

The Hunting of the SNARK*

Nir Bitansky[†] Ran Canetti[‡] Alessandro Chiesa[§] Shafi Goldwasser[§] Huijia Lin[¶]
Aviad Rubinfeld[†] Eran Tromer[†]

July 24, 2014

Abstract

The existence of succinct non-interactive arguments for NP (i.e., non-interactive computationally-sound proofs where the verifier’s work is essentially independent of the complexity of the NP nondeterministic verifier) has been an intriguing question for the past two decades. Other than CS proofs in the random oracle model [Micali, FOCS ’94], the only existing candidate construction is based on an elaborate assumption that is tailored to a specific protocol [Di Crescenzo and Lipmaa, CiE ’08].

We formulate a general and relatively natural notion of an *extractable collision-resistant hash function (ECRH)* and show that, if ECRHs exist, then a modified version of Di Crescenzo and Lipmaa’s protocol is a succinct non-interactive argument for NP. Furthermore, the modified protocol is actually a succinct non-interactive *adaptive argument of knowledge (SNARK)*. We then propose several candidate constructions for ECRHs and relaxations thereof.

We demonstrate the applicability of SNARKs to various forms of delegation of computation, to succinct non-interactive zero knowledge arguments, and to succinct two-party secure computation. Finally, we show that SNARKs essentially imply the existence of ECRHs, thus demonstrating the necessity of the assumption.

Going beyond ECRHs, we formulate the notion of *extractable one-way functions (EOWFs)*. Assuming the existence of a natural variant of EOWFs, we construct a 2-message selective-opening-attack secure commitment scheme and a 3-round zero-knowledge argument of knowledge. Furthermore, if the EOWFs are concurrently extractable, the 3-round zero-knowledge protocol is also concurrent zero-knowledge. Our constructions circumvent previous black-box impossibility results regarding these protocols by relying on EOWFs as the non-black-box component in the security reductions.

*This research was supported by the Check Point Institute for Information Security, by the Israeli Centers of Research Excellence (I-CORE) program (center No. 4/11), by the European Community’s Seventh Framework Programme (FP7/2007-2013) grant 240258, by a European Union Marie Curie grant, by the Israeli Science Foundation, and by the Israeli Ministry of Science and Technology.

This paper is a merge of [BCCT11] and [GLR11]. A preliminary version including part of the results appeared in ITCS 2012.

[†]Tel Aviv University, {nirbitan,tromer,aviadrub}@tau.ac.il

[‡]Boston University and Tel Aviv University, canetti@tau.ac.il

[§]MIT, {alexch,shafi}@csail.mit.edu

[¶]Boston University and MIT, huijia@csail.mit.edu

Contents

1	Introduction	1
1.1	Summary of Our Results	3
1.2	ECRHs, SNARKs, and Applications	4
1.3	ECRH Candidate Constructions and Sufficient Relaxations	6
1.4	High-Level Proof Strategy for Theorem 1	9
1.5	Discussion	12
1.6	Organization	12
2	Other Related Work	13
3	Preliminaries	14
3.1	Interactive Proofs, Zero Knowledge and Witness Indistinguishability	14
3.2	Proofs and Arguments of Knowledge	16
3.3	Commitment Scheme	16
3.4	Collision-Resistant Hashes	16
3.5	Merkle Trees	16
3.6	Private Information Retrieval	17
3.7	Probabilistically Checkable Proofs of Knowledge	18
4	SNARKs	18
4.1	The Universal Relation and NP Relations	18
4.2	Succinct Non-Interactive Arguments	19
5	From ECRHs to SNARKs	20
5.1	Construction Details	21
5.2	Proof of Security	23
5.3	Extension to Universal Arguments	27
6	ECRHs: a Closer Look	28
6.1	ECRHs	28
6.2	PECRHs	30
6.3	Weak PECRHs	30
7	From SNARKs to PECRHs (and More)	32
7.1	From SNARKs to PECRHs	33
7.2	From Leakage-Resilient Primitives and SNARKs to Extractable Primitives	34
8	Candidate ECRH and PECRH Constructions	36
8.1	ECRHs from t -Knowledge of Exponent	36
8.2	PECRHs from Knowledge of Knapsack	37
9	Zero-Knowledge SNARKs	42
9.1	SNARK on top of NIZK	42
9.2	NIZK on top of SNARK	43
10	Applications of SNARKs and zkSNARKs	44
10.1	Delegation of Computation	44
10.2	Succinct Non-Interactive Secure Computation	47

11 Extractable One-Way Functions and Their Applications	48
11.1 Definitions of sEOWFs and scEOWFs	49
11.2 A Special 3-round ZKAOK protocol	51
11.3 A 3-round Concurrent ZK Protocol	54
11.4 Two-message Selective Opening Attack Secure Commitments	55
11.5 Candidate Constructions of sEOWF and scEOWF	57
Acknowledgements	58
References	59

1 Introduction

For the Snark's a peculiar creature, that won't
Be caught in a commonplace way.
Do all that you know, and try all that you don't:
Not a chance must be wasted to-day!

*The Hunting of the Snark,
Lewis Carroll*

The notion of interactive proof systems [GMR89] is central to both modern cryptography and complexity theory. One extensively studied aspect of interactive proof systems is their expressiveness; this study culminated with the celebrated result that $IP = PSPACE$ [Sha92]. Another aspect of such systems, which is the focus of this work, is that proofs for rather complex NP-statements can potentially be verified much faster than by direct checking of an NP witness.

We know that if statistical soundness is required then any non-trivial savings would cause unlikely complexity-theoretic collapses (see, e.g., [BHZ87, GH98, GVW02, Wee05]). However, if we settle for proof systems with only computational soundness (also known as interactive arguments [BCC88]) then significant savings can be made. Indeed, using collision-resistant hash functions (CRHs), Kilian [Kil92] shows a four-message interactive argument for NP: the prover first uses a Merkle hash tree to bind itself to a polynomial-size PCP (Probabilistically Checkable Proof) string for the statement, and then answers the PCP verifier's queries while demonstrating consistency with the Merkle tree. This way, membership of an instance y in an NP language L can be verified in time that is bounded by $p(k, |y|, \log t)$, where t is the time to evaluate the NP verification relation for L on input y , p is a fixed polynomial independent of L , and k is a security parameter that determines the soundness error. Following tradition, we call such argument systems *succinct*.

Can we have succinct argument systems which are *non-interactive*? Having posed and motivated this question, Micali [Mic00] provides a one-message succinct non-interactive argument for NP, in the random oracle model, by applying the Fiat-Shamir paradigm [FS87] to Kilian's protocol. In the standard model, such "totally non-interactive" succinct arguments (against non-uniform provers) do not exist except for "quasi-trivial" languages (i.e., languages in $BPTIME(n^{\text{poly} \log n})$), because the impossibility results for statistical soundness can be directly extended to this case. Nonetheless, it may still be possible to obtain a slightly weaker notion of non-interactivity:

Definition 1.1. *A succinct non-interactive argument (SNARG) is a succinct argument where the verifier (or a trusted entity) generates ahead of time a succinct verifier-generated reference string (VGRS) and sends it to the prover. The prover can then use the VGRS to generate a succinct non-interactive proof π for a statement y of his choice. (The VGRS is thus independent of the statements to be proven later, and the definition requires "adaptive soundness", since y is chosen by the prover, potentially based on the VGRS.)*

In this paper, we consider the following questions:

Can SNARGs for NP exist in the standard model?
And if so, under what assumptions can we prove their existence?

Attempted solutions. To answer the above question, Aiello et al. [ABOR00] propose to avoid Kilian's hash-then-open paradigm, and instead use a polylogarithmic PIR (Private Information Retrieval) scheme

to access the PCP oracle as a long database. The verifier’s first message consists of the queries of the underlying PCP verifier, encrypted using the PIR chooser algorithm. The prover applies the PIR sender algorithm to the PCP oracle, and the verifier then runs the underlying PCP verifier on the values obtained from the PIR protocol. However, Dwork et al. [DLN⁺04] point out that this “PCP+PIR approach” is inherently problematic, because a cheating prover could “zigzag” and answer different queries according to different databases.¹ (Natural extensions that try to force consistency by using multiple PIR instances run into trouble due to potential PIR malleability.)

Di Crescenzo and Lipmaa [DCL08] propose to address this problem by further requiring the prover to bind itself (in the clear) to a specific database using a Merkle Tree (MT) as in Kilian’s protocol. Intuitively, the prover should now be forced to answer according to a single PCP string. In a sense, this “PCP+MT+PIR approach” squashes Kilian’s four-message protocol down to two messages “under the PIR”. However, while initially appealing, it is not a-priori clear how this intuition can be turned into a proof of security under some well defined properties of the Merkle tree hash. Indeed, to prove soundness of their protocol Di Crescenzo and Lipmaa use an assumption that is non-standard in two main ways: first, it is a “knowledge assumption,” in the sense that any adversary that generates a value of a certain form is assumed to “know” a corresponding preimage (see more discussion on such assumptions below). Furthermore, their assumption is very specific and intimately tied to the actual hash, PIR, and PCP schemes in use, as well as the language under consideration. Two other non-interactive arguments for NP, based on more concise knowledge assumptions, are due to Mie [Mie08] and Groth [Gro10]. However, neither of these protocols is succinct: in both protocols the verifier’s runtime is polynomially related to the time needed to directly verify the NP witness.

Recently, Gentry and Wichs [GW11] showed that some of the difficulty is indeed inherent by proving that no SNARG construction can be proved secure via a black-box reduction to an *efficiently falsifiable assumption* [Nao03]. For example, the assertion that one-way functions exist or that fully-homomorphic encryption exists are both falsifiable assumptions; in general, an assumption is efficiently falsifiable if it can be modeled as a game between an adversary and a challenger, where the challenger can *efficiently* decide whether the adversary has won the game. The impossibility result of Gentry and Wichs holds even for *designated-verifier* protocols, where the verifier needs secret randomness in order to verify. This suggests that non-standard assumptions, such as the knowledge (extractability) assumptions described next, may be inherent.

Knowledge (extractability) assumptions. Knowledge (or extractability) assumptions capture our belief that certain computational tasks can be achieved efficiently only by (essentially) going through specific intermediate stages and thereby obtaining, along the way, some specific intermediate values. Such an assumption asserts that, for any efficient algorithm that achieves the task, there exists a *knowledge extractor* algorithm that efficiently recovers the said intermediate values.

A number of different extractability assumptions exist in the literature, most of which are specific number theoretic assumptions (such as several variants of the *knowledge of exponent* assumption [Dam92]). It is indeed hard to gain assurance regarding their relative strengths. Abstracting from such specific assumptions, one can formulate general notions of extractability for one-way functions and other basic primitives (see [CD09, Dak09]). That is, say that a function family \mathcal{F} is *extractable* if, given a random $f \leftarrow \mathcal{F}$, it is infeasible to produce $y \in \text{Image}(f)$ without actually “knowing” x such that $f(x) = y$. This is expressed by saying that for any efficient adversary \mathcal{A} there is an efficient extractor $\mathcal{E}_{\mathcal{A}}$ such that, if $\mathcal{A}(f) = f(x)$ for some x , then $\mathcal{E}_{\mathcal{A}}(f)$ almost always outputs x' such that $f(x') = f(x)$. Typically, for such a family

¹The problem becomes evident when implementing the PIR using fully-homomorphic encryption; indeed, since any efficient adversarial strategy can be executed “under the encryption”, such a solution would be as insecure as sending the PCP queries in the clear.

to be interesting, \mathcal{F} should also have some sort of hardness property, e.g., one-wayness. Assuming that a certain function family is extractable does not typically fit in the mold of efficiently falsifiable assumptions. In particular, the impossibility result of Gentry and Wichs [GW11] does not apply to such assumptions.

1.1 Summary of Our Results

- We formulate a general and relatively-natural notion of *extractable collision-resistant hash functions* (ECRHs). We then show that if ECRHs exist then the modified construction is a SNARG for NP. Concretely, we revisit the PCP+MT+PIR approach of [DCL08], and show that it can be modified in a way that allows us to prove (adaptive) soundness when using an ECRH to construct the Merkle tree. Furthermore, we show that the construction is a *proof of knowledge*. We are thus able to obtain a *SNARG of knowledge* (SNARK) for NP.
- We propose a candidate ECRH construction, based on a Knowledge of Exponent assumption (that is a variant of the one found in [Gro10]) and the hardness of taking discrete logs. We also define *proximity ECRHs*, a weaker variant of ECRHs that are still sufficient for the construction of SNARKs; we propose two candidate PECRH constructions, based on knapsack (subset-sum) problems related to hard problems on lattices. Finally, we show how proximity ECRHs can be further weakened to obtain a (still sufficient) primitive for which it seems plausible to conjecture that “unstructured” hash functions such as (randomized) SHA-2 satisfy it.
- We show that existence of SNARKs for NP implies existence of (the weaker variants of) ECRHs, as well as extractable variants of some other cryptographic primitives; in other words, (proximity) ECRHs are not only sufficient but also necessary for the existence of SNARKs.
- We describe some applications of SNARKs: (a) to delegation of computations (even with long delegator input and with worker input), (b) to constructing *zero-knowledge* SNARKs, and (c) for constructing succinct non-interactive secure computation (in the common reference string model). See more below.
- We consider the related notion of extractable one-way functions (EOWFs). We formulate two natural variants of EOWFs and show that as ECRHs they help circumvent previous black-box impossibility results. Namely, we construct a two-message selective-opening-attack secure commitment scheme and a 3-round public-coin concurrent zero-knowledge (ZK) argument-of-knowledge, which in turn implies a 3-round simultaneously-resettable ZK protocol, (all assuming additionally the existence of enhanced trapdoor permutations). Previous works have shown that (under standard assumptions) these protocols can not be proven secure from standard assumptions using black-box reductions or simulation [Pas11, CKPR01].

The rest of the introduction describes these results in more detail. Before proceeding, though, we wish to highlight a prominent application of SNARGs that further motivates the need for proofs of knowledge.

Delegation of computation and adaptive arguments of knowledge. In the problem of delegation of computation, which has become ever more relevant with the advent of cloud computing, a client has some computational task (typically in P) and wishes to delegate the task to an untrusted worker, who responds with a claimed result along with a proof that the result is correct. Current delegation schemes, such as [KR06, GKR08, KR09, GGP10, CKV10], require either more than two messages of interaction between the client and the untrusted worker, or much work to be done by the verifier in a preprocessing stage.

A SNARG for NP could improve on these by minimizing both interaction and the verifier’s computational effort. (Note that adaptive soundness of the SNARG seems crucial for this application since a cheating worker might choose a bogus result of the computation based on the delegators first message.²)

However, the application to delegation schemes brings with it additional security concerns.

For example, the untrusted worker may store a long database z whose short Merkle hash $h = \text{MT}(z)$ is known to the delegator; the delegator may then ask the worker to compute $F(z)$ for some function F . However, from the delegator’s perspective, merely being convinced that “there exists \tilde{z} such that $h = \text{MT}(\tilde{z})$ and $F(\tilde{z}) = f$ ” is not enough. The delegator should also be convinced that the worker *knows* such a \tilde{z} , which implies due to collision resistance of MT that indeed $\tilde{z} = z$.

Thus, the delegator may not only be interested in establishing that a witness for a claimed theorem *exists*, but also want that such a witness can be *extracted* from a convincing prover. That is, we require *proof of knowledge* (or rather, an *argument of knowledge*) and thus SNARKs (rather than merely SNARGs) are needed. The ability to efficiently extract a witness for an adaptively-chosen theorem seems almost essential for making use of a delegation scheme when the untrusted worker is expected to contribute its own input (such as a database, as in the above example, or a signature, and so on) to a computation.

Another application where adaptive proofs of knowledge are crucial is *recursive proof composition*, a technique that has already been shown to enable desirable cryptographic tasks [Val08, CT10, BSW11, BCCT13].

1.2 ECRHs, SNARKs, and Applications

(i) Extractable collision-resistant hash functions. We start by defining a natural strengthening of collision-resistant hash functions (CRHs): a function ensemble $\mathcal{H} = \{\mathcal{H}_k\}_k$ is an extractable CRH (ECRH) if (a) it is collision-resistant in the standard sense, and (b) it is extractable in the sense that for any efficient adversary that is able to produce a valid evaluation of the function there is an extractor that is able to produce a corresponding preimage. More precisely, extractability is defined as follows:

Definition 1. A function ensemble $\mathcal{H} = \{\mathcal{H}_k\}_k$ mapping $\{0, 1\}^{\ell(k)}$ to $\{0, 1\}^k$ is extractable if for any polynomial-size adversary \mathcal{A} there exists a polynomial-size extractor $\mathcal{E}_{\mathcal{A}}$ such that for large enough security parameter $k \in \mathbb{N}$ and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$:

$$\Pr_{h \leftarrow \mathcal{H}_k} \left[\begin{array}{l} y \leftarrow \mathcal{A}(h, z) \\ \exists x : h(x) = y \end{array} \wedge \begin{array}{l} x' \leftarrow \mathcal{E}(h, z) \\ h(x') \neq y \end{array} \right] \leq \text{negl}(k) .$$

We do not require that there is an efficient way to tell whether a given string in $\{0, 1\}^k$ is in the image of a given $h \in \mathcal{H}_k$. We note that:

- For extractability and collision resistance (or one-wayness) to coexist, it should be hard to “obliviously sample images”; in particular, this implies that the image of almost any $h \in \mathcal{H}_k$ should be sparse in $\{0, 1\}^k$, i.e., with cardinality at most $2^{k-\omega(\log k)}$. (This remark is a bit over-simplified and not entirely accurate; see discussion in Section 6.1.)
- For simplicity of exposition, the above definition accounts for the most general case of arbitrary polynomial-size auxiliary-input; this is, in fact, too strong and cannot be achieved assuming indistinguishability obfuscation [BCPR14, BP13]. However, for our main result, we can actually settle for a

²More precisely, this seems to be the case when the worker contributes an input to the computation. In contrast, when the worker does not contribute any inputs, the delegator can use a SNARG with non-adaptive soundness by requesting two proofs for each bit of the claimed output.

relaxed definition that only considers a specific distribution over auxiliary inputs of a-priori bounded size. See further discussion in Section 6.1.

(ii) From ECRHs to adaptive succinct arguments of knowledge, and back again. We modify the “PCP+MT+PIR” construction of [DCL08] and show that the modified construction can be proven to be a SNARK based solely on the existence of ECRHs and PIR schemes with polylogarithmic complexity.³

Theorem 1 (informal). *If there exist ECRHs and (appropriate) PIRs then there exist SNARKs (for NP).*

A single VGRS in our construction suffices for only logarithmically many proofs; however, since the VGRS can be succinctly generated, the cost of occasionally resending a new one is limited.

