



Extractable Witness Encryption for KZG Commitments and Efficient Laconic OT

Nils Fleischhacker^{1*} , Mathias Hall-Andersen² , and Mark Simkin³ 

¹ Ruhr University Bochum
mail@nilsfleischhacker.de

² Aarhus University
ma@cs.au.dk

³ Ethereum Foundation
mark.simkin@ethereum.org

Abstract. We present a concretely efficient and simple extractable witness encryption scheme for KZG polynomial commitments. It allows to encrypt a message towards a triple $(\text{com}, \alpha, \beta)$, where com is a KZG commitment for some polynomial f . Anyone with an opening for the commitment attesting $f(\alpha) = \beta$ can decrypt, but without knowledge of a valid opening the message is computationally hidden. Our construction is simple and highly efficient. The ciphertext is only a single group element. Encryption and decryption both require a single pairing evaluation and a constant number of group operations.

Using our witness encryption scheme, we construct a simple and highly efficient laconic OT protocol, which significantly outperforms the state of the art in most important metrics.

1 Introduction

The polynomial commitment scheme of Kate, Zaverucha, and Goldberg (KZG) [KZG10] is a powerful tool that has allowed for constructing a variety of advanced cryptographic primitives. Many concretely efficient vector commitments [CF13] with all kinds of additional functionalities or security properties are extensions of the KZG polynomial commitment scheme [TAB⁺20, GRWZ20, LPR22, SCP⁺22, WUP23]. Many of the currently most efficient proof systems [MBKM19, GWC19, CHM⁺20] make crucial use of KZG commitments as part of their constructions.

In general, a polynomial commitment scheme allows for committing to a polynomial $f(X)$, such that one can later provide openings, attesting that $f(\alpha) = \beta$ for some chosen evaluation point α . The polynomial commitment scheme should ensure that the committed polynomial is at most of some degree d , that it is position binding in the sense that one cannot open the commitment to two different evaluations at the same point α , and in some cases it may also be desirable that the commitment itself hides the committed polynomial.

The KZG polynomial commitment scheme is highly efficient as both commitments and openings consists of a single group element, no matter how large the degree of the polynomial is. From a security perspective, the construction requires a trusted setup and can be proven secure in the algebraic group model (AGM) [FKL18] under a q -type variant of the standard discrete logarithm assumption.

Given the widespread usefulness of KZG commitments, it seems natural to ask the following somewhat abstract questions: Can we expand the toolkit surrounding KZG commitments in a fundamental way that would allow us to solve an even larger set of problems via these polynomial commitments?

* Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972.

1.1 Our Contribution

In this work, we present a simple, concretely efficient extractable witness encryption [GGSW13, GKP⁺13] scheme for KZG commitments and show how it can be used to construct concretely efficient laconic OT [CDG⁺17]. Concretely, we make the following contributions:

Witness Encryption for KZG Commitments. We present an encryption scheme that takes a KZG commitment com , evaluation point α , evaluation β , as well as a message m as input and produces ciphertext ct . Correctness of the encryption scheme allows anybody, who knows an opening for com that attests that the committed polynomial evaluates to β at point α , to decrypt the ciphertext ct . Informally, security of the encryption scheme ensures that anybody, who can decrypt ct , can also compute a valid opening with respect to $(\text{com}, \alpha, \beta)$. Now, if com is a commitment to polynomial $f(X)$ and ct is an encryption for $(\text{com}, \alpha, \beta)$, where $\beta \neq f(\alpha)$, then ct computationally hides the message, due to the evaluation binding property of the polynomial commitment scheme.

We note that for *any* tuple $(\text{com}, \alpha, \beta)$, there *exists* an opening and for this reason the vanilla notion of witness encryption as defined by Garg et al. [GGSW13] would not suffice to guarantee that the message is hidden. By focusing on extractable witness encryption as defined by Goldwasser et al. [GKP⁺13], we can ensure that the message is hidden, whenever computing the opening is assumed to be hard.

Our encryption scheme is highly efficient in terms of both ciphertext size and computational costs. One ciphertext for the scheme as described above is comprised of a single group element, encryption requires two group operations and one pairing evaluation and decryption requires a single pairing evaluation. The construction is generalized for multiple opening constraints. In that case the costs grow linearly in the number of constraints.

Efficient Laconic Oblivious Transfer. Using our witness encryption scheme for KZG commitments, we present a simple and highly-efficient protocol for laconic oblivious transfer [CDG⁺17]. Oblivious transfer [Rab81] is an interactive protocol between a sender holding messages m_0 and m_1 and a receiver holding a choice bit b . At the end of the protocol, the receiver should only learn m_b and the sender should learn nothing at all.

In laconic oblivious transfer, the receiver holds a database $D \in \{0, 1\}^n$ of n choice bits and publishes a digest $\text{digest} \leftarrow H(D)$, whose size is independent of the size of D . The sender can then repeatedly choose a message pair (m_0, m_1) , an index $i \in [n]$, and use the digest to compute a short message for the receiver, which allows them to obtain $m_{D[i]}$. As in regular OT, the receiver does not learn anything about the other message and the sender does not learn anything about the choice bit of the receiver.

Cho et al. [CDG⁺17] show that laconic OT is a useful building block for constructing secure computation protocols that operate on large inputs and for constructing multi-hop homomorphic encryption for RAM programs. Improving the efficiency of laconic OT, directly translates to faster protocols for these applications.

Implementation & Benchmarks. The full laconic OT construction was implemented⁴ in Rust. We provide benchmarks for various parameter regimes and show that our laconic OT construction outperforms the state of the art in most important metrics quite significantly.

⁴ <https://github.com/rot256/research-we-kzg>

1.2 Related Work

In the following, let us discuss related works from several research areas that intersect with our work here.

Witness Encryption. Witness encryption was originally introduced by Garg et al. [GGSW13], who presented a construction based on multilinear maps. Such encryption schemes take a statement x and a message m as input and produce a ciphertext ct . Correctness guarantees that given a witness w with $(x, w) \in \mathcal{R}$, one can decrypt ct and security ensures that the message is computationally hidden, if $x \notin \mathcal{L}$. It is important to note though that ct provides no explicit security guarantees, when $x \in \mathcal{L}$. In a different work, Garg et al. [GGH⁺13] showed that one can construct general-purpose witness encryption from indistinguishability obfuscation [BGI⁺01]. Two recent works by Tsabary [Tsa22] and Vaikuntanathan [VWW22] present construction based on new computational hardness assumptions. Extractable witness encryption is a strengthening of the original notion that was introduced by Goldwasser et al. [GKP⁺13]. It requires the existence of an extractor, which can recover a witness w from any adversary that can break the semantic security of the encryption scheme. How to construct general-purpose extractable witness encryption is currently unclear. Garg et al. [GGHW14] show that the existence of a type of special-purpose obfuscation would imply the non-existence of general-purpose extractable witness encryption. Benhamouda and Lin [BL18] introduce the notion of witness selectors, which lie somewhere between extractable and non-extractable witness encryption. They require semantic security of the encryption to hold if finding a witness for the relation is computationally hard.

While remotely practical general-purpose witness encryption from standard assumptions currently seems out of reach, there are several works [GS17, BL18, BL20, BJKL21, CFK22] that construct different types of witness encryption for specific languages. In the work of Garg and Srinivasan [GS17], they construct a primitive they call homomorphic proof commitments with encryption. Given a commitment com and a message m they allow for encrypting the message, such that a proof π that com is a commitment to 0 or 1 can be used to decrypt the message. Benhamouda and Lin [BL20] introduce witness encryption for NIZKs of commitments. Their scheme allows for encrypting towards a tuple (com, G, y) , where com is a commitment, G is a function, and y is an output. The witness for decryption is a non-interactive zero-knowledge proof for the statement “the message committed in com is m and $G(m) = y$ ”. In the presented construction, the commitment and proof size are at least linear in the statement size. Campanelli, Fiore, and Khoshakhlagh [CFK22] present succinct functional commitments with an associated witness encryption scheme with short commitments and short openings. They introduce a new computational hardness assumption and use it to prove their construction secure in the AGM.

In contrast to these works, we construct an extractable witness encryption scheme for a polynomial, rather than functional, commitment scheme from standard assumptions in the AGM. We rely on KZG commitments which have succinct commitments and openings and the ciphertexts of our construction are equally succinct. Our construction is surprisingly simple and in particular much simpler than the constructions discussed above. We believe this to be a significant advantage, when it comes to engineering and deploying a cryptographic primitive in the real world.

Laconic OT. The concept of laconic OT was originally introduced in a work by Cho et al. [CDG⁺17]. The authors presented a construction based on a special type of encryption scheme

in combination with garbled circuits. While asymptotically interesting, their construction makes non-blackbox use of cryptographic objects and is completely impractical when it comes to concrete performance.

Green, Jain, Van Laer [GJL23] focus on concretely efficient laconic OT protocols, but the security of their proposed construction relies on a new hardness assumption they introduce. In addition, they require a generic and somewhat expensive transformation via garbled circuits to achieve sender privacy. Döttling et al. [DKL⁺23] present another laconic OT protocol based on the learning with errors (LWE) with error leakage problem. They show that this problem reduces to standard LWE for certain parameter ranges.

From a concrete efficiency perspective, our construction is highly efficient in terms of bandwidth and computational overheads that are induced by the individual OT executions. The main drawback of our construction is a large crs, since we require the same setup as KZG commitments. In comparison the work of Green, Jain, Van Laer, our public parameters are smaller by a factor of 3-4x, our sender’s message size is smaller by a factor $\approx 5x$, and our receiver’s computation time is smaller by 1-6 *orders of magnitude* depending on the precise setting. In comparison to the work of Döttling et al., our sender’s OT message is smaller by 5 *orders of magnitude*. We provide a detailed discussion of the concrete efficiency of our construction in Section 6.

2 Preliminaries

In this section, we will recall some notation and standard definitions that we will use throughout the paper.

Notation. We denote by $\lambda \in \mathbb{N}$ the security parameter, by $\text{poly}(\lambda)$ any function that is bounded by a polynomial in λ and by $\text{negl}(\lambda)$ any negligible function in λ . An algorithm is PPT if it is modeled by a probabilistic Turing machine with a running time bounded by $\text{poly}(\lambda)$. Let S be a set. We write $x \leftarrow S$ for the process of sampling an element of S uniformly at random. For $n \in \mathbb{N}$, we write $[n]$ to denote the set $\{1, \dots, n\}$. For a vector $\mathbf{v} \in S^n$ and $i \in [n]$, we write v_i to denote its i -th component.

2.1 Algebraic Group Model

Some of our proofs will rely on the algebraic group model, which was introduced by Fuchsbauer, Kiltz, and Loss [FKL18]. In this model, one considers algebraic algorithms and, in particular, an algebraic adversary. Let \mathbb{G} of prime order q . Whenever an algorithm returns a value $Y \in \mathbb{G}$, it must also provide an algebraic representation of that element (e_1, \dots, e_n) , such that $Y = \prod_{i=1}^n X_i^{e_i}$, where (X_1, \dots, X_n) are group elements the algorithm has previously seen.

2.2 Pairings and Assumptions

Let GGen be the parameter generation algorithm that takes the security parameter λ as input and returns $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, g_1, g_2, q, e)$ as its output, where $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ are groups of prime order $p = p(\lambda)$, where $g_1 \in \mathbb{G}_1$ and $g_2 \in \mathbb{G}_2$ are generators and where $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is a bilinear map. For $i \in \{1, 2, T\}$, we write $[\alpha]_i$ as a shorthand notation for g_i^α and we write $[\alpha]_i + [\beta]_i = [\alpha + \beta]_i$ as a shorthand notation for the multiplication of two group elements. The type of pairing is irrelevant for our applications.

Definition 1 (ℓ -DLOG Assumption). *The ℓ -DLOG assumption holds with respect to GGen , if for any $\lambda \in \mathbb{N}$ and any PPT adversary \mathcal{A} , it holds that*

$$\Pr \left[\begin{array}{l} \tau = \tau' : \\ \tau' \leftarrow \mathcal{A}(\text{par}, ([\tau^0]_1, \dots, [\tau^\ell]_1, [1]_2, [\tau]_2)) \end{array} \middle| \begin{array}{l} \text{par} \leftarrow \text{GGen}(1^\lambda) \\ \tau \leftarrow \mathbb{F}_p \end{array} \right] \leq \text{negl}(\lambda),$$

where the probability is taken over the random coins of the group generation algorithm and the adversary and over the uniform choice of τ .

Remark 1. Throughout the paper we will omit explicitly spelling out the group parameter generation and we will assume all involved parties are always implicitly provided with the parameters.

