On Zero-Testable Homomorphic Encryption and Publicly Verifiable Non-interactive Arguments

Omer Paneth[1(✉)] and Guy N. Rothblum[2]

[1] MIT, Cambridge, USA
omerpa@gmail.com
[2] Weizmann Institute of Science, Rehovot, Israel

Abstract. We define and study *zero-testable homomorphic encryption* (ZTHE) – a *semantically secure*, somewhat homomorphic encryption scheme equipped with a *weak zero test* that can identify *trivial zeros*. These are ciphertexts that result from homomorphically evaluating an arithmetic circuit computing the zero polynomial over the integers. This is a relaxation of the (strong) zero test provided by the notion of graded encodings, which identifies all encodings of zero.

We show that ZTHE can suffice for powerful applications. Based on any ZTHE scheme that satisfies the additional properties of correctness on adversarial ciphertexts and multi-key homomorphism, we construct publicly verifiable non-interactive arguments for delegating computation. Such arguments were previously constructed from indistinguishability obfuscation or based on so-called knowledge assumptions. The arguments we construct are adaptively sound, based on an efficiently falsifiable assumption, and only make black-box use of the underlying cryptographic primitives.

We also show that a ZTHE scheme that is sufficient for our application can be constructed based on an efficiently-falsifiable assumption over so-called "clean" graded encodings.

1 Introduction

Recent breakthroughs in the study of fully homomorphic encryption [Gen09] and program obfuscation [GGH+13b] have revolutionized the foundations of cryptography. Fully homomorphic encryption (FHE) allows arbitrary polynomial-time computations to be performed "homomorphically" on encrypted data, while

This work subsumes an earlier report posted on the Cryptology ePrint Archive [PR14]. The current version features new results and addresses correctness issues with the earlier report.

O. Paneth—Research supported in part by NSF Grants CNS-1350619, CNS-1414119 and CNS-1413920, by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236 and by Simons Investigator Award Agreement Dated 6-5-12.

Y. Kalai and L. Reyzin (Eds.): TCC 2017, Part II, LNCS 10678, pp. 283–315, 2017.
https://doi.org/10.1007/978-3-319-70503-3_9

ensuring that semantic security is maintained and *nothing* about the data can be learned. While this powerful security guarantee enables important applications, other scenarios require more fine-grained control: allowing some information about the data to be exposed, while other information remains hidden. Multilinear maps [BS02] and graded encodings [GGH13a] are basic building blocks that have proven to be incredibly useful in such scenarios. Intuitively, a graded encoding scheme is a *somewhat homomorphic encryption*, supporting homomorphic evaluation of low-degree algebraic computations, with an additional capability: an efficient *zero test* procedure that publicly identifies encodings of zero. Graded encodings cannot be semantically secure: the zero test procedure leaks *partial information* on the encoded elements. Nevertheless, other information can remain hidden (in particular, inverting the encoding might still be hard). This balance between functionality and security makes the notion of graded encoding incredibly useful for computing on encrypted data, with applications such as indistinguishability obfuscation and functional encryption [GGH+13b, GGHZ16].

While homomorphic encryption can by based on the Learning with Errors assumption [BV11, GSW13], the situation for graded encodings is less clear. Analyzing the security of existing candidates and designing new ones are central challenges [GGH13a, CLT15, GGH15, CHL+15, HJ16, MSZ16, GMM+16].

Zero-testable homomorphic encryption. In this work we define and study a new relaxation of graded encodings that we call zero-testable (somewhat) homomorphic encryption (ZTHE). A ZTHE is a semantically secure somewhat homomorphic encryption scheme equipped with a *weak zero test* that can only identify *trivial zeros*. These are ciphertexts that result from homomorphically evaluating an arithmetic circuit computing the zero polynomial over $\mathbb{Z}$. The weak zero test should accept such trivial zeros, but reject ciphertexts that encrypt non-zero values.

Importantly, an efficient weak zero test poses no contradiction to semantic security, since it does not allow to distinguish between encryptions of two different values. Given a ciphertext c it is possible to homomorphically evaluate a circuit P on c and test if the result is a trivial zero. However, this does not give any information on the value encrypted in c, since the zero test only required to pass if P vanishes on *all* values. Intuitively, the zero test is giving information on the evaluated computation P rather then on the ciphertext c. Indeed, semantic security implies that if P only vanishes on some values, then even if the evaluated ciphertext encrypts zero it will not pass the weak zero test (except with negligible probability). Otherwise, the zero test would have revealed information on the original encrypted evaluation point.

From ZTHE to delegation. The main technical result in this work demonstrates that ZTHE can suffice for powerful applications. Based on any ZTHE scheme that satisfies the additional properties of correctness on adversarial ciphertexts and multi-key homomorphism (we elaborate on these additional properties below), we construct *publicly verifiable non-interactive arguments for delegating computation*. Such arguments were previously constructed from

indistinguishability obfuscation or based on so-called knowledge assumptions. Our construction follows a new approach and has important properties, such as adaptive soundness, reduction to an efficiently falsifiable assumption, and black-box use of the underlying cryptographic primitives. We note that the additional properties we assume (adversarial correctness and multi-key homomorphism) make ZTHE incomparable to "vanilla" graded encodings: the weak zero test assumption is more relaxed than the strong zero test of graded encodings schemes, but we require a stronger correctness property (namely correctness on adversarially generated ciphertexts).

ZTHE Candidate. We study the feasibility of constructing ZTHE. First, we observe that several existing somewhat homomorphic encryption schemes [Gen09, vDGHV10] admit a simple weak zero test. These schemes, however, do not satisfy the additional properties required for our non-interactive arguments. We construct ZTHE that is sufficient for our application based on an efficiently-falsifiable assumption over graded encodings with strong properties such as adversarial correctness. Our construction cannot be instantiated based on the existing graded encoding candidates (so-called "clean" graded encodings [Zim15, LV16] do guarantee these stronger properties). We leave the question of ZTHE instantiations as an important open problem and hope it will lead to new and improved deletion protocols based on weaker assumptions, as well as other applications.

Organization. In the rest of this introduction we elaborate on our results and techniques. Section 1.1 gives background on non-interactive arguments and discusses our main technical result, a construction of non-interactive arguments from ZTHE. In Sect. 1.2 we present our results in more detail. The construction of non-interactive arguments from ZTHE is described in Sect. 1.3. The construction of ZTHE from graded encodings is described in Sect. 1.4.

1.1 Non-interactive Arguments

Background. The power of efficiently verifiable proof systems is a foundational issue in the study of computation. A central goal is constructing proof systems that can be used by a powerful prover to convince a weak verifier of the correctness of a complex computational statement, usually framed as proving membership of an input x in a language $\mathcal{L}$. Beyond its foundational importance in the theory of computation, this question has real-world applications, such as *delegating computation*. In this setting, a powerful server (playing the role of the prover) can run a complex computation for a much weaker client (playing the role of the verifier), and provide a proof of the output's correctness.

A similar question was raised by Babai, Lund, Fortnow and Szegedy [BFLS91] in the PCP setting. Kilian [Kil92] and Micali [Mic94] gave the first candidate scheme for delegating computation. The question re-emerged in the theoretical literature in the work of Goldwasser, Kalai and Rothblum [GKR08], and became the focus of a rich body of research spanning theory and systems. See the recent survey by Walfish and Blumberg [WB13].

A "holy grail" for delegating computations is *fully non-interactive proofs*, comprised of a single message sent from the prover to the verifier with unconditional soundness, as in classic NP or Merlin-Arthur proofs. Unfortunately, there are serious barriers to constructing such proofs for delegating general deterministic computations (in particular, they imply Merlin-Arthur speedups for deterministic computations). Thus, a body of research has focused on computationally sound proofs in the common reference string model, where:

1. Soundness is only required to hold against *efficient* cheating provers. Computationally sound proof systems are commonly called *argument systems*.
2. There is a (public) *common reference string* (CRS), generated in advance by a trusted authority (or the verifier herself). This CRS can be used (repeatedly) by different parties to verify proofs. The prover and the verifier both have access to the CRS, but neither has access to the secret coins used to generate the CRS.

We focus on non-interactive argument systems for polynomial-time computations, where the verifier should be super-efficient (nearly-linear in the input length), and the honest prover should run in polynomial time. Non-interactive arguments are especially attractive for delegating computation, as any untrusted server can simply use the CRS to generate proofs and send them off (non-interactively and asynchronously), to be verified at the clients' convenience. We refer to such a system as a *publicly verifiable non-interactive argument for delegating computation*. For the remainder of this work, we use the term *non-interactive argument* as shorthand.

Prior works on non-interactive arguments. In his seminal work, Micali [Mic94] gave the first construction of non-interactive arguments in the random oracle model. However, instantiating random oracle model constructions in a provably secure way is notoriously difficult, and often impossible [CGH04, GW11]. A rich body of research has aimed to construct non-interactive arguments in the plain model led to a variety of beautiful constructions based on strong cryptographic assumptions.

One line of works based non-interactive arguments on non-falsifiable[1] *knowledge assumptions* such as the *knowledge of exponent assumption* in bilinear groups [Gro10, Lip12, DFH12, GGPR13, BCI+13, BCCT13]. A recent sequence of works [SW14, BGL+15, CHJV14, KLW14] show how to base non-interactive arguments on *indistinguishability obfuscation* (IO). Based on standard assumptions such as somewhat-homomorphic encryption or private information retrieval schemes, the works of [KRR13, KRR14, BHK16] achieve the weaker notion of *designated-verifier arguments*. These are two-message arguments where, in the first message, the verifier samples the CRS and sends it to the prover. The secret coins used to sample the CRS are required to verify the proof sent in the second message.

[1] A "falsifiable" assumption [Nao03] is one that can be efficiently refuted. Falsifiability is a basic "litmus test" for cryptographic assumptions.

This work. Our main technical result is a construction of non-interactive arguments from any ZTHE with the additional properties mentioned above (see Sect. 1.2). Our construction follows a different approach from previous works and leverages ideas and techniques that were previously used only in the context of designated-verifier arguments [KRR14,BHK16], such as efficient probabilistically checkable proofs and no-signaling soundness. As a result, our non-interactive arguments have some notable advantages compared to previous works:

- **Efficiently falsifiable assumptions.** Our arguments are based on the semantic security of the underlying ZTHE - an efficiently falsifiable assumption. Moreover, in our candidate construction of ZTHE from graded encodings, we further base semantic security of the ZTHE on a simple and efficiently falsifiable assumption on the graded encodings. Taken together, we can base soundness of the argument system on a falsifiable assumption on graded encodings.
  In contrast, the constructions of publicly verifiable non-interactive argument are based on assumptions that are not efficiently falsifiable. IO was recently constructed from simpler primitives such as multi-linear maps or functional encryption. However, these construction involve a sub-exponential security loss. While many applications of IO can be based directly on polynomially secure functional encryption, currently non-interactive arguments still require the full power of IO. For more information on this line of work, see [GPSZ17] and references therein.
  We note that for any particular non-interactive argument candidate, the assumption that the candidate is secure is efficiently falsifiable. Therefore, our focus will be on falsifiable assumptions that are elementary and natural compared to the tautological assumption that the candidate is secure.
- **Adaptive soundness.** The soundness of our non-interactive arguments is adaptive: it holds even when the statement proven is chosen as a function of the CRS. Adaptive soundness is required in many applications, and it is especially important in settings where the CRS is set "once and for all".
  We note that any sound argument can be turned into an adaptively sound one via "complexity leveraging". However, this reduction incurs an exponential loss in security, and therefore cannot be based on efficiently falsifiable assumptions.
- **Black-box construction.** In contrast to all previous construction of non-interactive arguments, our construction makes only black-box use of the underlying cryptographic primitives.[2] Understanding the feasibility and limitation of black-box constructions in cryptography is the subject of a rich body of work motivated both by theoretical interest as well as efficiency considerations.

[2] One exception is instantiating Micali's random oracle construction with a cryptographic hash function. However, beyond assuming this construction is secure, we do not know how to reduce its security to a simpler assumption.

1.2 Our Results in More Details

In this section we present our results in more details. We start by describing the basic notion of zero testable homomorphic encryption and the additional properties we consider.