We complement Theorem 1 by showing that ECRHs are in fact *essential* for SNARKs:

Theorem 2 (informal). *If there exist SNARKs and (standard) CRHs then there exist ECRHs.*

More accurately, we show that SNARKs and CRHs imply a slightly relaxed notion of ECRHs that we call *proximity ECRHs*, and which is still sufficient for our construction of SNARKs. To simplify the exposition of our main results we defer the discussion of the details of this relaxation to Section 1.3.

We also show that SNARKs can be used to construct extractable variants of other cryptographic primitives. A naïve strategy to obtain this may be to “add a succinct proof of knowledge of a preimage to the output”. While this strategy does not work as such because the proof may leak secret information, we show that in many cases this difficulty can be overcome by combining SNARKs with (non-extractable) *leakage-resilient* primitives. For example, since CRHs and subexponentially-hard OWFs are leakage-resilient, we obtain:

Theorem 3 (informal). *Assume SNARKs and (standard) CRHs exist. Then there exist extractable one-way functions and extractable computationally hiding and binding commitments. Alternatively, if there exist SNARKs and (standard) subexponentially-hard one-way functions then there exist extractable one-way functions. Furthermore, if these functions are one-to-one, then we can construct perfectly-binding computationally-hiding extractable commitments.*

We believe that this approach merits further investigation. One question, for example, is whether extractable pseudorandom generators and extractable pseudorandom functions can be constructed from generic extractable primitives (as was asked and left open in [CD09]). Seemingly, our SNARK-based approach can be used to obtain the weaker variants of extractable pseudo-entropy generators and pseudo-entropy functions, by relying on previous results regarding leakage-resilience of PRGs [DK08, RTTV08, GW11] and leakage-resilient pseudo-entropy functions [BHK11].

(iii) Applications of SNARKs. As discussed earlier, SNARKs directly enable non-interactive delegation of computation, including settings where the delegator has a very long input or where the worker supplies his own input to the computation. An important property of SNARK-based delegation is that it does not require expensive preprocessing and (as a result) soundness can be maintained even when the prover learns the verifier’s responses between successive delegation sessions because a fresh VGRS can simply be resent for each time.

In addition, SNARKs can be used to obtain zkSNARKs, that is, *zero-knowledge* succinct non-interactive arguments of knowledge in the common reference string (CRS) model. We provide two constructions to do so, depending on whether the succinct argument is “on top or below” the NIZK.

³More precisely, we shall require PIR schemes with polylogarithmic complexity where a fixed polynomial bound for the database size is not required by the query algorithm. See Section 1.4 for more details.

In an additional step, it is possible to obtain non-interactive two party secure computation (against malicious adversaries) between a powerful sender and a weak receiver. To do so, start with a non-interactive two-party computation protocol that is secure against “almost” honest-but-curious adversaries, who are allowed to choose arbitrary randomness; such protocols are known to exist, e.g., based on fully-homomorphic encryption [Gen09]. Then, to make the protocol resilient to malicious adversaries, let the sender attach to his message a zkSNARK that his computation was performed honestly. The succinct receiver can use a standard NIZK to prove knowledge of his inputs. As for the sender, if his input is short he can also prove knowledge using a standard NIZK (as also done in [DFH11]); otherwise, he must rely on the zkSNARK proof of knowledge. In particular, in the latter case, we only get non-concurrent security (rather than UC security) as the simulation relies on the non-black-box SNARK extractor.

In summary, SNARKs can be used for a number of applications:

Corollary 1.2 (informal). *If there exist SNARKs, then:*

1. *There exist non-interactive delegation schemes where the verifier’s response need not remain secret. (Furthermore, there are such schemes that allow the worker to contribute it own input, as well as ones allowing to delegate memory and data streams [CTY10, CKLR11].)*
2. *In the CRS model, there exist zero-knowledge SNARKs.*
3. *In the CRS model, there exist succinct non-interactive secure computation schemes (with UC security for short sender input and non-concurrent security for long sender input).*

For more details on the aforementioned applications see the respective sections Section 9, Section 10.1, and Section 10.2.

(iii) Applications of EOWFs. As demonstrated by the construction of SNARK from ECRHs, extractable functions can be used to circumvent impossibilities of constructing certain primitives using black-box security reductions. We explore this direction, and show that new feasibility results can be shown using EOWFs. Specifically, we first formalize a stronger variant of EOWFs—called strong EOWFs (sEOWF)—where every function in the ensemble is one-to-one and every sequence of functions is one-way. Using sEOWFs, we construct a two-message selective-opening-attack-secure (SOA-secure) commitment scheme. It was shown by Pass [Pas11] that such a commitment scheme cannot be proven secure via black-box security reductions from any assumptions that can be formulated as a game between a challenger and an adversary (a.k.a. falsifiable assumptions [Nao03]) of bounded rounds. In contrast, our scheme is proven secure using a non-black-box reduction from enhanced trapdoor permutations.

Furthermore, sEOWFs can be used to construct a 3-round zero-knowledge argument of knowledge (ZKAOK), and if the sEOWFs are “concurrently extractable” then the protocol is also concurrent ZK. Additionally, the 3-round protocol is public-coin; thus by applying standard transformations [BGGL01, DGS09], it yields a 3-round simultaneously resettable ZK protocol. It is known that these protocols do not admit black-box simulation [GK96, CKPR01]; our constructions indeed have non-black-box simulation utilizing the extractability of EOWFs.

For more details on the definitions and applications of (variants of) EOWFs see Section 11.

1.3 ECRH Candidate Constructions and Sufficient Relaxations

We propose a natural candidate ECRH based on a generalization of the knowledge of exponent assumption in large algebraic groups [Dam92]. The assumption, which we call *t-Knowledge-of-Exponent Assumption*

(or t -KEA for short) proceeds as follows. For any polynomial-size adversary, there exists a polynomial-size extractor such that, on input $g_1, \dots, g_t, g_1^\alpha, \dots, g_t^\alpha$ where each g_i is a random generator (of an appropriate group) and α is a random exponent: if the adversary outputs (f, f^α) , then the extractor finds a vector of “coefficients” (x_1, \dots, x_t) such that $f = \prod_{i \in [t]} g_i^{x_i}$. This assumption can be viewed as a simplified version of the assumption used by Groth in [Gro10] (the formal relation between the assumptions is discussed in Section 8.1). Similarly to Groth’s assumption, t -KEA holds in the *generic group model*.

Theorem 4 (informal). *If t -KEA holds in a group where taking discrete logs is hard, then there exists an ECRH whose compression is proportional to t .*

The construction is straightforward: the function family is parameterized by $(g_1, \dots, g_t, g_1^\alpha, \dots, g_t^\alpha)$. Given input (x_1, \dots, x_t) , the function outputs the two group elements $(\prod_{i \in [t]} g_i^{x_i}, \prod_{i \in [t]} g_i^{\alpha x_i})$. Extractability directly follows from t -KEA, while collision resistance is ensured by the hardness of taking discrete logs. See Section 8.1 for more details.

Next we proceed to propose ECRH candidates that are based on the subset sum problem in finite groups. Here, however, we are only able to construct candidates for somewhat weaker variants of ECRHs that are still sufficient for constructing SNARKs. While these variants are as not elegantly and concisely stated as the “vanilla” ECRH notion, they are still natural. Furthermore, we can show that these variants are necessary for SNARKs. We next proceed to formulate these weaker variants.

1.3.1 Proximity ECRH

We say that \mathcal{H} defined on domain D is a *proximity ECRH* (PECRH) if (for any $h \in \mathcal{H}$) there exist a reflexive “proximity” relation $\overset{h}{\approx}$ on values in the range and an extension of the hash to a larger domain $D_h \supseteq D$ fulfilling the following: (a) **proximity collision resistance**: given $h \leftarrow \mathcal{H}$, it is hard to find $x, x' \in D_h$ such that $h(x) \overset{h}{\approx} h(x')$, and (b) **proximity extraction**: for any poly-time adversary \mathcal{A} there exists an extractor \mathcal{E} such that, whenever \mathcal{A} outputs $y \in h(D)$, \mathcal{E} outputs $x \in D_h$ such that $h(x) \overset{h}{\approx} y$. (See Definition 6.2 for further details.)

Harder to find collisions, easier to extract. The notions of proximity extraction and proximity collision resistance are the same as standard extraction and collision resistance in the “strict” case, where $x \overset{h}{\approx} y$ is the equality relation and the domain is not extended ($D_h = \{0, 1\}^{\ell(k)}$, $\bar{h} = h$).

However, in general, proximity collision resistance is stronger than (standard) collision resistance, because even “near collisions” (i.e., $x \neq y$ such that $\bar{h}(x) \overset{h}{\approx} \bar{h}(y)$) must not be efficiently discoverable, not even over the extended domain D_h . Conversely, proximity extraction is weaker than (standard) extraction, since it suffices that the extractor finds a point mapping merely close to the adversary’s output (i.e., finds x' such that $\bar{h}(x') \overset{h}{\approx} y$); moreover, it suffices that the point is in the extended domain D_h . Thus, the notion of PECRH captures another, somewhat more flexible tradeoff between the requirements of extractability and collision resistance.

We show that any point on this tradeoff (i.e., any choice of $\overset{h}{\approx}$, D_h and \bar{h} fulfilling the conditions) suffices for the construction of SNARKs:

Theorem 5 (informal). *If there exist PECRHs and (appropriate) PIRs then there exist SNARKs for NP.*

Candidate PECRHs based on knapsack (subset sum) problems. A necessary property of ECRHs is that the image should be sparse; knapsack-based CRHs, which typically rely on a proper algebraic structure, can often be tweaked to obtain this essential property. For example, in the t -KEA-based ECRH that we

already discussed, we start from a standard knapsack hash $f = \prod_{i \in [t]} g_i^{x_i}$ and extend it to a “sparsified” knapsack hash (f, f^α) for a secret α . While for t -KEA this step is enough for plausibly assuming precise extractability (leading to a full fledged ECRH), for other knapsack-based CRHs this is not the case.

For example, let us consider the task of sparsifying modular subset-sum [Reg03]. Here, the hash function is given by random coefficients $l_1, \dots, l_t \in \mathbb{Z}_N$ and the hash of $x \in \{0, 1\}^t$ is simply the corresponding modular subset-sum $\sum_{i: x_i=1} l_i \bmod N$. A standard way to sparsify the function is, instead of drawing random coefficients, drawing them from a distribution of noisy multiples of some secret integer. However, by doing so, we lose the “algebraic structure” of the problem. Hence, now we also have to deal with new “oblivious image-sampling attacks” that exploit the noisy structure. For example, slightly perturbing an honestly computed subset-sum is likely to “hit” another image of the function. This is where the relaxed notion of proximity extraction comes into play: it allows the extractor to output the preimage of the nearby (honest) image and, more generally, to thwart “perturbation attacks”.

Sparsification of modular subset-sum in fact introduces additional problems. For instance, an attacker may take “small-norm” combinations of the coefficients that are not 0/1 and still obtain an element in the image (e.g., if there are two even coefficients); to account for this, we need to further relax the notion of extraction by allowing the extractor to output a preimage in an extended domain, while ensuring that (proximity) collision resistance still holds for the extended domain too. Additionally, in some cases a direct naïve sparsification is not sufficient and we also need to consider amplified knapsacks.

The relaxations of extractability discussed above have to be matched by a corresponding strengthening of collision resistance following the definition of PECRH. Fortunately, this can still be done under standard hardness assumptions.

A similar approach can be taken in order to sparsify the modular matrix subset-sum CRH [Ajt96, GGH96], resulting in a noisy inner-product knapsack hash based on the LWE assumption [Reg05]. Overall, we propose three candidate for PECRHs:

Theorem 6 (informal). *There exist PECRHs under any of the following assumptions:*

1. *A Knowledge of Knapsack of Exponent assumption (which in fact follows from t -KEA) and hardness of discrete logs.*
2. *A Knowledge of Knapsack of Noisy Multiples assumption and lattice assumptions.*
3. *A Knowledge of Knapsack of Noisy Inner Products assumption and learning with errors.*

1.3.2 Weak PECRHs⁴

Our second weakening is essentially orthogonal to the first one and relates to the condition that determines when the extractor has to “work”. The ECRH and PECRH definitions required extraction whenever the adversary outputs a valid image; here the sparseness of the image appears to be key. In particular, unstructured CRHs where one can sample elements in the image obliviously of their preimage have no hope to be either ECRH or PECRH. However, for our purposes it seems sufficient to only require the extractor to “work” when the adversary outputs an image y together with extra encoding of a preimage that can be verified given proper trapdoor information; oblivious image-sampling, on its own, is no longer sufficient for failing the extractor.

More formally, a family \mathcal{H} of functions is weakly extractable if for any efficient adversary \mathcal{A} there exists an efficient extractor $\mathcal{E}_{\mathcal{A}}^{\mathcal{H}}$ such that for any auxiliary input z and efficient decoder \mathcal{Y} , the probability of the

⁴This further weakening was inspired by private communication with Ivan Damgård.

event that \mathcal{A} outputs, given z and a random function h from the family, values y and e such that: (a) \mathcal{E}_A^H , given z and h , does not find a preimage of y , but (b) \mathcal{Y} , given e , does find a preimage of y , is negligible. See Definition 6.3 for further details.

We stress that the decoder circuit \mathcal{Y} is allowed to arbitrarily depend on the auxiliary input z . This is \mathcal{Y} 's advantage over the extractor \mathcal{E}_A^H (that must be efficient in z); in particular, the extractability condition is not trivially true. Another interpretation of the above definition is the following: The definition requires that there exists a single extractor \mathcal{E}_A^H that does as well as any other efficient extractor (that may depend on z); namely, any given decoder circuit \mathcal{Y} can be thought of as a candidate extractor and the extractor \mathcal{E}_A^H has to succeed whenever \mathcal{Y} does despite being efficient in the auxiliary input. In particular, whenever extraction is possible when given a trapdoor information related to the auxiliary input, it should also be possible without such trapdoor information. Indeed, the ability of the extractor to forgo the trapdoor is the key to a successful use of the above definition in our construction of SNARKs:

Theorem 7 (informal). *If there exist weak PECRHs and (appropriate) PIRs then there exist SNARKs for NP.*

Unlike ECRHs and PECRHs, weak ECRHs may in principle not require sparsity of the image or algebraic structure. For example, it may be plausible to assume that (a properly keyed) SHA-2 is a weak ECRHs.

We remark that the notion of weak ECRH is somewhat reminiscent of the notion of *preimage awareness* considered by Dodis et al. [DRS09]. Their notion, however, is set in an idealized model where the hash function can only be accessed as a black-box. When ported to such an ideal model, our definition can be viewed as a strengthening of the [DRS09] definition to account for auxiliary input. Similarly to the [DRS09] definition, the idealized version of our definition also holds in the (compressing) *random oracle model*. (We do not know if an idealized ECRH or PECRH can be constructed unconditionally in the random oracle model; however, assuming the existence of standard CRHs, so that a concise description for a CRH is available, it is possible to construct PECRHs in the random oracle model using CS proofs of knowledge.)

1.4 High-Level Proof Strategy for Theorem 1

In this section we provide some high level intuition for the proof of our main technical result: showing that the existence of ECRHs and (appropriate) PIR schemes imply the existence of SNARKs. At a very high-level, our construction follows the “PCP+MT+PIR approach” of [DCL08], which in turn is build upon Kilian’s 4-message succinct interactive argument system [Kil92]. So let’s start by reviewing them.

Kilian’s Construction: In [Kil92], Kilian presents a succinct special zero-knowledge argument for NP, (P_k, V_k) , where the prover P_k uses a Merkle tree in order to provide to V_k “virtual access” to a PCP proof. More precisely, to prove a statement $x \in L$, the verifier V_k starts by sending the prover a CRHF H . The prover on private input a witness w , constructs a PCP-proof π . In order to yield efficient verifiability, P_k cannot send V_k the witness w nor π . Rather, P_k builds a Merkle tree with the proof π as the leaf values (using the collision-free hash function H from the verifier) producing a root value rv . It then “commits” itself to π by sending rv to the verifier V_k . The verifier V_k tosses a fixed polynomial number of random coins r and sends them to the prover. Both the prover P_k and the verifier V_k computes the PCP queries by internally running the PCP verifier on input x and r . The prover P_k sends back answers to those queries, together with “proofs”—called authentication paths—that each answer is consistent with the root of the Merkle tree. Finally, the verifier accepts if all the answers are consistent with the root value rv , and convinces the PCP verifier.

Kilian’s protocol is succinct, because the verifier, invoking the PCP verifier, makes only a fixed polynomial number of queries and each query is answered with an authentication path of some fixed polynomial

length, all independent of the length of the witness. At a very high-level, the soundness follows from the fact that the Merkle tree provides the verifier “virtual access” to the PCP proof, in the sense that given the root value of the Merkle tree, for every query q , it is infeasible for a cheating prover to answer q differently depending on the queries. Therefore, interacting with the prover is “equivalent” to having access to a PCP proof oracle. Then it follows from the soundness of the PCP system that Kilian’s protocol is sound.

The “PCP+MT+PIR approach”: The work of [DCL08] proposed the “PCP+MT+PIR approach” to “squash” Kilian’s four-message protocol into a two-message protocol as follows. In Kilian’s protocol, the verifier obtains from the prover a Merkle hash to a PCP oracle and only then asks the prover to locally open the queries requested by the PCP verifier. In [DCL08]’s protocol, the verifier sends also in the first message, a PIR-encrypted version of the PCP queries (the first message of a PIR scheme can be viewed as an encryption to the queries); the prover then prepares the required PCP oracle, computes and sends a Merkle hash of it, and answers the verifier’s queries by replying to the PIR queries according to a database that contains the answer (as well as the authentication path with respect to the Merkle hash) to every possible verifier’s query.

[DCL08] proved the soundness of the above scheme based on the assumption that any convincing prover \mathcal{P}^* must essentially behave as an honest prover: Namely, if a proof is accepting, then the prover must have in mind a *full* PCP oracle π , which maps under the Merkle hash procedure to the claimed root, and such a proof π can be obtained by an efficient extractor $\mathcal{E}_{\mathcal{P}^*}$.⁵ [DCL08] then showed that, if this is the case, the extracted string π must be consistent with the answers the prover provides to the PCP queries, for otherwise the extractor can be used to obtain collisions of the hash function underlying the Merkle tree. Therefore, the extracted string π also passes the PCP test, where the queries are encrypted under PIR. Then, it follows from the privacy of the PIR scheme that, the string π is “computationally independent” of the query. Hence from the soundness of PCP, they conclude that the statement must be true.