2.3 Polynomial Commitments

A polynomial commitment allows for computing a short value com for a polynomial f of potentially high degree over a finite field \mathbb{F} . Later on, one can compute short openings that certify that the polynomial committed to by com evaluates to $\beta \in \mathbb{F}$ at some position $\alpha \in \mathbb{F}$. Polynomial commitment should be binding in the sense that it should be impossible to open the same point to two different values.

Definition 2 (Polynomial Commitments). *A polynomial commitment over \mathbb{F} is a tuple of PPT algorithms $\text{PC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Verify})$ defined as follows:*

$\text{ck} \leftarrow \text{Setup}(1^\lambda, 1^d)$: *The setup algorithm takes security parameter λ and degree upper bound d as input and returns commitment key ck .*

$\text{com} \leftarrow \text{Commit}(\text{ck}, f)$: *The commitment algorithm takes commitment key ck and polynomial $f \in \mathbb{F}[X]$ as input and returns commitment com .*

$\pi \leftarrow \text{Open}(\text{ck}, f, \alpha, \beta)$: *The opening algorithm takes commitment key ck , polynomial $f \in \mathbb{F}[X]$, point $\alpha \in \mathbb{F}$, and point $\beta \in \mathbb{F}$ as input and returns openings π .*

$b \leftarrow \text{Verify}(\text{ck}, \text{com}, \pi, \alpha, \beta)$: *The verification algorithm takes commitment key ck , commitment com , opening π , point $\alpha \in \mathbb{F}$, and point $\beta \in \mathbb{F}$ as input and returns a bit b .*

Definition 3 (Correctness). *A polynomial commitment scheme $\text{PC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Verify})$ over \mathbb{F} is correct, if for all $\lambda, d \in \mathbb{N}$, all $\text{ck} \leftarrow \text{Setup}(1^\lambda, 1^d)$, all polynomials $f \in \mathbb{F}[X]$ of degree at most d , all $\text{com} \leftarrow \text{Commit}(\text{ck}, f)$, all $\alpha, \beta \in \mathbb{F}$, and all $\pi \leftarrow \text{Open}(\text{ck}, f, \alpha, \beta)$ it holds that $\text{Verify}(\text{ck}, \text{com}, \pi, \alpha, \beta) = 1$.*

Definition 4 (Binding). *A polynomial commitment scheme $\text{PC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Verify})$ over \mathbb{F} is binding, if for all $\lambda, d \in \mathbb{N}$ and all PPT adversaries \mathcal{A} it holds that*

$$\Pr \left[\begin{array}{l} \beta_0 \neq \beta_1 \\ \wedge \text{Verify}(\text{ck}, \text{com}, \pi_0, \alpha, \beta_0) = 1 : \\ \wedge \text{Verify}(\text{ck}, \text{com}, \pi_1, \alpha, \beta_1) = 1 \end{array} \middle| \begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda, 1^d) \\ (\text{com}, \alpha, \beta_0, \beta_1, \pi_0, \pi_1) \leftarrow \mathcal{A}(\text{ck}) \end{array} \right] \leq \text{negl}(\lambda).$$

$\text{Setup}(1^\lambda, 1^d)$	$\text{Commit}(\text{ck}, f)$	$\text{Open}(\text{ck}, \text{com}, f, \alpha, \beta)$
$\tau \leftarrow \mathbb{F}_p$ return $([\tau^0]_1, \dots, [\tau^d]_1, [1]_2, [\tau]_2)$	$\text{com} := \sum_{i=0}^d f_i \cdot [\tau^i]_1$ return com	$q(X) := \frac{f(X) - \beta}{X - \alpha}$ $\pi := \sum_{i=0}^d q_i \cdot [\tau^i]_1$ return π
<hr style="width: 50%; margin: auto;"/> $\text{Verify}(\text{ck}, \text{com}, \pi, \alpha, \beta)$ <hr style="width: 50%; margin: auto;"/> return $e(\text{com} - [\beta]_1, [1]_2) \stackrel{?}{=} e(\pi, [\tau]_2 - [\alpha]_2)$		

Fig. 1. The KZG polynomial commitment scheme.

2.4 KZG Commitments

Let us recall the KZG polynomial commitment scheme. The scheme’s public parameters are powers of a secret point τ in the exponent of a group generator. Committing to a polynomial is done by evaluating it in the exponent at point τ , which can be done using the public parameters, but without knowledge of τ .

Theorem 1 ([CHM⁺20]). *If the d -DLOG assumption holds with respect to GGen , then KZG commitment scheme described in Figure 1 is a correct and binding polynomial commitment scheme in the AGM.*

It was shown by Feist and Khovratovich [FK23] that it is possible to open a KZG commitment in n positions much more efficiently than naively computing each opening separately. Specifically, it is possible to perform a batch opening at n points in time $\mathcal{O}(n \log^2(n))$. This can be further improved to $\mathcal{O}(n \log(n))$, if the points are powers of a root of unity. We refer to this batch opening algorithm as `BatchOpen`.

2.5 Weakly-Hiding Vector Commitments

A vector commitment allows for computing a short value `com` for a potentially long vector of messages (m_1, \dots, m_n) . Later on, one can compute short openings that certify that the vector committed to by `com` opens to m_i at some position $i \in [n]$. Vector commitment should be binding in the usual sense, meaning that no position can be opened to two different values. In addition, we will require a weak form of hiding from our vector commitments.

Definition 5. *A vector commitment with batch opening is a tuple of PPT algorithms $\text{VC} = (\text{Setup}, \text{Commit}, \text{BatchOpen}, \text{Verify})$ defined as follows:*

$\text{ck} \leftarrow \text{Setup}(1^\lambda, 1^n)$: *The setup algorithm takes security parameter λ and vector length n as input and returns commitment key ck .*

$(\text{com}, \text{aux}) \leftarrow \text{Commit}(\text{ck}, \mathbf{m})$: *The commitment algorithm takes commitment key ck and vector $\mathbf{m} \in \mathcal{M}^n$ as input and returns commitment com and auxiliary output aux .*

$(\pi_1, \dots, \pi_n) \leftarrow \text{BatchOpen}(\text{ck}, \text{com}, \text{aux})$: *The batch opening algorithm takes commitment key ck , a commitment com , and auxiliary input aux , as input and returns openings π_1, \dots, π_n .*

$b \leftarrow \text{Verify}(\text{ck}, \text{com}, \pi, i, m)$: The verification algorithm takes commitment key ck , commitment com , opening π , index i and message m as input and returns a bit b .

Definition 6 (Correctness). A vector commitment $\text{VC} = (\text{Setup}, \text{Commit}, \text{BatchOpen}, \text{Verify})$ is correct, if for all $\lambda, n \in \mathbb{N}$, all $\mathbf{m} \in \mathcal{M}^n$, all $(\text{com}, \text{aux}) \leftarrow \text{Commit}(\text{ck}, \mathbf{m})$, all $(\pi_1, \dots, \pi_n) \leftarrow \text{BatchOpen}(\text{ck}, \text{com}, \text{aux})$, and all $i \in [n]$ it holds that $\text{Verify}(\text{ck}, \text{com}, \pi_i, i, m) = 1$.

Definition 7 (Position-Binding). A vector commitment $\text{VC} = (\text{Setup}, \text{Commit}, \text{BatchOpen}, \text{Verify})$ is position binding if for all $\lambda, n \in \mathbb{N}$ and all PPT adversaries \mathcal{A} it holds that

$$\Pr \left[\begin{array}{l} m_0 \neq m_1 \\ \wedge \text{Verify}(\text{ck}, \text{com}, \pi_0, i, m_0) = 1 \\ \wedge \text{Verify}(\text{ck}, \text{com}, \pi_1, i, m_1) = 1 \end{array} : (\text{com}, i, m_0, m_1, \pi_0, \pi_1) \leftarrow \mathcal{A}(\text{ck}) \right] \leq \text{negl}(\lambda).$$

Definition 8 (Perfect Weak Hiding). A vector commitment $\text{VC} = (\text{Setup}, \text{Commit}, \text{BatchOpen}, \text{Verify})$ is perfectly weakly hiding if for all $\lambda, n \in \mathbb{N}$ and all $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}^n$ and all $\text{ck} \leftarrow \text{Setup}(1^\lambda, 1^n)$ it holds that over the random coins of the commitment algorithm, com_0 and com_1 computed as $(\text{com}_0, \text{aux}_0) \leftarrow \text{Commit}(\text{ck}, \mathbf{m}_0)$ and $(\text{com}_1, \text{aux}_1) \leftarrow \text{Commit}(\text{ck}, \mathbf{m}_1)$ are distributed identically.

Definition 9 (Efficiency). A vector commitment $\text{VC} = (\text{Setup}, \text{Commit}, \text{BatchOpen}, \text{Verify})$ is efficient, if commitments are of size independent of n , the commitment algorithm runs in time $n \cdot \text{poly}(\lambda)(\lambda)$ and the batch opening algorithm runs in time $n \cdot \text{poly}(\lambda)(\log n, \lambda)$.

$\text{Setup}(1^\lambda, 1^n)$ <hr/> $\text{pp} \leftarrow \text{KZG.Setup}(1^\lambda, 1^n)$ return pp	$\text{Commit}(\text{ck}, \mathbf{m})$ <hr/> $r \leftarrow \mathbb{F}_p$ $\mathbb{F}_p[X] \ni f(X) := \begin{cases} m_i & \text{if } X \in [n] \\ r & \text{if } X = 0 \end{cases}$ $\text{com} \leftarrow \text{KZG.Commit}(\text{pp}, f)$ return (com, f)
$\text{BatchOpen}(\text{ck}, \text{com}, f)$ <hr/> $\pi \leftarrow \text{KZG.BatchOpen}(\text{ck}, \text{com}, f, (i, f(i))_{i \in [n]})$ return π	$\text{Verify}(\text{ck}, \text{com}, \pi, i, m)$ <hr/> return $\text{KZG.Verify}(\text{ck}, \text{com}, \pi, i, m)$

Fig. 2. Weakly-hiding vector commitments from KZG polynomial commitments.

Given a polynomial commitment, it is easy to construct a vector commitment by encoding the message vector's entries into distinct polynomial evaluations. In the theorem statement below, we recall this transformation using KZG commitments but add a small twist that makes the construction weakly hiding.

Theorem 2. If KZG is a binding polynomial commitment scheme, then the construction specified in Figure 2 is a correct, efficient, computationally position binding, and perfectly weakly hiding commitment scheme.

Proof. Correctness is immediate from the correctness of KZG. Efficiency similarly follows from the definition of KZG and the batch opening algorithm of Feist and Khovratovich [FK23].

It is trivial to see, that any adversary breaking position binding of the vector commitment immediately also breaks binding of the KZG commitment scheme with the same probability. Position binding therefore follows immediately from the binding property of the KZG commitment. To see that the scheme is perfectly weakly hiding, consider any $\lambda, n \in \mathbb{N}$ and any $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}^n$ and any $\text{ck} \leftarrow \text{Setup}(1^\lambda, 1^n)$. Any commitment to either vector has the form $\text{com}_b = [\mu_b]_1 + [r]_1$, where μ_b is defined by m_b and independent of r . Since r is chosen uniformly from \mathbb{F}_p , com_b is always a uniformly distributed element of \mathbb{G}_1 . In particular, this means that com_0 and com_1 are distributed identically.

2.6 Symmetric Encryption

We quickly recall the definition of a symmetric encryption scheme.

Definition 10. *A symmetric encryption scheme with keyspace \mathcal{K} and message space \mathcal{M} is a pair of PPT algorithms $\text{SE} = (\text{Enc}^{\text{sym}}, \text{Dec}^{\text{sym}})$ defined as follows:*

$\text{ct} \leftarrow \text{Enc}^{\text{sym}}(\mathcal{K}, m)$: *The setup algorithm takes a key $\mathbf{k} \in \mathcal{K}$ and a message $m \in \mathcal{M}$ as input and returns ciphertext ct .*

$\text{ct} \leftarrow \text{Dec}^{\text{sym}}(\mathcal{K}, \text{ct})$: *The setup algorithm takes a key $\mathbf{k} \in \mathcal{K}$ and a ciphertext ct as input and returns message m .*

Definition 11 (Correctness). *A symmetric encryption scheme $\text{SE} = (\text{Enc}^{\text{sym}}, \text{Dec}^{\text{sym}})$ is correct, if for all $\mathbf{k} \in \mathcal{K}$, all $m \in \mathcal{M}$, and all $\text{ct} \leftarrow \text{Enc}^{\text{sym}}(\mathbf{k}, m)$ it holds that $\text{Dec}^{\text{sym}}(\mathbf{k}, \text{ct}) = m$.*

We will only require a very weak notion of security, namely indistinguishability against eavesdroppers, also called EAV-security. This security notion is e.g. satisfied by the one-time pad or a (nonce-less) stream cipher.