Zero-testable homomorphic encryption. A homomorphic encryption is a semantically secure public key encryption equipped with a public evaluation algorithm that adds, subtracts and multiplies values homomorphically "under the encryption". We focus on somewhat homomorphic encryption that only supports homomorphic evaluation of polynomial-size arithmetic circuits of logarithmic degree. That is, of degree $c \cdot \log \lambda$ for any constant c, where λ is the security parameter. We require that ciphertexts are succinct: their size is bounded by some fixed polynomial in λ that is independent of c.

A zero-testable somewhat homomorphic encryption (ZTHE) has an additional zero test procedure that takes a ciphertext and tests if it is a trivial zero. In more detail, we consider the homomorphic evaluation of a circuit P over freshly encrypted ciphertexts $c_1, \ldots, c_n$, resulting in the evaluated ciphertext c. If the polynomial computed by P is identically zero over $\mathbb{Z}$, then we require that c passes the zero test. We also require that a ciphertext c' that decrypts to a non-zero value does not pass the zero-test. If c decrypts to zero, but it is not a trivial zero, we make no requirement on the outcome of the zero test. However, as discussed above, it follows from the semantic security of the encryption that such a ciphertext should not pass the zero test. Moreover, we note that even if P vanishes on all boolean inputs, but it is not identically zero as a polynomial over $\mathbb{Z}$, we still expect the zero test to fail. Otherwise, the zero test can be used to efficiently decide the satisfiability of P.

We further study the following additional properties of ZTHE, which we use in our construction of non-interactive arguments:

Multi-key evaluation. In multi-key homomorphic encryption, introduced by López-Alt et al. [LTV12], homomorphic computation can be executed over ciphertexts encrypted under different keys. To ensure semantic security, decrypting the result requires all secret keys. We use ZTHE for three keys. That is, it is possible to homomorphically compute over ciphertexts encrypted under at most three different keys, and to run a weak zero test on the result. Importantly, a system can generate ciphertext under an unbounded number of keys and any three of them can be combined in a homomorphic computation. The encryption may also use shared public parameters to generate all keys.

Correctness for adversarially generated ciphertexts. We require that an efficient adversary, given the public key, cannot generate a pair of ciphertexts that result in an evaluation error. A pair of ciphertexts c_1, c_2 cause an evaluation error if computing a homomorphic operation $\star$ over c_1, c_2 and decrypting the evaluated ciphertext c give a different result than decrypting c_1 and c_2 and computing $\star$ on the decrypted values. If c_1 and c_2 are generated honestly, this follows from the standard correctness guarantee of the encryption. However, we

require correctness even when the ciphertext are not generated honestly. Note that the zero test is only required to accept honest ciphertexts that are trivially zero. However, even a malformed ciphertext that decrypts to a non-zero value should make the zero test reject.

In known constructions of somewhat homomorphic encryption, there exist invalid ciphertexts that do not represent an encryption of any value. To account for such candidates, we allow the decryption algorithm to fail. If c_1 or c_2 are invalid (fail to decrypt) we require that the evaluated ciphertext c is invalid as well. If both c_1 and c_2 are valid, we require that c is either invalid or it decrypts to the correct value.

Theorem 1.1 (Informal). *Assuming a 3-key zero-testable somewhat homomorphic encryption scheme with correctness for adversarially-generated ciphertexts, there exists an adaptively-secure publicly-verifiable non-interactive argument for delegating all polynomial time computations. The non-interactive argument uses the encryption scheme as a black box.*

Instantiations: discussion. We observe that existing constructions of somewhat homomorphic encryption, such as the ones in [Gen09, vDGHV10], already support zero testing: simply test if the ciphertext is zero in the ring of ciphertexts. More generally, in any encryption scheme where ciphertexts are elements of some ring, and the homomorphic operations on ciphertext identify with the ciphertext-ring operations, every trivial zero is represented by the zero of the ciphertext ring. While these construction satisfy the weak zero test requirement, they do not seem to support the additional properties stated above.

Following the observations in [LTV12, GHV10, HRSV11], any homomorphic encryption scheme that supports homomorphic computations of sufficiently large degree can be generically modified to satisfy both multi-key evaluation for a constant number of keys and correctness for adversarially generated ciphertexts. This transformation, however, may not preserve the weak zero test property. Roughly speaking, the generic transformation is based on the idea of bootstrapping [Gen09], where the evaluated circuit is modified to include the decryption circuit of the scheme itself. Now, even if we evaluate a circuit computing the zero polynomial, the modified circuit, which now runs the scheme's decryption circuit, will not be identically zero.

We show that ZTHE satisfying both additional properties can be constructed from graded encodings with additional properties described below.

Graded encoding. A graded encoding is an encoding scheme for elements of a ring. We consider a symmetric graded encoding that supports homomorphic computations of bounded degree Δ. The encoding scheme also features a (strong) zero test that identifies encodings of zero (even non-trivial ones). In Sect. 1.4 we describe the interface of a graded encoding scheme in more detail.

We consider graded encodings that satisfy a simple and natural decisional assumption.

Assumption 1.2 (Informal). *Given encoded coefficients $\alpha_0, \ldots, \alpha_\Delta$ of a random degree Δ polynomial, it is hard to distinguish an encoding of a root from an encoding of a random element.*

Intuitively, this problem should be hard since testing if the given encoding is a root requires a homomorphic computation of degree $\Delta + 1$.

To reduce the semantic security of the ZTHE to the above assumption on the graded encoding, we need the graded encodings to support a re-randomization operation. Intuitively, re-randomizing an encoding results in a new encoding of the same value that is otherwise independent of the original encodings. As in many other applications of graded encoding (for example [GLSW15]), the re-randomization operation is only needed in the reduction and not in the construction. We note that it is possible to avoid the use of randomization, but this requires making a more complicated and less natural (though still efficiently falsifiable) hardness assumption.

Correctness for adversarially generated encodings. In order to construct a ZTHE scheme with correctness for adversarially generated ciphertexts we need to require that the graded encoding themselves have correctness for adversarially generated ciphertexts. This is a non-standard requirement for graded encoding schemes, and it is not required in other applications such as obfuscation (where all encodings are generated by an honest party).

The correctness requirement for adversarially generated encodings is somewhat stronger than in the context of encryption. We require that it is hard to find a pair of valid encodings such that a homomorphic operation on them results in an invalid encoding. In order to support "noisy" candidates, where such an evaluation error always occurs after a large enough number of homomorphic evaluations, we also consider a relaxed requirement. Intuitively, it should be possible to publicly test that the level of noise in an adversarially generated encoding is low. If we determine that an encoding has low noise, it should support a large number of homomorphic operation without an error.

Theorem 1.3 (Informal). *Assuming a graded encoding scheme satisfying Assumption 1.2, there exists a $O(1)$-key zero-testable somewhat homomorphic encryption scheme. Moreover, if the graded encoding scheme is correct for adversarially generated encodings, then the encryption scheme is correct for adversarially generated ciphertexts.*

Instantiations: discussion. The existing constructions of graded encodings [GGH13a, CLT15, GGH15] that support re-randomization do not satisfy our hardness assumption [GGH13a, CHL+15, HJ16]. We don't know if in existing constructions of graded encodings it is possible to publicly test for low noise level. One potential strategy to implement such a test would be to combine the re-randomization and zero test operations. We note that so-called "clean" graded encoding schemes (see for example [Zim15, LV16]), where every element has a unique encoding, trivially satisfy correctness for adversarially generated encodings, and support re-randomization.

1.3 Non-interactive Arguments from Zero-Testable Homomorphic Encryption

Our construction is based on ideas developed in the context of designated-verifier arguments.

Designated-verifier arguments. Aiello *et al.* [ABOR00] suggested the following approach to constructing designated verifier arguments: The prover computes a probabilistically checkable proof (PCP) for the statement. The verifier's message contains PCP queries, encrypted using an FHE scheme, where each query is encrypted under a different key. The prover computes the PCP answers homomorphically, and the verifier decrypts and verifies the answers. The hope was that since a cheating prover couldn't tailor its answer to one query depending on other queries's values, the argument would inherit the PCP's soundness. Dwork *et al.* [DLN+04, DNR16] showed obstacles to proving this construction's soundness. Nonetheless, Kalai, Raz and Rothblum [KRR14] proved that when the underlying PCP satisfies a strong notion of soundness called *no-signaling* soundness, the suggested arguments are in fact sound.

Leaking information on queries: a failed attempt. A naive attempt to turn the above designated-verifier protocol into a publicly verifiable non-interactive argument would be to place the verifier's encrypted queries in the CRS, and provide some leakage on encrypted queries that allows verifying the evaluated answers, but (somehow) does not compromise the soundness of the protocol. We argue, however, that any such leakage must (inherently) compromise soundness. A cheating prover can begin with an accepting PCP proof, changing it into a rejecting proof one symbol at a time. By observing which of the intermediate proofs makes the verifier reject, the prover can recover the encrypted queries and break soundness.

Our approach: intuition. Our protocol follows the blueprint described above: the CRS contains encrypted queries, and the prover homomorphically evaluates the PCP and sends the evaluated queries as the proof. However, to make the proof publicly verifiable we *do not leak any information about the encrypted queries or their answers*. The main idea is to encrypt the queries with a ZTHE. By executing a sequence of homomorphic evaluations and zero tests on the evaluated ciphertexts in the proof, the verifier *learns information about the PCP proof computed by the prover*, which is sufficient to verify its validity.

Next we elaborate on this idea. We start by giving some background on the PCP system we use.

The BFLS PCP. The PCP of Babai et al. [BFLS91] proves that a given computation accepts its input. The tableau of the computation is translated into a multi-variate low-degree polynomial P_0 and the PCP proof contains all the evaluations of P_0 over some finite field. Testing the validity of the tableau is reduced to testing that P_0 is indeed a low-degree polynomial and that it vanishes on all *boolean* inputs. The proof that P_0 vanishes on all boolean inputs is based on the well-known sum-check protocol. The sum-check proof contains auxiliary polynomials $P_1, \ldots, P_m$ and the verifier tests that these polynomials satisfy some local

low-degree relations of the form $R(P_i, P_{i+1}) \equiv 0$. These tests are carried out by probing the polynomials on a small number of random inputs and testing that the relations are satisfied.

A sketch of our protocol. As described above, the CRS contains encryptions $c_1, \ldots, c_m$ that specify queries to the PCP. Each triplet c_j, c_k, c_ℓ specifies an evaluation point for the polynomials $P_1, \ldots, P_m$. For every such triplet, and for every polynomial P_i, the proof contains the homomorphically evaluated answer $d_i = P_i(c_j, c_k, c_\ell)$. To verify the relation $R(P_i, P_{i+1}) \equiv 0$, the verifier homomorphically evaluates $R(d_i, d_{i+1})$ and tests that the evaluation results in a trivial zero. Since the different queries are encrypted under different keys, we use a *multi-key* homomorphic encryption scheme. While the CRS contains encryptions under m different keys, the verifier only computes homomorphically on three keys at a time, therefore we only need 3-key homomorphism.

The proof strategy. Intuitively, if the prover is cheating and $R(P_i, P_{i+1}) \not\equiv 0$ it follows from sematic security that the verifier's zero test fails. Alas, this intuition is fundamentally flawed. A cheating prover may not derive its answers by homomorphically evaluating the low degree polynomials $P_1, \ldots, P_m$, or any other polynomial for that matter. Our actual proof strategy is inspired by that of Kalai, Raz and Rothblum [KRR14] and consists of the following steps.

1. Since the encryption is semantically secure, the prover's answers are *no-signaling*, meaning that the decrypted answer to one query gives no information on the other queries values.
2. In the BFLS PCP, it is possible to reconstruct any small subset of entries L of the computation's tableau based on PCP values in some small set of locations $q(L)$. We show that our proof satisfies the following local soundness guarantee: if the verifier's encrypted queries include the locations $q(L)$ and if the verifier accepts the prover's encrypted answers then the reconstructed subset of the tableau is *locally consistent*. That is, it obeys the computation's local constrains. To show that this is the case even when the prover sends malformed answers we use the fact that the encryption scheme is correct for *adversarially generated ciphertext*.
3. By the semantic security of encrypted queries, and by the fact that the protocol is publicly verifiable, we deduce that if the verifier accepts the answers to *any* queries encrypted in the CRS (say the all-0 queries), it would also accept the answers to the to queries $q(L)$, for *every* subset L.
4. It follows that we can turn any convincing prover in our protocol into an algorithm that samples local assignments for any subset L of the computation's tableau that are guaranteed to be *both no-signaling and locally consistent*.
5. Based on the *augmented circuit technique* of [KRR14], we show how to use such a *local-assignment generator* to reconstruct a complete and valid tableau.