1.4.1 The main challenges and our solutions.

Our goal is to obtain the stronger notion of SNARK, based on the more restricted assumption that ECRHs exist. At a very high-level, our construction follows the “PCP+MT+PIR approach” but replacing the CRH underlying the Merkle tree in their construction with an ECRH. Unlike [DCL08] which directly assumed the “global extraction” guarantees from a Merkle tree, we show that we can *lift* the “local extraction” guarantee provided by the ECRH, to the “global extraction” guarantee on the entire Merkle tree. More precisely, by relying on the (local) ECRH extraction guarantee, we show that for every convincing prover, there is an extractor that can efficiently extract a PCP proof $\tilde{\pi}$ that is “sufficiently satisfying” (i.e., a PCP verifier given $\tilde{\pi}$ accepts with high probability). Furthermore, to obtain the argument of knowledge property, we instantiate the underlying PCP system with PCPs of knowledge, which allows for extracting a witness from any sufficiently-satisfying proof oracle. (See details for the requisite PCP system in Section 3.7.) We next describe the high-level ideas for achieving global extraction guarantees from the local extraction of ECRHs. Full details are contained in Section 5, and the construction is summarized in Figure 1.

From local to global extraction. The main technical challenge lies in establishing a “global” knowledge feature (i.e., informally, extraction of a sufficiently satisfying proof $\tilde{\pi}$ from a convincing prover given that it outputs the root of a Merkle tree) from a very “local” one (namely, extraction of a preimage from a machine that outputs a valid hash value of the ECRH h). A natural attempt is to start from the root of the Merkle tree

⁵ Note that, as originally formulated, the assumption of [DCL08] seems to be false; indeed, a malicious prover can always start from a good PCP oracle π for a true statement and compute an “almost full” Merkle hash on π , skipping very few branches — so one should at least formulate an analogous but more plausible assumption by only requiring “sufficient consistency” with the claimed root.

and extract the values of the intermediate nodes of the Merkle tree layer by layer towards the leaves; that is, to extract a candidate proof $\tilde{\pi}$ by recursively applying the ECRH-extractor to extract the entire Merkle tree \widetilde{MT} , where the leaves should correspond to $\tilde{\pi}$.

However, recursively composing ECRH-extractors encounters a difficulty: each time applying the ECRH extraction incurs a polynomial blowup in extraction time. Therefore, (without making a very strong assumption on the amount of “blowup” incurred by the extractor,) we can only apply the ECRH extraction recursively a constant number of times; as a result, we can only extract from Merkle trees of a constant depth. We address this problem by opting to use a “squashed” Merkle tree, where the fan-in of each intermediate node is polynomial rather than binary as in the traditional case. Consequently, for every NP-statement, since its PCP proof is of polynomial length, the depth ℓ of the Merkle tree built over the PCP proof is a constant. Then, we can apply the ECRH extraction procedure recursively to extract the whole Merkle tree, while overall incurring polynomial overhead in the size of the extractor.

A technical problem that arises when applying the above recursive extraction procedure is that we might not be able to extract out a full Merkle tree (where all paths have the same length). More precisely, after applying the ECRH extraction recursively ℓ times, we obtain the values for the intermediate nodes up to level ℓ (thinking about the root as level zero). However, at each intermediate node, when applying the ECRH extractor, extraction could have failed to extract a preimage if the given intermediate node is not a valid hash image under the ECRH. Hence, the extracted tree might be at some points “disconnected”.⁶ Nevertheless, we show (relying solely on ECRH extraction) that the leaves in the extracted (perhaps partial) tree connected to the root are sufficiently satisfying for witness-extraction.

Proof at high level. Given the foregoing discussion, we show the correctness of the extraction procedure in two steps:

- *Step 1: “local consistency”.* We first show that whenever the verifier is convinced, the recursively extracted string $\tilde{\pi}$ contains valid answers to the verifier’s PCP queries specified in its PIR queries. Otherwise, it is possible to find collisions within the ECRH h as follows. A collision finder could simulate the PIR-encryption on its own, invoke both the extraction procedure and the prover, and obtain two paths that map to the same root but must differ somewhere (as one is satisfying and the other is not) and therefore obtain a collision.
- *Step 2: “from local to global consistency”.* Next, using the privacy guarantees of the PIR scheme, we show that, whenever we extract a set of leaves that are satisfying with respect to the PIR-encrypted queries, the same set of leaves must also be satisfying for almost all other possible PCP queries and is thus sufficient for witness-extraction. Indeed, if this was not the case then we would be able to use the polynomial-size extraction circuit to break the semantic security of the PIR.

For further details of the proof we refer the reader to Section 5.2.

Our construction ensures that the communication complexity and the verifier’s time complexity are bounded by a polynomial in the security parameter, the size of the instance, and the logarithm of the time it takes to verify a valid witness for the instance; furthermore, this polynomial is *independent* of the specific NP language at hand.

On Adaptive Soundness: The soundness requirement of SNARK prevents a cheating prover from proving any false statement even if it is allowed to choose the statement *after* seeing the verifier’s first message (or, rather, the verifier-generated reference string). In the above construction, the verifier sends in the first message the PIR-encrypted PCP queries. However, in general, the PCP queries may depend on the statement

⁶This captures for example the behavior of the prover violating the [DCL08] assumption described above.

to be proven, and hence limit the construction to be only statically sound. To resolve this problem, we modify the above protocol as follows: In the first message, the verifier PIR-encrypt the PCP verifier’s random coins rather than its actual queries (as the former are independent of the statement), and require the prover to prepare an appropriate database (containing all the possible answers for each setting of the coins, rather than a proof oracle). One subtle problem with this approach is that in general, the number of random coins tossed by the PIR verifier may depend on the size of the database (or a known upper bound on the size of the database). Therefore, to handle the case where no a-priori bound on the time of verifying the witness is given, and thus no a-priori bound on the size of the corresponding PCP oracle (which is the database), we require that the PIR supports random queries with respect to a-priori unknown database size.

On Universal Soundness: As for soundness, our main construction is not universal, in the sense that the verifier needs to know a constant c such that the verification time of an instance y does not exceed $|y|^c$. Fortunately, this very weak dependence on the specific NP language at hand (weak because it does not even depend on the Turing machine verifying witnesses) does not affect the application to delegation of computation, because the delegator, having in mind a specific poly-time task to delegate, *knows* c . Furthermore, we also show how, by assuming the existence of exponentially-hard one-way functions, our main construction can be extended to be a *universal* SNARK, that is, a single protocol that can simultaneously work with all NP languages.

1.5 Discussion

We conclude the introduction with an attempt to briefly motivate the premise of this work. Our main contribution can be seen as providing additional understanding of the security of a construction that has frustrated researchers. Towards this goal we prove strong security properties of the scheme based on a new cryptographic primitive, ECRHs, that, while not fitting into the mold of “standard cryptographic primitives or assumptions”, can be defined concisely and investigated separately.

Furthermore, we investigate a number of relaxations of ECRHs as well as a number of candidate constructions that have quite different underlying properties. Looking beyond the context of our particular protocol, this work can be seen as another step towards understanding the nature of extractability assumptions and their power in cryptography.

1.6 Organization

In Section 2, we discuss more related work. In Section 3, we give basic definitions for the cryptographic primitives that we use (along with any non-standard properties that we may need). In Section 4, we give the definitions for SNARKs. In Section 5, we give our main construction showing that ECRHs, along with appropriate PIRs, imply SNARKs. In Section 6, we discuss two relaxations of ECRHs, namely PECRHs and weak PECRHs, that still suffice for SNARKs. In Section 7, we explain how SNARKs imply proximity PECRHs, thus showing an almost tight relation between the existence of SNARKs and PECRHs. In Section 8, we propose candidate constructions for ECRHs and PECRHs. In Section 9, we show how to obtain zero-knowledge SNARKs. In Section 10, we provide further discussion for how SNARKs can be used in the setting of delegation of computation and secure computation. Finally, in Section 11 we define two notions of extractable one-way functions and show how to use them to construct, respectively, non-interactive (two-message) selective-opening-attack secure commitments and 3-round concurrent zero-knowledge protocols.

2 Other Related Work

Knowledge assumptions. A popular class of knowledge assumptions, which have been successfully used to solve a number of (at times notoriously open) cryptographic problems, is that of *Knowledge of Exponent* assumptions. These have the following flavor: if an efficient circuit, given the description of a finite group along with some other public information, computes a list of group elements that satisfies a certain algebraic relation, then there exists a knowledge extractor that outputs some related values that “explain” how the public information was put together to satisfy the relation. Most such assumptions have been proven secure against generic algorithms (see Nechaev [Nec94], Shoup [Sho97], and Dent [Den06]), thus offering some evidence for their truth. In the following we briefly survey prior works which, like ours, relied on Knowledge of Exponent assumptions.

Damgård [Dam92] first introduced a Knowledge of Exponent assumption to construct a CCA-secure encryption scheme. Later, Hada and Tanaka [HT98] showed how to use two Knowledge of Exponent assumptions to construct the first three-round zero-knowledge argument. Bellare and Palacio [BP04] then showed that one of the assumptions of [HT98] was likely to be false, and proposed a modified assumption, using which they constructed a three-round zero-knowledge argument.

More recently, Abe and Fehr [AF07] extended the assumption of [BP04] to construct the first perfect NIZK for NP with “full” adaptive soundness. Prabhakaran and Xue [PX09] constructed statistically-hiding sets for trapdoor DDH groups [DG06] using a new Knowledge of Exponent assumption. Gennaro et al. [GKR10] used another Knowledge of Exponent assumption (with an interactive flavor) to prove that a modified version of the Okamoto-Tanaka key-agreement protocol [OT89] satisfies perfect forward secrecy against fully active attackers.

In a different direction, Canetti and Dakdouk [CD08, CD09, Dak09] study *extractable functions*. Roughly, a function f is extractable if finding a value x in the image of f implies knowledge of a preimage of x . The motivation of Canetti and Dakdouk for introducing extractable functions is to capture the abstract essence of prior knowledge assumptions, and to formalize the “knowledge of query” property that is sometimes used in proofs in the Random Oracle Model. They also study which security reductions are “knowledge-preserving” (e.g., whether it possible to obtain extractable commitment schemes from extractable one-way functions).

[BCPR14, BP13] show that, assuming indistinguishability obfuscation [BGI⁺12], extractable one-way functions (and thus also ECRHs) cannot be constructed against adversaries with arbitrary polynomial-size auxiliary-input if the (efficient) extractor is universally fixed before the adversary’s auxiliary input. On the other hand, they show that, under standard assumptions, extractable one-way functions are achievable against adversaries with a-priori bounded auxiliary input. (It is still not known whether such ECRHs can also be constructed under standard assumptions).

Prior (somewhat) succinct arguments from Knowledge of Exponent assumptions. Knowledge of Exponent assumptions have been used to obtain somewhat succinct arguments, in the sense the non-interactive proof is short, but the verifier’s running time is long.

Recently, Groth [Gro10] introduced a family of knowledge of exponent assumptions, generalizing those of [AF07], and used them to construct extractable length-reducing commitments, as a building block for short non-interactive perfect zero-knowledge arguments system for circuit satisfiability. These arguments have very succinct proofs (independent of the circuit size), though the public key is large: quadratic in the size of the circuit. Groth’s assumption holds in the generic group model. For a comparison between our t -KEA assumption and Groth’s assumptions see Section 8.1.

Mie [Mie08] observes that the PCP+MT+PIR approach works as long as the PIR scheme is *database aware* — essentially, a prover that is able to provide valid answers to PIR queries must “know” their de-

encrypted values, or, equivalently, must “know” a database consistent with those answers (by arbitrarily setting the rest of the database). Mie then shows how to make the PIR scheme of Gentry and Ramzan [GR05] PIR-aware, based on Damgård’s Knowledge of Exponent assumption [Dam92]; unfortunately, while the communication complexity is very low, the receiver in [GR05], and thus also the verifier in [Mie08], are inefficient relative to the database size. We note that PIR schemes with database awareness can be constructed directly from ECRHs (without going through the PCPs of a SNARK construction); moreover, if one is willing to use PCPs to obtain a SNARK, one would then be able to obtain various stronger notions of database awareness.

Delegation of computation. An important application of succinct arguments is *delegation of computation* schemes, where one usually also cares about privacy, and not only soundness, guarantees. Specifically, a succinct argument can be usually combined in a trivial way with fully-homomorphic encryption [Gen09] (in order to ensure privacy) to obtain a delegation scheme where the delegator runs in time polylogarithmic in the running time of the computation (see Section 10.1).

Within the setting of delegation, however, where the same weak delegator may be asking a powerful untrusted worker to evaluate an expensive function on many different inputs, a weaker preprocessing approach may still be meaningful. In such a setting, the delegator performs a one-time function-specific expensive setup phase, followed by inexpensive input-specific delegations to amortize the initial expensive phase. Indeed, in the preprocessing setting a number of prior works have already achieved constructions where the online stage is only two messages [GGP10, CKV10, AIK10]. These constructions do not allow for an untrusted worker to contribute his own input to the computation, namely they are “P-delegation schemes” rather than “NP-delegation schemes”. Note that all of these works do not rely on any knowledge assumption; indeed, the impossibility results of [GW11] only apply for NP and not for P.

However, even given that the preprocessing model is very strong, all of the mentioned works maintain soundness over many delegations only as long as the verifier’s answers remain secret. (A notable exception is the work of Benabbas et al. [BGV11], though their constructions are not generic, and are only for specific functionalities such as polynomial functions.)

Goldwasser et al. [GKR08] construct interactive proofs for log-space uniform NC where the verifier running time is quasi-linear. When combining [GKR08] with the PIR-based squashing technique of Kalai and Raz [KR06], one can obtain a succinct two-message delegation scheme. Canetti et al. [CRR11] introduce an alternative way of squashing [GKR08], in the preprocessing setting; their scheme is of the public coin type and hence the verifier’s answers need not remain secret (another bonus is that the preprocessing state is publicly verifiable and can thus be used by anyone).

3 Preliminaries

In this section we give basic definitions for the cryptographic primitives that we use (along with any non-standard properties that we may need). Throughout, $\text{negl}(k)$ is any negligible function in k .

3.1 Interactive Proofs, Zero Knowledge and Witness Indistinguishability

We use the standard definitions of interactive proofs (and interactive Turing machines) [GMR89] and arguments (also known as computationally-sound proofs) [BCC88]. Given a pair of interactive Turing machines, P and V , we denote by $\langle P(w), V \rangle(x)$ the random variable representing the final (local) output of V , on common input x , when interacting with machine P with private input w , when the random input to each machine is uniformly and independently chosen.

Definition 3.1 (Interactive Proof System). A pair (P, V) of interactive machines is called an interactive proof system for a language L if there is a negligible function $\nu(\cdot)$ such that the following two conditions hold :

- Completeness: For every $x \in L$ and $w \in R_L(x)$, $\Pr[\langle P(w), V \rangle(x) = 1] = 1$.
- Soundness: For every $x \in \{0, 1\}^n \setminus L$, and every interactive machine B , $\Pr[\langle B, V \rangle(x) = 1] \leq \nu(n)$

In case that the soundness condition is required to hold only with respect to a computationally bounded prover, the pair $\langle P, V \rangle$ is called an interactive argument system.

Zero-Knowledge. An interactive proof is said to be *zero-knowledge* (ZK) if a verifier V learns nothing beyond the validity of the assertion being proved. Because “feasible computation” is typically captured through the notion of probabilistic polynomial-time, the ZK property is formalized by requiring that the output of every (possibly malicious) verifier interacting with the honest prover P can be *simulated* by a probabilistic (expected) polynomial-time machine S , known as the the *simulator*. The idea behind this definition is that the anyone can learn by himself, by running the simulator S , whatever V^* learns by interacting with P .

The notion of ZK was introduced and formalized by Goldwasser, Micali and Rackoff in [GMR89]. We present their definition below.

Definition 3.2 (ZK). Let L be a language in NP, R_L a witness relation for L , (P, V) an interactive proof (argument) system for L . We say that (P, V) is statistical/computational ZK, if for every probabilistic polynomial-time interactive machine V there exists a probabilistic algorithm S whose expected running-time is polynomial in the length of its first input, such that the following ensembles are statistically close/computationally indistinguishable over L .

- $\left\{ \text{View}_V[\langle P(y), V(z) \rangle(x)] \right\}_{n \in \mathbb{N}, x \in \{0, 1\}^n \cap L, y \in R_L(x), z \in \{0, 1\}^*}$
- $\left\{ S(x, z) \right\}_{n \in \mathbb{N}, x \in \{0, 1\}^n \cap L, y \in R_L(x), z \in \{0, 1\}^*}$

where $\text{View}[\langle P(y), V(z) \rangle(x)]$ denote the random variable describing the view of V in interaction with P on common input x and private inputs y and z respectively.

Witness-Indistinguishable Proofs. An interactive proof is said to be *witness indistinguishable* (WI) if the verifier’s output is “computationally independent” of the witness used by the prover for proving the statement. In this context, we focus on languages $L \in \mathcal{NP}$ with a corresponding witness relation R_L . Namely, we consider interactions in which, on common input x , the prover is given a witness in $R_L(x)$. By saying that the output is computationally independent of the witness, we mean that for any two possible NP-witnesses that could be used by the prover to prove the statement $x \in L$, the corresponding outputs are computationally indistinguishable.

Definition 3.3 (Witness indistinguishability). Let (P, V) be an interactive proof system for a language $L \in \mathcal{NP}$. We say that (P, V) is witness-indistinguishable for R_L , if for every probabilistic polynomial-time interactive machine V^* and for every two sequences $\{w_{n,x}^1\}_{n \in \mathbb{N}, x \in L}$ and $\{w_{n,x}^2\}_{n \in \mathbb{N}, x \in L}$, such that $w_{n,x}^1, w_{n,x}^2 \in R_L(x)$ for every $x \in L \cap \{0, 1\}^n$, the following probability ensembles are computationally indistinguishable over $n \in \mathbb{N}$.

- $\left\{ \text{View}_V[\langle P(w_{n,x}^1), V^*(z) \rangle(x)] \right\}_{n \in \mathbb{N}, x \in L \cap \{0, 1\}^n, z \in \{0, 1\}^*}$
- $\left\{ \text{View}_V[\langle P(w_{n,x}^2), V^*(z) \rangle(x)] \right\}_{n \in \mathbb{N}, x \in L \cap \{0, 1\}^n, z \in \{0, 1\}^*}$

3.2 Proofs and Arguments of Knowledge

Given a language $L \in \text{NP}$ and an instance x , a *proof or argument of knowledge* (POK or AOK) not only convinces the verifier that $x \in L$, but also to demonstrate that the prover possesses an NP-witness for x . This is formalized by the existence of an *extractor*: given black-box access to a machine that can successfully complete the proof or argument of knowledge on input x , the extractor can compute a witness for x .