Definition 12 (EAV-Security). *A symmetric encryption scheme $\text{SE} = (\text{Enc}^{\text{sym}}, \text{Dec}^{\text{sym}})$ has indistinguishable encryptions in the presence of an eavesdropper, or is EAV-secure, if for any PPT adversary \mathcal{A} it holds that*

$$\Pr[\text{Expt}_{\text{SE}, \mathcal{A}}^{\text{EAV}}(1^\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where the experiment is defined as follows.

$$\begin{array}{l} \text{Expt}_{\text{SE}, \mathcal{A}}^{\text{EAV}}(1^\lambda) \\ \hline \mathbf{k} \leftarrow \mathcal{K} \\ (m_0, m_1) \leftarrow \mathcal{A}(1^\lambda) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{Enc}^{\text{sym}}(\mathbf{k}, m_b) \\ b' \leftarrow \mathcal{A}(\text{ct}) \\ \text{return } \begin{cases} 1 & \text{if } b' = b \\ 0 & \text{otherwise} \end{cases} \end{array}$$

3 Extractable Witness KEMs

As a building block for our main constructions, we first define and instantiate the notion of an extractable witness KEM. This notion, with minor differences, has previously been defined by Choi and Vaudenay [CV21]. In their work, the authors present an instantiation of this notion for some class of problems using a new non-falsifiable hardness assumption in a new restricted non-standard computational model. In our work, we will focus on instantiating the notion of an extractable witness KEM for specific relations using a standard assumption in the AGM.

We first define the notions of an indexed family of NP relations. We will later define extractable witness KEMs relative to such a family. This is more general than defining them for an individual relation and allows some part of the relation to be fixed by system parameters.

Definition 13. Let $\mathcal{I} \subseteq \{0, 1\}^*$ be a set. A set $\mathcal{F} = \{\mathcal{R}_I\}_{I \in \mathcal{I}}$ a family of NP relations with index set \mathcal{I} if for all $I \in \mathcal{I}$, \mathcal{R}_I is an NP relation. We call I the index of \mathcal{R}_I and \mathcal{R}_I the relation identified by I . We use \mathcal{L}_I to refer to the corresponding NP language.

Relative to these families we can now define a witness KEM. The general idea of a witness KEM is that it works like a regular key encapsulation mechanism, but uses the pairs of statement and witness in an NP relation as a key-pair.

Definition 14. A witness key encapsulation mechanism for a family of NP relations \mathcal{F} and a keyspace \mathcal{K} is a pair of PPT algorithms $\text{WKEM} = (\text{Encap}, \text{Decap})$, defined as follows:

$(\text{ct}, \mathbf{k}) \leftarrow \text{Encap}(I, \mathbf{x})$: The encapsulation algorithm takes as input an index I identifying a relation $\mathcal{R}_I \in \mathcal{F}$ and a statement \mathbf{x} and returns as output a ciphertext ct and a key $\mathbf{k} \in \mathcal{K}$.

$\mathbf{k} \leftarrow \text{Decap}(I, \mathbf{w}, \text{ct})$: The deterministic decapsulation algorithm takes as input an index I identifying a relation $\mathcal{R}_I \in \mathcal{F}$, a witness \mathbf{w} , and a ciphertext ct and returns a key $\mathbf{k} \in \mathcal{K}$.

Definition 15 (Correctness). A witness KEM $\text{WKEM} = (\text{Encap}, \text{Decap})$ for a family of NP relations \mathcal{F} is correct, if for any $\mathcal{R}_I \in \mathcal{F}$, any $\lambda \in \mathbb{N}$, any $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}_I$, and any $(\text{ct}, \mathbf{k}) \leftarrow \text{Encap}(I, \mathbf{x})$ it holds that $\text{Decap}(I, \mathbf{w}, \text{ct}) = \mathbf{k}$.

The above definitions by themselves do not guarantee any kind of security. Security will be derived from the extractability of the scheme. Extractability essentially says that if any efficient adversary can distinguish between a key encapsulated under some statement \mathbf{x} and a random key, then this adversary can also be used to extract a witness \mathbf{w} for \mathbf{x} .

Definition 16 (Extractability). A witness KEM $\text{WKEM} = (\text{Encap}, \text{Decap})$ for a family of NP-relations \mathcal{F} is extractable, if there exists a PPT algorithm Ext such that for any stateful PPT adversary \mathcal{A} and any relation $\mathcal{R}_I \in \mathcal{F}$ such that

$$\Pr[\text{Expt}_{\text{WKEM}, \mathcal{A}}^{\text{KEM-CPA}}(1^\lambda, I) = 1] \geq \frac{1}{2} + \epsilon(\lambda)$$

for some non-negligible function $\epsilon(\lambda)$, it holds that

$$\Pr \left[(\mathbf{x}, \mathbf{w}) \in \mathcal{R}_I : \begin{array}{l} \mathbf{x} \leftarrow \mathcal{A}(1^\lambda, I) \\ \mathbf{w} \leftarrow \text{Ext}^{\mathcal{A}(\cdot, \cdot)}(I, \mathbf{x}) \end{array} \right] \geq \delta(\lambda),$$

for some non-negligible function $\delta(\lambda)$. The latter probability is taken over the random coins of the adversary and the extractor and the experiment $\text{Expt}_{\text{WKEM}, \mathcal{A}}^{\text{KEM-CPA}}(1^\lambda)$ is defined as follows.

$$\begin{array}{l}
\text{Expt}_{\text{WKEM}, \mathcal{A}}^{\text{KEM-CPA}}(1^\lambda, I) \\
\hline
\mathbf{x} \leftarrow \mathcal{A}(1^\lambda, I) \\
b \leftarrow \{0, 1\} \\
(\mathbf{ct}, \mathbf{k}_0) \leftarrow \text{Encap}(I, \mathbf{x}) \\
\mathbf{k}_1 \leftarrow \mathcal{K} \\
b' \leftarrow \mathcal{A}(\mathbf{ct}, \mathbf{k}_b) \\
\mathbf{return} \begin{cases} 1 & \text{if } b' = b \\ 0 & \text{otherwise} \end{cases}
\end{array}$$

3.1 An Extractable Witness KEM for KZG Openings

We now proceed with constructing an extractable witness KEM as just defined for a very specific family of relations, specifically those describing valid opening of KZG commitments. Let $(\text{Setup}, \text{Commit}, \text{Open}, \text{Verify})$ be the KZG commitment relative to the bilinear group $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ of prime order p as specified in Figure 1. Let $\text{CK} = \{\text{ck} \in \mathbb{G}_1^{d+1} \times \mathbb{G}_2^2 \mid d \in \mathbb{N} \wedge \text{ck} \in \text{Setup}(1^\lambda, 1^d)\}$ be the set of valid KZG commitment keys. We can then define the family of NP relations of valid KZG openings as

$$\mathcal{F}_{KZG} := \{\mathcal{R}_{\text{ck}}\}_{\text{ck} \in \text{CK}}$$

where

$$\mathcal{R}_{\text{ck}} = \{((\text{com}_j, \alpha_j, \beta_j)_{j \in [\ell]}, (\pi_j)_{j \in [\ell]}) \mid \ell \in \mathbb{N} \wedge \forall j \in [\ell]. \text{Verify}(\text{ck}, \text{com}_j, \pi_j, \alpha_j, \beta_j) = 1\}$$

for any $\text{ck} \in \text{CK}$.

$\text{Encap}^H(\text{ck}, (\text{com}, \alpha, \beta))$	$\text{Decap}^H(\text{ck}, (\pi_1, \dots, \pi_\ell), \text{ct})$
for $1 \leq j \leq \ell$ $r_j \leftarrow \mathbb{F}_p$ $s_j := e(r_j \cdot (\text{com} - [\beta_j]_1), [1]_2)$ $\text{ct}_j \leftarrow r_j \cdot ([\tau]_2 - [\alpha_j]_2)$ $\text{ct} := (\text{ct}_1, \dots, \text{ct}_\ell)$ $\mathbf{k} := H(s_1, \dots, s_\ell)$ return (ct, \mathbf{k})	parse ct as $(\text{ct}_1, \dots, \text{ct}_\ell)$ for $1 \leq j \leq \ell$ $s_j := e(\pi_j, \text{ct}_j)$ $\mathbf{k} := H(s_1, \dots, s_\ell)$ return \mathbf{k}

Fig. 3. Construction of an Extractable Witness KEM for \mathcal{F}_{KZG} with keyspace $\{0, 1\}^\lambda$ in the combined Algebraic Group and Random Oracle Model.

Theorem 3. *Let $H : \mathbb{G}_T^* \rightarrow \{0, 1\}^\lambda$ be a hash function modeled as a random oracle. If the d -DLOG assumption holds with respect to GGen , then the construction described in Figure 3 is an extractable witness KEM for \mathcal{F}_{KZG} in the algebraic group model.*

Proof. Let \mathcal{A} be an arbitrary algebraic PPT adversary with non-negligible advantage $\epsilon(\lambda)$ for the extractability of the construction described in Figure 3.

We construct an extractor Ext as follows. The extractor receives as input an index $\text{ck} = ([\tau^0]_1, \dots, [\tau^d]_1, [1]_2, [\tau]_2) \in \text{CK}$ as well as a statement $((\text{com}_1, \alpha_1, \beta_1), \dots, (\text{com}_\ell, \alpha_\ell, \beta_\ell)) \in (\mathbb{G}_1 \times \mathbb{F}_p^2)^\ell$. Further, since \mathcal{A} is an algebraic algorithm, the extractor also receives an algebraic representations of each com_j in the form of coefficients $f_{j,0}, \dots, f_{j,d}$ such that

$$\text{com}_j := \sum_{i=0}^d f_{j,i} \cdot [\tau^i]_1.$$

The extractor chooses $r_1, \dots, r_\ell \leftarrow \mathbb{F}_p$, computes $\text{ct}_j := r_j \cdot ([\tau]_2 - [\alpha_j]_2)$ for each $j \in [\ell]$, sets $\text{ct} := (\text{ct}_1, \dots, \text{ct}_\ell)$, chooses a random key $\mathbf{k} \leftarrow \{0, 1\}^\lambda$, initializes an empty list $\Gamma := \emptyset$ and invokes $\mathcal{A}(\text{ct}, \mathbf{k})$. The adversary \mathcal{A} expects access to a random oracle that Ext simulates as follows. For any query $\mathbf{s} \in \mathbb{G}_T^*$ such that either $|\mathbf{s}| \neq \ell$ or for at least one $j \in [\ell]$, $s_j \neq e(r_j \cdot (\text{com} - [\beta_j]_1), [1]_2)$, Ext continues to simulate the random oracle via lazy sampling. That is, if there exists an entry $(\mathbf{s}, \mathbf{k}_\mathbf{s}) \in \Gamma$ it returns $\mathbf{k}_\mathbf{s}$. If such an entry does not yet exist, Ext samples $\mathbf{k}_\mathbf{s} \leftarrow \{0, 1\}^n$, adds $(\mathbf{s}, \mathbf{k}_\mathbf{s})$ to Γ , and returns $\mathbf{k}_\mathbf{s}$.