We note that our soundness proof is significantly simpler than that of [KRR14]. In particular we only use a striped down version of the BFLS PCP without any low-degree tests, and we do not argue that this PCP has no-signaling

soundness. Intuitively, what enables this simplification is that in the publicly-verifiable setting we can move from local consistency for one subset to local consistency on all subsets using semantic security (see Step 3 above) and without using global properties of the PCP.

Proving *adaptive* soundness presents additional challenges. To argue adaptive soundness, we use ideas inspired by the recent work of Brakerski et al. [BHK16], who constructed an adaptively sound arguments in the designated-verifier setting. Roughly, they show how to reconstruct a tableau from any local-assignment generator that can chose the statement adaptively as a function of the subset L.

On the notion of local-assignment generator. The augmented circuit technique as well as the technique of reconstructing the computation's tableau by reading subsets that are no-signaling and locally-consistent originates from the analysis of [KRR14]. The notion of local-assignment generator and the generic transformation from a local-assignment generator to global soundness first appeared in an earlier version of this work [PR14]. Since then the local-assignment generator abstraction played a key role in achieving stronger designated-verifier arguments for RAM computations [KP16] and Batch-NP computations [BHK16], as well as in achieving adaptive soundness [BHK16]. In the current version of this work we use the adaptive local-assignment generator of [BHK16].

1.4 Zero-Testable Homomorphic Encryption from Graded Encodings

We start by describing the interface of a graded encoding scheme in more details. The scheme has public parameters that define a ring R and a maximal degree Δ. The scheme encodes elements in R and supports homomorphic computations up to degree Δ. Every encoding has a level. Freshly generated encodings are of level 1 and level-δ encodings are the result of a degree-δ homomorphic computation. We also refer to the elements of R as level-0 encoding. Following the standard formulation of graded encodings, we do not assume that the ring R is public. Instead, there is a public interface for sampling random level-0 encodings and evaluating the ring operations. We also assume that the public parameters include encodings of the constants 0 and 1 in every level.

The graded encoding supports a (strong) zero test that can publicly identify encodings of zero in any level. It also supports a re-randomization operation that, given an encoding, generates a new random encoding of the same element. For example, re-randomizing an encoding can be used to hide the homomorphic computation that generated it.

The ZTHE scheme. We construct multi-key ZTHE from graded encoding as follows. The scheme's public parameters are the parameters of a graded encoding scheme with degree bound Δ. The secret key is a random ring element $t \in R$ and the corresponding public key is a level-1 encoding of t.

An encryption c of a message $m \in \{0,1\}$ is given by a random degree-Δ univariate polynomial P such that $P(t) = m$. The ciphertext c consists of level-1

encodings of the $\Delta+1$ coefficients $\alpha_0, \ldots, \alpha_\Delta$ of P. The semantic security of this encryption follows from Assumption 1.2 that states that even given the public key encoding of t, the encodings in c are indistinguishable from encodings of random elements, independent of m.

Encryption. We need to sample such an encryption using only the public parameters and the public key encoding of t. A naive approach would be to sample all the coefficients of P except for the free coefficient α_0 randomly and then homomorphically compute an encoding of α_0. However, this would result in an encoding in level Δ instead of level 1. Instead we can sample all the coefficients of P as linear functions of t. We sample random ring elements $r_1, \ldots, r_\Delta$ and homomorphically compute encodings of the coefficients

$$\alpha_0 = m - r_1 \cdot t, \quad \ldots \quad , \alpha_i = r_i - r_{i+1} \cdot t, \quad \ldots \quad , \alpha_\Delta = r_\Delta.$$

Note that $\alpha_0, \ldots, \alpha_\Delta$ are indeed random subject to $\sum \alpha_i \cdot t^i = m$. Finally, we re-randomize the encoded coefficient to hide the process in which they where sampled (which depends on m).

We note that the re-randomization operation is only used during encryption. In our non-interactive argument the ZTHE encryption procedure is only used to generate the CRS and in the security proof. As noted above, we could avoid the use of re-randomization at the cost of making a more complicated assumption on the graded encoding that implies the CPA security of our encryption scheme in the secret key setting.

Same-Key homomorphic evaluation. Let c_1 and c_2 be ciphertexts encrypting messages m_1 and m_2 respectively under the same secret key t. Let P_1 and P_2 be the polynomials encoded by c_1 and c_2, where

$$P_1(t) = m_1, \quad P_2(t) = m_2.$$

To evaluate a homomorphic operation $\star \in \{+, -, \times\}$ we homomorphically compute the encoded coefficients of the polynomial $P_1 \star P_2$. Correctness follows since

$$(P_1 \star P_2)(t) = P_1(t) \star P_1(t) = m_1 \star m_2.$$

For addition and subtraction, the homomorphic computation of the new coefficients is a linear operation (over the input coefficients), and the degree of the resulting polynomial is the maximal degree of the two input polynomials. For multiplication, we homomorphically compute a convolution of the input coefficients, and the degree of the resulting polynomial is the sum of the degrees of the input polynomials. Thus, the evaluation of a degree-δ homomorphic computation yields coefficients that are encoded in level-δ of the graded encoding scheme, and the resulting (univariate) polynomial has degree $(\delta \cdot \Delta)$. It follows that the encryption supports degree-Δ homomorphic computations, before the level of encoded coefficient exceeds the degree bound.

Multi-key homomorphic evaluation. To compute a homomorphic operation $\star$ over ciphertexts c_1, c_2 encrypted under different secret keys t_1, t_2, we

homomorphically compute the coefficients of the *bivariate* polynomial $P(x, y) \equiv P_1(x) \star P_2(y)$, where P_1 and P_2 are the polynomials encoded by c_1 and c_2 respectively. In general, a homomorphic computation involving ciphertexts under d different keys will result in a ciphertext encoding a d-variate polynomial. Since the number of coefficients grows exponentially with d, we only support homomorphic computation involving a constant number of keys.

Decryption. To decrypt a ciphertext c, we homomorphically evaluate the polynomial P it encodes on the secret key t. Since the secret key is a level-0 encoding, this homomorphic evaluation does not exceed the degree bound Δ. We then use the graded encoding zero test to compare the evaluated encoding to an encoding of 0 or of 1. If none of the tests succeed decryption fails.

Note that in homomorphic evaluation, the algebraic operation on the plaintexts are evaluated over the ring R. However, since our decryption only obtains *an encoding of* the plaintext, we can only decrypt messages in $\{0, 1\}$ (or more generally, messages taken from a small plaintext space). This is analogous to the behaviour of the additively-homomorphic ElGamal encryption and other schemes [BGN05]. Such decryption is sufficient for our application, where we evaluate arithmetic circuits (over $\mathbb{Z}$) whose outputs are expected to be boolean.

Zero Test. A ciphertext c that results from a homomorphic evaluation of a polynomial that is identically zero always encodes a polynomial $P \equiv 0$. We can test this by using the zero test procedure of the graded encoding, testing that all the encoded coefficient of P are zero. It is also the case that a ciphertext that passes the zero test must encode a polynomial $P \equiv 0$ and therefore it must decrypt to zero.

Correctness for adversarially generated ciphertexts. If the graded encoding scheme is correct even on adversarially generated encodings, we inherit this strong correctness guarantee also for the ciphertext. Note, however, that even a ciphertext that consists of valid encodings may encode a polynomial P such that $P(t) \notin \{0, 1\}$, and therefore fail to decrypt. To deal with this case, we consider an alternative decryption algorithm that is inefficient and can decrypt any value in R. The correctness requirement for adversarially generated ciphertexts is therefore defined with respect to this inefficient decryption procedure. The weaker correctness requirement suffices for proving the computational soundness of the non-interactive argument, even though it considers an inefficient decryption algorithm: once the correctness requirement is guaranteed, the remainder of the soundness proof is information theoretic.

1.5 Organization

The definition of non-interactive arguments and other preliminaries are given in Sect. 2. In Sect. 3 we define the notion of ZTHE and the additional properties we use. Section 4 describes the construction of non-interactive argument from ZTHE. The analysis of the non-interactive argument and the construction of ZTHE from graded encodings appear in the full version of this work.

2 Preliminaries

For a sequence $\mathbf{x} = (x_1, \ldots, x_n)$, we denote by $\mathbf{x}_{-i}$ the sequence with the i-th elements removed

$$\mathbf{x}_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, x_n)\,.$$

For a pair of sequences $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{y} = (y_1, \ldots, y_{n'})$ we denote by $\mathbf{x} \,|\, \mathbf{y}$ the concatenated sequence

$$\mathbf{x} \,|\, \mathbf{y} = (x_1, \ldots, x_n, y_1, \ldots, y_{n'})\,.$$

2.1 Arithmetic Circuits

We consider arithmetic circuits with binary addition, subtraction and multiplication gates. We only allow use of the constants $\{0, 1\}$.

Degree. For an arithmetic circuit C, the degree (resp. total degree) of C is the individual (resp. total) degree of the formal polynomial computed by C. A degree-1 circuit is said to be multi-linear.

Equivalence. An arithmetic circuit C is said to be identically zero (denoted by $C \equiv 0$) if the formal polynomial computed by C is identically zero over $\mathbb{Z}$. Two arithmetic circuits C_1, C_2 are said to be equivalent (denoted by $C_1 \equiv C_2$) if $C_1 - C_2 \equiv 0$.

Computing boolean functions. An arithmetic circuit C is said to compute a boolean function f if C agrees with f when evaluated over $\mathbb{Z}$. That is, if f takes n inputs, then for every $x \in \{0,1\}^n$ we have that $f(x) = C(x)$ when C is evaluated over $\mathbb{Z}$.

Fact 2.1. *Let C_1 and C_2 be arithmetic circuits with n inputs wires computing boolean functions f_1 and f_2 respectively.*

1. *The circuit $1 - C_1$ computes the boolean function $1 - f_1$.*
2. *The circuit $C_1 \cdot C_2$ computes the boolean function $f_1 \cdot f_2$.*
3. *If for every $x \in \{0,1\}^n$, at most one of the values $C_1(x)$ and $C_2(x)$ is non-zero, then the circuit $C_1 + C_2$ computes the boolean function $f_1 + f_2$.*

Circuit restrictions. Let C be an arithmetic circuit with n inputs wires and individual degree δ. For $i \in [n]$ let $C|_{i,0}, \ldots, C|_{i,\delta}$ be the arithmetic circuits with $n-1$ inputs wires and individual degree δ such that

$$C(x_1, \ldots, x_n) \equiv \sum_{j \in [0,\delta]} C|_{i,j}(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) \cdot x_i^j. \tag{1}$$

For $j > \delta$ let $C|_{i,j}$ denote the identically 0 circuit.

Fact 2.2. *There is an procedure that given an arithmetic circuit C with n inputs wires and individual degree δ and given an index $i \in [n]$ computes $C|_{i,0}, \ldots, C|_{i,\delta}$ in time* $\mathrm{poly}(|C|, \delta)$.

2.2 Multi-linear Extension

A multi-linear extension of a boolean function f is a multi-linear arithmetic circuit C computing f. Next we describe a multi-linear extension circuit of an arbitrary boolean function f.

Let β_n be the multi-linear arithmetic circuit with $2n$ inputs computing the boolean identity function. That is, for every $\mathbf{x}, \mathbf{y} \in \{0,1\}^n$, $\beta_n(\mathbf{x}, \mathbf{y}) = 1$ if and only if $\mathbf{x} = \mathbf{y}$. The arithmetic circuit β_n is given by the expression

$$\beta_n(x_1, \ldots, x_n, y_1, \ldots, y_n) = \prod_{i \in [n]} x_i y_i + (1 - x_i)(1 - y_i). \tag{2}$$

We sometimes omit the subscript n when it is clear from the context.

The multi-linear extension of a boolean function f with n inputs is given by the arithmetic circuit

$$C(\mathbf{x}) = \sum_{\mathbf{y} \in \{0,1\}^n} \beta_n(\mathbf{x}, \mathbf{y}) \cdot f(\mathbf{y}). \tag{3}$$

Since for every $\mathbf{x} \in \{0,1\}^n$ there exist only one value of $\mathbf{y} \in \{0,1\}^n$ such that $\beta_n(\mathbf{x}, \mathbf{y}) \neq 0$, it follows by Fact 2.1 that C computes the boolean function f.