Definition 3.4 (Proofs and arguments of knowledge). *An interactive protocol $\Pi = (P, V)$ is a proof of knowledge (resp. argument of knowledge) of NP-language L with respect to witness relation R_L if Π is indeed an interactive proof (resp. argument) for L . Additionally, there exists a polynomial q , a negligible function ν , and a PPT oracle machine E , such that for every interactive machine P^* (resp. for every polynomially-sized machine P^*), every $x \in L$ and every auxiliary input $z \in \{0, 1\}^*$, the following holds: On input x and oracle access to $P^*(x, z)$, machine E outputs a string from the $R_L(x)$ with probability at least $q(\Pr[\langle P^*(z), V \rangle(x) = 1]) - \nu(|x|)$. The machine E is called the knowledge extractor.*

3.3 Commitment Scheme

A commitment scheme (C, R) consists of a pair of PPT ITMs C and R that interact in a commit stage and a reveal stage. In this work, we consider commitment schemes that are computationally binding and computationally hiding. The computational-binding property asserts that the probability that a malicious commitment, after the interaction with an honest receiver, can decommit to two different values is negligible. The computational-hiding property guarantees that commitments to any two different values are computationally indistinguishable. (See [Gol01] for a formal definition.) Furthermore, we restrict our attention to commitment schemes where the reveal phase is non-interactive—the committer decommits to value v by simply sending a decommitment pair (v, d) . One-message statistically-binding commitment schemes can be constructed using any one-to-one one-way function (see Section 4.4.1 of [Gol01]). Allowing some minimal interaction (in which the receiver first sends a single random initialization message), statistically-binding commitment schemes can be obtained from any one-way function [Nao91, HILL99].

3.4 Collision-Resistant Hashes

A *collision-resistant hash* (CRH) is a function ensemble for which it is hard to find two inputs that map to the same output. Formally:

Definition 3.5. *A function ensemble \mathcal{H} is a CRH if it is collision-resistant in the following sense: for every polynomial-size adversary \mathcal{A} ,*

$$\Pr_{h \leftarrow \mathcal{H}_k} \left[\begin{array}{l} x \neq y \\ h(x) = h(y) \end{array} : (x, y) \leftarrow \mathcal{A}(h) \right] \leq \text{negl}(k) .$$

We say that a function ensemble \mathcal{H} is $(\ell(k), k)$ -*compressing* if each $h \in \mathcal{H}_k$ maps strings of length $\ell(k)$ to strings of length $k < \ell(k)$.

3.5 Merkle Trees

Merkle tree (MT) hashing [Mer89] enables a party to use a CRH to compute a succinct commitment to a long string π and later to locally open to any bit of π (again in a succinct manner). Specifically, given a function $h: \{0, 1\}^{\ell(k)} \rightarrow \{0, 1\}^k$ randomly drawn from a CRH ensemble, the committer divides π into

$|\pi|/\ell(k)$ parts (padding with 0's if needed) and evaluates h on each of these; the same operation is applied to the resulting string, and so on, until one reaches the single k -bit root. For $|\pi| = (\ell/k)^{d+1}$, this results in a tree of depth d , whose nodes are all the intermediate k -bit hash images. An opening to a leaf in π (or any bit within it) includes all the nodes and their siblings along the path from the root to the leaf, and is of size ℓd . Typically, $\ell(k) = 2k$, resulting in a binary tree of depth $\log \pi$. In this work, we shall also be interested in “wide trees” with polynomial fan-in (relying on CRHs with polynomial compression); see Section 5.1.

3.6 Private Information Retrieval

A (single-server) polylogarithmic *private information retrieval* (PIR) scheme [CMS99] consists of a triple of algorithms (PEnc, PEval, PDec) where:

- $\text{PEnc}_R(1^k, i)$ outputs an encryption C_i of query i to a database DB using randomness R ,
- $\text{PEval}(\text{DB}, C_i)$ outputs a succinct blob e_i “containing” the answer $\text{DB}[i]$, and
- $\text{PDec}_R(e_i)$ decrypts the blob e_i to an answer $\text{DB}[i]$.

Formally:

Definition 3.6. *A triple of algorithms (PEnc, PEval, PDec) is a PIR if it has the following properties:*

1. **Correctness.** *For any database DB, any query $i \in \{1, \dots, |\text{DB}|\}$, and security parameter $k \in \mathbb{N}$,*

$$\Pr_R \left[\text{PDec}_R(e_i) = \text{DB}[i] : \begin{array}{l} C_i \leftarrow \text{PEnc}_R(1^k, i) \\ e_i \leftarrow \text{PEval}(\text{DB}, C_i) \end{array} \right] = 1 ,$$

where $\text{PEval}(\text{DB}, C_i)$ runs in $\text{poly}(k, |\text{DB}|)$ time.

2. **Succinctness.** *The running time of both $\text{PEnc}_R(1^k, i)$ and $\text{PEval}(\text{DB}, C_i)$ is bounded by*

$$\text{poly}(k, \log |\text{DB}|) .$$

In particular, the sizes of the two messages C_i and e_i are also so bounded.

3. **Semantic security.** *The query encryption is semantically secure for multiple queries, i.e., for any polynomial-size \mathcal{A} , all large enough security parameter $k \in \mathbb{N}$ and any two tuples of queries $\mathbf{i} = (i_1 \dots i_q), \mathbf{i}' = (i'_1 \dots i'_q) \in \{0, 1\}^{\text{poly}(k)}$,*

$$\Pr \left[\mathcal{A}(\text{PEnc}_R(1^k, \mathbf{i})) = 1 \right] - \Pr \left[\mathcal{A}(\text{PEnc}_R(1^k, \mathbf{i}')) = 1 \right] \leq \text{negl}(k) ,$$

where $\text{PEnc}_R(1^k, \mathbf{i})$ the coordinate-wise encryption the tuple \mathbf{i} .

PIR schemes with the above properties have been constructed under various hardness assumptions such as Φ HA [CMS99] or LWE [BV11].

A-priori unknown database size. In certain cases, we want the server to be able to specify the database DB only *after* receiving the query. In such cases, the client might not know the database's size when generating his query, but will only know some superpolynomial bound, such as $|\text{DB}| \leq 2^{\log^2 k}$. In this case we require that the PIR scheme allows the server to deduce from an encrypted (long) query $r \in \{0, 1\}^{\log^2 k}$ a shorter ρ -bit query $\hat{r} \in \{0, 1\}^\rho$ that works for $|\text{DB}| = 2^\rho$. In FHE-based schemes, such as the one in [BV11] (which is in turn based on LWE), this extra property can be easily supported. In other cases (as in delegation of computation), even if an adversary is adaptive, an a-priori bound on the database size is still available; whenever this is the case, then no additional properties are required of the PIR scheme.

3.7 Probabilistically Checkable Proofs of Knowledge

A (verifier-efficient) *probabilistically checkable proof* (PCP) of knowledge for the universal relation $\mathcal{R}_{\mathcal{U}}$ is a triple of algorithms $(P_{\text{pcp}}, V_{\text{pcp}}, E_{\text{pcp}})$, where P_{pcp} is the prover, V_{pcp} is the (randomized) verifier, and E_{pcp} is the knowledge extractor.

Given $(y, w) \in \mathcal{R}_{\mathcal{U}}$, $P_{\text{pcp}}(y, w)$ generates a proof π of length $\text{poly}(t)$ and runs in time $\text{poly}(|y|, t)$. The verifier $V_{\text{pcp}}^\pi(y, r)$ runs in time $\text{poly}(|y|) = \text{poly}(|M| + |x| + \log t)$, while accessing $\text{polylog}(t)$ locations in the proof π and using $\rho = O(\log t)$ random bits. We require:

1. **Completeness.** For every $(y, w) = ((M, x, t), w) \in \mathcal{R}_{\mathcal{U}}$, $\pi \leftarrow P_{\text{pcp}}(y, w)$:

$$\Pr_{r \leftarrow \{0,1\}^{\rho(t)}} [V_{\text{pcp}}^\pi(y, r) = 1] = 1 .$$

2. **Proof of knowledge (PoK).** There is a constant ε such that for any $y = (M, x, t)$ if

$$\Pr_{r \leftarrow \{0,1\}^{\rho(t)}} [V_{\text{pcp}}^\pi(y, r) = 1] \geq 1 - \varepsilon ,$$

then $E_{\text{pcp}}(y, \pi)$ outputs a witness w such that $(y, w) \in \mathcal{R}_{\mathcal{U}}$, and runs in time $\text{poly}(|y|, t)$.

(Note that proof of knowledge in particular implies that the soundness error is at most ε .)

PCPs of knowledge as defined above can be based on the efficient-verifier PCPs of [BSS08, BSGH⁺05]. (See [Val08] for a simple example of how a PCP of proximity can yield a PCP with a proof of knowledge.) Moreover, the latter PCPs' proof length is quasi-linear in t ; for simplicity, we shall settle for a bound of t^2 .

We shall typically invoke the verifier V_{pcp} for $q(k)$ times to reduce the proof-of-knowledge threshold to $(1 - \varepsilon)^{q(k)}$, where k is the security parameter and $q(k) = \omega(\log k)$. Namely, extraction should succeed whenever $\Pr_{\mathbf{r}} [V_{\text{pcp}}^\pi(y, \mathbf{r}) = 1] \geq (1 - \varepsilon)^q$, where $\mathbf{r} = (r_i)_{i \in [q]}$ and $V_{\text{pcp}}^\pi(y, \mathbf{r}) = \bigwedge_{i \in [q]} V_{\text{pcp}}^\pi(y, r_i)$.

4 SNARKS

In this section we formally introduce the main cryptographic primitive studied in this paper — the SNARK.

4.1 The Universal Relation and NP Relations

The *universal relation* [BG08] is defined to be the set $\mathcal{R}_{\mathcal{U}}$ of instance-witness pairs (y, w) , where $y = (M, x, t)$, $|w| \leq t$, and M is a Turing machine, such that M accepts (x, w) after at most t steps. While the witness w for each instance $y = (M, x, t)$ is of size at most t , there is *no a-priori* polynomial bounding t in terms of $|x|$.

Also, for any $c \in \mathbb{N}$, we denote by \mathcal{R}_c the subset of $\mathcal{R}_{\mathcal{U}}$ consisting of those pairs $(y, w) = ((M, x, t), w)$ for which $t \leq |x|^c$. This is a “generalized” NP relation, where we do not insist on the same Turing machine accepting different instances, but only insist on a fixed polynomial bounding the running time in terms of the instance size.

4.2 Succinct Non-Interactive Arguments

A *succinct non-interactive argument* (SNARG) is a triple of algorithms $(\mathcal{P}, \mathcal{G}_V, \mathcal{V})$. For a security parameter k , the verifier runs $\mathcal{G}_V(1^k)$ to generate $(\text{vgrs}, \text{priv})$, where vgrs is a (public) verifier-generated reference string and priv are corresponding private verification coins; the honest prover $\mathcal{P}(y, w, \text{vgrs})$ produces a proof Π for the statement $y = (M, x, t)$ given a witness w ; then $\mathcal{V}(\text{priv}, y, \Pi)$ verifies the validity of Π . The SNARG is *adaptive* if the prover may choose the statement *after* seeing the vgrs , otherwise, it is *non-adaptive*.

Definition 4.1. A triple of algorithms $(\mathcal{P}, \mathcal{G}_V, \mathcal{V})$ is a SNARG for the relation $\mathcal{R} \subseteq \mathcal{R}_U$ if the following conditions are satisfied:

1. **Completeness.** For any $(y, w) \in \mathcal{R}$,

$$\Pr \left[\mathcal{V}(\text{priv}, y, \Pi) = 1 : \begin{array}{l} (\text{vgrs}, \text{priv}) \leftarrow \mathcal{G}_V(1^k) \\ \Pi \leftarrow \mathcal{P}(y, w, \text{vgrs}) \end{array} \right] = 1 .$$

In addition, $\mathcal{P}(y, w, \text{vgrs})$ runs in time $\text{poly}(k, |y|, t)$.

2. **Succinctness.** The length of the proof Π that $\mathcal{P}(y, w, \text{vgrs})$ outputs, as well as the running time of $\mathcal{V}(\text{priv}, y, \Pi)$, is bounded by

$$p(k + |y|) = p(k + |M| + |x| + \log t) ,$$

where p is a universal polynomial that does not depend on \mathcal{R} . In addition, $\mathcal{G}_V(1^k)$ runs in time $\text{poly}(k)$; in particular, $(\text{vgrs}, \text{priv})$ are of length $\text{poly}(k)$.

3. **Soundness.** Depending on the notion of adaptivity:

- **Non-adaptive soundness.** For all polynomial-size prover \mathcal{P}^* , large enough $k \in \mathbb{N}$, and $y \notin \mathcal{L}_{\mathcal{R}}$,

$$\Pr \left[\mathcal{V}(\text{priv}, y, \Pi) = 1 : \begin{array}{l} (\text{vgrs}, \text{priv}) \leftarrow \mathcal{G}_V(1^k) \\ \Pi \leftarrow \mathcal{P}^*(y, \text{vgrs}) \end{array} \right] \leq \text{negl}(k) .$$

- **Adaptive soundness.** For all polynomial-size prover \mathcal{P}^* and large enough $k \in \mathbb{N}$,

$$\Pr \left[\mathcal{V}(\text{priv}, y, \Pi) = 1 : \begin{array}{l} (\text{vgrs}, \text{priv}) \leftarrow \mathcal{G}_V(1^k) \\ (y, \Pi) \leftarrow \mathcal{P}^*(\text{vgrs}) \\ y \notin \mathcal{L}_{\mathcal{R}} \end{array} \right] \leq \text{negl}(k) .$$

A *SNARG of knowledge*, or *SNARK* for short, is a SNARG where soundness is strengthened as follows:

Definition 4.2. A triple of algorithms $(\mathcal{P}, \mathcal{G}_V, \mathcal{V})$ is a SNARK if it is a SNARG where adaptive soundness is replaced by the following stronger requirement:

- **Adaptive proof of knowledge.**⁷ For any polynomial-size prover \mathcal{P}^* there exists a polynomial-size extractor $\mathcal{E}_{\mathcal{P}^*}$ such that for all large enough $k \in \mathbb{N}$ and all auxiliary inputs $z \in \{0, 1\}^{\text{poly}(k)}$,

$$\Pr \left[\begin{array}{l} (\text{vgrs}, \text{priv}) \leftarrow \mathcal{G}_V(1^k) \\ (y, \Pi) \leftarrow \mathcal{P}^*(z, \text{vgrs}) \\ \mathcal{V}(\text{priv}, y, \Pi) = 1 \end{array} \wedge \begin{array}{l} (y, w) \leftarrow \mathcal{E}_{\mathcal{P}^*}(z, \text{vgrs}) \\ w \notin \mathcal{R}(y) \end{array} \right] \leq \text{negl}(k) .$$

⁷One can also formulate weaker proof-of-knowledge notions; in this work we focus on the above strong notion.

Universal arguments vs. weaker notions. A SNARG for the relation $\mathcal{R} = \mathcal{R}_{\mathcal{U}}$ is called a *universal argument*.⁸ A weaker notion that we will also consider is a SNARG for the relation $\mathcal{R} = \mathcal{R}_c$ for a constant $c \in \mathbb{N}$. In this case, soundness is only required to hold with respect to \mathcal{R}_c ; in particular, the verifier algorithm may depend on c . Nevertheless, we require (and achieve) *universal succinctness*, where a universal polynomial p , independent of c , upper bounds the length of every proof and the verification time.

Designated verifiers vs. public verification. In a *publicly-verifiable* SNARG the verifier does not require a private state priv . In this work, however, we concentrate on *designated verifier* SNARGs, where priv must remain secret for soundness to hold.

The verifier-generated reference string. A very desirable property is the ability to generate the verifier-generated reference string vgrs once and for all and then reuse it in polynomially-many proofs (potentially by different provers). In publicly verifiable SNARGs, this *multi-theorem soundness* is automatically guaranteed; in designated verifier SNARGs, however, multi-theorem soundness needs to be required explicitly as an additional property. Usually, this is achieved by ensuring that the verifier’s response “leaks” only a negligible amount of information about priv . (Note that $O(\log k)$ -theorem soundness always holds; the “non-trivial” case is whenever $\omega(\log k)$. Weaker solutions to support more theorems include assuming that the verifier’s responses remain secret, or re-generating vgrs every logarithmically-many rejections, e.g., as in [KR06, Mie08, GKR08, KR09, GGP10, CKV10].)

The SNARK extractor \mathcal{E} and auxiliary input. Above, we require that any polynomial-size family of circuits \mathcal{P}^* has a specific polynomial-size family of extractors $\mathcal{E}_{\mathcal{P}^*}$; in particular, we allow the extractor to be of arbitrary polynomial-size and to be “more non-uniform” than \mathcal{P}^* . In addition, we require that for any prover auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$, the polynomial-size extractor manages to perform its witness-extraction task given the same auxiliary input z . As shown in [BCPR14, BP13], this aspect of the definition is too strong assuming indistinguishability obfuscation [BGI⁺12]. The definition can be naturally relaxed to consider only specific distributions of auxiliary inputs according to the required application. (In our setting, the restrictions on the auxiliary-input handled by the knowledge extractor will be related to the auxiliary-input that the underlying ECRH extractor can handle. See further discussion in Section 6.1.)

5 From ECRHs to SNARKs

In this section we describe and analyze our construction of adaptive SNARKs for NP based on ECRHs. (Recall that an ECRH is a CRH as in Definition 3.5 that is extractable as in Definition 1).

In Section 5.3 we discuss the universal features of our constructions, and the difficulties in extending it to a full-fledged universal argument; we propose a solution that can overcome the difficulties based on exponentially hard one-way functions.

Our modified approach. We modify the PCP+MT+PIR approach of [DCL08] and show that the knowledge assumption of [DCL08] (which involves the entire PIR+MT structure) can be replaced by the simpler generic assumption that ECRHs exist. Furthermore, our modification enables us to improve the result from a two-message succinct argument with non-adaptive soundness to an adaptive SNARG of knowledge (SNARK) — this improvement is crucial cryptographic applications. Specifically, we perform two modifications:

1. We instantiate the Merkle tree hash using an ECRH and, unlike the traditional construction where a

⁸Barak and Goldreich [BG08] define universal arguments for \mathcal{R} with a black-box “weak proof-of-knowledge” property; in contrast, our proof of knowledge property is not restricted to black-box extractors, and does not require the extractor to be an implicit representation of a witness.

$(2k, k)$ -CRH is used, we use a polynomially-compressing (k^2, k) -ECRH; in particular, for k^{d+1} -long strings the resulting tree will be of depth d (rather than $d \log k$).⁹ As we shall see later, doing so allows us to avoid superpolynomial blowup of the final knowledge extractor that will be built via composition of ECRH extractors. The initial construction we present will be specialized for “generalized” NP-relations \mathcal{R}_c ; after presenting and analyzing it, we propose an extension to the universal relation $\mathcal{R}_{\mathcal{U}}$ by further assuming the existence of exponentially-hard one-way functions.