If, however, $|\mathbf{s}| = \ell$ and for all $j \in [\ell]$, $s_j = e(r_j \cdot (\text{com} - [\beta_j]_1), [1]_2)$, then the extractor aborts \mathcal{A} and continues as follows. Since \mathcal{A} is an algebraic algorithm it also provides an algebraic description of each s_j in the form of coefficients⁵

$$\tilde{w}_{j,0}, \dots, \tilde{w}_{j,2d}, \tilde{q}_{j,1,0}, \dots, \tilde{q}_{j,1,d}, \tilde{q}_{j,2,0}, \dots, \tilde{q}_{j,\ell,d}, \tilde{h}_{j,1}, \dots, \tilde{h}_{j,\ell}$$

such that

$$\begin{aligned} s_j := & \sum_{i=0}^{2d} \tilde{w}_{j,i} \cdot [\tau^i]_T + \sum_{k=1}^{\ell} \sum_{i=0}^d \tilde{q}_{j,k,i} \cdot [r_k \cdot (\tau - \alpha_k) \cdot \tau^i]_T \\ & + \sum_{k=1}^{\ell} \tilde{h}_{j,k} \cdot [r_j \cdot (\tau - \alpha_j) \cdot r_k \cdot (\tau - \alpha_k)]_T. \end{aligned}$$

Since Ext computed each ct_k , for $k \neq j$ algebraically, however, we can simplify this representation. Specifically Ext can compute an algebraic description of each s_j in the form of coefficients $w_{j,0}, \dots, w_{j,2d}, q_{j,0}, \dots, q_{j,d}, h_j$ defined as

$$\begin{aligned} w_{j,i} := & \begin{cases} \tilde{w}_{j,i} + \sum_{k \in [\ell] \setminus \{j\}} \tilde{q}_{j,k,i} r_k \alpha_k & \text{if } i = 0 \\ \tilde{w}_{j,i} + \sum_{k \in [\ell] \setminus \{j\}} \tilde{q}_{j,k,i} r_k \alpha_k + \tilde{q}_{j,k,i-1} r_k & \text{if } 0 < i \leq d \\ \tilde{w}_{j,i} + \sum_{k \in [\ell] \setminus \{j\}} \tilde{q}_{j,k,i-1} r_k & \text{if } i = d + 1 \\ \tilde{w}_{j,i} & \text{if } i > d + 1 \end{cases} \\ q_{j,i} := & \begin{cases} \tilde{q}_{j,j,i} + \sum_{k \in [\ell] \setminus \{j\}} \tilde{h}_{j,k} r_k & \text{if } i = 0 \\ \tilde{q}_{j,j,i} + \sum_{k \in [\ell] \setminus \{j\}} \tilde{h}_{j,k} r_k \alpha_k & \text{if } i = 1 \\ \tilde{q}_{j,j,i} & \text{if } i > 1 \end{cases} \\ h_j := & \tilde{h}_{j,j} \end{aligned}$$

⁵ Note that we allow the algebraic adversary the maximum amount of freedom here that they would only have in the context of a Type-1 pairing. In the context of Type-2 or Type-3 pairings, many of these coefficients would necessarily be 0.

such that

$$\begin{aligned} s_j &:= \sum_{i=0}^{2d} w_{j,i} \cdot [\tau^i]_T + \sum_{i=0}^d q_{j,i} \cdot [r_j \cdot (\tau - \alpha_j) \cdot \tau^i]_T + h_j \cdot [r^2 \cdot (\tau - \alpha_j)^2]_T \\ &= \sum_{i=0}^{2d} w_{j,i} \cdot [\tau^i]_T + r_j \cdot (\tau - \alpha_j) \cdot \sum_{i=0}^d q_{j,i} \cdot [\tau^i]_T + h_j r_j^2 \cdot [(\tau - \alpha_j)^2]_T \end{aligned}$$

For each $j \in [\ell]$ we define the two polynomials

$$F_j(X) := \sum_{i=1}^d f_{j,i} \cdot X^i$$

and

$$G_j(X, Y) := \sum_{i=0}^{2d} w_{j,i} \cdot X^i + Y \cdot (X - \alpha_j) \cdot \sum_{i=0}^d q_{j,i} \cdot X^i + h_j Y^2 \cdot (X - \alpha_j)^2.$$

It is easy to see that $\text{com}_j = [F_j(\tau)]_T$ and $s_j = [G_j(\tau, r_j)]_T$.

If for all $j \in [\ell]$, $G_j(X, Y) = Y(F_j(X) - \beta_j)$ the extractor computes $\pi_j := \sum_{i=1}^d q_{j,i} \cdot [\tau^i]_1$ and outputs $\boldsymbol{\pi} := (\pi_1, \dots, \pi_\ell)$. Otherwise the extractor outputs \perp . If the adversary terminates without querying an \mathbf{s} as described above, then the extractor also outputs \perp .

We now prove two claims about the extractor.

Claim 4. *If Ext outputs $\boldsymbol{\pi} \neq \perp$, then for all $j \in [\ell]$ it holds that $\text{Verify}(\text{ck}, \text{com}_j, \pi_j, \alpha_j, \beta_j) = 1$.*

Proof. The extractor only outputs $\boldsymbol{\pi} \neq \perp$, if for all $j \in [\ell]$, $G_j(X, Y) = Y(F_j(X) - \beta_j)$. We observe that, this can only be true if $h_j = 0$ and $w_{j,i} = 0$ for all $j \in [\ell]$ and $0 \leq i \leq 2d$. Therefore,

$$G_j(X, Y) := Y \cdot (X - \alpha_j) \cdot \sum_{i=0}^d q_{j,i} \cdot X^i.$$

and thus

$$\begin{aligned} G_j(X, Y) &= Y(F_j(X) - \beta_j) \\ \iff Y \cdot (X - \alpha_j) \cdot \sum_{i=0}^d q_{j,i} \cdot X^i &= Y(F_j(X) - \beta_j) \\ \iff \sum_{i=0}^d q_{j,i} \cdot X^i &= \frac{F_j(X) - \beta_j}{X - \alpha_j}. \end{aligned}$$

It follows that the extractor's output

$$\pi_j = \left[\frac{F_j(\tau) - \beta_j}{\tau - \alpha_j} \right]_1$$

and thus the verification equation

$$e(\pi_j, [\tau]_2 - [\alpha_j]_2) = e\left(\left[\frac{F_j(\tau) - \beta_j}{\tau - \alpha_j} \right]_1, [\tau - \alpha_j]_2\right)$$

$$\begin{aligned}
&= [F_j(\tau) - \beta_j]_T \\
&= e([F_j(\tau) - \beta_j]_1, [1]_2) = e(\mathbf{com}_j - [\beta_j]_1, [1]_2)
\end{aligned}$$

holds as required. \square

Claim 5. Ext outputs $\pi = \perp$ with probability at most $1 - 2\epsilon(\lambda) + \text{negl}(\lambda)$.

Proof. Let **Hit** denote the event that \mathbf{s} with $|\mathbf{s}| = \ell$ and $s_j = e(r_j \cdot (\mathbf{com}_j - [\beta_j]_1), [1]_2)$ for all $j \in [\ell]$ is queried to the random oracle. The extractor outputs \perp , either if **Hit** does not occur, or if **Hit** occurs but for at least one $j \in [\ell]$ $G_j(X, Y) \neq Y(F_j(X) - \beta_j)$. Since those two events are mutually exclusive, we thus have

$$\Pr[\pi = \perp] = \Pr[\overline{\mathbf{Hit}}] + \Pr[\mathbf{Hit} \wedge \exists j \in [\ell]. G_j(X, Y) \neq Y(F_j(X) - \beta_j)]. \quad (1)$$

We bound the two probabilities separately.

Let us first bound $\Pr[\overline{\mathbf{Hit}}]$. Recall, that \mathcal{A} is an adversary with advantage $\epsilon(\lambda)$. Whenever **Hit** does not occur, the view of \mathcal{A} remains independent of the bit b in the KEM-CPA experiment. This is the case, since unless **Hit** occurs, both \mathbf{k}_0 and \mathbf{k}_1 are uniformly distributed. Therefore we have that

$$\begin{aligned}
\frac{1}{2} + \epsilon(\lambda) &= \Pr[\text{Expt}_{\text{WKEM}_{KZG}, \mathcal{A}}^{\text{KEM-CPA}}(1^\lambda) = 1] \\
&= \Pr[b' = b] \\
&= \underbrace{\Pr[b' = b \mid \mathbf{Hit}]}_{\leq 1} \Pr[\mathbf{Hit}] + \underbrace{\Pr[b' = b \mid \overline{\mathbf{Hit}}]}_{=1/2} \Pr[\overline{\mathbf{Hit}}] \\
&\leq 1 - \Pr[\overline{\mathbf{Hit}}] + \frac{1}{2} \Pr[\overline{\mathbf{Hit}}] \\
&= 1 - \frac{1}{2} \Pr[\overline{\mathbf{Hit}}]
\end{aligned}$$

and thus, we have that

$$\Pr[\overline{\mathbf{Hit}}] \leq 1 - 2\epsilon(\lambda). \quad (2)$$

Next we bound $\Pr[\mathbf{Hit} \wedge \exists j \in [\ell]. G_j(X, Y) \neq Y(F_j(X) - \beta_j)]$. For each $j \in [\ell]$, define the bivariate polynomial $Q_j(X, Y) = G_j(X, Y) - Y(F_j(X) - \beta_j)$. Since **Hit** happened, it must hold that $s_j = e(r_j \cdot (\mathbf{com}_j - [\beta_j]_1), [1]_2) = [r_j \cdot (F_j(\tau) - \beta_j)]_T$ and since $s_j = [G_j(\tau, r_j)]_T$, it must hold that $Q_j(\tau, r_j) = 0$. However, since $G_j(X, Y) \neq Y(F_j(X) - \beta_j)$, it also holds that $Q_j(X, Y) \neq 0$, meaning $Q_j(X, Y)$ is a non-zero polynomial with a root at (τ, r_j) .

We can rewrite $Q_j(X, Y)$ as a polynomial of the form $Q_j(X, Y) = C_{j,0}(X) + C_{j,1}(X) \cdot Y + C_{j,2}(X) \cdot Y^2$, where each $C_{j,i}(X) \in \mathbb{F}_p[X]$ is a univariate polynomial of degree at most $2d$. Since $Q_j(X, Y)$ is non-zero, at least one of $C_{j,0}(X), C_{j,1}(X), C_{j,2}(X)$ must also be non-zero.

Now consider the univariate Polynomial $P_j(Y) := Q_j(\tau, Y)$. We can consider two cases, either $P_j(Y) = 0$, or $P_j(Y) \neq 0$. In the case that $P_j(Y) = 0$, it would need to hold that $C_{j,0}(\tau) = C_{j,1}(\tau) = C_{j,2}(\tau) = 0$, but at least one of these is a non-zero polynomial. We can therefore find the roots of one of the non-zero $C_{j,0}(Y), C_{j,1}(Y), C_{j,2}(Y)$ to recover τ . On the other hand, if $P_j(Y) \neq 0$, we can find the roots of $P_j(Y)$ to recover r_j .

We use this to specify the following reduction to d -DLOG. On input $[\eta^0]_1, \dots, [\eta^d]_1, [\eta^0]_2, [\eta^1]_2$, the reduction \mathcal{B} flips a bit $b \leftarrow \{0, 1\}$. If $b = 0$, then \mathcal{B} defines $\text{ck} = ([\eta^0]_1, \dots, [\eta^d]_1, [\eta^0]_2, [\eta^1]_2)$.

If $b = 1$, then \mathcal{B} chooses $\tau \leftarrow \mathbb{F}_q$ and defines $\text{ck} = ([\tau^0]_1, \dots, [\tau^d]_1, [1]_2, [\tau]_2)$. In both cases, \mathcal{B} then invokes $\mathcal{A}(1^\lambda, \text{ck})$. Eventually \mathcal{A} will output $((\text{com}_1, \alpha_1, \beta_1), \dots, (\text{com}_\ell, \alpha_\ell, \beta_\ell))$ together with the algebraic explanation of each com_j . If $b = 0$, then for $j \in [\ell]$, the reduction chooses $r_j \leftarrow \mathbb{F}_p$ and computes $\text{ct}_j := r_j \cdot ([\eta]_2 - [\alpha_j]_2)$. If $b = 1$, then for $j \in [\ell]$, the reduction chooses $z_j \leftarrow \mathbb{F}_p$ and computes $\text{ct}_j := (\tau - \alpha_j) \cdot ([\eta]_2 + z_j \cdot [1]_2)$. Note, that, if we define $r_j := z_j + \eta$ it holds that $(\tau - \alpha_j) \cdot ([\eta]_2 + z_j \cdot [1]_2) = (z_j + \eta) \cdot ([\tau]_2 - [\alpha_j]_2) = r_j \cdot ([\tau]_2 - [\alpha_j]_2)$. Therefore, in the case of $b = 1$ the ciphertexts are computed correctly for implicitly defined *but uniformly distributed* r_j . From this point, \mathcal{B} proceeds exactly as `Ext` until `Hit` occurs. It is important to verify that \mathcal{B} can actually check whether this query occurs, even in the case $b = 1$. However, due to the bilinearity of the pairing, \mathcal{B} can compute the relevant values as

$$\begin{aligned} e(\text{com}_j - [\beta_j]_1, z_j \cdot [1]_2 + [\eta]_2) &= e(\text{com}_j - [\beta_j]_1, (z_j + \eta) \cdot [1]_2) \\ &= e(\text{com}_j - [\beta_j]_1, r_j \cdot [1]_2) \\ &= e(r_j \cdot (\text{com}_j - [\beta_j]_1), [1]_2) \end{aligned}$$

and thus can check each query against these values. If `Hit` does not occur or if for all $j \in [\ell]$, it holds that $G_j(X, Y) = Y(F_j(X) - \beta_j)$, then \mathcal{B} aborts. Otherwise, we consider two cases.

