cryptology – CRYPTO '02*, pages 61–76, 2002.

[COS+17] M. Ciampi, R. Ostrovsky, L. Siniscalchi, and I. Visconti. Four-round concurrent non-malleable commitments from one-way functions. In *Advances in Cryptology – CRYPTO '17*, pages 127–157, 2017.

[CRF+11] D. Catalano, M. D. Raimondo, D. Fiore, and M. Messina. Zero-knowledge sets with short proofs. *IEEE Transaction on Information Theory*, 57(4):2488–2502, 2011.

[DDN00] D. Dolev, C. Dwork, and M. Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000.

[DG03] I. Damgard and J. Groth. Non-interactive and reusable non-malleable commitment schemes. In *Proceedings of the 35th Annual ACM Symposium on the Theory of Computing*, pages 426–437, 2003.

[DKN+04] Y. Dodis, A. Kiayias, A. Nicolosi, and V. Shoup. Anonymous identification in ad hoc groups. In *Advances in Cryptology – EUROCRYPT '04*, pages 609–626, 2004.

[FF00] M. Fischlin and R. Fischlin. Efficient non-malleable commitment schemes. In *Advances in Cryptology – CRYPTO '00*, pages 413–431, 2000.

[Fis01] M. Fischlin. Trapdoor commitment schemes and their applications. PhD Thesis, University of Frankfurt (available at `https://www.math.uni-frankfurt.de/~dmst/research/phdtheses/mfischlin.dissertation.2001.html`), 2001.

[FVY14] C. Fromknecht, D. Velicanu, and S. Yakoubov. A decentralized public key infrastructure with identity retention. Cryptology ePrint Archive, Report 2014/803, 2014.

[GLO+12] V. Goyal, C.-K. Lee, R. Ostrovsky, and I. Visconti. Constructing non-malleable commitments: A black-box approach. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, pages 51–60, 2012.

[GM06] R. Gennaro and S. Micali. Independent zero-knowledge sets. In *Proceedings of the 33th International Colloquium on Automata, Languages and Programming*, pages 34–45, 2006.

[GMY03] J. A. Garay, P. MacKenzie, and K. Yang. Strengthening zero-knowledge protocols using signatures. In *Advances in Cryptology – EUROCRYPT '03*, pages 177–194, 2003.

[GPR16] V. Goyal, O. Pandey, and S. Richelson. Textbook non-malleable commitments. In *Proceedings of the 48th annual ACM Symposium on Theory of Computing*, pages 1128–1141, 2016.

[GR97] R. Gennaro and P. Rohatgi. How to sign digital streams. In *Advances in Cryptology – CRYPTO '97*, pages 180–197, 1997.

[GRW+20] S. Gorbunov, L. Reyzin, H. Wee, and Z. Zhang. Pointproofs: Aggregating proofs for multiple vector commitments. In *Proceedings of the 27th ACM Conference on Computer and Communications Security*, pages 2007–2023, 2020.

[HIL+99] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

[Khu17] D. Khurana. Round optimal concurrent non-malleability from polynomial hardness. In *Proceedings of the 15th Theory of Cryptography Conference*, pages 139–171, 2017.

[Kil92] J. Kilian. A note on efficient zero-knowledge proofs and arguments. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 723–732, 1992.

[KSS+16] J. Krupp, D. Schröder, M. Simkin, D. Fiore, G. Ateniese, and S. Nürnberger. Nearly optimal verifiable data streaming. In *Proceedings of the 19th International Conference on Practice and Theory in Public-Key Cryptography*, pages 417–445, 2016.

[Lam79] L. Lamport. Constructing digital signatures from a one way function. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory, 1979.

[LM19] R. W. F. Lai and G. Malavolta. Subvector commitments with application to succinct arguments. In *Advances in Cryptology – CRYPTO '19*, pages 530–560, 2019.

[LP09] H. Lin and R. Pass. Non-malleability amplification. In *Proceedings of the 41st annual ACM Symposium on Theory of Computing*, pages 189–198, 2009.