2. In order to ensure that the first message of the verifier does not depend on the theorem being proved, the database that we use does not consist of (authenticated) bits of π ; rather, the r -th entry of the database corresponds to the authenticated answers to the queries of $V_{\text{pcp}}^\pi(y, r)$ where y is chosen by the prover and, of course, the authentication is relative to a single string π (to avoid cheating provers claiming one value for a particular location of π in one entry of the database, and another value for the same location of π in another entry of the database). An additional requirement for this part is the use of a PIR scheme that can support databases where the exact size is a-priori unknown (and only a superpolynomial bound is known).

5.1 Construction Details

We start by providing a short description of our MT and then present the detailed construction of the protocol in Figure 1.

The shallow Merkle tree. By padding when required, we assume without loss of generality that the compressed string π is of size k^{d+1} (where d is typically unknown to the verifier). A node in the MT of distance j from the root can be represented by a string $\mathbf{i} = i_1 \cdots i_j \in [k]^j$ containing the path traversal indices (and the root is represented by the empty string). We then label the nodes with k -bit strings according to π and the hash $h : \{0, 1\}^{k^2} \rightarrow \{0, 1\}^k$ as follows:

- The leaf associated with $\mathbf{i} = i_1 \cdots i_d \in [k]^d \cong [k^d]$ is labeled by the \mathbf{i} -th k -bit block of π , denoted by $\ell_{\mathbf{i}}$ (here \mathbf{i} is interpreted as number in $[k^d]$).
- An internal node associated with $\mathbf{i} = i_1 \cdots i_j \in [k]^j$ is labeled by $h(\ell_{i_1} \ell_{i_2} \cdots \ell_{i_k})$, denoted by $\ell_{\mathbf{i}}$.
- Thus, the label of the root is $h(\ell_1 \ell_2 \cdots \ell_k)$, which we denote by ℓ_ϵ .

The MT commitment is the pair (d, ℓ_ϵ) . An opening $\text{dcom}_{\mathbf{i}}$ to a leaf \mathbf{i} consists of all the labels of all the nodes and their siblings along the corresponding path. To verify the opening information, evaluate the hash h from the leaves upwards. Specifically, for each node $\mathbf{i}' = \mathbf{i}j$ along the opening path labeled by $\ell_{\mathbf{i}'} = \ell_{\mathbf{i}j}$ and his siblings' labels $\ell_{i_1}, \ell_{i_2}, \dots, \ell_{i_{(j-1)}}, \ell_{i_{(j+1)}}, \dots, \ell_{i_k}$, verify that $h(\ell_{i_1}, \dots, \ell_{i_k}) = \ell_{\mathbf{i}}$.

We shall prove the following theorem:

Theorem 5.1. *For any NP-relation \mathcal{R}_c , the construction in Figure 1 yields a SNARK that is secure against adaptive provers. Moreover, the construction is universally succinct: the proof's length and verifier's running-time are bounded by the same universal polynomials for all $\mathcal{R}_c \subseteq \mathcal{R}_{\mathcal{U}}$.*

The completeness of the construction follows directly from the completeness of the PCP and PIR. In Section 5.2, we give a security reduction establishing the PoK property (against adaptive provers). In Section 5.3, we discuss *universal succinctness* and possible extensions of our construction to a full-fledged universal argument.

⁹We note that any $(k^\epsilon, k^{\epsilon'})$ -compressing ECRH would have sufficed (for any constants $\epsilon > \epsilon'$); for the sake of simplicity, we stick with (k^2, k) -compression.

Ingredients.

- A universal efficient-verifier PCP of knowledge $(P_{\text{pcp}}, V_{\text{pcp}}, E_{\text{pcp}})$ for $\mathcal{R}_{\mathcal{U}}$; for $((M, x, t), w) \in \mathcal{R}_{\mathcal{U}}$, a proof π is s.t. $|\pi| \leq t^2$ and the non-repeated verifier uses $\rho = O(\log t)$ coins and $O(1)$ queries.
- A succinct PIR (PEnc, PEval, PDec) that supports an a-priori unknown database size.^a
- An (k^2, k) -ECRH.

Setup $\mathcal{G}_{\mathcal{V}}(1^k)$.

- Generate private verification state:
 - Sample coins for $q = \omega(\log k)$ repetitions of V_{pcp} : $\mathbf{r} = (r_1, \dots, r_q) \xleftarrow{\mathcal{U}} \{0, 1\}^{(\log^2 k) \times q}$.
 - Sample coins for encrypting q PIR-queries: $R \xleftarrow{\mathcal{U}} \{0, 1\}^{\text{poly}(k)}$.
 - Sample an ECRH: $h \leftarrow \mathcal{H}_k$.
 - Set $\text{priv} := (h, \mathbf{r}, R)$.
- Generate corresponding verifier-generated reference string:
 - Compute $C_{\mathbf{r}} \leftarrow \text{PEnc}_R(1^k, \mathbf{r})$.
 - Set $\text{vgrs} := (h, C_{\mathbf{r}})$.

Proof generation by \mathcal{P} .

- Input: $1^k, \text{vgrs}, (y, w) \in \mathcal{R}_c$ where $y = (M, x, t)$ and $t \leq |x|^c$.
- Proof generation:
 - Compute a PCP proof $\pi \leftarrow P_{\text{pcp}}(y, w)$ of size $|\pi| = k^{d+1} \leq t^2$.
 - Compute an MT commitment for π : $\ell_{\epsilon} = \text{MT}_h(\pi)$ of depth d .
 - Let $\rho = O(\log t) < \log^2 k$ be the amount of coins required by V_{pcp} . Compute a database DB with 2^{ρ} entries; in each entry $\hat{r} \in \{0, 1\}^{\rho}$ store the openings $\text{dcom}_{\hat{r}}$ for all the locations of π that are queried by $V_{\text{pcp}}^{\pi}(y, \hat{r})$.^b
 - Compute $C_{\text{dcom}_{\hat{r}}} \leftarrow \text{PEval}(\text{DB}, C_{\mathbf{r}})$. Here, each coordinate $r_j \in \{0, 1\}^{\log^2 k}$ of \mathbf{r} is interpreted by the PIR as a shorter query $\hat{r}_j \in \{0, 1\}^{\rho}$.
 - The proof is set to be $\Pi := (d, \ell_{\epsilon}, C_{\text{dcom}_{\hat{r}}})$.

Proof verification by \mathcal{V} .

- Input: $1^k, \text{priv}, y, \Pi$, where $y = (M, x, t), \Pi = (d, \ell_{\epsilon}, C_{\text{dcom}_{\hat{r}}})$.
- Proof verification:
 - Verify^c that $k^{d+1} \leq t^2 \leq |x|^{2c}$.
 - Decrypt PIR answers $\text{dcom}_{\hat{r}} = \text{PDec}_R(C_{\text{dcom}_{\hat{r}}})$, and verify opened paths (against h and ℓ_{ϵ}).
 - Let $\pi|_{\hat{r}}$ be the opened values of π in the locations corresponding to \hat{r} (where again \hat{r} is the interpretation of $r \in \{0, 1\}^{\log^2 k}$ as $\hat{r} \in \{0, 1\}^{\rho}$). Check whether $V_{\text{pcp}}^{\pi|_{\hat{r}}}(y, \hat{r})$ accepts.
 - In case any of the above fail, reject.

^aSuch a PIR interprets a “long” query $r \in \{0, 1\}^{\log^2 k}$ as a shorter one $\hat{r} \in \{0, 1\}^{\rho}$, the ρ -bit prefix of r (see Section 3.6).

^b V_{pcp} ’s queries might be adaptive; such behavior can be simulated by the prover.

^cThis is the single place where the verification algorithm depends on c . See further discussion in Section 5.3.

Figure 1: A SNARK for the relation \mathcal{R}_c .

5.2 Proof of Security

A high-level overview of the proof and main technical challenges are presented in Section 1.4. We now turn to the detailed proof, which concentrates on establishing and proving the validity of the knowledge extractor.

Proposition 5.2 (adaptive proof of knowledge). *For any polynomial-size \mathcal{P}^* there exists a polynomial-size extractor $\mathcal{E}_{\mathcal{P}^*}$, such that for all large enough $k \in \mathbb{N}$ and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$:*

$$\Pr_{h, \mathbf{r}, R} \left[\begin{array}{l} (y, \Pi) \leftarrow \mathcal{P}^*(z, h, \text{PEnc}_R(\mathbf{r})) \\ \mathcal{V}((1^k, h, R, \mathbf{r}), y, \Pi) = 1 \end{array} \wedge \begin{array}{l} (y, w) \leftarrow \mathcal{E}_{\mathcal{P}^*}(1^k, z, h, \text{PEnc}_R(\mathbf{r})) \\ w \notin \mathcal{R}_c(y) \end{array} \right] \leq \text{negl}(k) ,$$

where h is an ECRH, \mathbf{r} are the PCP coins and R are the PIR coins.

We start by describing how the extraction circuit is constructed and then prove that it satisfies Proposition 5.2. In order to simplify notation, we will address provers \mathcal{P}^* that get as input only $(h, C_{\mathbf{r}})$, where $C_{\mathbf{r}} = \text{PEnc}_R(\mathbf{r})$; the analysis can be extended to the case that \mathcal{P}^* also gets additional auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$. We note that, formally, a prover \mathcal{P}^* is a family $\{\mathcal{P}_k^*\}$ of polynomial-size circuits, and so is its corresponding extractor $\mathcal{E}_{\mathcal{P}^*}$; for notational convenience, we omit the subscript k .

The extraction procedure. As discussed in Section 1.4, extraction is done in two phases: first, we recursively extract along all the paths of the Merkle tree (MT); doing so results in a string (of leaf labels) $\tilde{\pi}$; then, we apply to $\tilde{\pi}$ the PCP witness-extractor E_{pcp} . As we shall see, $\tilde{\pi}$ will exceed the knowledge-threshold ε of the PCP and hence E_{pcp} will produce a valid witness.

We now describe the recursive extraction procedure of the of the ECRH-based MT. Given a polynomial-size prover \mathcal{P}^* , let d be such that $|\mathcal{P}^*| \leq k^d$. We derive $2cd$ circuit families of extractors, one for each potential level of the MT. Define $\mathcal{E}_1 := \mathcal{E}_{\mathcal{P}^*}^{\mathcal{H}}$ to be the ECRH extractor for \mathcal{P}^* ; like \mathcal{P}^* , \mathcal{E}_1 also expects input $(h, C_{\mathbf{r}}) \in \{0, 1\}^{\text{poly}(k)}$ and returns a string $(\tilde{\ell}_1, \dots, \tilde{\ell}_k) \in \{0, 1\}^{k \times k}$ (which will be a preimage in case \mathcal{P}^* produces a valid image). We can interpret the string output by \mathcal{E}_1 as k elements in the range of the hash. Since the ECRH extraction guarantee only considers a single image, we define an augmented family of circuits \mathcal{E}'_1 that expects input $(h, C_{\mathbf{r}}, i_1)$, where $i_1 \in [k]$, and returns ℓ_{i_1} , which is the i_1 -th k -bit block of $\mathcal{E}_1(h, C_{\mathbf{r}})$.

Next, we define $\mathcal{E}_2 := \mathcal{E}'_1$ to be the extractor for \mathcal{E}'_1 . In general, we define $\mathcal{E}_{j+1} := \mathcal{E}'_j$ to be the extractor for \mathcal{E}'_j , and \mathcal{E}'_j expects an input $(h, C_{\mathbf{r}}, \mathbf{i})$, where $\mathbf{i} \in [k]^j$. For each $\mathbf{i} \in [k]^j$, $\mathcal{E}_{j+1}(h, C_{\mathbf{r}}, \mathbf{i})$ is meant to extract the labels $\tilde{\ell}_{i_1}, \dots, \tilde{\ell}_{i_k}$.

Recall, however, that the ECRH extractor \mathcal{E}_{j+1} is only guaranteed to output a preimage whenever the corresponding circuit \mathcal{E}'_j outputs a true image. For simplicity, we assume that in case \mathcal{E}'_j doesn't output a true image, \mathcal{E}_{j+1} still outputs some string of length k^2 (without any guarantee on this string).

Overall, the witness extractor $\mathcal{E}_{\mathcal{P}^*}$ operates as follows. Given input $(1^k, h, C_{\mathbf{r}})$, (a) first invoke $\mathcal{P}^*(h, C_{\mathbf{r}})$ and obtain (y, Π) ; (b) obtain the depth \tilde{d} from Π , and abort if $\tilde{d} > 2cd$; (c) otherwise, for each $\mathbf{i} \in [k]^{\tilde{d}-1}$, extract the labels $(\tilde{\ell}_{i_1}, \dots, \tilde{\ell}_{i_k}) \leftarrow \mathcal{E}'_{\tilde{d}}(h, C_{\mathbf{r}}, \mathbf{i})$; (d) letting $\tilde{\pi}$ be the extracted PCP-proof given by the leaves, apply the PCP witness extractor $\tilde{w} \leftarrow E_{\text{pcp}}(y, \tilde{\pi})$ and output \tilde{w} .

We now turn to prove that (except with negligible probability), whenever the verifier is convinced, the extractor $\mathcal{E}_{\mathcal{P}^*}$ outputs a valid witness. The proof is divided into two main claims as outlined in Section 1.4.

A reminder and some notation. Recall that prior to the prover's message, the randomness for the PCP verifier is of the form $\mathbf{r} = (r_i)_{i \in [q]} \in \{0, 1\}^{(\log^2 k) \times q}$ (and $q = \omega(\log k)$ is some fixed function). Once

the verifier receives (y, Π) , where $y = (M, x, t)$ and $\Pi = (\tilde{d}, \ell_\epsilon, C_{\text{dcom}})$, he computes the amount of coins required $\rho = O(\log t) < \log^2 k$ and interprets each r_j as a shorter $\hat{r}_j \in \{0, 1\}^\rho$. (Recall that \hat{r}_j is the ρ -bit prefix of r_j ; in particular, when r_j is uniformly-random, so is \hat{r}_j .) The corresponding PCP proof π (or the extracted $\tilde{\pi}$) is of size $k^{\tilde{d}+1}$. We shall denote by $\tilde{\pi} = \mathcal{E}_{\tilde{d}}(h, \text{PEnc}_R(\mathbf{r})) = \cup_{\mathbf{i} \in [k]^{\tilde{d}-1}} \mathcal{E}_{\tilde{d}}(h, \text{PEnc}_R(\mathbf{r}), \mathbf{i})$ the extraction of the full set of leaves.

Using collision resistance and ECRH extraction, we show that (almost) whenever the verifier is convinced, we extract a proof $\tilde{\pi}$ that locally satisfies the queries induced by the encrypted $\text{PEnc}_R(\mathbf{r})$.

Claim 5.3 (local consistency). *Let \mathcal{P}^* be a polynomial-size prover strategy, where $|\mathcal{P}| \leq k^d$, and let $(\mathcal{E}_1, \dots, \mathcal{E}_{2cd})$ be its ECRH extractors as defined above. Then for all large enough $k \in \mathbb{N}$,*

$$\Pr_{(h, R, \mathbf{r}) \leftarrow \mathcal{G}_V(1^k)} \left[\begin{array}{l} (y, \Pi) \leftarrow \mathcal{P}^*(h, \text{PEnc}_R(\mathbf{r})) \\ y = (M, x, t), \Pi = (\tilde{d}, \ell_\epsilon, C_{\text{dcom}}) \\ \mathcal{V}(1^k, (h, R, \mathbf{r}), y, \Pi) = 1 \end{array} \wedge \begin{array}{l} \tilde{\pi} \leftarrow \mathcal{E}_{\tilde{d}}(1^k, h, \text{PEnc}_R(\mathbf{r})) \\ V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}) = 0 \end{array} \right] \leq \text{negl}(k) ,$$

where $\hat{\mathbf{r}} \in \{0, 1\}^{\rho \times q}$ is the interpretation of $\mathbf{r} \in \{0, 1\}^{(\log^2 k) \times q}$ as (a vector of) shorter random strings (as detailed above.)

Proof. Let us say that a tuple (h, R, \mathbf{r}) is “bad” if it leads to the above event. Assume towards contradiction that for infinitely many $k \in \mathbb{N}$, there is a noticeable fraction $\varepsilon(k)$ of bad tuples (h, R, \mathbf{r}) . We show how to find collisions in \mathcal{H} .

Given $h \leftarrow \mathcal{H}$, sample coins R for the PIR encryption and coins \mathbf{r} for the PCP verifier to simulate $\text{PEnc}_R(\mathbf{r})$. Given that the resulting (h, R, \mathbf{r}) is bad, let us show how to produce a collision in \mathcal{H} .

First, invoke $\mathcal{P}^*(h, \text{PEnc}_R(\mathbf{r}))$ to obtain (y, Π) , where $y = (M, x, t)$ and $\Pi = (\tilde{d}, \ell_\epsilon, C_{\text{dcom}})$. Next, decrypt C_{dcom} (using R) and obtain a set S of $O(q)$ opened paths (each r_j in $\mathbf{r} = (r_i)_{i \in [q]}$ induces a constant amount of queries). Each path corresponds to some leaf $\mathbf{i} \in [k]^{\tilde{d}}$ and contains \tilde{d} k^2 -bit strings $\mathbf{l}_1^{\mathbf{i}}, \dots, \mathbf{l}_{\tilde{d}}^{\mathbf{i}} \in \{0, 1\}^{k^2 \times \tilde{d}}$; each string $\mathbf{l}_j^{\mathbf{i}}$ contains the label of the j -th node along the path and the labels of all its siblings.

Next, note that $\tilde{d} \leq 2cd$. Indeed, if the verifier accepts then: $k^{\tilde{d}} \leq |x|^{2c}$, and in our case $|x| \leq |\mathcal{P}^*| \leq k^d$. Accordingly, we can use our extractors as follows: for each opened path in $\mathbf{i} \in S$, where $\mathbf{i} = i_1 \dots i_{\tilde{d}} \in [k]^{\tilde{d}}$, invoke:

$$\begin{aligned} & \mathcal{E}_1(h, \text{PEnc}_R(\mathbf{r})) \\ & \mathcal{E}_2(h, \text{PEnc}_R(\mathbf{r}), i_1) \\ & \quad \vdots \\ & \mathcal{E}_{\tilde{d}}(h, \text{PEnc}_R(\mathbf{r}), i_1 \dots i_{\tilde{d}-1}) \end{aligned}$$

and obtain $\tilde{\mathbf{l}}_1^{\mathbf{i}}, \dots, \tilde{\mathbf{l}}_{\tilde{d}}^{\mathbf{i}} \in \{0, 1\}^{k^2 \times \tilde{d}}$.