If $b = 0$ and there exists an index j^* such that $G_{j^*}(X, Y) \neq Y(F_{j^*}(X) - \beta_{j^*})$ and $P_{j^*}(Y) = 0$, then, as discussed above, there exists a $C_{j^*}(X) \in \{C_{j^*,0}(X), C_{j^*,1}(X), C_{j^*,2}(X)\}$ such that $C_{j^*}(X)$ is non-zero. \mathcal{B} factors $C_{j^*}(X)$ to find all of its at most $2d$ roots. For each root ξ , \mathcal{B} checks whether $\xi \cdot [1]_1 = [\eta]_1$ and returns ξ if that's the case. Since η must be one of the roots, \mathcal{B} will always correctly identify η in this case.

If $b = 1$ and there exists an index j^* such that $G_{j^*}(X, Y) = Y(F_{j^*}(X) - \beta_{j^*})$ and $P_{j^*}(Y) \neq 0$, then \mathcal{B} factors $P_{j^*}(Y)$ to find all of its at most 2 roots. For each root ξ , \mathcal{B} checks whether $\xi \cdot [1]_1 = z_{j^*} \cdot [1]_1 + [\eta]_1$ and returns $\xi - z_{j^*}$ if that is the case. Since $r_{j^*} = z_{j^*} + \eta$ must be one of the roots, \mathcal{B} will always correctly identify $\eta = r_{j^*} - z_{j^*}$ in this case.

If neither of those two cases happens, then \mathcal{B} also aborts.

Since the view of \mathcal{A} is independent of the value of b and for each j^* such that $G_{j^*}(X, Y) \neq Y(F_{j^*}(X) - \beta_{j^*})$, it must either be the case that $P_{j^*}(Y) = 0$ or that $P_{j^*}(Y) \neq 0$, it follows under the d -DLOG assumption, that

$$\begin{aligned} \text{negl}(\lambda) &\geq \Pr \left[\begin{array}{l} \text{par} \leftarrow \text{GGen}(1^\lambda) \\ \tau = \tau' : \\ \tau \leftarrow \mathbb{F}_p \\ \tau' \leftarrow \mathcal{B}(\text{par}, ([\tau^0]_1, \dots, [\tau^d]_1, [1]_2, [\tau]_2)) \end{array} \right] \\ &= \Pr[b = 0] \cdot \Pr[\text{Hit} \wedge (\exists j^* \in [\ell]. G_{j^*}(X, Y) \neq Y(F_{j^*}(X) - \beta_{j^*}) \wedge P_{j^*}(Y) = 0)] \\ &\quad + \Pr[b = 1] \cdot \Pr[\text{Hit} \wedge (\exists j^* \in [\ell]. G_{j^*}(X, Y) \neq Y(F_{j^*}(X) - \beta_{j^*}) \wedge P_{j^*}(Y) \neq 0)] \\ &= \frac{1}{2} \cdot \left(\Pr[\text{Hit} \wedge (\exists j^* \in [\ell]. G_{j^*}(X, Y) \neq Y(F_{j^*}(X) - \beta_{j^*}) \wedge P_{j^*}(Y) = 0)] \right. \\ &\quad \left. + \Pr[\text{Hit} \wedge (\exists j^* \in [\ell]. G_{j^*}(X, Y) \neq Y(F_{j^*}(X) - \beta_{j^*}) \wedge P_{j^*}(Y) \neq 0)] \right) \\ &\geq \frac{1}{2} \Pr[\text{Hit} \wedge \exists j^* \in [\ell]. G_{j^*}(X, Y) \neq Y(F_{j^*}(X) - \beta_{j^*})]. \end{aligned}$$

The claim thus follows. \square

By combining the two claims we finally get

$$\Pr \left[\forall j \in [\ell]. \text{Verify}(\text{ck}, \text{com}_j, \pi_j, \alpha_j, \beta_j) = 1 : \begin{array}{l} (\text{com}_j, \alpha_j, \beta_j)_{j \in [\ell]} \leftarrow \mathcal{A}(1^\lambda, \text{ck}) \\ \pi \leftarrow \text{Ext}^{\mathcal{A}(\cdot, \cdot)}(\text{ck}, (\text{com}_j, \alpha_j, \beta_j)_{j \in [\ell]}) \end{array} \right]$$

$$= \Pr[\pi \neq \perp] = 1 - \Pr[\pi = \perp] = 2\epsilon(\lambda) - \text{negl}(\lambda).$$

which is non-negligible for any non-negligible function $\epsilon(\lambda)$, as required. \square

We will later require an extractable witness KEM not just for plain KZG, but in fact for the derived weakly hiding vector commitment described in Figure 3. Let VC be the vector commitment and let $\text{CK}_{\text{VC}} = \{\text{ck} \mid \text{ck} \in \text{VC.Setup}(1^\lambda, 1^n)\}$ be the set of valid commitment keys. We can then define the family of NP relations of valid VC openings as

$$\mathcal{F}_{\text{VC}} := \{\mathcal{R}_{\text{ck}}\}_{\text{ck} \in \text{CK}_{\text{VC}}}$$

where

$$\mathcal{R}_{\text{ck}} = \left\{ \left((\text{com}_j, i_j, m_j)_{j \in [\ell]}, (\pi_j)_{j \in [\ell]} \right) \mid \ell \in \mathbb{N} \wedge \forall j \in [\ell]. \text{VC.Verify}(\text{ck}, \text{com}_j, \pi_j, i_j, m_j) = 1 \right\}$$

for any $\text{ck} \in \text{CK}_{\text{VC}}$.

Since VC commitments are simply KZG commitments, openings are KZG openings, and verification is KZG verification, it follows immediately that $\mathcal{F}_{\text{VC}} = \mathcal{F}_{\text{KZG}}$ and we thus get the following corollary from Theorem 3.

Corollary 6. *Let $\text{H} : \mathbb{G}_T^* \rightarrow \{0, 1\}^\lambda$ be a hash function modeled as a random oracle. If the d -DLOG assumption holds with respect to GGen , then the construction described in Figure 3 is an extractable witness KEM for \mathcal{F}_{VC} in the algebraic group model.*

4 Extractable Witness Encryption

In this section, we recall the definition of extractable witness encryption following broadly the definitions of [GKP⁺13], with the extension to *families* of relations. We then go on to show that one can generically construct it from any extractable witness KEM using the standard KEM/DEM paradigm.

As is the case for regular encryption and KEMs, witness encryption and witness KEMs are very similar, with the only difference being that a witness encryption scheme is capable of encrypting a freely chosen message instead of a random one. The definitions are otherwise very similar to the definitions from the previous section.

Definition 17. *A witness encryption scheme for a family of NP relations \mathcal{F} and a messagespace \mathcal{M} is a pair of PPT algorithms $\text{WE} = (\text{Enc}, \text{Dec})$, defined as follows:*

$\text{ct} \leftarrow \text{Enc}(I, \mathbf{x}, m)$: *The encryption algorithm takes as input an index I identifying a relation $\mathcal{R}_I \in \mathcal{F}$, a statement \mathbf{x} , and a message $m \in \mathcal{M}$ and returns as output a ciphertext ct .*

$m/\perp \leftarrow \text{Dec}(I, \mathbf{w}, \text{ct})$: *The deterministic decryption algorithm takes as input an index I identifying a relation $\mathcal{R}_I \in \mathcal{F}$, a ciphertext ct , and a witness \mathbf{w} and returns a message $m \in \mathcal{M}$ or a error symbol \perp .*

Definition 18 (Correctness). *A witness encryption scheme $\text{WE} = (\text{Enc}, \text{Dec})$ for a family of NP relations \mathcal{F} is correct, if for any $\mathcal{R}_I \in \mathcal{F}$, any $\lambda \in \mathbb{N}$, any $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}_I$, any $m \in \mathcal{M}$, and any $\text{ct} \leftarrow \text{Enc}(\mathbf{x})$ it holds that $\text{Dec}(I, \mathbf{w}, \text{ct}) = m$.*

Definition 19 (Extractability). A witness encryption scheme $\text{WE} = (\text{Enc}, \text{Dec})$ for a family of NP-relations \mathcal{F} is extractable, if there exists a PPT algorithm Ext such that for any stateful PPT adversary \mathcal{A} and any relation $\mathcal{R}_I \in \mathcal{F}$ such that

$$\Pr[\text{Expt}_{\text{WE}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I) = 1] \geq \frac{1}{2} + \epsilon(\lambda),$$

for some non-negligible function $\epsilon(\lambda)$ it holds that

$$\Pr \left[(\mathbf{x}, \mathbf{w}) \in \mathcal{R}_{\mathcal{L}} : \begin{array}{l} (\mathbf{x}, m_0, m_1) \leftarrow \mathcal{A}(1^\lambda, I) \\ \mathbf{w} \leftarrow \text{Ext}^{\mathcal{A}(\cdot)}(I, \mathbf{x}, m_0, m_1) \end{array} \right] \geq \delta(\lambda),$$

for some non-negligible function $\delta(\lambda)$. The latter probability is taken over the random coins of the adversary and the extractor and the experiment $\text{Expt}_{\mathcal{A}}^{\text{CPA}}$ is defined as follows.

$$\begin{array}{l} \text{Expt}_{\text{WE}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I) \\ \hline (\mathbf{x}, m_0, m_1) \leftarrow \mathcal{A}(1^\lambda, I) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{Enc}(I, \mathbf{x}, m_b) \\ b' \leftarrow \mathcal{A}(\text{ct}) \\ \text{return} \begin{cases} 1 & \text{if } b' = b \\ 0 & \text{otherwise} \end{cases} \end{array}$$

4.1 Extractable Witness Encryption from Extractable Witness KEMs

We can construct extractable witness encryption for a family of NP-relations \mathcal{F} and a message space \mathcal{M} from any extractable witness KEM for \mathcal{F} and any EAV secure symmetric encryption scheme with message space \mathcal{M} , as long as the two schemes share a compatible key space \mathcal{K} . The construction, shown in Figure 4 essentially follows the standard KEM/DEM paradigm instantiated with an extractable witness KEM. Even though the security of the KEM/DEM paradigm has been

$\text{Enc}(I, \mathbf{x}, m)$	$\text{Dec}(I, \mathbf{w}, (\text{ct}_1, \text{ct}_2))$
$(\text{ct}_1, \mathbf{k}) \leftarrow \text{Encap}(I, \mathbf{x})$	$\mathbf{k} := \text{Decap}(I, \mathbf{w}, \text{ct}_1)$
$\text{ct}_2 \leftarrow \text{Enc}^{\text{sym}}(\mathbf{k}, m)$	$m := \text{Dec}^{\text{sym}}(\mathbf{k}, \text{ct}_2)$
return $(\text{ct}_1, \text{ct}_2)$	return m

Fig. 4. Construction of an extractable witness encryption Scheme for \mathcal{F} based on an extractable witness KEM and an EAV secure symmetric encryption scheme.

proven ad nauseam, we will not skip the proof here. Since we are not proving indistinguishability, but *extractability* we need to be a bit more careful. Note that encapsulated keys are in general *not* indistinguishable from random keys. In fact, since the *adversary* chooses the statement, they may very well *know the witness* and thus be capable of distinguishing with overwhelming probability. It is merely the case in this case *we are capable of extracting the witness*. We need to be careful in our proof that this capability is preserved in the KEM/DEM paradigm. This may not always be the case, depending on how one executes the standard hybrid argument.

Theorem 7. Let $\text{WKEM} = (\text{Encap}, \text{Decap})$ be an extractable witness KEM for \mathcal{F} and key space \mathcal{K} . Let $\text{SE} = (\text{Enc}^{\text{sym}}, \text{Dec}^{\text{sym}})$ be an EAV secure symmetric encryption scheme with key space \mathcal{K} , message space \mathcal{M} . Then $\text{WE} = (\text{Enc}, \text{Dec})$ as specified in Figure 4 is an extractable witness encryption scheme for \mathcal{F} and message space \mathcal{M} .

Proof. We first define a *modified* witness encryption scheme $\widetilde{\text{WE}} = (\widetilde{\text{Enc}}, \cdot)$ that, instead of using the encapsulated key \mathbf{k} to perform the symmetric encryption, chooses a fresh key $\mathbf{k}' \leftarrow \mathcal{K}$ to do so. This scheme has no well-defined decryption algorithm. Nevertheless, for any PPT adversary \mathcal{A} , the probability $\Pr[\text{Expt}_{\widetilde{\text{WE}}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I) = 1]$ is still well defined. Let \mathcal{A} be an arbitrary PPT adversary such that

$$\Pr[\text{Expt}_{\widetilde{\text{WE}}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I) = 1] = \frac{1}{2} + \epsilon(\lambda). \quad (3)$$

for some non-negligible function $\epsilon(\lambda)$. We first prove the following claim.