2.3 Publicly-Verifiable Non-interactive Arguments

In this section we define publicly verifiable non-interactive arguments.

Let $\mathcal{U}$ be the universal language such that $(x, \mathsf{T}) \in \mathcal{U}$ for $x = (M, y)$ if and only if the Turing machine M accepts the input y within at most T steps.

Syntax. A publicly verifiable non-interactive argument scheme for the universal language $\mathcal{U}$ consists of PPT algorithms $(\mathsf{Del.Gen}, \mathsf{Del.P}, \mathsf{Del.V})$ with the following syntax.

Del.Gen: Given the security parameter 1^λ, outputs a common reference string CRS.

Del.P: Given the common reference string, a time bound 1^T in unary representation and an instance $x \in \{0,1\}^*$, outputs a proof Π.

Del.V: Given the common reference string, a time bound T in binary representation, an instance $x \in \{0,1\}^*$ and a proof Π, outputs a bit.

Definition 2.1. *A publicly verifiable non-interactive argument scheme* $(\mathsf{Del.Gen}, \mathsf{Del.P}, \mathsf{Del.V})$ *for the universal language* $\mathcal{U}$ *satisfies the following requirements*

Completeness: *For every* $\lambda \in \mathbb{N}$ *and every* $(x, \mathsf{T}) \in \mathcal{U}$

$$Pr\left[\mathsf{Del.V}(\mathsf{CRS}, \mathsf{T}, x, \Pi) = 1 \;\middle|\; \begin{array}{l} \mathsf{CRS} \leftarrow \mathsf{Del.Gen}(1^\lambda) \\ \Pi \leftarrow \mathsf{Del.P}(\mathsf{CRS}, 1^\mathsf{T}, x) \end{array}\right] = 1.$$

Efficiency: *In the above (honest) experiment the size of the proof* Π *is* $\mathrm{poly}(\lambda, \log \mathsf{T})$. *The running time of* $\mathsf{Del.V}$ *is* $|x| \cdot \mathrm{poly}(|\mathsf{CRS}|, |\Pi|, \log \mathsf{T})$.

Adaptive Soundness: *For every polynomial* T *and for every poly-size cheating prover* P^* *there exists a negligible function* μ *such that for every* $\lambda \in \mathbb{N}$

$$Pr\left[\begin{array}{l|l} (x^*, \mathsf{T}(\lambda)) \notin \mathcal{U} & \mathsf{CRS} \leftarrow \mathsf{Del.Gen}(1^\lambda) \\ \mathsf{Del.V}(\mathsf{CRS}, \mathsf{T}, x^*, \Pi^*) = 1 & (x^*, \Pi^*) \leftarrow \mathsf{P}^*(\mathsf{CRS}) \end{array}\right] \leq \mu(\lambda),$$

3 Zero-Testable Homomorphic Encryption

In this section we define the notion of zero-testable homomorphic encryption. We also define a multi-key variant [LTV12].

3.1 Homomorphic Encryption

We start by recalling the notion of homomorphic encryption.

Syntax. A homomorphic encryption scheme consists of PPT algorithms

$$(\mathsf{HE.KeyGen}, \mathsf{HE.Enc}, \mathsf{HE.Dec}, \mathsf{HE.Eval})$$

with the following syntax.

HE.KeyGen: Given the security parameter 1^λ, outputs a secret key sk, a public key pk and a description of a ring R.

HE.Enc: Given the public key pk and a message $m \in \{0,1\}$, outputs a ciphertext c.

HE.Dec: Given the secret key sk and a ciphertext c, outputs a ring element $\alpha \in R$ or a special symbol $\perp$.

HE.Eval: Given e public key pk, an operation $\star \in \{+, -, \times\}$, and a pair of ciphertexts c_1, c_2, outputs a ciphertext c or a special symbol $\perp$.

Evaluating circuits. Some formulations of homomorphic encryption only consider an evaluation algorithm for circuits and not individual gates. By explicitly requiring that the evaluation is performed gate by gate, we ensure correctness for a "multi-hop" evaluation [GHV10] where ciphertexts that result from a homomorphic computation support further homomorphic operations.

Homomorphic evaluation of an arithmetic circuit C is implemented by iteratively applying the basic evaluation algorithm $\mathsf{HE.Eval}$ for every gate in C. This process is described formally below.

We only consider arithmetic circuits containing constants from $\{0,1\}$, which can be evaluated over any ring. When evaluating a gate that takes a constant $b \in \{0,1\}$ we do not generate a fresh random encryption of b. Instead, we assume that the public key includes ciphertexts $\hat{0}$ and $\hat{1}$ of 0 and 1 respectively. This evaluation strategy guarantees that all occurrences of a constant in C are replaced with the same ciphertext. This will be crucial later when we introduce the notion of zero-testable homomorphic encryption.

For an arithmetic circuit C, and ciphertexts $(c_1, \ldots, c_n)$ encrypted under public key pk we denote by $\langle C\,(c_1, \ldots, c_n)\rangle$ the evaluated ciphertext c computed as follows.

- If C is the constant 0 then $c = \hat{0}$.
- If C is the constant 1 then $c = \hat{1}$.
- If C is the i-th input wire then $c = c_i$.
- If C is of the form $C = C_1 \star C_2$ then

$$c = \mathsf{HE.Eval}\left(\mathsf{pk}, \star, \left(\langle C_1\,(c_1, \ldots, c_n)\rangle, \langle C_2\,(c_1, \ldots, c_n)\rangle\right)\right).$$

Definition 3.1 (Homomorphic Encryption). *Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda\in\mathbb{N}}$ be an ensemble of circuits. A homomorphic encryption scheme* $(\mathsf{HE.KeyGen}, \mathsf{HE.Enc}, \mathsf{HE.Dec}, \mathsf{HE.Eval})$ *for $\mathcal{C}$ satisfies the following requirements.*

Correctness: *For every $\lambda \in \mathbb{N}$, every $C \in \mathcal{C}_\lambda$ with n inputs wires, and every $m_1, \ldots, m_n \in \{0,1\}$*

$$Pr\left[C(m_1, \ldots, m_n) = \alpha \;\middle|\; \begin{array}{l} (\mathsf{sk}, \mathsf{pk}, R) \leftarrow \mathsf{HE.KeyGen}(1^\lambda) \\ \forall i \in [n] : c_i \leftarrow \mathsf{HE.Enc}(\mathsf{pk}, m_i) \\ c \leftarrow \langle C(c_1, \ldots, c_n)\rangle \\ \alpha \leftarrow \mathsf{HE.Dec}(\mathsf{sk}, c) \end{array}\right] = 1,$$

where C is evaluated over R.

Compactness: *There exists a polynomial L such that in the above honest experiment $|c| \le L(\lambda)$ (independently of $|C|$).*

Semantic Security: *For every poly-size adversary* Adv *there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$*

$$Pr\left[m = m' \;\middle|\; \begin{array}{l} m \leftarrow \{0,1\} \\ (\mathsf{sk}, \mathsf{pk}, R) \leftarrow \mathsf{HE.KeyGen}(1^\lambda) \\ c \leftarrow \mathsf{HE.Enc}(\mathsf{pk}, m) \\ m' \leftarrow \mathsf{Adv}(\mathsf{pk}, c) \end{array}\right] \le \frac{1}{2} + \mu(\lambda).$$

Definition 3.2 (Somewhat Homomorphic Encryption). *For $B, \Delta \in \mathbb{N}$ let $\mathcal{C}_{B,\Delta}$ be the set of arithmetic circuits of size at most B and total degree at most Δ. Let $B = B(\lambda), \Delta = \Delta(\lambda)$ be polynomially bounded functions. A homomorphic encryption scheme is (B, Δ)-somewhat homomorphic if it satisfies Definition 3.1 for the circuit ensemble $\left\{\mathcal{C}_{B(\lambda),\Delta(\lambda)}\right\}_{\lambda\in\mathbb{N}}$. A scheme is Δ-somewhat homomorphic if it is (B, Δ)-somewhat homomorphic for every polynomial B.*

3.2 Correctness for Adversarial Ciphertexts

We formulate an additional correctness requirement that considers evaluation of adversatively generated ciphertexts. Informally, we require that an efficient adversary cannot generate a pair of ciphertexts that cause en evaluation error. A homomorphic evaluation $\langle c_1 \star c_2\rangle$ is erroneous if the following two experiments have different outputs

1. Homomorphically evaluate $\langle c_1 \star c_2 \rangle$ and output the decryption of the evaluated ciphertext.
2. Decrypt c_1, c_s. If one of the ciphertexts fails to decrypt (decryption output $\perp$), then output $\perp$. Otherwise output the evaluation of $\star$ on the decrypted elements.

Many existing homomorphic encryption candidates only support a polynomially bounded number of homomorphic operations before the noise in the ciphertexts becomes too large and causes an evaluation error. Therefore, in such candidates, ciphertexts that cause en evaluation error are easy to generate. To support candidate of this nature we allow the output of the first experiment above to be $\perp$ even if the output of the second experiment is different than $\perp$.

Correctness for Adversarial Ciphertexts: For every poly-size adversary Adv there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and for every operation $\star \in \{+, -, \times\}$

$$\Pr\left[\alpha \notin \{\alpha_1 \star \alpha_2, \perp\} \;\middle|\; \begin{array}{l} (\mathsf{sk}, \mathsf{pk}, R) \leftarrow \mathsf{HE.KeyGen}(1^\lambda) \\ c_1, c_2 \leftarrow \mathsf{Adv}(\mathsf{pk}) \\ c \leftarrow \mathsf{HE.Eval}(\mathsf{pk}, \star, (c_1, c_2)) \\ \forall i \in \{1,2\} : \alpha_i \leftarrow \mathsf{HE.Dec}(\mathsf{sk}, c_i) \\ \alpha \leftarrow \mathsf{HE.Dec}(\mathsf{sk}, c) \end{array}\right] \leq \mu(\lambda),$$

where in the probability above, if $\alpha_1, \alpha_2 \in R$, the expression $\alpha_1 \star \alpha_2$ is evaluated over R. If either $\alpha_1 = \perp$ or $\alpha_1 = \perp$ then $\alpha_1 \star \alpha_2 = \perp$.

3.3 Zero Test

A zero test for a homomorphic encryption scheme is a PPT algorithm $\mathsf{HE.ZT}$ that can identify trivial encryptions of 0. These are ciphertexts that result from homomorphically evaluating an arithmetic circuit that is identically zero. We additionally require that the zero test never incorrectly identifies encryptions of non-zero values. This holds even for adversatively generated ciphertexts.

Given the public key pk and a ciphertext c, the zero test $\mathsf{HE.ZT}$ outputs a bit. The zero test satisfies the following requirements.

Zero-Test Completeness: For every $\lambda \in \mathbb{N}$, every $C \in \mathcal{C}_\lambda$ with n inputs wires such that C is identically zero, and every $m_1, \ldots, m_n \in \{0,1\}$

$$\Pr\left[\mathsf{HE.ZT}(\mathsf{pk}, c) = 1 \;\middle|\; \begin{array}{l} (\mathsf{sk}, \mathsf{pk}, R) \leftarrow \mathsf{HE.KeyGen}(1^\lambda) \\ \forall i \in [n] : c_i \leftarrow \mathsf{HE.Enc}(\mathsf{pk}, m_i) \\ c \leftarrow \langle C\,(c_1, \ldots, c_n) \rangle \end{array}\right] = 1.$$

Zero-Test Soundness: For every poly-size adversary Adv there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$

$$\Pr\left[\begin{array}{l} \mathsf{HE.ZT}(\mathsf{pk}, c) = 1 \\ \alpha \neq 0 \end{array} \;\middle|\; \begin{array}{l} (\mathsf{sk}, \mathsf{pk}, R) \leftarrow \mathsf{HE.KeyGen}(1^\lambda) \\ c \leftarrow \mathsf{Adv}(\mathsf{pk}) \\ \alpha \leftarrow \mathsf{HE.Dec}(\mathsf{sk}, c) \end{array}\right] \leq \mu(\lambda).$$

3.4 Weak Decryption

We define a relaxation of homomorphic encryption where

- The decryption procedure $\mathsf{HE.Dec}$ is not required to be PPT.
- Instead we require that there exists a weak decryption procedure $\mathsf{HE.WeakDec}$ which is PPT but does not decrypt messages outside $\{0,1\}$.
- The weak decryption result should be consistent with the inefficient decryption result even for adversarially generated ciphertexts.