Let $\pi|_S = \left(\mathbf{l}_d^{\mathbf{i}} \right)_{\mathbf{i} \in S}$ be the leaves \mathcal{P}^* opened and let $\tilde{\pi}|_S = \left(\tilde{\mathbf{l}}_d^{\mathbf{i}} \right)_{\mathbf{i} \in S}$ be the extracted leaves. Since (h, R, \mathbf{r}) are bad, it holds that $V_{\text{pcp}}^{\pi|_S}(x, \hat{\mathbf{r}}) = 1$ while $V_{\text{pcp}}^{\tilde{\pi}|_S}(x, \hat{\mathbf{r}}) = 0$; in particular, there exist some $\mathbf{i} \in S$ such that $\mathbf{l}_d^{\mathbf{i}} \neq \tilde{\mathbf{l}}_d^{\mathbf{i}}$. We focus from hereon on this specific \mathbf{i} . Let $j \in [\tilde{d}]$ be the smallest index such that $\mathbf{l}_j^{\mathbf{i}} \neq \tilde{\mathbf{l}}_j^{\mathbf{i}}$ (we just established that such an index exists); then it holds that $\mathbf{l}_{j-1}^{\mathbf{i}} = \tilde{\mathbf{l}}_{j-1}^{\mathbf{i}}$. Furthermore, since (h, R, \mathbf{r}) are bad, \mathcal{V} accepts; this in turn implies that h compresses $\mathbf{l}_j^{\mathbf{i}}$ to the i_{j-1} -th block of $\mathbf{l}_{j-1}^{\mathbf{i}} = \tilde{\mathbf{l}}_{j-1}^{\mathbf{i}}$,

which we will denote by ℓ^* . However, the latter is also the output of $\mathcal{E}'_{j-1}(h, \text{PEnc}_R(\mathbf{r}), i_1 \cdots i_{j-1})$, which in turn implies that $\mathcal{E}_j(h, \text{PEnc}_R(\mathbf{r}), i_1 \cdots i_{j-1}) = \tilde{\mathbf{I}}_j^i$ is also compressed by h to the same ℓ^* (except when extraction fails, which occurs with negligible probability). It follows that $\mathbf{I}_j^i \neq \tilde{\mathbf{I}}_j^i$ form a collision in h . \square

The second step in the proof of Proposition 5.2, is to show that if the aforementioned extractor outputs a proof $\tilde{\pi}$ that convinces the PCP verifier with respect to the encrypted randomness, then the same proof $\tilde{\pi}$ must be globally satisfying (at least for witness extraction); otherwise, the polynomial-size extractor can be used to break the semantic security of the PIR.

Claim 5.4 (from local satisfaction to extraction). *Let \mathcal{P}^* be a polynomial-size prover and let $\mathcal{E}_{\mathcal{P}^*}$ be its polynomial-size extractor.¹⁰ Then for all large enough $k \in \mathbb{N}$,*

$$\Pr_{(h, R, \mathbf{r}) \leftarrow \mathcal{G}_V(1^k)} \left[\begin{array}{l} t \leq |x|^c \quad (y, \Pi) \leftarrow \mathcal{P}^*(h, \text{PEnc}_R(\mathbf{r})) \\ V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}) = 1 \quad : \quad y = (M, x, t), \Pi = (\tilde{d}, \ell_e, C_{\text{dcom}}) \\ E_{\text{pcp}}(y, \tilde{\pi}) \notin \mathcal{R}_c(y) \quad \tilde{\pi} \leftarrow \mathcal{E}_{\mathcal{P}^*}(1^k, h, \text{PEnc}_R(\mathbf{r})) \end{array} \right] \leq \text{negl}(k) ,$$

where $\hat{\mathbf{r}} \in \{0, 1\}^{\rho \times q}$ is the interpretation of $\mathbf{r} \in \{0, 1\}^{(\log^2 k) \times q}$ as (a vector of) shorter random strings (as detailed above.)

Proof. Assume towards contradiction that for infinitely many $k \in \mathbb{N}$ the above event occurs with noticeable probability $\delta = \delta(k)$; we show how to break the semantic security of PEnc . First note that whenever the event occurs, it holds that $\Pr_{\hat{\mathbf{r}} \leftarrow \{0, 1\}^{\rho \times q}} [V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}) = 1] \leq (1 - \varepsilon)^q$, where ε is the (constant) knowledge threshold of the PCP (see Section 3.7), and $q = \omega(\log k)$ is the number of repetitions.

To break semantic security, we consider the following CPA game, where a breaker \mathcal{B} hands its challenger two independent strings of PCP randomness, $(\mathbf{r}_0, \mathbf{r}_1) \in \{0, 1\}^{(\log^2 k) \times 2}$, and gets back $\text{PEnc}_R(\mathbf{r}_b)$ for a random $b \in \{0, 1\}$. Now, \mathcal{B} samples a random h , and runs $\mathcal{P}^*(h, \text{PEnc}_R(\mathbf{r}_b))$ and $\mathcal{E}_{\mathcal{P}^*}(1^k, h, \text{PEnc}_R(\mathbf{r}_b))$ to obtain an instance $y = (M, x, t)$ from the prover and an extracted PCP proof $\tilde{\pi}$ from the extractor. Then, \mathcal{B} computes the amount of coins required for V_{pcp} , $\rho = \rho(t)$, and derives the corresponding ρ -prefixes $(\hat{\mathbf{r}}_0, \hat{\mathbf{r}}_1)$ of $(\mathbf{r}_0, \mathbf{r}_1)$.

The breaker now runs the PCP extractor E_{pcp} on input $(y, \tilde{\pi})$ to obtain a string \tilde{w} and verifies whether it is a valid witness for y (which can be done in $\text{poly}(|x|) = \text{poly}(k)$ time). In case the witness \tilde{w} is valid or $V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}_0) = V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}_1)$, the breaker \mathcal{B} outputs a random guess for the bit b . Otherwise, the breaker outputs the single b' such that $V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}_{b'}) = 1$.

We now analyze the success probability of \mathcal{B} . We define two events F and E over a random choice of $(h, R, \hat{\mathbf{r}}_0, \hat{\mathbf{r}}_1, b)$; note that any choice of $(h, R, \hat{\mathbf{r}}_0, \hat{\mathbf{r}}_1, b)$ induces a choice of $y = (M, x, t)$ and $\tilde{\pi}$. Define F to be the event that $t \leq |x|^c$ and $E_{\text{pcp}}(y, \tilde{\pi})$ fails to output a valid witness \tilde{w} ; next, define E to be the event that $V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}_0) \neq V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}_1)$.

First, since we have assumed by way of contradiction that the event in the statement of the claim occurs with probability δ , we know that

$$\Pr [V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}_b) = 1 \mid F] = \frac{\delta}{\Pr[F]} .$$

Second, since E_{pcp} cannot extract a valid witness from $\tilde{\pi}$ and $\hat{\mathbf{r}}_{1-b}$ are random coins independent of $(\tilde{\pi}, y)$, we also know that

$$\Pr_{\hat{\mathbf{r}}_{1-b} \leftarrow \{0, 1\}^{\rho \times q}} [V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{\mathbf{r}}_{1-b}) = 1 \mid F] \leq (1 - \varepsilon)^q .$$

¹⁰The claim actually holds for any circuit family \mathcal{E} , but we'll be interested in the extractor of \mathcal{P}^*

Combining these two facts, we deduce that

$$\begin{aligned}
& \Pr [E \mid F] \\
& \geq \Pr [V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{r}_b) = 1 \wedge V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{r}_{1-b}) = 0 \mid F] \\
& \geq \Pr [V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{r}_b) = 1 \mid F] - \Pr [V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{r}_{1-b}) = 1 \mid F] \\
& \geq \Pr [V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{r}_b) = 1 \mid F] - (1 - \varepsilon)^q = \frac{\delta}{\Pr[F]} - \text{negl}(k) ,
\end{aligned}$$

so that, in particular, we can also deduce that

$$\Pr [F \wedge E] \geq \delta - \text{negl}(k) .$$

Therefore,

$$\begin{aligned}
& \Pr [\mathcal{B} \text{ guesses } b \mid F \wedge E] \\
& \geq 1 - \Pr [V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{r}_{1-b}) = 1 \mid F \wedge E] \\
& = 1 - \frac{\Pr [V_{\text{pcp}}^{\tilde{\pi}}(y, \hat{r}_{1-b}) = 1 \wedge E \mid F]}{\Pr [E \mid F]} \\
& \geq 1 - \frac{(1 - \varepsilon)^q}{\delta / \Pr [F]} \\
& \geq 1 - \text{negl}(k) .
\end{aligned}$$

We now deduce that the breaker \mathcal{B} guesses b with a noticeable advantage; indeed,

$$\begin{aligned}
\Pr [\mathcal{B} \text{ guesses } b] &= \Pr [F \wedge E] \Pr [\mathcal{B} \text{ guesses } b \mid F \wedge E] + (1 - \Pr [F \wedge E]) \Pr [\mathcal{B} \text{ guesses } b \mid \bar{F} \vee \bar{E}] \\
&= \Pr [F \wedge E] \Pr [\mathcal{B} \text{ guesses } b \mid F \wedge E] + (1 - \Pr [F \wedge E]) \cdot \frac{1}{2} \\
&= \frac{1}{2} + \Pr [F \wedge E] \left(\Pr [\mathcal{B} \text{ guesses } b \mid F \wedge E] - \frac{1}{2} \right) \\
&\geq \frac{1}{2} + (\delta - \text{negl}(k)) \left(\frac{1}{2} - \text{negl}(k) \right) \\
&\geq \frac{1}{2} + \frac{\delta - \text{negl}(k)}{2} ,
\end{aligned}$$

thus completing the proof of the claim. \square

Putting it all together. By Claim 5.3 we conclude that whenever the verifier accepts, $\mathcal{E}_{\mathcal{P}^*}$ almost always extracts a proof $\tilde{\pi}$ which locally satisfies the PCP verifier on the encrypted randomness. By Claim 5.4, we deduce that whenever this occurs, $\tilde{\pi}$ must satisfy sufficiently many queries for PCP witness-extraction. This completes the proof of Proposition 5.2 and thus of Theorem 5.1.

Efficiency: “universal succinctness”. For input $y = (M, x, t)$ (where $t < k^{\log k}$ for a security parameter k), the proof $\Pi = (\ell_\varepsilon, d, C_{\text{dcom}_r})$ is essentially dominated by the PIR answers C_{dcom_r} ; this includes $q = \text{polylog}(k)$ PIR answers for entries of size $\tilde{O}(k^2)$.¹¹ In the PIR scheme of [BV11] the size of each PIR-answer is bounded by $E \cdot k \cdot \text{polylog}(k) + \log D$, where E is the size of an entry and D is the size of the entire

¹¹Recall that $d = \log_k t < \log k$.

DB. Hence, the overall length of the proof is bounded by a fixed polynomial $\tilde{O}(k^2)$, independently of $|x|$, $|w|$ or c . The verifier’s and prover’s running time are bounded respectively by fixed universal polynomials $\text{poly}(|y|, k)$, $\text{poly}(k, t)$, again independently of c .

Parameter scaling. In Kilian’s original protocol, succinctness of the proof could be improved by making stronger hardness assumptions. For example, for security parameter k , if one is willing to assume collision-resistant hash functions with a $\text{polylog}(k)$ -long output, the proof length would be $\text{polylog}(k)$, rather than $\text{poly}(k)$. Unfortunately, in our construction we use a Merkle tree with $\text{poly}(k)$ fan-in; therefore, we cannot afford the same scaling. Specifically, even if we assume that our hash and PIR scheme have polylogarithmic-size output, each node in the Merkle tree still has $\text{poly}(k)$ siblings.¹² Nonetheless, scaling can be performed if we make a stronger extractability assumption, such as the interactive one of [DFH11], because in such a case there is no need to consider Merkle trees with polynomial fan-in as binary Merkle trees suffice for the security reduction.

5.3 Extension to Universal Arguments

We now discuss the possibility of extending our construction to a full-fledged universal argument, namely an argument for the universal relation \mathcal{R}_U as defined in Section 4.1.

Indeed, Theorem 5.1 tells us that for every $c \in \mathbb{N}$ we obtain a specific protocol that is sound with respect to \mathcal{R}_c . The dependence on c , however, only appears in the first step of \mathcal{V} , where it is checked that $k^{d+1} \leq |x|^{2c}$. In particular, as we already discussed, the running time of both the prover and verifier, as well as the proof-length, are universal and do *not* depend on c .

Towards a full-fledged universal argument. To obtain a full-fledged universal argument we might try to omit the above c -dependent size check. However, now we encounter the following difficulty: for the proof of knowledge to go through, we must ensure that the number of recursive extractions is a-priori fixed to some constant \tilde{d} (that may depend on the prover). In particular, we need to prevent the prover \mathcal{P}^* from convincing the verifier of statements $y = (M, x, t)$ with $t > k^{\tilde{d}}$. The natural candidate for \tilde{d} is typically related to the polynomial-size bound on the size of \mathcal{P}^* . Indeed, any prover that actually “writes down” a proof of size t should be of size at least t ; intuitively, one could hope that being able to convince the verifier should essentially amount to writing down the proof and computing a Merkle hash of it. However, we have not been able to prove this. Instead, we propose the following modification to the protocol to make it a universal argument.

Proofs of work. For the relation \mathcal{R}_c , the above problem of \mathcal{P}^* claiming an artificially large t can be avoided by ensuring that the size of a convincing proof t can only be as large as $|x|^c$, where $|x|$ is a lower-bound on the prover’s size. More generally, to obtain a *universal argument*, we can omit the verifier’s check (thus collapsing the family of protocols to a *single* protocol) and enhance the protocol of Figure 1 with a *proof of work* attesting that the prover has size at least t^ε for some constant $\varepsilon > 0$. Concretely, if we are willing to make an additional (though admittedly quite strong) assumption we can obtain such proofs of work:

Theorem 5.5. *If there exist $2^{\varepsilon n}$ -hard one-way functions (where n is the input size), then, under the same assumptions as Theorem 5.1, we can modify the protocol in Figure 1 to obtain a universal argument.*

Proof sketch. Let $f: \{0, 1\}^* \rightarrow \{0, 1\}^*$ a $2^{\varepsilon n}$ -hard one-way function. Modify $\mathcal{G}_V(1^k)$ to also output z_1, \dots, z_ℓ , with $\ell := \log^2 k$ and $z_i := f(s_i)$, where each s_i is drawn at random from $\{0, 1\}^i$. Then, when claiming a proof Π for an instance $y = (M, x, t)$, the prover must also present s'_i such that $f(s'_i) = z_i$

¹²We thank Kai-Min Chung and the anonymous referees of ITCS for pointing out the scaling problem.

where $i > \log t$.¹³ The verifier \mathcal{V} can easily check that this is the case by evaluating f . (Note also that the honest prover will have to pay at most an additive factor of $\tilde{O}(t)$ in its running time when further requested to present this challenge.) Then, by the hardness of f , we know that if the prover has size k^d , then it must be that $k^d > 2^{\varepsilon i} > t^\varepsilon$, so that we conclude that $k^{d/\varepsilon} > t$. Therefore, in the proof of security, we know that the claimed depth of the prover is a constant depending only on d and ε , and thus the same proof of security as that of Theorem 5.1 applies. \square

Admittedly, assuming exponentially-hard one-way functions is unsatisfying, and we hope to remove the assumption with further analysis; in the meantime, we would like to emphasize that this assumption has already been made, e.g., in natural proofs [RR94] or in works that improve PRG constructions [HHR06].

6 ECRHs: a Closer Look

In this section we take a closer look at the notion of ECRH, and propose relaxations of this notion that still suffice for constructing SNARKs. These relaxations are crucial to expand our set of candidate constructions; for more details on the constructions, see Section 8.

6.1 ECRHs

We begin by discussing several technical aspects regarding the definition of ECRH. Recall that an ECRH is a collision-resistant function ensemble \mathcal{H} that is extractable in the sense of Definition 1, which we reproduce:

Definition 6.1. *An efficiently-samplable function ensemble $\mathcal{H} = \{\mathcal{H}_k\}_k$ is an $(\ell(k), k)$ -compressing ECRH if it is $(\ell(k), k)$ -compressing, collision-resistant, and moreover extractable: for any polynomial-size adversary \mathcal{A} , there exists a polynomial-size extractor $\mathcal{E}_\mathcal{A}^\mathcal{H}$, such that for all large enough $k \in \mathbb{N}$ and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$:*

$$\Pr_{h \leftarrow \mathcal{H}_k} \left[\begin{array}{l} y \leftarrow \mathcal{A}(h, z) \\ \exists x : h(x) = y \end{array} \wedge \begin{array}{l} x' \leftarrow \mathcal{E}_\mathcal{A}^\mathcal{H}(h, z) \\ h(x') \neq y \end{array} \right] \leq \text{negl}(k) . \quad (1)$$

In other words, the only way an adversary \mathcal{A} can sample elements in the image of the hash is by knowing a corresponding preimage (which an extractor $\mathcal{E}_\mathcal{A}^\mathcal{H}$ could in principle find).

Image verification. In known applications of extractable primitives (e.g., 3-round zero knowledge [HT98, BP04, CD09]), an extra image-verifiability feature is required. Namely, given $y \in \{0, 1\}^k$ and h , one should be able to efficiently test whether $y \in \text{Image}(h)$. Here, there are two flavors to consider: (a) public verifiability, where to verify an image all that is required is the (public) seed h ; and (b) private verifiability; that is, the seed h is generated together with private verification parameters priv , so that anyone in hold of priv may perform image verification. We emphasize that our main ECRH-based construction (presented in Section 5.1) *does not require any verifiability features*.

Necessity of sparseness. For \mathcal{H} to be collision-resistant, it must also be one-way; namely, the image distribution

$$\mathcal{I}_h = \left\{ h(x) : x \xleftarrow{U} \{0, 1\}^{\ell(k)} \right\}$$

¹³ Note that in any case the verifier will reject any claim for t above the superpolynomial universal bound $2^{\log^2 k}$; hence, $\ell = \log^2 k$ challenges are sufficient for any polynomial-size prover.

should be hard to invert (except with negligible probability over h). In particular, \mathcal{I}_h must be very far from the uniform distribution over $\{0, 1\}^k$ (for almost all h).

Indeed, suppose that the statistical distance between \mathcal{I}_h and uniform is $1 - 1/\text{poly}(k)$ and consider an adversary \mathcal{A} that simply outputs range elements $y \in \{0, 1\}^k$ uniformly at random, and any $\mathcal{E}_{\mathcal{H}}^{\mathcal{A}}$. In this case, there is no “knowledge” to extract from \mathcal{A} , so $\mathcal{E}_{\mathcal{H}}^{\mathcal{A}}$ has to invert uniformly random elements of the range $\{0, 1\}^k$. Thus, the success probability of $\mathcal{E}_{\mathcal{H}}^{\mathcal{A}}$ will differ by at most $1 - 1/\text{poly}(k)$ from its success probability had the distribution been \mathcal{I}_h , which is negligible (by one-wayness); hence $\mathcal{E}_{\mathcal{H}}^{\mathcal{A}}$ will still fail with probability $1 - 1/\text{poly}(k)$ often, thereby violating Equation (1).