Claim 8. For any $I \in \mathcal{I}$, it holds that

$$\Pr[\text{Expt}_{\widetilde{\text{WE}}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I) = 1] \leq \frac{1}{2} + \text{negl}(\lambda).$$

Proof. We construct a PPT adversary \mathcal{B} against the EAV security of SE as follows. Upon input 1^λ , \mathcal{B} invokes $\mathcal{A}(1^\lambda, I)$, receiving \mathbf{x}, m_0, m_1 in response, and outputs m_0, m_1 . After receiving as input the challenge ciphertext ct_2 , it computes $(\text{ct}_1, \mathbf{k}) \leftarrow \text{Encap}(I, \mathbf{x})$ and invokes $\mathcal{A}((\text{ct}_1, \text{ct}_2))$. In response, \mathcal{A} will output a bit b' , which \mathcal{B} also outputs.

It is easy to see that \mathcal{B} perfectly simulates the $\text{Expt}_{\widetilde{\text{WE}}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I)$ experiment for \mathcal{A} . Further, whenever \mathcal{A} would be successful, so is \mathcal{B} . It thus follows from the EAV security of SE that

$$\frac{1}{2} + \text{negl}(\lambda) \geq \Pr[\text{Expt}_{\text{SE}, \mathcal{B}}^{\text{EAV}}(1^\lambda) = 1] = \Pr[\text{Expt}_{\widetilde{\text{WE}}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I) = 1]$$

as claimed. \square

Claim 9. There exists a PPT adversary \mathcal{B} such that for any $I \in \mathcal{I}$, it holds that

$$\Pr[\text{Expt}_{\text{WKEM}, \mathcal{B}}^{\text{KEM-CPA}}(1^\lambda, I) = 1] = \frac{1}{2} + \frac{1}{2}\epsilon(\lambda) - \text{negl}(\lambda).$$

Proof. Upon input 1^λ and I , \mathcal{B} invokes $\mathcal{A}(1^\lambda, I)$, receiving \mathbf{x}, m_0, m_1 in response and outputs \mathbf{x} . After receiving ct_1, \mathbf{k} , \mathcal{B} samples $b' \leftarrow \{0, 1\}$, computes $\text{ct}_2 \leftarrow \text{Enc}^{\text{sym}}(\mathbf{k}, m_{b'})$, and invokes $\mathcal{A}((\text{ct}_1, \text{ct}_2))$. Eventually \mathcal{A} will output a bit b'' and \mathcal{B} will output 0, if $b'' = b'$ and 1 otherwise.

Let b denote the random bit of the experiment $\text{Expt}_{\text{WKEM}, \mathcal{B}}^{\text{KEM-CPA}}(1^\lambda, I)$. It is then easy to see that, if $b = 0$, then \mathcal{B} perfectly simulates the experiment $\text{Expt}_{\widetilde{\text{WE}}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I)$ and outputs 0 iff the experiment outputs 1. Similarly, if $b = 1$, then \mathcal{B} perfectly simulates the experiment $\text{Expt}_{\widetilde{\text{WE}}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I)$ and outputs 0 iff the experiment outputs 1. It then follows from Claim 8 and Equation 3 that

$$\begin{aligned} & \Pr[\text{Expt}_{\text{WKEM}, \mathcal{B}}^{\text{KEM-CPA}}(1^\lambda, I) = 1] \\ &= \Pr[b = 0] \cdot \Pr[\text{Expt}_{\text{WKEM}, \mathcal{B}}^{\text{KEM-CPA}}(1^\lambda, I) = 1 \mid b = 0] + \Pr[b = 1] \cdot \Pr[\text{Expt}_{\text{WKEM}, \mathcal{B}}^{\text{KEM-CPA}}(1^\lambda, I) = 1 \mid b = 1] \\ &= \frac{1}{2} \cdot (\Pr[\text{Expt}_{\widetilde{\text{WE}}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I) = 1] + 1 - \Pr[\text{Expt}_{\widetilde{\text{WE}}, \mathcal{A}}^{\text{CPA}}(1^\lambda, I) = 1]) \\ &\geq \frac{1}{2} \cdot \left(\frac{1}{2} + \epsilon(\lambda) + 1 - \frac{1}{2} - \text{negl}(\lambda) \right) \\ &= \frac{1}{2} + \frac{1}{2}\epsilon(\lambda) - \text{negl}(\lambda) \end{aligned} \quad \square$$

Finally, since WKEM is known to be extractable and $\frac{1}{2}\epsilon(\lambda) - \text{negl}(\lambda)$ is a non-negligible function whenever $\epsilon(\lambda)$ is non-negligible, there exists a PPT extractor $\widetilde{\text{Ext}}$, such that

$$\Pr \left[(\mathbf{x}, \mathbf{w}) \in \mathcal{R}_{\mathcal{L}} : \begin{array}{l} \mathbf{x} \leftarrow \mathcal{B}(1^\lambda, I) \\ \mathbf{w} \leftarrow \widetilde{\text{Ext}}^{\mathcal{B}(\cdot)}(I, \mathbf{x}) \end{array} \right] \geq \delta(\lambda)$$

for some non-negligible function $\delta(\lambda)$. We can thus construct an extractor Ext for WE as follows. The extractor receives as input $(I, \mathbf{x}, m_0, m_1)$ and is given oracle access to \mathcal{A} . It then invokes $\widetilde{\text{Ext}}^{\mathcal{B}}(\cdot)(I, \mathbf{x})$. Whenever $\widetilde{\text{Ext}}$ queries ct_1 to its oracle, Ext samples $b' \leftarrow \{0, 1\}$, computes $\text{ct}_2 \leftarrow \text{Enc}^{\text{sym}}(\mathbf{k}, m_{b'})$, and queries $(\text{ct}_1, \text{ct}_2)$ to its own oracle, forwarding the reply to $\widetilde{\text{Ext}}$. It is easy to see that this perfectly simulates oracle access to \mathcal{B} . Therefore, it holds that

$$\Pr \left[(\mathbf{x}, \mathbf{w}) \in \mathcal{R}_{\mathcal{L}} : \begin{array}{l} (\mathbf{x}, m_0, m_1) \leftarrow \mathcal{A}(1^\lambda, I) \\ \mathbf{w} \leftarrow \text{Ext}^{\mathcal{A}(\cdot)}(I, \mathbf{x}, m_0, m_1) \end{array} \right] \geq \delta(\lambda)$$

as required. \square

5 Laconic OT

As discussed in the introduction, we will now show how to construct a concretely efficient laconic OT protocol from our extractable witness encryption scheme for KZG commitments and openings.

Definition 20 ([CDG⁺17]). *A laconic oblivious transfer scheme is a tuple of PPT algorithms $\text{LOT} = (\text{Setup}, \text{H}, \text{Send}, \text{Receive})$ defined as follows:*

$\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^n)$: *The setup algorithm takes security parameter λ and database length n as input and returns public parameters pp .*

$(\text{digest}, \text{aux}) \leftarrow \text{H}(\text{pp}, D)$: *The hashing algorithm takes public parameters pp and database D as input and returns a public hash digest and some secret auxiliary information aux .*

$c \leftarrow \text{Send}(\text{pp}, \text{digest}, i, m_0, m_1)$: *The sender algorithm takes public parameters pp , database hash digest , index i , and message m_0 and m_1 as input and returns message c .*

$m \leftarrow \text{Receive}(\text{pp}, \text{aux}, c, i)$: *The receiver algorithm takes public parameters pp , auxiliary information aux , sender message c , and index i as input and returns message m .*

Definition 21 (Correctness). *A laconic oblivious transfer scheme $\text{LOT} = (\text{Setup}, \text{H}, \text{Send}, \text{Receive})$ is correct, if for all $\lambda, n \in \mathbb{N}$ with $n = \text{poly}(\lambda)$, all $\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^n)$, all databases $D \in \{0, 1\}^n$, all $(\text{digest}, \text{aux}) \leftarrow \text{H}(\text{pp}, D)$, all $i \in [n]$, all $m_0, m_1 \in \mathcal{M}$, and all $c \leftarrow \text{Send}(\text{pp}, \text{digest}, i, m_0, m_1)$ it holds that $\text{Receive}(\text{pp}, \text{aux}, c, i) = m_{D[i]}$.*

Sender privacy requires the sender's message to the receiver to hide the message that was not selected by the receiver's choice bit.

Definition 22 (Sender Privacy). *A laconic oblivious transfer scheme $\text{LOT} = (\text{Setup}, \text{H}, \text{Send}, \text{Receive})$ is sender private against semi-honest adversaries, if for all $\lambda, n \in \mathbb{N}$ with $n = \text{poly}(\lambda)$, any PPT adversary \mathcal{A} , any database $D \in \{0, 1\}^n$, any $i \in [n]$, and any pair of messages $m_0, m_1 \in \mathcal{M}$, it holds that*

$$\left| \begin{array}{l} \Pr[\text{Expt}_{\mathcal{A}}^{\text{OT-S-Real}}(1^\lambda, 1^n, D, m_0, m_1, i) = 1] \\ - \Pr[\text{Expt}_{\mathcal{A}}^{\text{OT-S-Sim}}(1^\lambda, 1^n, D, m_0, m_1, i) = 1] \end{array} \right| \leq \text{negl}(\lambda),$$

where $\text{Expt}_{\mathcal{A}}^{\text{OT-S-Real}}$ and $\text{Expt}_{\mathcal{A}}^{\text{OT-S-Sim}}$ are defined as follows

$\frac{\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Real}}(1^\lambda, 1^n, D, m_0, m_1, i)}{\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^n)}$ $(\text{digest}, \text{aux}) \leftarrow \text{H}(\text{pp}, D)$ $c \leftarrow \text{Send}(\text{pp}, \text{digest}, i, m_0, m_1)$ $b \leftarrow \mathcal{A}(\text{pp}, c, \text{aux})$ $\text{return } b$	$\frac{\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Sim}}(1^\lambda, 1^n, D, m_0, m_1, i)}{\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^n)}$ $c \leftarrow \text{Sim}(\text{pp}, D, i, m_{D[i]})$ $b \leftarrow \mathcal{A}(\text{pp}, c, \text{aux})$ $\text{return } b$
--	--

In the original work of Cho et al. [CDG⁺17], no explicit definition for receiver privacy was stated. The authors informally argued that any protocol that is not receiver privacy can be transformed into one that is, via the use of generic secure computation. In our work, we do define receiver privacy and we focus on the arguably strongest possible notion, namely that of perfect receiver privacy, where the sender learns no information about the receivers choice bits in the information-theoretic sense. We do not make use of generic secure computation and instead our construction will directly satisfy this security notion.

Definition 23 (Receiver Privacy). *A laconic oblivious transfer scheme $\text{LOT} = (\text{Setup}, \text{H}, \text{Send}, \text{Receive})$ is receiver private against semi-honest adversaries, if for all $\lambda, n \in \mathbb{N}$ with $n = \text{poly}(\lambda)$, any PPT adversary \mathcal{A} , all databases $D \in \{0, 1\}^n$, it holds that*

$$\left| \Pr[\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-R-Real}}(1^\lambda, 1^n, D) = 1] - \Pr[\text{Expt}_{\mathcal{A}}^{\text{OT-R-Sim}}(1^\lambda, 1^n, D) = 1] \right|,$$

where $\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-R-Real}}$ and $\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-R-Sim}}$ are defined as follows.

$\frac{\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-R-Real}}(1^\lambda, 1^n, D)}{\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^n)}$ $(\text{digest}, \text{aux}) \leftarrow \text{H}(\text{pp}, D)$ $b \leftarrow \mathcal{A}(\text{pp}, \text{digest})$ $\text{return } b$	$\frac{\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-R-Sim}}(1^\lambda, 1^n, D)}{\text{pp} \leftarrow \text{Setup}(1^\lambda, 1^n)}$ $\text{digest} \leftarrow \text{Sim}(\text{pp})$ $b \leftarrow \mathcal{A}(\text{pp}, \text{digest})$ $\text{return } b$
---	---

The main property that makes laconic OT non-trivial and interesting is its efficiency. The digest is required to be independent of the original database size, the hashing should run in quasi-linear time in the database size, and both the computational complexity of sending and receiving should have at most a polylogarithmic dependency on the size of the database.

Definition 24 (Efficiency). *A laconic oblivious transfer scheme $\text{LOT} = (\text{Setup}, \text{H}, \text{Send}, \text{Receive})$ is efficient, if $|\text{digest}| \in \text{poly}(\lambda)$ and in particular independent of $|D|$, the hashing algorithm H runs in time $|D| \cdot \text{poly}(\log|D|, \lambda)$, and both Send and Receive run in time $\text{poly}(\log|D|, \lambda)$.*

5.1 Constructing Laconic OT

Our construction is conceptually very simple. The receiver will use a hiding vector commitment to compute a commitment com to their database D . Upon the i -th invocation of the OT, the sender will use our extractable witness encryption scheme to separately encrypt m_0 and m_1 , such that they can be decrypted, if com can be opened to 0 and 1, respectively. The sender, who will receive the two ciphertexts, can then use their opening π_i for commitment com to decrypt the ciphertext containing $m_{D[i]}$. Sender privacy effectively follows from the security guarantees of the extractable witness encryption scheme, whereas receiver privacy follows from the hiding properties of the vector commitment.