The encryption scheme we construct from graded encodings will only satisfy this relaxation which is sufficient for our application.

Given the secret key sk and a ciphertext c, the weak decryption procedure $\mathsf{HE.WeakDec}$ outputs a message $m \in \{0,1\}$ or a special symbol $\perp$. The weak decryption procedure satisfies the following requirement.

Weak Decryption: For every poly-size adversary Adv there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$

$$\Pr\left[m \neq \alpha' \;\middle|\; \begin{array}{l} (\mathsf{sk}, \mathsf{pk}, R) \leftarrow \mathsf{HE.KeyGen}(1^\lambda) \\ c \leftarrow \mathsf{Adv}(\mathsf{pk}) \\ \alpha \leftarrow \mathsf{HE.Dec}(\mathsf{sk}, c) \\ m \leftarrow \mathsf{HE.WeakDec}(\mathsf{sk}, c) \end{array} \right] \leq \mu(\lambda),$$

where in the above probability, $\alpha' = \alpha$ if $\alpha \in \{0,1\}$ and $\alpha' = \perp$ otherwise.

3.5 Multi-key Zero-Testable Homomorphic Encryption

In this section we define a multi-key variant of homomorphic encryption that also satisfies the other requirements defined above. In multi-key homomorphic encryption, introduced by López-Alt et al. [LTV12] homomorphic computation can be executed over ciphertexts encrypted under d different keys. To ensure semantic security, decrypting the result requires all secret keys. Importantly, a system can generate ciphertext under an unbounded number of keys and any d of them can be combined in a homomorphic computation. We assume that the number of different keys d is constant. We also allow for common public parameters used to generate all keys.

Syntax. A d-key zero-testable homomorphic encryption scheme consists of PPT algorithms

$$(\mathsf{MHE.ParamGen}, \mathsf{MHE.KeyGen}, \mathsf{MHE.Enc}, \mathsf{MHE.WeakDec}, \mathsf{MHE.Eval}, \mathsf{MHE.ZT})$$

and an unbounded algorithm $\mathsf{MHE.Dec}$ with the following syntax.

$\mathsf{MHE.ParamGen}$**:** Given the security parameter 1^λ, outputs public parameters pp and a description of a ring R.

$\mathsf{MHE.KeyGen}$**:** Given the public parameters pp, outputs a secret key sk and a public key pk.

$\mathsf{MHE.Enc}$**:** Given public parameters pp, a public key pk and a message $m \in \{0,1\}$, outputs a ciphertext c.

$\mathsf{MHE.Dec}$**:** Given public parameters pp, d secret keys $\mathsf{sk}_1, \ldots, \mathsf{sk}_d$ and a ciphertext c, outputs a ring element $\alpha \in R$ or a special symbol $\perp$.

$\mathsf{MHE.WeakDec}$**:** Given public parameters pp, d secret keys $\mathsf{sk}_1, \ldots, \mathsf{sk}_d$ and a ciphertext c, outputs a message $m \in \{0,1\}$ or a special symbol $\perp$.

$\mathsf{MHE.Eval}$**:** Given public parameters pp, a pair of public keys $\mathsf{pk}_1, \mathsf{pk}_2$, an operation $\star \in \{+, -, \times\}$ and a pair c_1, c_2, outputs a ciphertext c or a special symbol $\perp$.

$\mathsf{MHE.ZT}$**:** Given public parameters pp, d public keys $\mathsf{pk}_1, \ldots, \mathsf{pk}_d$ and a ciphertext c, outputs a bit.

Remark 3.1 (Superfluous keys). The decryption and zero test algorithms take d keys, even if the input ciphertext results from a computation involving less keys. We assume without loss of generality that adding superfluous keys does not affect the procedures functionality.

Definition 3.3 (Multi-key Zero-Testable Homomorphic Encryption). *Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of circuits. A d-key zero-testable homomorphic encryption scheme*

$$(\mathsf{MHE.ParamGen}, \mathsf{MHE.KeyGen}, \mathsf{MHE.Enc}, \mathsf{MHE.Dec}, \mathsf{MHE.WeakDec},$$
$$\mathsf{MHE.Eval}, \mathsf{MHE.ZT})$$

for $\mathcal{C}$ satisfies the following requirements.

Correctness: *There exists a negligible function μ such that for every $\lambda \in \mathbb{N}$, every $C \in \mathcal{C}_\lambda$ with n inputs wires, every $m_1, \ldots, m_n \in \{0,1\}$ and every indices $j_1, \ldots, j_n \in [d]$*

$$Pr\left[C(m_1, \ldots, m_n) = \alpha \;\middle|\; \begin{array}{l} (\mathsf{pp}, R) \leftarrow \mathsf{MHE.ParamGen}(1^\lambda) \\ \forall j \in [d] : (\mathsf{pk}_j, \mathsf{sk}_j) \leftarrow \mathsf{MHE.KeyGen}(\mathsf{pp}) \\ \forall i \in [n] : c_i \leftarrow \mathsf{MHE.Enc}(\mathsf{pp}, \mathsf{pk}_{j_i}, m_i) \\ c \leftarrow \langle C(c_1, \ldots, c_n) \rangle \\ \alpha \leftarrow \mathsf{MHE.Dec}(\mathsf{pp}, (\mathsf{sk}_1, \ldots, \mathsf{sk}_d), c) \end{array} \right] \geq 1 - \mu(\lambda),$$

where C is evaluated over R.

Compactness: *There exists a polynomial L (that may depend on d) such that in the above honest experiment $|c| \leq L(\lambda)$ (independently of $|C|$).*

Correctness for Adversarial Ciphertexts: *For every poly-size adversary* Adv *there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and for every operation $\star \in \{+, -, \times\}$*

$$Pr\left[\alpha \notin \{\alpha_1 \star \alpha_2, \perp\} \;\middle|\; \begin{array}{l} (\mathsf{pp}, R) \leftarrow \mathsf{MHE.ParamGen}(1^\lambda) \\ \forall j \in [d] : (\mathsf{pk}_j, \mathsf{sk}_j) \leftarrow \mathsf{MHE.KeyGen}(\mathsf{pp}) \\ c_1, c_2 \leftarrow \mathsf{Adv}(\mathsf{pp}, \mathsf{pk}_1, \ldots, \mathsf{pk}_d) \\ c \leftarrow \mathsf{MHE.Eval}(\mathsf{pp}, (\mathsf{pk}_1, \ldots, \mathsf{pk}_d), \star, (c_1, c_2)) \\ \forall i \in \{1,2\} : \alpha_i \leftarrow \mathsf{MHE.Dec}(\mathsf{pp}, (\mathsf{sk}_1, \ldots, \mathsf{sk}_d), c_i) \\ \alpha \leftarrow \mathsf{MHE.Dec}(\mathsf{pp}, (\mathsf{sk}_1, \ldots, \mathsf{sk}_d), c) \end{array} \right] \leq \mu(\lambda),$$

where in the probability above, if $\alpha_1, \alpha_2 \in R$, the expression $\alpha_1 \star \alpha_2$ is evaluated over R. If either $\alpha_1 = \perp$ or $\alpha_1 = \perp$ then $\alpha_1 \star \alpha_2 = \perp$.

Zero Test Completeness: *There exists a negligible function μ such that for every $\lambda \in \mathbb{N}$, every $C \in \mathcal{C}_\lambda$ with n inputs wires that is identically zero, every $m_1, \ldots, m_n \in \{0,1\}$, and every indices $j_1, \ldots, j_n \in [d]$*

$$Pr\left[b = 1 \;\middle|\; \begin{array}{l} (\mathsf{pp}, R) \leftarrow \mathsf{MHE.ParamGen}(1^\lambda) \\ \forall j \in [d] : (\mathsf{pk}_j, \mathsf{sk}_j) \leftarrow \mathsf{MHE.KeyGen}(\mathsf{pp}) \\ \forall i \in [n] : c_i \leftarrow \mathsf{MHE.Enc}(\mathsf{pp}, \mathsf{pk}_{j_i}, m_i) \\ c \leftarrow \langle C\,(c_1, \ldots, c_n)\rangle \\ b \leftarrow \mathsf{MHE.ZT}(\mathsf{pp}, (\mathsf{pk}_1, \ldots, \mathsf{pk}_d), c) \end{array}\right] \geq 1 - \mu(\lambda).$$

Zero-Test Soundness: *For every poly-size adversary* Adv *there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$*

$$Pr\left[\begin{array}{l} b = 1 \\ \alpha \neq 0 \end{array} \;\middle|\; \begin{array}{l} (\mathsf{pp}, R) \leftarrow \mathsf{MHE.ParamGen}(1^\lambda) \\ \forall j \in [d] : (\mathsf{pk}_j, \mathsf{sk}_j) \leftarrow \mathsf{MHE.KeyGen}(\mathsf{pp}) \\ c \leftarrow \mathsf{Adv}(\mathsf{pp}, \mathsf{pk}_1, \ldots, \mathsf{pk}_d) \\ \alpha \leftarrow \mathsf{MHE.Dec}(\mathsf{pp}, (\mathsf{sk}_1, \ldots, \mathsf{sk}_d), c) \\ b \leftarrow \mathsf{MHE.ZT}(\mathsf{pp}, (\mathsf{pk}_1, \ldots, \mathsf{pk}_d), c) \end{array}\right] \leq \mu(\lambda).$$

Weak Decryption: *For every poly-size adversary* Adv *there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$*

$$Pr\left[m \neq \alpha' \;\middle|\; \begin{array}{l} (\mathsf{pp}, R) \leftarrow \mathsf{MHE.ParamGen}(1^\lambda) \\ \forall j \in [d] : (\mathsf{pk}_j, \mathsf{sk}_j) \leftarrow \mathsf{MHE.KeyGen}(\mathsf{pp}) \\ c \leftarrow \mathsf{Adv}(\mathsf{pp}, \mathsf{pk}_1, \ldots, \mathsf{pk}_d) \\ \alpha \leftarrow \mathsf{MHE.Dec}(\mathsf{pp}, (\mathsf{sk}_1, \ldots, \mathsf{sk}_d), c) \\ m \leftarrow \mathsf{MHE.WeakDec}(\mathsf{pp}, (\mathsf{sk}_1, \ldots, \mathsf{sk}_d), c) \end{array}\right] \geq 1 - \mu(\lambda),$$

where in the above probability, $\alpha' = \alpha$ if $\alpha \in \{0,1\}$ and $\alpha' = \perp$ otherwise.

Semantic Security: *For every poly-size adversary* Adv *there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$*

$$Pr\left[m = m' \;\middle|\; \begin{array}{l} m \leftarrow \{0,1\} \\ (\mathsf{pp}, R) \leftarrow \mathsf{MHE.ParamGen}(1^\lambda) \\ (\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{MHE.KeyGen}(1^\lambda) \\ c \leftarrow \mathsf{MHE.Enc}(\mathsf{pp}, \mathsf{pk}, m) \\ m' \leftarrow \mathsf{Adv}(\mathsf{pk}, c) \end{array}\right] \leq \frac{1}{2} + \mu(\lambda).$$

4 A Non-interactive Argument

This section describes our publicly-verifiable non-interactive arguments. We start with an overview of the construction.

4.1 Overview

We construct a non-interactive argument system for the universal language $\mathcal{U}$. Given an instance $x = (M, y) \in \{0,1\}^n$ and a time bound T the verifier wants to

ascertain that $(x, \mathsf{T}) \in \mathcal{U}$, that is, that the Turing machine M accepts the input y within T steps. The protocol should be *adaptively sound*: even an adaptive cheating prover, who first sees the CRS and then picks an instance $(x, \mathsf{T}) \notin \mathcal{U}$ adaptively, should not be able to generate am accepting proof.