A simple way to ensure that the image distribution \mathcal{I}_h is indeed far from uniform is to make the support of \mathcal{I}_h sparse. We will take this approach when constructing candidates, making sure that all $h(x)$ fall into a superpolynomially sparse subset of $\{0, 1\}^k$: $\text{Image}(h) < 2^{k-\omega(\log k)}$ (except with negligible probability over $h \leftarrow \mathcal{H}_k$).

Of course, this merely satisfies one necessary condition, and is a long way off from implying extractability. Still, this rules out one of the few generic attacks about which we can reason without venturing into the poorly-charted territory of non-black-box extraction. Moreover, the sparseness (or more generally, statistical distance) requirement rules out many natural constructions; for example, traditional cryptographic CRH ensembles, and heuristic constructions such as the SHA family, have an image distribution \mathcal{I}_h that is close to uniform (by design) and are thus not extractable.

On auxiliary input. The ECRH definition requires that for any adversary auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$, the polynomial-size extractor manages to perform its extraction task given the same auxiliary input z . As observed by Hada and Tanaka [HT98], and by Goldreich [Gol04], this requirement is rather strong considering the fact that z could potentially encode arbitrary circuits. Indeed, [BCPR14, BP13] show that such a definition cannot be satisfied, assuming indistinguishability obfuscation [BGI⁺12]. The mentioned impossibility results strongly rely on the fact that the auxiliary input can be of arbitrary polynomial size, and can be taken from an arbitrary distribution.

While for presentational purposes the above definition may be simple and convenient, for our main application (i.e., SNARKs) we can actually settle for definition that is weaker in both aspects—it can be restricted to a specific “benign distribution” on auxiliary inputs, which is also of a-priori bounded polynomial size, a setting in which no negative results are known. (Naturally, when relying on ECRHs with respect to such restricted auxiliary-input distributions, the resulting SNARK will account for the same auxiliary-input distributions.)

Specifically, in our setting the extractor is required to handle an auxiliary input (C, I) consisting of (honestly-generated) PIR-encryptions C of random strings, whose length is a-priori bounded by a fixed polynomial $\text{poly}(k)$ in the security parameter k , and a short path index I of length $O(\log k)$. The logarithmically-short auxiliary input I can be shown not to give any additional power to the adversary. Specifically, an ECRH for auxiliary input distribution C is also an ECRH for auxiliary input (C, I) as long as $|I| = O(\log k)$. (Roughly, the extractor of an adversary \mathcal{A} for auxiliary input (C, I) can be constructed from a polynomial number of extractors, obtained by considering the extractor of $\mathcal{A}_I(C) := \mathcal{A}(C, I)$ for every possible value of I ; for the formal argument one has to adapt this intuition to circuit families.)

As for the “main” part C of the auxiliary input, we note that by using a PIR scheme where honestly generated ciphers are pseudo-random, we can actually restrict the “main” part C of the auxiliary input to be a random string (for example, the LWE-based PIR scheme of [BV11] has this property), as long as the ECRH is has efficient (even private) image verification. Indeed, an extractor that does noticeably better for random auxiliary-input than for pseudo-random auxiliary-input, can be used to distinguish the two cases. For this we should be able to efficiently check when that adversary produces an actual image (whereas the

extractor fails to produce a corresponding preimage), which is why efficient image verification is needed.

6.2 PECRHs

As discussed in Section 1.3.1, our first weakening of ECRH, namely proximity ECRH, captures a more flexible tradeoff between the requirements of extractability and collision resistance. Formally:

Definition 6.2. *An efficiently-samplable function ensemble $\mathcal{H} = \{\mathcal{H}_k\}_k$ is an $(\ell(k), k)$ -compressing PECRH if it is $(\ell(k), k)$ -compressing and, for every h in the support of \mathcal{H}_k , there exist a reflexive “proximity relation” $\overset{h}{\approx}$ over pairs in $\{0, 1\}^k \times \{0, 1\}^k$, an “extended domain” $D_h \supseteq \{0, 1\}^{\ell(k)}$, and an extension $\bar{h} : D_h \rightarrow \{0, 1\}^k$ consistent with h (i.e., $\forall x \in \{0, 1\}^{\ell(k)}$ it holds that $h(x) = \bar{h}(x)$), such that:*

1. \mathcal{H} is proximity-extractable in the following weakened sense: for any polynomial-size adversary \mathcal{A} there exists a polynomial-size extractor $\mathcal{E}_{\mathcal{A}}^{\mathcal{H}}$ such that for large enough security parameter $k \in \mathbb{N}$ and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$:

$$\Pr_{h \leftarrow \mathcal{H}_k} \left[\begin{array}{c} y \leftarrow \mathcal{A}(h, z) \\ \exists x \in \{0, 1\}^{\ell(k)} : y = h(x) \end{array} \wedge \neg \left(x' \in D_h \wedge \bar{h}(x') \overset{h}{\approx} y \right) \right] < \text{negl}(k) .$$

2. \mathcal{H} is proximity-collision-resistant in the following strengthened sense: for any polynomial-size adversary \mathcal{A} ,

$$\Pr_{h \leftarrow \mathcal{H}_k} \left[(x, x') \leftarrow \mathcal{A}(h) \wedge x, x' \in D_h \wedge x \neq x' \wedge \bar{h}(x) \overset{h}{\approx} \bar{h}(x') \right] < \text{negl}(k) .$$

We now discuss why any point on the tradeoff (i.e., any choice of $\overset{h}{\approx}$, D_h and \bar{h} fulfilling the conditions) suffices for the construction of SNARKs as claimed in Theorem 5 in Section 1.3.1.

Proof sketch for Theorem 5. We argue that the *same construction* the we use in the proof of Theorem 1 to construct SNARKs from ECRHs still suffices even when the underlying hash function is only a PECRH.

First, observe that moving from ECRHs to PECRHs only affects the “local consistency” step of our proof (as described in our high-level description in Section 1.4 and then formally as Claim 5.3). Indeed, in the proof based on ECRHs, the local-consistency step is where we employ collision resistance to claim that the Merkle tree output by the extractor locally agrees with the opened paths (except with negligible probability).

The same argument holds. By the proximity extraction guarantee, it must be that the hash image of every node label that appears in an opened path is “close” to the image of the corresponding node label in the extracted tree. By the proximity collision resistance, however, these two node labels must in fact *be the same*; for if they were not, then we could utilize the prover and extractor for finding “proximity collisions”. The rest of the proof of Theorem 1 remains unchanged.

We emphasize that the proximity relation $\overset{h}{\approx}$ need not be efficiently computable for the above to hold. \square

6.3 Weak PECRHs

As discussed in Section 1.3.2, our second weakening of ECRH, namely weak PECRH, relaxes the condition that determines when the extractor has to work. Formally:

Definition 6.3. An efficiently-samplable function ensemble $\mathcal{H} = \{\mathcal{H}_k\}_k$ is a $(\ell(k), k)$ -compressing weak PECRH if it satisfies Definition 6.2 with the following modified first condition:

1. \mathcal{H} is weak proximity-extractable in the following sense: for any polynomial-size adversary \mathcal{A} there exists a polynomial-size extractor $\mathcal{E}_A^{\mathcal{H}}$ such that for large enough security parameter $k \in \mathbb{N}$ and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$ and polynomial-size decoder circuit \mathcal{Y} :

$$\Pr_{h \leftarrow \mathcal{H}_k} \left[\begin{array}{l} (y, e) \leftarrow \mathcal{A}(h, z) \\ h(\mathcal{Y}(e)) = y \end{array} \wedge \neg \left(x' \leftarrow \mathcal{E}_A^{\mathcal{H}}(h, z) \wedge x' \in D_h \wedge \bar{h}(x') \stackrel{h}{\approx} y \right) \right] < \text{negl}(k) .$$

We show that even weak PECRHs are enough for the construction of SNARKs as claimed in Theorem 7.

Proof sketch of Theorem 7. We argue that the *same construction* that we use in the proof of Theorem 1 to construct SNARKs from ECRHs also obtains SNARKs even when the underlying hash function is only a weak PECRH. As was the case when moving from ECRHs to PECRHs, moving from PECRHs to weak PECRHs only affects the “local consistency” step of our proof (as described in our high-level description in Section 1.4 and then formally as Claim 5.3); specifically, we must still be able to guarantee local consistency even when the condition under which the extractor is guaranteed to work is weakened to the case where the adversary outputs an encoding of a preimage (as opposed to when the adversary merely outputs a value in the image).

In the construction, this is always the case, because preimages along the opened path are provided as encrypted authentication paths, which can be “decoded” by a decoder that knows encryption the secret key (i.e., the PIR private coins). Therefore, we are still able to show local consistency. \square

Unlike PECRHs, weak PECRHs may in principle not require sparsity of the image or special algebraic structure. While an attacker trying to fool a PECRH extractor only has to obviously sample an image of the function, now, to fool a weak PECRH extractor, it needs to simultaneously obviously sample an image of the function and an encoding of a corresponding preimage. This raises the following natural question: “is any CRH also a weak PECRH?” We believe this to be unlikely; indeed, the following example shows that (assuming the existence of one-way permutations) there exists a CRH that is not a weak PECRH when the proximity relation is forced to be equality. (To extend this to an actual counter-example, one would have to rule out all proximity relations.)

Let $\mathcal{H} = \{\mathcal{H}_k\}$ be any CRH mapping $\{0, 1\}^{2k}$ to $\{0, 1\}^k$ and P any one-way permutation mapping $\{0, 1\}^k$ to itself. Consider the (contrived) new CRH \mathcal{H}' , mapping $\{0, 1\}^{2k}$ to $\{0, 1\}^{k+1}$, that is defined as follows. A seed $h' \in \mathcal{H}'_k$ corresponds to a seed $h \in \mathcal{H}_k$; each $(x_1 || 0^k) \in \{0, 1\}^k \times \{0^k\}$ is mapped by h' to $0 || P(x_1)$ and each $(x_1 || x_2) \in \{0, 1\}^k \times (\{0, 1\}^k \setminus \{0^k\})$ is mapped by h' to $1 || h(x_1 || x_2)$.

Since P is one-to-one and the intersection of the image sets $h'(\{0, 1\}^k \times \{0^k\})$ and $h'(\{0, 1\}^k \times (\{0, 1\}^k \setminus \{0^k\}))$ is empty, any collision in h' implies a corresponding collision in h and, therefore, \mathcal{H}' is also a CRH (that is, a proximity CRH relative to the equality proximity relation). However, \mathcal{H}' is not weakly proximity extractable relative to the equality proximity relation; indeed, consider the auxiliary input distribution $z = P(x_1)$ for a random $x_1 \leftarrow \{0, 1\}^k$, with a corresponding (correlated) decoder \mathcal{Y} that always outputs $x_1 || 0^k$. In addition, consider a dummy adversary that given z, h simply outputs $0 || z = 0 || P(x_1) = h(x_1 || 0^k)$. Note that since the decoder always outputs a valid preimage, any extractor would have to do the same. However, any efficient extractor that manages to do so has to invert the one-way permutation P .

7 From SNARKs to PECRHs (and More)

In this section we provide more details about Theorem 2 and Theorem 3, which were only informally stated in the introductory discussion of Section 1.2. That is, we show that (proximity) extractable collision-resistant hash functions (PECRHs) are in fact not only sufficient (together with appropriate polylog PIRs) but also necessary for SNARKs (assuming standard CRHs). We then describe a general method for obtaining additional (non-interactive) extractable primitives.

Extractability and proximity extractability. We say that a function ensemble $\mathcal{F} = \{\mathcal{F}_k\}_k$ is *extractable* if, given a random $f \leftarrow \mathcal{F}_k$, it is infeasible to produce $y \in \text{Image}(f)$ without actually “knowing” $x \in \text{Domain}(f)$ such that $f(x) = y$. This is formalized by the requirement that for any polynomial-size adversary \mathcal{A} there is a polynomial-size extractor $\mathcal{E}_{\mathcal{A}}$ such that for any auxiliary input z and randomly chosen f : if $\mathcal{A}(z, f)$ outputs a proper image, $\mathcal{E}_{\mathcal{A}}(z, f)$ outputs a corresponding preimage. Typically, for such a family to be interesting, it is required that \mathcal{F} also has some hardness property, e.g., one-wayness; in particular, for the two features (of hardness and extractability) to coexist, $\text{Image}(f)$ must be sparse in the (recognizable) range of the function.

As explained (for ECRHs) in Section 1.3.1 and later in Section 6.1, the above notion of extraction can be relaxed to consider proximity extraction, according to which it is infeasible to produce $y \in \text{Image}(f)$ without knowing an x such that $f(x)$ is proximate to y relative to some given reflexive proximity relation (e.g., relative hamming distance at most $1/2$). For such a notion of extraction to be useful, the corresponding hardness of f should be upgraded accordingly. For example, a *proximity extractable one-way function* (f, \approx) , where \approx is some proximity relation, is such that, given $f(x)$ for a random x , it is infeasible to find an x' for which $f(x') \approx f(x)$ (and it is proximity extractable in the sense that we just described).

In Section 6 we discussed even further relaxations: in one, the extractor also had the freedom to output elements x from an extended domain $D_f \supseteq \text{Domain}(f)$; in another, the extractor only had to work when the adversary manages to output not only an image but also an encoding of a corresponding preimage. In this section, however, we do not consider these further relaxations.

Note that, naturally, known cryptographic schemes (e.g., 3-round zero-knowledge) have relied on standard extraction rather than proximity extraction; however, to the best of our knowledge we can safely replace standard extraction in these schemes with proximity extraction (as we did for the purpose of constructing SNARKs).

Verification and proximity verification. Another extraction-related issue is *image verification*; here, there are two notions that can be considered:

- *Public verification:* Given f and $y \in \text{Range}(f)$, one can efficiently test whether $y \in \text{Image}(f)$.
- *Private verification:* Together with the function f , \mathcal{F}_k also generates a private verification state priv_f . Given f , priv_f and $y \in \text{Range}(f)$, one can efficiently test whether $y \in \text{Image}(f)$.

In addition, when considering proximity extractability, we can consider a corresponding notion of *proximity verifiability*, where the verifier should only check whether $y \in_{\approx} \text{Image}(f)$, namely there is some element $y' \in \text{Image}(f)$ for which $y \approx y'$. Again we note that for the known applications of (verifiable) extractable primitives, proximity verification is sufficient.

Also note that the weaker notion of *extractability with no efficient verification* might also be meaningful in certain scenarios. Indeed, for our main ECRH-based construction of SNARKs (presented in Section 5.1), this weak notion of extractability with no efficient verification suffices, and in fact ultimately allows us to deduce, through the SNARK construction, many extractable primitives with efficient private (proximity) verification.

7.1 From SNARKs to PECRHs

We now present the implications of SNARKs to the existence of extractable primitives, starting with the necessity of PECRHs:

Proposition 7.1. *If there exist SNARKs and (standard) CRHs, then there exist proximity verifiable PECRHs.*

Proof sketch. We show that designated-verifier SNARKs imply PECRHs with private proximity verification. The proof can be easily extended to the case of public verifiability (where publicly-verifiable SNARKs imply PECRHs with public proximity verification). Let \mathcal{H} be a (tk, k) -compressing CRH, where $t = t(k) > 1$ is a compression parameter. Let $(\mathcal{P}, \mathcal{G}_\nu, \mathcal{V})$ be an (adaptive) SNARK such that, given security parameter \hat{k} , the length of any proof is bounded by \hat{k}^c .¹⁴ We define a $(tk, 2k)$ -compressing ECRH, $\tilde{\mathcal{H}} = \{\tilde{\mathcal{H}}_k\}_k$. A function \tilde{h} and private verification state $\text{priv}_{\tilde{h}}$ are sampled by $\tilde{\mathcal{H}}_k$ as follows:

1. Draw a function $h \leftarrow \mathcal{H}_k$,
2. Draw public and private parameters $(\text{vgrs}, \text{priv}) \leftarrow \mathcal{G}_\nu(k^{1/c})$, and
3. Set $\tilde{h} = (h, \text{vgrs})$, $\text{priv}_{\tilde{h}} = \text{priv}$.

Then, for an input x and defining $y = h(x)$, we define $\tilde{h}(x) = (y, \Pi)$ where $\Pi = \mathcal{P}(\text{vgrs}, \text{thm}, x)$ is a proof of knowledge for the NP-statement $\text{thm} = \text{“there exists an } x \in \{0, 1\}^{tk} \text{ such that } h(x) = y\text{”}$.

We now show that the above is a PECRH with respect to the relation \approx , where $(y, \Pi) \approx (y', \Pi')$ if and only if $y = y'$.

The proximity collision resistance of $\tilde{\mathcal{H}}$ follows directly from the collision resistance of \mathcal{H} , because any proximity-colliding pair (x, x') for $\tilde{\mathcal{H}}$ is a colliding pair for \mathcal{H} . The proximity extractability property of $\tilde{\mathcal{H}}$ follows from the (adaptive) proof of knowledge of the SNARK $(\mathcal{P}, \mathcal{G}_\nu, \mathcal{V})$; that is, for any image-computing polynomial-size adversary \mathcal{A} , the ECRH extractor is set to be the SNARK witness-extractor $\mathcal{E}_\mathcal{A}$. In addition, an image can be proximity-verified (with respect to \approx) by invoking the SNARK verifier \mathcal{V} with the private verification state priv , the proof Π , and the corresponding statement. We note that, for the proposition to go through, it is crucial for the SNARK to hold against adaptive provers; indeed, the adversary gets to choose on which inputs to compute the hash function, and these may very well depend on the public parameters. \square

Why PECRH and not ECRH? At first glance, it is not clear why the above does not imply an (“exact”) ECRH rather than a PECRH. The reason lies in the fact that the extractor is only guaranteed to output one of many possible witnesses (preimages). In particular, given an honest image $(h(x), \Pi_x)$ (corresponding to some preimage x), the extractor may output x' such that $h(x') = h(x)$ but applying the honest prover to x' (or, more precisely, the NP-statement proving knowledge of x') results in a proof $\Pi_{x'} \neq \Pi_x$.

We now can immediately deduce that SNARKs also imply *proximity extractable one-way functions* (PEOWFs) and *proximity extractable computationally binding and hiding commitments* (PECOMs):

Corollary 7.2. *If there exist SNARKs and (standard) CRHs, then there exist PEOWFs and PECOMs. Moreover, the verifiability features of the SNARK carry over to the implied primitives.*

¹⁴More precisely, the length of any proof is bounded by $(\hat{k} + \log t)^c$, where t is the computation time; however, we only address statements where the computation is poly-time and in particular $\log t < \hat{k}$.