$\text{Setup}(1^\lambda, 1^n)$ <hr/> $\text{pp} \leftarrow \text{VC.Setup}(1^\lambda, 1^n)$	$\text{H}(\text{pp}, D)$ <hr/> $(\text{com}, \overline{\text{aux}}) \leftarrow \text{VC.Commit}(\text{pp}, D)$ $(\pi_1, \dots, \pi_n) \leftarrow \text{BatchOpen}(\text{pp}, \overline{\text{aux}})$ return $(\text{com}, (D, \pi_1, \dots, \pi_n))$
$\text{Send}(\text{pp}, \text{digest}, i, m_0, m_1)$ <hr/> $\text{ct}_0 \leftarrow \text{WE.Enc}(\text{pp}, (\text{digest}, i, 0), m_0)$ $\text{ct}_1 \leftarrow \text{WE.Enc}(\text{pp}, (\text{digest}, i, 1), m_1)$ return $(\text{ct}_0, \text{ct}_1)$	$\text{Receive}(\text{pp}, \text{aux}, (\text{ct}_0, \text{ct}_1), i)$ <hr/> parse aux as (D, π_1, \dots, π_n) $b := D_i$ $m_b \leftarrow \text{WE.Dec}(\text{pp}, \pi_i, \text{ct}_b)$ return m_b

Fig. 5. Laconic OT construction.

Theorem 10. *Let $\lambda, n \in \mathbb{N}$ with $n = \text{poly}(\lambda)$. Let $\text{VC} = (\text{Setup}, \text{Commit}, \text{BatchOpen}, \text{Verify})$ be a correct, efficient, position binding, and perfectly weakly-hiding vector commitment. Let $\text{WE} = (\text{Enc}, \text{Dec})$ be an extractable witness encryption scheme for \mathcal{F}_{VC} . Then the construction in Figure 5 is an efficient, sender-private, and perfectly receiver-private laconic OT.*

Proof. Efficiency follows immediately from the efficiency of VC. We will prove the other two properties separately.

Lemma 11. *Let $\text{VC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Verify})$ be a correct and position binding vector commitment. Let $\text{WE} = (\text{Enc}, \text{Dec})$ be an extractable witness encryption scheme for \mathcal{F}_{VC} . Then the construction in Figure 5 is sender private.*

Proof. We specify a simulator as follows. On input $\text{pp}, D, i, m_{D[i]}$ the simulator simply computes

$$\begin{aligned} \text{ct}_{D[i]} &\leftarrow \text{WE.Enc}(\text{pp}, (\text{digest}, i, D[i]), m_{D[i]}) \\ \text{and } \text{ct}_{1-D[i]} &\leftarrow \text{WE.Enc}(\text{pp}, (\text{digest}, i, 1 - D[i]), 0) \end{aligned}$$

and outputs $(\text{ct}_0, \text{ct}_1)$.

Let $D \in \{0, 1\}^n$ be an arbitrary database, $i \in [n]$ an arbitrary index, $m_0, m_1 \in \mathbb{F}$ an arbitrary pair of messages, and \mathcal{A} an arbitrary PPT adversary against sender privacy with

$$\begin{aligned} \epsilon(\lambda) &= \Pr[\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Real}}(1^\lambda, 1^n, D, m_0, m_1, i) = 1] \\ &\quad - \Pr[\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Sim}}(1^\lambda, 1^n, D, m_0, m_1, i) = 1]. \end{aligned} \tag{4}$$

Assume wlog, that $\epsilon(\lambda) \geq 0$.

We construct a PPT adversary \mathcal{B} against the extractability of WE as follows. On input $1^\lambda, \text{ck}$, the adversary \mathcal{B} computes $(\text{com}, \overline{\text{aux}}) \leftarrow \text{VC.Commit}(\text{pp}, D)$ and outputs

$$\left(\underbrace{(\text{com}, i, 1 - D[i])}_{=:x}, 0, m_{1-D[i]} \right).$$

Upon receiving $\text{ct}_{1-D[i]}$, \mathcal{B} further computes $\text{ct}_{D[i]} \leftarrow \text{WE.Enc}(\text{ck}, (\text{digest}, i, D[i]), m_{D[i]})$, sets $\text{aux} := (D, \text{com}, \overline{\text{aux}})$ and invokes $\mathcal{A}(\text{ck}, (\text{ct}_0, \text{ct}_1), \text{aux})$. Eventually \mathcal{A} will output a bit, which \mathcal{B} will also output.

Let $b \in \{0, 1\}$ be the bit in the experiment $\text{Expt}_{\text{WE}, \mathcal{B}}^{\text{CPA}}(1^\lambda)$. It is easy to see that, if $b = 1$, then \mathcal{B} perfectly simulates $\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Real}}(1^\lambda, 1^n, D, m_0, m_1, i)$. On the other hand, if $b = 0$, then \mathcal{B} perfectly simulates $\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Sim}}(1^\lambda, 1^n, D, m_0, m_1, i)$. It thus follows that

$$\begin{aligned}
& \Pr[\text{Expt}_{\text{WE}, \mathcal{B}}^{\text{CPA}}(1^\lambda, \text{ck}) = 1] \\
&= \frac{1}{2} (\Pr[\text{Expt}_{\text{WE}, \mathcal{A}}^{\text{CPA}}(1^\lambda) = 1 \mid b = 1] + \Pr[\text{Expt}_{\text{WE}, \mathcal{A}}^{\text{CPA}}(1^\lambda) = 1 \mid b = 0]) \\
&= \frac{1}{2} (\Pr[\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Real}}(1^\lambda, 1^n, D, m_0, m_1, i) = 1] + \Pr[\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Sim}}(1^\lambda, 1^n, D, m_0, m_1, i) = 0]) \\
&= \frac{1}{2} (\Pr[\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Real}}(1^\lambda, 1^n, D, m_0, m_1, i) = 1] + 1 - \Pr[\text{Expt}_{\text{LOT}, \mathcal{A}}^{\text{OT-S-Sim}}(1^\lambda, 1^n, D, m_0, m_1, i) = 1]) \\
&= \frac{1}{2} + \frac{1}{2} \epsilon(\lambda).
\end{aligned}$$

Since WE is extractable, it follows that there exists a PPT algorithm Ext, such that, if $\epsilon(\lambda)$, and thereby also $\epsilon(\lambda)/2$ were non-negligible, it would hold that

$$\Pr \left[\text{Verify}(\text{ck}, \text{com}, \pi, i, 1 - D[i]) = 1 : \begin{array}{l} ((\text{com}, i, 1 - D[i]), 0, m_{1-D[i]}) \leftarrow \mathcal{B}(1^\lambda, \text{ck}) \\ \pi \leftarrow \text{Ext}^{\mathcal{B}(\cdot)}(\text{ck}, (\text{com}, i, 1 - D[i]), 0, m_{1-D[i]}) \end{array} \right] \geq \delta(\lambda)$$

for some non-negligible function $\delta(\lambda)$.

To show that this can't be the case, we construct an adversary \mathcal{C} against the position binding of VC as follows. On input ck , the adversary \mathcal{C} computes $(\text{com}, \overline{\text{aux}}) \leftarrow \text{VC.Commit}(\text{pp}, D)$, and invokes $\text{Ext}^{\mathcal{B}(\cdot)}((\text{com}, i, 1 - D[i]), 0, m_{1-D[i]})$. Whenever Ext queries its oracle with $\text{ct}_{1-D[i]}$, \mathcal{C} further computes $\text{ct}_{D[i]} \leftarrow \text{WE.Enc}(\text{ck}, (\text{digest}, i, D[i]), m_{D[i]})$, sets $\text{aux} := (D, \text{com}, \overline{\text{aux}})$, invokes $b \leftarrow \mathcal{A}(\text{ck}, (\text{ct}_0, \text{ct}_1), \text{aux})$ and replies with b . Eventually, Ext will output π . \mathcal{C} will then compute $\pi' \leftarrow \text{VC.Open}(\text{ck}, \overline{\text{aux}}, i)$ and output $\text{com}, i, 1 - D[i], D[i], \pi, \pi'$.

It is easy to see, that \mathcal{C} perfectly simulates \mathcal{B} for the extractor. Therefore, it holds with probability at least $\delta(\lambda)$, that $\text{Verify}(\text{ck}, \text{com}, \pi, i, 1 - D[i]) = 1$. The correctness of the vector commitment guarantees that $\text{Verify}(\text{ck}, \text{com}, \pi', i, D[i]) = 1$. Since further $D[i] \neq 1 - D[i]$, it thus holds that

$$\text{negl}(\lambda) \geq \Pr \left[\begin{array}{l} m_0 \neq m_1 \\ \wedge \text{Verify}(\text{ck}, \text{com}, \pi_0, i, m_0) = 1 : \\ \wedge \text{Verify}(\text{ck}, \text{com}, \pi_1, i, m_1) = 1 \end{array} : \begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda, 1^n) \\ (\text{com}, i, m_0, m_1, \pi_0, \pi_1) \leftarrow \mathcal{C}(\text{ck}) \end{array} \right] = \delta(\lambda)$$

Which immediately implies that $\epsilon(\lambda) \leq \text{negl}(\lambda)$ as required. \square

Lemma 12. *Let $\text{VC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Verify})$ be a perfectly weakly hiding vector commitment. Then the construction in Figure 5 is perfectly receiver private.*

Proof. The simulator takes as input the public parameters $\text{pp} = \text{ck}$, computes $(\text{com}, \overline{\text{aux}}) \leftarrow \text{VC.Commit}(\text{ck}, 0^n)$ and outputs com . The only difference between the two experiments from the point of view of an adversary is that in one case they receive a commitment to D , whereas in the other they receive a commitment to 0^n . However, since VC is perfectly weakly hiding, those two commitments are distributed identically. \square

Now Theorem 10 follows directly by combining Lemmas 11 and 12. \square

6 Benchmarks

The laconic OT protocol from Section 5 was instantiated with the KZG based vector commitment from Section 2.5 and the extractable witness commitment derived by combining the extractable witness KEM from Section 3 with the generic transformation from Section 4. A simple one-time pad as the symmetric encryption scheme in the transformation. The full construction was implemented⁶ in Rust using `arkworks` [ac22]. The BLS12-381 [Bow17] curve was used as the pairing-friendly elliptic curve throughout the implementation. All involved computational costs and bandwidth overheads were measured for various parameters. All benchmarks were run on a personal laptop with an i7-11800H @ 2.30GHz CPU and 64 GB of RAM.

As already explained in Section 2.5, the batch opening technique of Feist and Khovratovich was used [FK23] to precompute all openings of the used vector commitment during the hashing of the database of choice indices. For concreteness the sender’s OT inputs are assumed to be 256-bit messages. A single witness encryption for such a message consists of a 96 byte witness KEM ciphertext and a 32 bytes large one-time pad encryption of the message itself, thus leading to a 128 byte ciphertext. The public parameters for a database of size n consists of $n + 1$ many \mathbb{G}_1 elements and two \mathbb{G}_2 elements with \mathbb{G}_1 and \mathbb{G}_2 for BLS12-381 clocking in at 48 and 96 bytes respectively. The benchmark results can be found in Table 1.

$ D $	Sizes			Times		
	pp	digest	Sender Msg.	Hash	Send	Receive
2^6	3.2 KB	48 B	256 B	173 ms	4 ms	1 ms
2^8	12.2 KB	48 B	256 B	723 ms	4 ms	1 ms
2^{10}	48.2 KB	48 B	256 B	3 s	4 ms	1 ms
2^{12}	192.2 KB	48 B	256 B	10 s	4 ms	1 ms
2^{14}	768.2 KB	48 B	256 B	43 s	4 ms	1 ms
2^{16}	3.0 MB	48 B	256 B	3 min	5 ms	1 ms
2^{18}	12.0 MB	48 B	256 B	8 min	5 ms	1 ms
2^{31}	96.0 GB	48 B	256 B	—	5 ms	1 ms

Table 1. Runtimes and bandwidth overheads of our laconic OT protocol for varying sizes of the receiver’s database D . The computation times for $|D| = 2^{31}$ are extrapolated to compare with [GJL23].

Our construction is highly efficient in most parameters, in particular database digest and the sender’s message are both not just constant in size but concretely very small. The time to compute the sender’s message and to decode on the receiver’s side in an OT invocation is below 5 ms. The main drawbacks of our construction the initial one-time cost of computing the hash of the receiver’s database and the size of the public parameters.