In the protocol, the prover and verifier translate the instance (x, T) into a 3CNF formula φ over $\text{poly}(n, \mathsf{T})$ variables, which is satisfiable if and only if $(x, \mathsf{T}) \in \mathcal{U}$. φ has a "short" implicit description via an arithmetic circuit Φ of small size and degree that, given the labels of three literals, determines whether their disjunction is a clause in φ. Note that given φ, the formula Φ and the original instance (x, T) can be efficiently reconstructed. More over, if $(x, \mathsf{T}) \in \mathcal{U}$, a satisfying assignment for φ can be efficiently computed. With this formula in mind, the argument system has two main ingredients:

Ingredient 1: the core protocol. The first ingredient is a publicly-verifiable non-interactive "core protocol". The prover in the core protocol is presented with a CRS, a circuit Φ describing a 3CNF φ (as above), and a satisfying assignment σ to φ. It generates a proof Π that will convince the verifier that the 3CNF described by Φ is satisfiable.

The core protocol has a relaxed soundness property: it is *not* guaranteed that an adaptive cheating prover P^* cannot generate a circuit Φ describing an unsatisfiable 3CNF together with a proof Π^* that makes the verifier accept. Rather, the soundness guarantee is that any adaptive cheating prover for the core protocol can be used to derive a *no-signalling adaptive local assignment generator* Assign. The adaptive assignment generator Assign is a randomized algorithm that gets as input a small set S of variables, and outputs a pair (Φ, σ), where $\sigma : S \to \{0,1\}$ is a local assignment to the variables in S. The algorithm Assign satisfies the following properties:

1. **No-signalling.** Given a set S of variables, Assign outputs a pair (Φ, σ). Intuitively, the joint distribution of Φ and the values assigned to any subset of the variables in S are independent of the other variables in S. More precisely, for every two sets of variables S_1, S_2 both containing a subset T, the distributions obtained by executing Assign on S_1 and on S_2 to obtain (Φ, σ), and then restricting σ to the variables in T, are computationally indistinguishable.
2. **Adaptive local soundness.** We consider an execution of the cheating prover P^* in the core protocol that generates a pair (Φ, Π^*). Additively, for every small subset S of variables, we consider an execution of Assign on the set S that generates a pair (Φ', σ'). We require that Φ' is indistinguishable from Φ, and moreover, if the proof Π^* is accepting, then the assignment σ' is *locally-consistent* with the 3CNF φ' described by Φ'. We say that the assignment $\sigma' : S \to \{0,1\}$ is locally-consistent with φ' if σ' satisfies all clauses of φ' that are comprised entirely of variables in S.
   In particular, we have that if P^* has a noticeable probability of generating a pair (Φ, Π^*) such that Φ describes an unsatisfiable 3CNF, but the verifier accepts Π^*. Then for every small subset S of variables, running Assign on the set S has a noticeable probability of producing a pair (Φ', σ') where Φ' describes an unsatisfiable 3CNF φ', but σ is a locally-consistent with Φ'.

Some remarks are in order. First, we note that the relaxed soundness property has a flavor of "knowledge extraction": while we do not claim that any cheating prover for the core protocol must "know" a satisfying assignment to the 3CNF (indeed, the 3CNF might not be satisfiable, in which case no such assignment exists), a cheating prover *can* be used to generate "locally consistent" assignments on any set of variables. This extraction property is slightly more involved because it is concerned with *adaptive* cheating provers: the 3CNF is not fixed in advance. Rather, an adaptive cheating prover for the core protocol can be used to adaptively generate, given a set S of variables, an unsatisfiable 3CNF together with a locally-consistent assignment for those variables in S. The distribution of 3CNFs generated by the core protocol cheating prover (together with the bit indicating whether the verifier accepts the jointly-generated proof) is computationally indistinguishable from the distribution of 3CNFs generated by the assignment generator (together with the bit indicating whether the jointly-generated assignment is locally satisfiable). We note further that the no-signalling property implies that for any two sets S and S', the distributions of the circuit Φ generated by Assign are themselves computationally indistinguishable.

While the core protocol's soundness guarantee is robust to adaptive provers, it is weak in the sense that it only guarantees *local* consistency of the assignment generator. Even for a fixed 3CNF (let alone for an adaptively-generated one) the existence of no-signalling locally-consistent assignments does not imply that the 3CNF is satisfiable! As in prior works, we provide a "circuit-augmentation" procedure that encodes a Turing Machine computation as a 3CNF with a particular structure. The existence of a (no-signalling) locally-consistent assignment generator for the augmented 3CNF guarantees that the Turing Machine accepts its input. Here too, we need to take care to handle adaptive adversaries. This is the second main ingredient of our delegation protocol.

Ingredient 2: adaptive augmented circuit. To build an adaptively-sound delegation protocol we need an adaptive variant of the *augmented circuit construction* from [KRR14]. We describe this as a circuit-augmentation algorithm that transforms an instance (x, T) for $\mathcal{U}$ into an arithmetic circuit Φ of small size and degree, which describes a 3CNF φ. The 3CNF φ should be satisfiable if and only if $(x, \mathsf{T}) \in \mathcal{U}$. This property alone, of course, is not sufficient, since the core protocol does not prove the 3CNF's global satisfiability. Prior work showed a transformation where if $(x, \mathsf{T}) \notin \mathcal{U}$, then it is not possible to generate even locally-consistent assignments in a no-signalling manner. Since we want an adaptively-sound delegation protocol, we need an even stronger property: let Assign be a no-signalling adaptive assignment generator as above. We assume that Assign generates the circuit Φ by applying the adaptive circuit-augmentation procedure to an instance (x, T). Then for some small set S^* of variables the probability that $(x, \mathsf{T}) \notin \mathcal{U}$ but Assign generates a locally-consistent assignment for S^* is negligible. The transformation and its proof are based on [KRR14, PR14, BHK16].

There is a (slight) gap between the soundness we consider in the augmented-circuit transformation and in the core protocol: the core protocol is simply con-

cerned with 3CNFs described by small circuits. The augmented-circuit transformation, on the other hand, considers (and relies on) the procedure used to derive these 3CNFs from a computation described by a Turing Machine. This gap makes the presentation of the core protocol considerably simpler and more modular (in particular, there is no need to consider Turing Machines in the core protocol). We bridge the gap by noting that the augmentation procedure Aug is easy to invert: given a circuit Φ, it is easy to recover the instance (x, T) from which it was derived (or to output $\perp$ if Φ is not an output of Aug). This allows us to argue that for two computationally indistinguishable distributions on Φ, if the first distribution is over outputs of Aug, then the second must be over such outputs too (except with negligible probability). Moreover, given a circuit Φ produced by Aug, we can determine whether it describes a satisfiable 3CNF by recovering the original instance for $\mathcal{U}$ and testing (in polynomial time) whether the Turing Machine accepts or rejects.

Putting it together. To derive a delegation protocol, we use the core protocol's CRS. Given an instance (x, T), the prover and verifier both use the augmented-circuit transformation to derive Φ and execute the core protocol on Φ. A prover P^* that cheats with noticeable probability can be used to derive a no-signalling adaptive local assignment generator Assign^*. By the core protocol's soundness we conclude that for every set S of variables, with noticeable probability Assign^* generates pairs (Φ, σ) where Φ describes an unsatisfiable 3CNF, but σ is locally consistent. Moreover, Φ is derived by running the augmented circuit construction on an instance $(x, \mathsf{T}) \notin \mathcal{U}$ (this is true for the execution of the core protocol, by computational indistinguishability it holds also for the outputs of Assign^*). However, the the augmented circuit construction guarantees that no such assignment generator exists, leading to a contradiction.

Organization. We define adaptive no-signalling local assignment generators in Sect. 4.2. The core protocol is given in Sect. 4.3. The properties of the augmented-circuit transformation are discussed in Sect. 4.4. The analysis of the core protocol, the augmented-circuit transformation, and the full delegation protocol appear in the full version of this work.

4.2 Adaptive No-Signaling Local-Assignment Generator

Before stating the properties of the core protocol, we introduce some notation and formalize the notion of an adaptive no-signaling local-assignment generator.

Succinct formula representation $\mathcal{I}_\varphi$. Let φ be a 3CNF boolean formula with variables $\alpha_1, \ldots, \alpha_B$. Let $B = 2^m$ and identify the indices in $[B]$ with strings in $\{0,1\}^m$. We define a boolean *indicator function* $\mathcal{I}_\varphi : \{0,1\}^{3m+3} \to \{0,1\}$ of φ as follows. For every indices $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3 \in \{0,1\}^m$ and for every bits $b_1, b_2, b_3 \in \{0,1\}^3$, we have that

$$\mathcal{I}_\varphi(\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, b_1, b_2, b_3) = 1,$$

if and only if φ contains the clause:

$$(\alpha_{\mathbf{u}_1} = b_1) \vee (\alpha_{\mathbf{u}_2} = b_2) \vee (\alpha_{\mathbf{u}_3} = b_3).$$

The locally consistency verifier $\mathsf{V}_{\mathsf{local}}$. We denote by $\mathsf{V}_{\mathsf{local}}$ the verification algorithm for local assignments to φ. The algorithm is given as input

- An arithmetic circuit Φ computing a boolean function with $3m + 3$ inputs (we think of Φ as computing the indicator function $\mathcal{I}_\varphi$ for some formula φ).
- A partial assignments $\sigma : S \to \{0,1\}$ for a set $S \subseteq \{0,1\}^m$.

$\mathsf{V}_{\mathsf{local}}(\Phi, \sigma)$ accepts if an only if the assignment σ is locally consistent with the formula described by Φ. That is, for every $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3 \in S$ and every $b_1, b_2, b_3 \in \{0,1\}$

$$\Phi(\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, b_1, b_2, b_3) = 1 \quad \Rightarrow \quad (\sigma(\mathbf{u}_1) = b_1) \vee (\sigma(\mathbf{u}_2) = b_2) \vee (\sigma(\mathbf{u}_3) = b_3).$$

Adaptive local-assignment generator. Let $Q = Q(\lambda), B = B(\lambda)$ be functions and let $B = 2^m$. An adaptive Q-local-assignment generator Assign for B-variate formulas is a probabilistic algorithm with the following syntax: given the the security parameter 1^λ and a set of indices $S \subseteq \{0,1\}^m$ of size at most Q, Assign outputs

- An arithmetic circuit Φ computing a boolean function with $3m + 3$ inputs.
- A partial assignment $\sigma : S \to \{0,1\}$.

We define a no-signaling adaptive local-assignment generator

Definition 4.1 (No-Signaling Adaptive Local-Assignment Generator). *A Q-local-assignment generator* Assign *for $B = 2^m$-variate formulas is (computationally) no-signalling if for every polynomial-size distinguisher D there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and every subsets $S \subseteq S' \subseteq \{0,1\}^m$ of size at most Q*

$$\left| \Pr_{(\Phi,\sigma) \leftarrow \mathsf{Assign}(1^\lambda, S)} [D(\Phi, \sigma(S)) = 1] - \Pr_{(\Phi,\sigma') \leftarrow \mathsf{Assign}(1^\lambda, S')} [D(\Phi, \sigma'(S)) = 1] \right| \leq \mu(\lambda).$$

4.3 The Core Protocol

In this section we describe the syntax and the properties of the core delegation protocol. The protocol itself is given in Sect. 4.3.

Syntax. Let $\Delta = \Delta(\lambda)$ be a polynomially bounded function. The core protocol with degree bound Δ consists of PPT algorithms $(\mathsf{Core.Gen}, \mathsf{Core.P}, \mathsf{Core.V})$ with the following syntax. Let φ be a B-variate 3CNF boolean formula where $B = 2^m$ and let Φ be an arithmetic circuit of total degree $\delta \leq \Delta$ computing the function $\mathcal{I}_\varphi$.

$\mathsf{Core.Gen}$**:** Given the security parameter 1^λ and a locality parameter 1^Q outputs a common reference string CRS.

$\mathsf{Core.P}$**:** Given the common reference string CRS, the circuit Φ and an assignment $\sigma : \{0,1\}^m \to \{0,1\}$, outputs a proof Π.

$\mathsf{Core.V}$**:** Given the common reference string CRS, the circuit Φ and the proof Π outputs a bit.

The protocol satisfies the following requirements.