[LP11] H. Lin and R. Pass. Constant-round non-malleable commitments from any one-way function. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, pages 705–714, 2011.

[LPV08] H. Lin, R. Pass, and M. Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In *Proceedings of the 5th Theory of Cryptography Conference*, pages 571–588, 2008.

[LY10] B. Libert and M. Yung. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In *Proceedings of the 7th Theory of Cryptography Conference*, pages 499–517, 2010.

[Mer87] R. C. Merkle. A digital signature based on a conventional encryption function. In *Advances in Cryptology – CRYPTO '87*, pages 369–378, 1987.

[MGG+13] I. Miers, C. Garman, M. Green, and A. D. Rubin. Zerocoin: Anonymous distributed e-cash from bitcoin. In *IEEE Symposium on Security and Privacy*, pages 397–411, 2013.

[Mic94] S. Micali. CS proofs. In *Proceedings of the 35th Annual IEEE Symposium on the Foundations of Computer Science*, pages 436–453, 1994.

[MND+04] C. Martel, G. Nuckolls, P. Devanbu, M. Gertz, A. Kwong, and S. G. Stubblebine. A general model for authenticated data structures. *Algorithmica*, 39(1):21–24, 2004.

[MRK03] S. Micali, M. O. Rabin, and J. Kilian. Zero-knowledge sets. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 80–91, 2003.

[Nak08] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Available at `https://bitcoin.org/bitcoin.pdf`, 2008.

[Nao91] M. Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, 1991.

[Ngu05] L. Nguyen. Accumulators from bilinear pairings and applications. In *Topics in Cryptology – CT-RSA '05*, pages 275–292, 2005.

[NN98] M. Naor and K. Nissim. Certificate revocation and certificate update. In *Proceedings of the 7th USENIX Security Symposium*, pages 217–228, 1998.

[NY89] M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 33–43, 1989.

[OWW+20] A. Ozdemir, R. Wahby, B. Whitehat, and D. Boneh. Scaling verifiable computation using efficient set accumulators. In *Proceedings of the 29th USENIX Security Symposium*, pages 2075–2092, 2020.

[PPV08] O. Pandey, R. Pass, and V. Vaikuntanathan. Adaptive one-way functions and applications. In *Advances in Cryptology – CRYPTO '08*, pages 57–74, 2008.

[PR05] R. Pass and A. Rosen. Concurrent non-malleable commitments. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 563–572, 2005.

[PR08] R. Pass and A. Rosen. New and improved constructions of nonmalleable cryptographic protocols. *SIAM Journal on Computing*, 38(2):702–752, 2008.

[PW10] R. Pass and H. Wee. Constant-round non-malleable commitments from subexponential one-way functions. In *Advances in Cryptology – EUROCRYPT '10*, pages 638–655, 2010.

[Rey01] L. Reyzin. Zero-knowledge with public keys. PhD Thesis, Massachusetts Institute of Technology (available at `https://www.cs.bu.edu/~reyzin/phd-thesis.html`), 2001.

[Rom90] J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, pages 387–394, 1990.

[STS99] T. Sander and A. Ta-Shma. Auditable, anonymous electronic cash. In *Advances in Cryptology – CRYPTO '99*, pages 555–572, 1999.

[SvDJ+12] E. Stefanov, M. van Dijk, A. Jules, and A. Opera. Iris: a scalable cloud file system with efficient integrity checks. In *Proceedings of the 28th Annual Computer Security Applications Conference*, pages 229–238, 2012.

[TAB+20] A. Tomescu, I. Abraham, V. Buterin, J. Drake, D. Feist, and D. Khovratovich. Aggregatable subvector commitments for stateless cryptocurrencies. In *Proceedings of the 12th International Conference on Security and Cryptography for Networks*, pages 45–64, 2020.

[Tod16] P. Todd. Making UTXO set growth irrelevant with low-latency delayed TXO commitments, 2016. Available at `https://petertodd.org/2016/delayed-txo-commitments`.

[Wee10] H. Wee. Black-box, round-efficient secure computation via non-malleability amplification. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 531–540, 2010.