Weakly Laconic OT. We also examine the efficiency of a variant of our construction, which does not strictly satisfy the formal efficiency requirements for laconic OT protocols from Definition 24 but has significantly smaller public parameters at the cost of somewhat larger digests. The idea for achieving this trade-off is to simply partition the receiver’s database D into $\sqrt{|D|}$ smaller

⁶ <https://github.com/rot256/research-we-kzg>

$ D $	Sizes		Sender Msg.	Times		
	pp	digest		Hash	Send	Receive
2^6	624 B	384 B	256 B	194 ms	4 ms	1 ms
2^8	1008 B	768 B	256 B	743 ms	4 ms	1 ms
2^{10}	1.7 KB	1.5 KB	256 B	3 s	4 ms	1 ms
2^{12}	3.2 KB	3.0 KB	256 B	11 s	4 ms	1 ms
2^{14}	6.2 KB	6.0 KB	256 B	44 s	4 ms	1 ms
2^{16}	12.2 KB	12.0 KB	256 B	3 min	4 ms	1 ms
2^{18}	24.2 KB	24.0 KB	256 B	12 min	4 ms	1 ms
2^{31}	3 MB	1.5 MB	256 B	69 days	4 ms	1 ms

Table 2. Runtimes and bandwidth overheads for the variant of our laconic OT protocol, which splits the database D into $\sqrt{|D|}$ many smaller databases. The numbers are extrapolated from our experimental results for laconic OT.

databases and to then hash each of them individually. The hash returned by the receiver is simply the concatenation of all $\sqrt{|D|}$ hashes. This construction does not formally satisfy the asymptotic efficiency requirement of a laconic OT because the digest size now (sublinearly) depends on the size of the input database. The main observation here is that all of those hashes can be computed using the same trusted setup for instances of size $\sqrt{|D|}$. This simple trick was already observed by Green, Jain, and Van Laer in their work [GJL23]. We calculate the sizes of objects and we extrapolate the timing benchmarks for this construction from our benchmarks for regular laconic OT. The results can be found in Table 2.

Comparison to Related Works. Equipped with the concrete performance numbers for our constructions, we can compare its efficiency to that of existing protocols. Green, Jain, Van Laer [GJL23] provide benchmarks for a single data point for their laconic OT construction, which does not achieve sender privacy. They argue that one can generically obtain sender privacy by using garbled circuits, which would incur a significant overhead on top of their provided benchmarks. In their favour, we shall ignore this and simply compare it to our construction that achieves both sender and receiver privacy.

For a database size of 2^{31} , their construction’s public parameters 412.3 GB, their digest is 48 B, their sender’s message size is 1.34 KB and receiving takes 27.7 *minutes*. That means their public parameters are $\approx 4x$ larger, their digests have the same size, their sender’s message is $\approx 5x$ larger, and their receiving time is larger by 6 *orders of magnitude*. No further data points were provided in their benchmarks.

The authors also provide benchmark results for a weakly laconic version of their construction, which splits the database in $\sqrt{|D|}$ chunks as described above. Comparing it to our weakly laconic OT construction, they have $\approx 3x$ larger public parameters, comparable digest size, and $\approx 5x$ larger sender’s message. Their receiver’s computational time is still $\approx 38x$ larger.

Döttling et al. [DKL⁺23] only sketch how to construct laconic OT from another cryptographic primitive, which they do provide some benchmarks for. The main advantage of their work is that the public parameters do not grow with the receiver’s database size. Taking the benchmarks for their lowest security levels, their public parameters are 8 MB, their digest is 2 KB, and their sender’s message size is 98 MB. The constant size of their public parameters makes them concretely smaller

than ours once the databases size reaches 2^{18} . Their digests are always $\approx 8x$ larger and their sender’s message size is 5 orders of magnitude larger.

Comparing with our weakly laconic OT protocol, our public parameters would only become concretely larger than theirs, when the database size reaches at least 2^{35} . For a database size of 2^{18} , our digest is $\approx 12x$ larger, but our sender’s message is still 5 orders of magnitude smaller. We believe that a slightly larger digest size in exchange for much smaller sender messages is a valuable trade-off. The digest is sent once, but the sender’s messages need to be sent for each OT invocation.

Acknowledgements

We would like to thank Nathan Xiong for pointing out a flaw in an earlier version of the proof of Claim 5.

References

- ac22. arkworks contributors. arkworks zksnark ecosystem, 2022. URL: <https://arkworks.rs>. 6
- BGI⁺01. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Heidelberg, Germany. doi:10.1007/3-540-44647-8_1. 1.2
- BJKL21. Fabrice Benhamouda, Aayush Jain, Ilan Komargodski, and Huijia Lin. Multiparty reusable non-interactive secure computation from LWE. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 724–753, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-77886-6_25. 1.2
- BL18. Fabrice Benhamouda and Huijia Lin. k-round multiparty computation from k-round oblivious transfer via garbled interactive circuits. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 500–532, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany. doi:10.1007/978-3-319-78375-8_17. 1.2
- BL20. Fabrice Benhamouda and Huijia Lin. Mr NISC: Multiparty reusable non-interactive secure computation. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020: 18th Theory of Cryptography Conference, Part II*, volume 12551 of *Lecture Notes in Computer Science*, pages 349–378, Durham, NC, USA, November 16–19, 2020. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-64378-2_13. 1.2
- Bow17. Sean Bowe. Bls12-381: New zk-snark elliptic curve construction, March 2017. URL: <https://electriccoin.co/blog/new-snark-curve/>. 6
- CDG⁺17. Chongwon Cho, Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Antigoni Polychroniadou. Laconic oblivious transfer and its applications. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part II*, volume 10402 of *Lecture Notes in Computer Science*, pages 33–65, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany. doi:10.1007/978-3-319-63715-0_2. 1.1, 1.1, 1.2, 20, 5
- CF13. Dario Catalano and Dario Fiore. Vector commitments and their applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013: 16th International Conference on Theory and Practice of Public Key Cryptography*, volume 7778 of *Lecture Notes in Computer Science*, pages 55–72, Nara, Japan, February 26 – March 1, 2013. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-36362-7_5. 1
- CFK22. Matteo Campanelli, Dario Fiore, and Hamidreza Khoshakhlagh. Witness encryption for succinct functional commitments and applications. *Cryptology ePrint Archive*, Report 2022/1510, 2022. <https://eprint.iacr.org/2022/1510>. 1.2
- CHM⁺20. Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Psi Vesely, and Nicholas P. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 738–768, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-45721-1_26. 1, 1

- CV21. Gwangbae Choi and Serge Vaudenay. Towards witness encryption without multilinear maps - extractable witness encryption for multi-subset sum instances with no small solution to the homogeneous problem. In Jong Hwan Park and Seung-Hyun Seo, editors, *ICISC 21: 24th*, volume 13218, pages 28–47, Seoul, South Korea, December 1–3 2021. doi:10.1007/978-3-031-08896-4_2. 3
- DKL⁺23. Nico Döttling, Dimitris Kolonelos, Russell W. F. Lai, Chuanwei Lin, Giulio Malavolta, and Ahmadreza Rahimi. Efficient laconic cryptography from learning with errors. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part III*, volume 14006 of *Lecture Notes in Computer Science*, pages 417–446, Lyon, France, April 23–27, 2023. Springer, Heidelberg, Germany. doi:10.1007/978-3-031-30620-4_14. 1.2, 6
- FK23. Dankrad Feist and Dmitry Khovratovich. Fast amortized KZG proofs. Cryptology ePrint Archive, Report 2023/033, 2023. https://eprint.iacr.org/2023/033. 2.4, 2.5, 6
- FKL18. Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 33–62, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany. doi:10.1007/978-3-319-96881-0_2. 1, 2.1
- GGH⁺13. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual Symposium on Foundations of Computer Science*, pages 40–49, Berkeley, CA, USA, October 26–29, 2013. IEEE Computer Society Press. doi:10.1109/FOCS.2013.13. 1.2
- GGHW14. Sanjam Garg, Craig Gentry, Shai Halevi, and Daniel Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 518–535, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany. doi:10.1007/978-3-662-44371-2_29. 1.2
- GGSW13. Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th Annual ACM Symposium on Theory of Computing*, pages 467–476, Palo Alto, CA, USA, June 1–4, 2013. ACM Press. doi:10.1145/2488608.2488667. 1.1, 1.1, 1.2
- GJL23. Matthew Green, Abhishek Jain, and Gijs Van Laer. Efficient set membership encryption and applications. pages 1080–1092, Copenhagen, Denmark, November 26–30, 2023. doi:10.1145/3576915.3623131. 1.2, 1, 6, 6
- GKP⁺13. Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. How to run Turing machines on encrypted data. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 536–553, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-40084-1_30. 1.1, 1.1, 1.2, 4
- GRWZ20. Sergey Gorbunov, Leonid Reyzin, Hoeteck Wee, and Zhenfei Zhang. Pointproofs: Aggregating proofs for multiple vector commitments. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020: 27th Conference on Computer and Communications Security*, pages 2007–2023, Virtual Event, USA, November 9–13, 2020. ACM Press. doi:10.1145/3372297.3417244. 1
- GS17. Sanjam Garg and Akshayaram Srinivasan. Garbled protocols and two-round MPC from bilinear maps. In Chris Umans, editor, *58th Annual Symposium on Foundations of Computer Science*, pages 588–599, Berkeley, CA, USA, October 15–17, 2017. IEEE Computer Society Press. doi:10.1109/FOCS.2017.60. 1.2
- GWC19. Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. https://eprint.iacr.org/2019/953. 1
- KZG10. Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *Advances in Cryptology – ASIACRYPT 2010*, volume 6477 of *Lecture Notes in Computer Science*, pages 177–194, Singapore, December 5–9, 2010. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-17373-8_11. 1
- LPR22. Benoît Libert, Alain Passetègue, and Mahshid Riahinia. PointProofs, revisited. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part IV*, volume 13794 of *Lecture Notes in Computer Science*, pages 220–246, Taipei, Taiwan, December 5–9, 2022. Springer, Heidelberg, Germany. doi:10.1007/978-3-031-22972-5_8. 1
- MBKM19. Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings. In Lorenzo Cavallaro, Johannes

- Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019: 26th Conference on Computer and Communications Security*, pages 2111–2128, London, UK, November 11–15, 2019. ACM Press. doi: [10.1145/3319535.3339817](https://doi.org/10.1145/3319535.3339817). 1
- Rab81. Michael O. Rabin. How to exchange secrets with oblivious transfer. *Technical Report TR-81, Aiken Computation Lab, Harvard University*, 1981. URL: <http://eprint.iacr.org/2005/187>. 1.1
- SCP⁺22. Shravan Srinivasan, Alexander Chepur, Charalampos Papamanthou, Alin Tomescu, and Yupeng Zhang. Hyperproofs: Aggregating and maintaining proofs in vector commitments. In Kevin R. B. Butler and Kurt Thomas, editors, *USENIX Security 2022: 31st USENIX Security Symposium*, pages 3001–3018, Boston, MA, USA, August 10–12, 2022. USENIX Association. 1
- TAB⁺20. Alin Tomescu, Ittai Abraham, Vitalik Buterin, Justin Drake, Dankrad Feist, and Dmitry Khovratovich. Aggregatable subvector commitments for stateless cryptocurrencies. In Clemente Galdi and Vladimir Kolesnikov, editors, *SCN 20: 12th International Conference on Security in Communication Networks*, volume 12238 of *Lecture Notes in Computer Science*, pages 45–64, Amalfi, Italy, September 14–16, 2020. Springer, Heidelberg, Germany. doi: [10.1007/978-3-030-57990-6_3](https://doi.org/10.1007/978-3-030-57990-6_3). 1
- Tsa22. Rotem Tsabary. Candidate witness encryption from lattice techniques. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part I*, volume 13507 of *Lecture Notes in Computer Science*, pages 535–559, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Heidelberg, Germany. doi: [10.1007/978-3-031-15802-5_19](https://doi.org/10.1007/978-3-031-15802-5_19). 1.2
- VWW22. Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and null-IO from evasive LWE. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part I*, volume 13791 of *Lecture Notes in Computer Science*, pages 195–221, Taipei, Taiwan, December 5–9, 2022. Springer, Heidelberg, Germany. doi: [10.1007/978-3-031-22963-3_7](https://doi.org/10.1007/978-3-031-22963-3_7). 1.2
- WUP23. Weijie Wang, Annie Ulichney, and Charalampos Papamanthou. Balanceproofs: Maintainable vector commitments with fast aggregation. pages 4409–4426, Anaheim, CA, USA, August 9–11, 2023. 1