Completeness. For every security parameter $\lambda \in \mathbb{N}$, every 3CNF boolean formula φ with B variables, every satisfying assignment σ, every arithmetic circuit Φ of individual degree $\delta \leq \Delta$ computing the function $\mathcal{I}_\varphi$, and every locality parameter $Q \in [B]$

$$\Pr\left[\mathsf{Core.V}(\mathsf{CRS}, \Phi, \Pi) = 1 \;\middle|\; \begin{array}{l} \mathsf{CRS} \leftarrow \mathsf{Core.Gen}(1^\lambda, 1^Q) \\ \Pi \leftarrow \mathsf{Core.P}(\mathsf{CRS}, \Phi, \sigma) \end{array}\right] = 1.$$

Efficiency. There exists a polynomial L such that in the above honest experiment $|\Pi| \leq L(\lambda) \cdot Q \cdot \delta$ where δ is the individual degree of the circuit Φ. Additionally the verifier's running time is bounded by $L(|\mathsf{CRS}|) \cdot (|\Phi| + |\Pi|)$.

No-Signaling adaptive local soundness. For every polynomially bounded functions $Q = Q(\lambda), B = B(\lambda)$ there exists an algorithm Assign such that for every poly-size cheating prover P^* the following holds

- $\mathsf{Assign}^{\mathsf{P}^*}$ is a no-signaling adaptive Q-local-assignment generator for B-variate formulas.
- For every polynomial-size distinguisher D there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$, letting $B = 2^m$, for every set of indices $S \subseteq \{0,1\}^m$ of size at most Q

$$\left| \begin{array}{l} \Pr\left[D(\Phi, b) = 1 \;\middle|\; \begin{array}{l} \mathsf{CRS} \leftarrow \mathsf{Core.Gen}(1^\lambda, 1^Q) \\ (\Phi, \Pi^*) \leftarrow \mathsf{P}^*(\mathsf{CRS}) \\ b \leftarrow \mathsf{Core.V}(\mathsf{CRS}, \Phi, \Pi^*) \end{array}\right] \\ \\ - \Pr\left[D(\Phi, b) = 1 \;\middle|\; \begin{array}{l} \\ (\Phi, \sigma) \leftarrow \mathsf{Assign}^{\mathsf{P}^*}(1^\lambda, S) \\ b \leftarrow \mathsf{V}_{\mathsf{local}}(\Phi, \sigma) \end{array}\right] \end{array} \right| \leq \mu(\lambda).$$

Construction. Let $\Delta = \Delta(\lambda)$ be the function bounding the total degree of the circuit Φ. The core protocol makes use of a 3-key zero-testable 2Δ-somewhat homomorphic encryption scheme

$$(\mathsf{MHE.ParamGen}, \mathsf{MHE.KeyGen}, \mathsf{MHE.Enc}, \mathsf{MHE.Dec}, \mathsf{MHE.WeakDec},$$
$$\mathsf{MHE.Eval}, \mathsf{MHE.ZT}).$$

The CRS generator. The CRS generation algorithm $\mathsf{Core.Gen}$ is given as input the security parameter 1^λ and a locality parameter 1^Q. It outputs a common reference string CRS as follows.

1. Sample public parameters for the encryption scheme

$$(\mathsf{pp}, R) \leftarrow \mathsf{MHE.ParamGen}(1^\lambda).$$

2. For every $q \in [Q]$, generate a key pair

$$(\mathsf{sk}^q, \mathsf{pk}^q) \leftarrow \mathsf{MHE.KeyGen}(\mathsf{pp}),$$

and λ encryptions of 0

$$\{c_i^q \leftarrow \mathsf{MHE.Enc}(\mathsf{pp}, \mathsf{pk}^q, 0)\}_{i\in[\lambda]}.$$

3. Output a reference string containing the public parameters and all the public keys and ciphers

$$\mathsf{CRS} = \left(\mathsf{pp}, \{\mathsf{pk}^q, (c_1^q, \ldots, c_\lambda^q)\}_{q\in[Q]}\right).$$

The prover. The prover algorithm $\mathsf{Core.P}$ is given as input

- The common reference string

$$\mathsf{CRS} = \left(\mathsf{pp}, \{\mathsf{pk}^q, (c_1^q, \ldots, c_\lambda^q)\}_{q\in[Q]}\right).$$

- An (individual) degree δ arithmetic circuit Φ computing a boolean function with $3m + 3$ inputs.
- An assignment $\sigma : \{0,1\}^m \to \{0,1\}$.

We start by introducing some notation.

1. For every query $q \in [Q]$, let $\mathbf{c}^q = (c_1^q, \ldots, c_m^q)$. We refer to the ciphertext vector $\mathbf{c}^q$ as an encryption of the q-th CRS index (in an honestly generated CRS the index value is always 0^m).
2. Let Σ be a multi-linear extension of σ (See Sect. 2.2).
3. For every triplet of bits $\mathbf{b} = (b_1, b_2, b_3) \in \{0,1\}^3$ let $P_0^{\mathbf{b}}$ be the degree $\delta + 1$ arithmetic circuit taking $3m$ inputs

$$P_0^{\mathbf{b}}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) = \Phi(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{b}) \cdot \prod_{k\in[3]} (1 - \beta(b_k, \Sigma(\mathbf{x}_k))). \tag{4}$$

(See Sect. 2.2 for the definition of the circuit β.)

4. For every $i \in [3m]$, let $P_i^{\mathbf{b}}$ be the linearization of the first i variables of the circuit $P_0^{\mathbf{b}}$. That is, $P_i^{\mathbf{b}}$ is the following arithmetic circuit taking $3m$ inputs which is multilinear in its first i variables, and of degree at most $\delta + 1$ in its other variables.

$$P_i^{\mathbf{b}}(x_1, \ldots, x_{3m}) = \sum_{y_1,\ldots y_i\in\{0,1\}} \beta(y_1, \ldots y_i, x_1, \ldots x_i) \cdot P_0^{\mathbf{b}}(y_1, \ldots, y_i, x_{i+1}, \ldots, x_{3m}). \tag{5}$$

$\mathsf{Core.P}$ outputs a proof Π as follows.

1. For every $q \in [Q]$ obtain an encryption of the assignment Σ evaluated on the q-th CRS index. That is, homomorphically obtain the ciphertext $d^q = \langle \Sigma(\mathbf{c}^q) \rangle$.
2. For every triplet of bits $\mathbf{b} \in \{0,1\}^3$, triplet of queries $\mathbf{q} = (q_1, q_2, q_3) \in [Q]^3$, and $i \in [3m]$ obtain the encrypted coefficients of the circuit $P^{\mathbf{b}}_{i-1}$ evaluated on the CRS indices $\mathbf{q}$ and restricted to its i-th input variable (see Sect. 2.1). Since the individual degree of $P^{\mathbf{b}}_{i-1}$ is at most $\delta+1$, the restricted polynomial will have at most $\delta+2$ coefficients. That is, homomorphically obtain the sequence of $\delta+2$ ciphertexts $\mathbf{e}^{\mathbf{q},\mathbf{b}}_{i-1}$

$$\mathbf{e}^{\mathbf{q},\mathbf{b}}_{i-1} = \left(\left\langle P^{\mathbf{b}}_{i-1}\big|_{i,j} \left((\mathbf{c}^{q_1} \mid \mathbf{c}^{q_2} \mid \mathbf{c}^{q_3})_{-i} \right) \right\rangle \right)_{j \in [0,\delta+1]}.$$

3. Output a proof Π that contains all the ciphertexts

$$\Pi = \left(\{d^q\}_{q \in [Q]}, \quad \left\{ \mathbf{e}^{\mathbf{q},\mathbf{b}}_{i-1} \right\}_{\mathbf{b} \in \{0,1\}^3, \mathbf{q} \in [Q]^3, i \in [3m]} \right).$$

The verifier. The verifier algorithm $\mathsf{Core.V}$ is given as input

– The common reference string

$$\mathsf{CRS} = \left(\mathsf{pp}, \{\mathsf{pk}^q, (c^q_1, \ldots, c^q_\lambda)\}_{q \in [Q]} \right).$$

– A degree δ arithmetic circuit Φ computing a boolean function with $3m+3$ inputs.
– An proof

$$\Pi = \left(\{d^q\}_{q \in [Q]}, \quad \left\{ \mathbf{e}^{\mathbf{q},\mathbf{b}}_{i-1} \right\}_{\mathbf{b} \in \{0,1\}^3, \mathbf{q} \in [Q]^3, i \in [3m]} \right).$$

$\mathsf{Core.V}$ performs the following tests for every triplet of bits $\mathbf{b} = (b_1, b_2, b_3) \in \{0,1\}^3$ and triplet of queries $\mathbf{q} = (q_1, q_2, q_3) \in [Q]^3$. $\mathsf{Core.V}$ accepts only if all tests pass.

First, $\mathsf{Core.V}$ homomorphically evaluates the following ciphertexts

– Let $\tilde{P}^{\mathbf{b}}$ be the following arithmetic circuit taking $3m+3$ inputs

$$\tilde{P}^{\mathbf{b}}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, y_1, y_3, y_3) = \Phi(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{b}) \cdot \prod_{k \in [3]} (1 - \beta(b_k, y_k)). \tag{6}$$

Evaluate the ciphertext

$$f'_0 = \left\langle \tilde{P}^{\mathbf{b}} \left(\mathbf{c}^{q_1}, \mathbf{c}^{q_2}, \mathbf{c}^{q_3}, d^{q_1}, d^{q_2}, d^{q_3} \right) \right\rangle.$$

– Let F be the following arithmetic circuit taking $\delta+3$ inputs

$$F(x, y_0, \ldots, y_{\delta+1}) = \sum_{j \in [0,\delta+1]} y_j \cdot x^j. \tag{7}$$

For $i \in [3m]$, evaluate the ciphertext f_{i-1} that encrypts the evaluation of the univariate polynomial with encrypted coefficient $\mathbf{e}_{i-1}^{\mathbf{q},\mathbf{b}}$ on the i-th input bit of the concatenated CRS indices $\mathbf{q}$. Recall that $\mathbf{e}_{i-1}^{\mathbf{q},\mathbf{b}}$ are supposedly the encrypted coefficients of the circuit $P_{i-1}^{\mathbf{b}}$ evaluated on the CRS indices $\mathbf{q}$ and restricted to its i-th input variable. Therefore, f_{i-1} is suppose to encrypt the evaluation of $P_{i-1}^{\mathbf{b}}$ on the CRS indices $\mathbf{q}$.

$$f_{i-1} = \left\langle F\left(\left(\mathbf{c}^{q_1} \mid \mathbf{c}^{q_2} \mid \mathbf{c}^{q_3}\right)_i, \mathbf{e}_{i-1}^{\mathbf{q},\mathbf{b}}\right)\right\rangle.$$

– Let F' be the following arithmetic circuit taking $\delta + 3$ inputs

$$F'(x, y_0, \ldots, y_{\delta+1}) = \sum_{z \in \{0,1\}} \beta(z, x) \cdot F(z, y_0, \ldots, y_{\delta+1}).$$

For $i \in [3m]$, evaluate the ciphertext f_i' that encrypts the linearization of the univariate polynomial with encrypted coefficient $\mathbf{e}_{i-1}^{\mathbf{q},\mathbf{b}}$ evaluated on the on the i-th input bit of the concatenated CRS indices $\mathbf{q}$. Therefore, f_{i-1} is suppose to encrypt the evaluation of the circuit $P_{i-1}^{\mathbf{b}}$ with its i-th variable linearized on the CRS indices $\mathbf{q}$.

$$f_i' = \left\langle F'\left(\left(\mathbf{c}^{q_1} \mid \mathbf{c}^{q_2} \mid \mathbf{c}^{q_3}\right)_i, \mathbf{e}_{i-1}^{\mathbf{q},\mathbf{b}}\right)\right\rangle.$$

– Let $f_{3m} = \hat{0}$.
For every $i \in [0, 3m]$, $\mathsf{Core.V}$ tests that

$$\mathsf{MHE.ZT}\left(\mathsf{pp}, \left(\mathsf{pk}^{q_1}, \mathsf{pk}^{q_2}, \mathsf{pk}^{q_3}\right), \langle f_i - f_i' \rangle\right) = 1.$$

4.4 The Augmented Circuit

Syntax. Let $\mathcal{U}$ be the universal language (see Sect. 2.3). The augmented-circuit transformation consists of deterministic polynomial time algorithms $(\mathsf{Aug}, \mathsf{Aug}^{-1}, \mathsf{Trans})$ with the following syntax.

Aug: the circuit-augmentation procedure takes as input an instance $x = (M, y)$ and a time bound T for $\mathcal{U}$. It outputs an arithmetic circuit Φ computing the indicator function $\mathcal{I}_\varphi$ of the "augmented formula" φ (see Sect. 4.2)). We say that Φ *represents* φ.

Aug^{-1}: the inversion procedure takes as input an arithmetic circuit Φ. It either outputs (x, T) or fails and outputs $\perp$.

Trans: the assignment generation procedure takes as input an instance x and a time bound T for $\mathcal{U}$. It outputs an assignment σ for φ.

These procedures satisfy the following properties:

Efficiency. For $x \in \{0,1\}^n$

– $\mathsf{Aug}(x, \mathsf{T})$ runs in time $n \cdot \mathrm{polylog}(\mathsf{T})$ and outputs an arithmetic circuit Φ such that

- Φ is of size $n \cdot \text{polylog}(\mathsf{T})$.
- Φ is of total degree $\delta = \delta(n, \mathsf{T}) = \text{polylog}(n, \mathsf{T})$.
- Φ represents a formula φ on $B = B(n, \mathsf{T}) = \text{poly}(n, \mathsf{T})$ variables.

– $\mathsf{Aug}(x, \mathsf{T})$ and $\mathsf{Aug}^{-1}(\Phi)$ run in time $n \cdot \text{polylog}(\mathsf{T})$.
– $\mathsf{Trans}(x, \mathsf{T})$ runs in time $\text{poly}(n, \mathsf{T})$.

Inversion. For every $(x, \mathsf{T}) \in \{0,1\}^*$

$$\mathsf{Aug}^{-1}(\mathsf{Aug}(x, \mathsf{T})) = (x, \mathsf{T}).$$

Completeness. For every $(x, \mathsf{T}) \in \mathcal{U}$, $\mathsf{Trans}(x, \mathsf{T})$ outputs a satisfying assignment σ for the formula φ represented by the output of $\mathsf{Aug}(x, \mathsf{T})$.

Soundness. At a high level, the soundness guarantees that there does not exist an adaptive no-signalling local-assignment generator (see Sect. 4.2) that for *every* small set of indices S generates a circuit $\Phi = \mathsf{Aug}(x, \mathsf{T})$, such that $(x, \mathsf{T}) \notin \mathcal{U}$, together with partial assignment $\sigma : S \to \{0,1\}$ that is locally consistent with the formula represented by Φ.

Lemma 4.1 (Augmented Circuit Soundness). *There exists a function* $Q = \text{polylog}(\lambda)$ *such that for every polynomially bounded function* $B = B(\lambda)$*, and every polynomial-time no-signaling* Q*-local-assignment generator* Assign *for* B*-variate formulas there exists a negligible function* μ *such that for every* $\lambda \in \mathbb{N}$*, letting* $B = 2^m$*, there exists a set* $S^* \subseteq \{0,1\}^m$ *of size at most* Q *such that*

$$Pr\left[\begin{matrix}\mathsf{Aug}^{-1}(\Phi) \notin \mathcal{U} \cup \{\bot\} \\ \mathsf{V}_{\mathsf{local}}(\Phi, \sigma) = 1\end{matrix} \;\middle|\; (\Phi, \sigma) \leftarrow \mathsf{Assign}(1^\lambda, S^*)\right] \leq \mu(\lambda).$$

Acknowledgements. We thank Zvika Brakerski, Yael Kalai, Ron Rothblum and Nir Bitansky for many helpful and illuminating conversations.

References

[ABOR00] Aiello, W., Bhatt, S., Ostrovsky, R., Rajagopalan, S.R.: Fast verification of any remote procedure call: short witness-indistinguishable one-round proofs for NP. In: Montanari, U., Rolim, J.D.P., Welzl, E. (eds.) ICALP 2000. LNCS, vol. 1853, pp. 463–474. Springer, Heidelberg (2000). doi:10.1007/3-540-45022-X_39

[BCCT13] Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: Recursive composition and bootstrapping for SNARKS and proof-carrying data. In: STOC, pp. 111–120 (2013)

[BCI+13] Bitansky, N., Chiesa, A., Ishai, Y., Paneth, O., Ostrovsky, R.: Succinct non-interactive arguments via linear interactive proofs. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 315–333. Springer, Heidelberg (2013). doi:10.1007/978-3-642-36594-2_18

[BFLS91] Babai, L., Fortnow, L., Levin, L.A., Szegedy, M.: Checking computations in polylogarithmic time. In: Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, 5–8 May 1991, New Orleans, Louisiana, USA, pp. 21–31 (1991)

[BGL+15] Bitansky, N., Garg, S., Lin, H., Pass, R., Telang, S.: Succinct randomized encodings and their applications. In: Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, 14–17 June 2015, pp. 439–448 (2015)

[BGN05] Boneh, D., Goh, E.-J., Nissim, K.: Evaluating 2-DNF formulas on ciphertexts. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 325–341. Springer, Heidelberg (2005). doi:10.1007/978-3-540-30576-7_18

[BHK16] Brakerski, Z., Holmgren, J., Kalai, Y.T.: Non-interactive RAM and batch NP delegation from any PIR. IACR Cryptology ePrint Archive, 2016:459 (2016)

[BS02] Boneh, D., Silverberg, A.: Applications of multilinear forms to cryptography. IACR Cryptology ePrint Archive, 2002:80 (2002)

[BV11] Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. IACR Cryptology ePrint Archive, 2011:344 (2011)

[CGH04] Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited. J. ACM **51**(4), 557–594 (2004)

[CHJV14] Canetti, R., Holmgren, J., Jain, A., Vaikuntanathan, V.: Indistinguishability obfuscation of iterated circuits and ram programs. Cryptology ePrint Archive, Report 2014/769 (2014). http://eprint.iacr.org/

[CHL+15] Cheon, J.H., Han, K., Lee, C., Ryu, H., Stehlé, D.: Cryptanalysis of the multilinear map over the integers. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 3–12. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46800-5_1

[CLT15] Coron, J.-S., Lepoint, T., Tibouchi, M.: New multilinear maps over the integers. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 267–286. Springer, Heidelberg (2015). doi:10.1007/978-3-662-47989-6_13

[DFH12] Damgård, I., Faust, S., Hazay, C.: Secure two-party computation with low communication. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 54–74. Springer, Heidelberg (2012). doi:10.1007/978-3-642-28914-9_4

[DLN+04] Dwork, C., Langberg, M., Naor, M., Nissim, K., Reingold, O.: Succinct proofs for NP and spooky interactions (2004, unpublished manuscript). http://www.cs.bgu.ac.il/~kobbi/papers/spooky_sub_crypto.pdf

[DNR16] Dwork, C., Naor, M., Rothblum, G.N.: Spooky interaction and its discontents: compilers for succinct two-message argument systems. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9816, pp. 123–145. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53015-3_5

[Gen09] Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, 31 May–2 June 2009, pp. 169–178 (2009)

[GGH13a] Garg, S., Gentry, C., Halevi, S.: Candidate multilinear maps from ideal lattices. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 1–17. Springer, Heidelberg (2013). doi:10.1007/978-3-642-38348-9_1

[GGH+13b] Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: FOCS (2013)

[GGH15] Gentry, C., Gorbunov, S., Halevi, S.: Graph-induced multilinear maps from lattices. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015. LNCS, vol. 9015, pp. 498–527. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46497-7_20

[GGHZ16] Garg, S., Gentry, C., Halevi, S., Zhandry, M.: Functional encryption without obfuscation. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016. LNCS, vol. 9563, pp. 480–511. Springer, Heidelberg (2016). doi:10.1007/978-3-662-49099-0_18

[GGPR13] Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct NIZKs without PCPs. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 626–645. Springer, Heidelberg (2013). doi:10.1007/978-3-642-38348-9_37

[GHV10] Gentry, C., Halevi, S., Vaikuntanathan, V.: i-hop homomorphic encryption and rerandomizable yao circuits. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 155–172. Springer, Heidelberg (2010). doi:10.1007/978-3-642-14623-7_9

[GKR08] Goldwasser, S., Kalai, Y.T., Rothblum, G.N.: Delegating computation: interactive proofs for muggles. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, 17–20 May 2008, pp. 113–122 (2008)

[GLSW15] Gentry, C., Lewko, A.B., Sahai, A., Waters, B.: Indistinguishability obfuscation from the multilinear subgroup elimination assumption. In: IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17–20 October 2015, pp. 151–170 (2015)

[GMM+16] Garg, S., Miles, E., Mukherjee, P., Sahai, A., Srinivasan, A., Zhandry, M.: Secure obfuscation in a weak multilinear map model. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9986, pp. 241–268. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53644-5_10

[GPSZ17] Garg, S., Pandey, O., Srinivasan, A., Zhandry, M.: Breaking the subexponential barrier in obfustopia. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10212, pp. 156–181. Springer, Cham (2017). doi:10.1007/978-3-319-56617-7_6

[Gro10] Groth, J.: Short pairing-based non-interactive zero-knowledge arguments. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 321–340. Springer, Heidelberg (2010). doi:10.1007/978-3-642-17373-8_19

[GSW13] Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: conceptually-simpler, asymptotically-faster, attribute-based. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8042, pp. 75–92. Springer, Heidelberg (2013). doi:10.1007/978-3-642-40041-4_5

[GW11] Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, pp. 99–108 (2011)

[HJ16] Hu, Y., Jia, H.: Cryptanalysis of GGH map. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9665, pp. 537–565. Springer, Heidelberg (2016). doi:10.1007/978-3-662-49890-3_21

[HRSV11] Hohenberger, S., Rothblum, G.N., Shelat, A., Vaikuntanathan, V.: Securely obfuscating re-encryption. J. Cryptol. **24**(4), 694–719 (2011)

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

[KLW14] Koppula, V., Lewko, A.B., Waters, B.: Indistinguishability obfuscation for turing machines with unbounded memory. Cryptology ePrint Archive, Report 2014/925 (2014). http://eprint.iacr.org/

[KP16] Kalai, Y., Paneth, O.: Delegating RAM computations. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9986, pp. 91–118. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53644-5_4

[KRR13] Kalai, Y.T., Raz, R., Rothblum, R.D.: Delegation for bounded space. In: STOC, pp. 565–574 (2013)

[KRR14] Kalai, Y.T., Raz, R., Rothblum, R.D.: How to delegate computations: the power of no-signaling proofs. In: Symposium on Theory of Computing, STOC 2014, New York, NY, USA, 31 May–03 June 2014, pp. 485–494 (2014)

[Lip12] Lipmaa, H.: Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 169–189. Springer, Heidelberg (2012). doi:10.1007/978-3-642-28914-9_10

[LTV12] López-Alt, A., Tromer, E., Vaikuntanathan, V.: On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In: Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, 19–22 May 2012, pp. 1219–1234 (2012)

[LV16] Lin, H., Vaikuntanathan, V.: Indistinguishability obfuscation from DDH-like assumptions on constant-degree graded encodings. In: IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9–11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA, pp. 11–20 (2016)

[Mic94] Micali, S.: CS proofs (extended abstracts). In: 35th Annual Symposium on Foundations of Computer Science, Santa Fe, New Mexico, USA, 20–22 November 1994, pp. 436–453 (1994)

[MSZ16] Miles, E., Sahai, A., Zhandry, M.: Annihilation attacks for multilinear maps: cryptanalysis of indistinguishability obfuscation over GGH13. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9815, pp. 629–658. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53008-5_22

[Nao03] Naor, M.: On cryptographic assumptions and challenges. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 96–109. Springer, Heidelberg (2003). doi:10.1007/978-3-540-45146-4_6

[PR14] Paneth, O., Rothblum, G.N.: Publicly verifiable non-interactive arguments for delegating computation. Cryptology ePrint Archive, Report 2014/981 (2014). http://eprint.iacr.org/2014/981

[SW14] Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: STOC (2014)

[vDGHV10] van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully homomorphic encryption over the integers. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 24–43. Springer, Heidelberg (2010). doi:10.1007/978-3-642-13190-5_2

[WB13] Walfish, M., Blumberg, A.J.: Verifying computations without reexecuting them: from theoretical possibility to near-practicality. Electronic Colloquium on Computational Complexity (ECCC), 20:165 (2013)

[Zim15] Zimmerman, J.: How to obfuscate programs directly. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 439–467. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46803-6_15