Proof sketch. First, note that any $(\ell(k), k)$ -compressing PECRH is also a (keyed) PEOF (with respect to the same image proximity relation). Indeed, it is a proximity OWF since it is a proximity CRH and independently of that it is also proximity extractable (and verifiable).

Second, to get an extractable bit-commitment scheme, one can use the classic CRH plus hardcore bit construction of Blum [Blu81]. Specifically, the commitment scheme is keyed by a seed h for the PECRH and a commitment to a bit b is obtained by sampling $r, \hat{r} \xleftarrow{U} \{0, 1\}^{\ell(k)}$ and computing

$$\text{Eval}_{\text{Com}}(h; b; r, \hat{r}) := (h(r), \hat{r}, b \oplus \langle r, \hat{r} \rangle) .$$

The fact that this is a computationally binding and hiding commitment holds for any CRH (note that any proximity CRH is in particular a CRH). Moreover, any adversary that computes a valid commitment $c = (y, \hat{r}, b)$ (under the random seed h) also computes a valid image y under h ; hence, we can use the PECRH extractor to extract the commitment randomness r , such that $y \approx h(r)$ and $c \approx \text{Eval}_{\text{Com}}(h; b \oplus \langle r, \hat{r} \rangle; r, \hat{r})$, where $c = (y, \hat{r}, b) \approx c' = (y', \hat{r}', b')$ if and only if $y \approx y', \hat{r} = \hat{r}', b = b'$.

In addition, verifying a proximity commitment is done by verifying that y is proximate to an image under h . \square

7.2 From Leakage-Resilient Primitives and SNARKs to Extractable Primitives

Given the results in the previous section, naïvely, it seems that non-interactive adaptive arguments of knowledge offer a generic approach towards constructing extractable primitives: “simply add a non-interactive proof of preimage knowledge” (which might seem to be applicable even without succinctness when compression is not needed). However, this approach may actually compromise privacy because the attached proofs may leak too much information about the preimage.

One may try to overcome this problem by using non-interactive *zero-knowledge* proofs of knowledge; however, this can only be done in the common reference string model, which typically trivializes the notion of extractable functions. (Nonetheless, in later sections, we will still discuss zero-knowledge SNARKs and some of their meaningful applications in the CRS model.)

In this section, we consider a different approach towards overcoming the problem of proof-induced preimage leakage: we suggest to consider stronger (non-extractable) primitives that are resilient to bounded amounts of leakage on the preimage. Then, we can leverage the succinctness of SNARKs to claim that proving knowledge of a preimage does not leak too much and hence does not compromise the security of the primitive. Indeed, CRHs are in a sense optimally leakage-resilient OWFs; hence, the first part of Corollary 7.2 can be viewed as an application of this paradigm. Moreover, in this approach, there is no need to assume a trusted third party to set up a common reference string (as would be the case if we were to use zero-knowledge techniques).

Applying this paradigm to one-to-one subexponentially hard OWFs, yields:

Proposition 7.3. *Given any $2^{|x|^\epsilon}$ -hard OWF $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ and SNARKs, there exist extractable PEOFs (against polynomial-size adversaries). Moreover, the verifiability features of the SNARK carry over to the implied PEOF.*

Proof sketch. Let $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be any $2^{|x|^\epsilon}$ -hard OWF, where n is the size of the input. As in Proposition 5.2, we define an extractable function $\mathcal{F} = \{\mathcal{F}_k\}_k$. Let c be the constant such that any SNARK proof is bounded by \hat{k}^c for security parameter \hat{k}^c . The functions generated by \mathcal{F}_k are defined on the domain $\{0, 1\}^{k^{2c/\epsilon}}$ and are indexed by a $\text{vgrs} \leftarrow \mathcal{G}_{\mathcal{V}}(1^k)$. For $x \in \{0, 1\}^{k^{2c/\epsilon}}$, $f_{\text{vgrs}}(x) = (f(x), \Pi)$, where Π is a SNARK for the statement that “there exists an $x \in \{0, 1\}^{k^{2c/\epsilon}}$ such that $f(x) = y$ ”. As for PEGRHs, we

define the proximity relation to be $(y, \Pi) \approx (y', \Pi')$ if and only if $y = y'$. As in the proof of Proposition 7.1, proximity extraction and verifiability follow directly from the extraction and verifiability of the SNARK. We claim that \mathcal{F} is one-way with respect to \approx ; namely, given the image (y, Π) of a random x in the domain, it is infeasible to come up with an x' such that $f(x') = y$. This follows from the fact that f is $2^{|x|^\epsilon}$ -hard, and the proof Π is of length at most $|x|^{\epsilon/2}$. In particular, any polynomial-size adversary which proximity-inverts \mathcal{F} , can be transformed to a $2^{|x|^\epsilon}$ -size adversary that inverts f by simply enumerating all the (short) proofs π . \square

Note that, given SNARK with proof size $\text{polylog}(k)$, one can start from f that is only hard against quasipolynomial-size adversaries. (As noted in Section 5.3 our SNARK security proof does not scale in this way, but given stronger extractability assumptions it would.) We also note that the above reduction essentially preserves the structure of the original OWF f ; in particular, if f is one-to-one so is \mathcal{F} . Moreover, in such a case in the NP-language corresponding to f , any theorem claiming that $y = f(x)$ is a proper image has a single witness x . In this case we would get (exact) EOWF rather than PEOWF. We thus get:

Corollary 7.4. *Given any $2^{|x|^\epsilon}$ -hard one-to-one OWF and SNARKs, there exist (exactly) extractable commitments that are perfectly binding and computationally hiding (against polynomial-size adversaries).*

Proof sketch. Indeed, the EOWFs given by Proposition 7.3 would now be one-to-one, which in turn imply perfectly-binding ECOMs, using the hardcore bit construction as in Corollary 7.2 instantiated with a one-to-one EOWF. (The fact that one-to-one EOWFs imply perfectly binding ECOMs was already noted in [CD09].) \square

More extractable primitives based on SNARKs and leakage-resilience. We believe there is room to further investigate the above approach towards obtaining more powerful extractable primitives. In this context, one question that was raised by [CD09] is whether extractable *pseudorandom generators* and *pseudorandom functions* can be constructed from generic extractable primitives, e.g., EOWFs. (They show that the generic constructions of [HILL99] are not knowledge preserving.)

Our SNARK-based approach can seemingly be used to obtain two weaker variants; namely, extractable *pseudo-entropy generators* and *pseudo-entropy functions*. Specifically, the results of [DK08, RTTV08, GW11] imply that any strong enough PRG is inherently also leakage-resilient, in the sense that, even given leakage on the seed, the PRG's output still has high pseudo-entropy (specifically, *HILL entropy*). The results of Braverman et. al. [BHK11] show how to obtain the more general notion of leakage-resilient pseudo-entropy functions. We leave the investigation of these possibilities and their applicability for future work.

Non-verifiable extractable primitives. Perfectly-binding ECOMs (as given by Corollary 7.4) provide a generic way of obtaining limited extractable primitives that do not admit efficient verification (and if compression is needed SNARKs can be used on top). Specifically, one can transform a function \mathcal{F} to an extractable $\tilde{\mathcal{F}}$ as follows. The seed $\tilde{f}_{(f,g)}$ generated by $\tilde{\mathcal{F}}_k$ includes $f \leftarrow \mathcal{F}_k$ and a seed for the perfectly binding ECOM $g \leftarrow \text{Gen}_{\text{Com}}(1^k)$. To apply the sampled function on x , sample extra randomness r for the commitment, and define $\tilde{f}_{(f,g)}(x; r) = (f(x), \text{Eval}_{\text{Com}}(g; x; r))$. That is, add to $f(x)$ a perfectly-binding commitment to a preimage. The hiding property of the commitment clearly prevents the problem of leakage on x . The fact that the commitment is perfectly-binding and extractable implies that $\tilde{\mathcal{F}}$ is also extractable. Indeed, any adversary that produces a valid image, also produces a valid perfectly-binding commitment to a valid preimage; hence, using the extractor for the commitment, we obtain a valid preimage. A major caveat of this approach is that the resulting $\tilde{\mathcal{F}}$ does not support efficient image-verification; indeed, the commitment is never opened, and the seed generator does not have any trapdoor on it. At this time we are not

aware of applications for non-verifiable extractable primitives other than our non-verifiable ECRH-based SNARK construction. We leave an investigation of other possible applications of non-verifiable extractable primitives for future work.

8 Candidate ECRH and PECRH Constructions

In this section we discuss:

- a candidate construction for an ECRH, based on a Knowledge of Exponent assumption and the hardness of discrete logs; and
- a generic technique for obtaining candidate constructions for PECRHs, which we instantiate in three different ways.

As already discussed, the relaxation of ECRHs to PECRHs is crucial for (a) obtaining more candidate constructions, and (b) arguing the necessity of PECRHs to the construction of SNARKs.

8.1 ECRHs from t -Knowledge of Exponent

Recall that ECRHs are formally discussed in Definition 6.1. The *Knowledge of Exponent assumption* (KEA) [Dam92] states that any adversary that, given a generator g and a random group element g^α , manages to produce $g^x, g^{\alpha x}$, must “know” the exponent x . The assumption was later extended in [HT98, BP04], by requiring that given $g^{r_1}, g^{r_1\alpha}, g^{r_2}, g^{r_2\alpha}$ it is infeasible to produce f, f^α without “knowing” x_1, x_2 such that $f = g^{x_1 r_1} g^{x_2 r_2} = g^{x_1 r_1 + x_2 r_2}$. The t -KEA assumption is a natural extension to $t = \text{poly}(k)$ pairs $g^{r_i}, g^{\alpha r_i}$.

Assumption 8.1 (t -KEA). *There exists an efficiently-samplable ensemble $\mathcal{G} = \{\mathcal{G}_k\}$ where each $(\mathbb{G}, g) \in \mathcal{G}_k$ consists of a group of prime order $p \in (2^{k-1}, 2^k)$ and a generator $g \in \mathbb{G}$, such that the following holds. For any polynomial-size adversary \mathcal{A} there exists a polynomial-size extractor $\mathcal{E}_{\mathcal{A}}$ such that for all large enough $k \in \mathbb{N}$ and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$,*

$$\Pr_{\substack{(\mathbb{G}, g) \leftarrow \mathcal{G}_k \\ (\alpha, \mathbf{r}) \leftarrow \mathbb{Z}_p \times \mathbb{Z}_p^t}} \left[\begin{array}{l} (f, f') \leftarrow \mathcal{A}(g^{\mathbf{r}}, g^{\alpha \mathbf{r}}, z) \\ f' = f^\alpha \end{array} \wedge \begin{array}{l} \mathbf{x} \leftarrow \mathcal{E}_{\mathcal{A}}(g^{\mathbf{r}}, g^{\alpha \mathbf{r}}, z) \\ g^{\langle \mathbf{x}, \mathbf{r} \rangle} \neq f \end{array} \right] \leq \text{negl}(k) ,$$

where $|\mathbb{G}| = p$, $\mathbf{r} = (r_1, \dots, r_t)$, $g^{\mathbf{r}} = (g^{r_1}, \dots, g^{r_t})$, $\mathbf{x} = (x_1, \dots, x_t)$, and $\langle \cdot, \cdot \rangle$ denotes inner product.

A related assumption was made by Groth [Gro10]; there, instead of random r_1, \dots, r_t , the exponents are powers of the same random element, i.e., $r_i = r^i$. (As formalized in [Gro10], the assumption does not account for auxiliary inputs, but it could naturally be strengthened to do so.)

Our assumption can be viewed as a simplified version of Groth’s assumption; in particular, we could use Groth’s assumption directly to get ECRHs. Furthermore, Groth’s assumption is formally stated in bilinear groups, while in our setting bilinearity is not necessary. When considered in (non bilinear) groups where t -DDH is assumed to hold, the two assumptions are actually equivalent.¹⁵ Therefore, as Groth shows that his assumption holds in the generic group model (independently of the bilinear structure) and as t -DDH is also known to hold in this model, our assumption holds in the generic group model as well.

A candidate ECRH from t -KEA. A $(k \cdot t(k), 2k)$ -compressing ECRH \mathcal{H} can now be constructed in the natural way:

¹⁵ t -DDH asserts that, over suitable groups, tuples of the form $g^x, g^{x^2}, \dots, g^{x^t}$ are indistinguishable from random tuples.

- To sample from \mathcal{H}_k : sample $(\mathbb{G}, g) \leftarrow \mathcal{G}_k$ and $(\alpha, \mathbf{r}) \xleftarrow{U} \mathbb{Z}_p \times \mathbb{Z}_p^t$, and output $h := (\mathbb{G}, g^\alpha, g^{\alpha\mathbf{r}})$.
- To compute $h(x_1, \dots, x_t)$: output the pair $(g^{\langle \mathbf{r}, \mathbf{x} \rangle}, g^{\langle \alpha\mathbf{r}, \mathbf{x} \rangle}) = \left(\prod_{i \in [t]} g^{r_i x_i}, \prod_{i \in [t]} g^{\alpha r_i x_i} \right)$.

The extractability of \mathcal{H} easily follows from the t -KEA assumption. We show that \mathcal{H} is collision-resistant based on the hardness of computing discrete logarithms in \mathcal{G} .

Claim 8.2. *If \mathcal{C} finds a collision within \mathcal{H} w.p. ε , then we can compute discrete logarithms w.p. ε/t .*

Proof sketch. Given g^r , where $r \xleftarrow{U} \mathbb{Z}_p$, choose a random $i \in [t]$ and sample $\alpha, r_1, \dots, r_{i-1}, r_{i+1}, \dots, r_t$. Denote $r_i = r$ and $\mathbf{r} = (r_1, \dots, r_t)$. Feed \mathcal{C} with $g^\alpha, g^{\alpha\mathbf{r}}$. By our initial assumption and the independent choice of i , \mathcal{C} outputs \mathbf{x}, \mathbf{x}' such that $x_i \neq x'_i$ and $g^{\langle \mathbf{x}, \mathbf{r} \rangle} = g^{\langle \mathbf{x}', \mathbf{r} \rangle}$ w.p. at least ε/t . It follows that $r_i = (x_i - x'_i)^{-1} \sum_{j \in [k] - \{i\}} (x_j - x'_j) r_j$. \square

8.2 PECRHs from Knowledge of Knapsack

In Section 8.1 we presented a candidate ECRH based on a generalization of the Knowledge of Exponent assumption in large algebraic groups. We are now going to introduce a class of knowledge assumptions with a “lattice flavor”, which we call *Knowledge of Knapsack*, to construct candidates for the weaker notion of a proximity ECRH (PECRH). Recall that PECRHs are formally discussed in Definition 6.2.

Indeed, we are not able to achieve the strict notion of ECRH from “lattice-flavor” Knowledge of Knapsack assumptions; instead, we only obtain the “noisy” notion of ECRH that we have formalized as a PECRH (which yet is still sufficient, and essentially necessary, for constructing SNARKs, as discussed in Section 6.2). This might not be surprising, given that problems about lattices tend to involve statements about noise distributions, rather than about exact algebraic relations as in the case of t -KEA.

At high level, we define a candidate PECRH family based on knowledge assumptions of the following form: given a set of elements l_1, \dots, l_t in some group, the only way to compute a subset sum is (essentially) to pick a subset $S \subseteq [t]$ and output the subset sum $\sum_{i \in S} l_i$. As before, this is expressed by saying that for any adversary there exists an extractor such that whenever the adversary outputs a value y which happens to be a subset sum, the extractor “explains” this y by outputting a corresponding subset.

For convenience of exposition, we first define a very general “Knowledge of Knapsack” template, where the set size t , the group, and the distribution of l_i are left as parameters, along with an amplification factor λ (saying how many such subset-sum instances are to be solved simultaneously).

Hashes from knapsacks. A *knapsack* is a tuple $K = (\mathbb{H}, l_1, \dots, l_t)$, such that \mathbb{H} is (the description of) an additive finite group and $l_1, \dots, l_t \in \mathbb{H}$.

We construct hash function ensembles out of knapsack ensembles in a natural way. Given a size parameter $t = t(k)$, amplification parameter $\lambda = \lambda(k)$, and an ensemble of knapsacks $\mathcal{K} = \{\mathcal{K}_k\}_k$, we define the hash function ensemble $\mathcal{H}^{t, \lambda, \mathcal{K}} = \left\{ \mathcal{H}_k^{t, \lambda, \mathcal{K}} \right\}_k$ as follows. For $K = (\mathbb{H}, l_1, \dots, l_t) \leftarrow \mathcal{K}_k$, let $h^{t, K} : \{0, 1\}^t \rightarrow \mathbb{H}$ be given by $h^{t, K}(\vec{s}) := \sum_{i: s_i=1} l_i$ represented in $\{0, 1\}^{\lceil \log |\mathbb{H}| \rceil}$, where the summation is over \mathbb{H} . Then to sample $\mathcal{H}_k^{t, \lambda, \mathcal{K}}$, draw $K^1, \dots, K^\lambda \leftarrow \mathcal{K}_k$ and output the hash function $h(x) := (h^{t, K^1}(x), \dots, h^{t, K^\lambda}(x))$. (That is, h is the λ -wise repetition of $h^{t, K}$.)

Knowledge of knapsack. The *Knowledge of Knapsack* assumption with respect to $(t, \lambda, \mathcal{K}, \overset{h}{\approx}, D_h)$ asserts that the function ensemble $\mathcal{H}^{t, \lambda, \mathcal{K}}$ is proximity-extractable with respect to some proximity relation $\overset{h}{\approx}$, some extended domain $D_h \subseteq \mathbb{Z}^t$, and extended function $\bar{h} : D_h \rightarrow \mathbb{H}$ defined by taking a linear combinations with coefficients in D_h (rather than just subset sums). Explicitly:

Definition 8.3 (Knowledge of Knapsack). Let $t = t(k) \in \mathbb{N}$ (size parameter) and let $\lambda = \lambda(k) \in \mathbb{N}$ (amplification parameter). Let $\mathcal{K} = \{\mathcal{K}_k\}_k$ be an efficiently-samplable ensemble of knapsacks. For each h in the support of \mathcal{K}_k , let \approx^h be a relation on the image of h and let D_h be an extended domain $D_h \subseteq \mathbb{Z}^t$ where $D_h \supseteq \{0, 1\}$.

The Knowledge of Knapsack assumption with respect to $(t, \lambda, \mathcal{K}, \approx^h, D_h)$ states the following: for any polynomial-size adversary \mathcal{A} there exists a polynomial-size extractor $\mathcal{E}_{\mathcal{A}}$ which outputs subsets of $[t]$ such that for all large enough $k \in \mathbb{N}$ and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$,

$$\Pr_{(\mathbb{H}^j, l_1^j, \dots, l_t^j)_{j=1}^{\lambda} \leftarrow \mathcal{K}_k} \left[\begin{array}{l} (y^1, \dots, y^{\lambda}) \leftarrow \mathcal{A}(K^1, \dots, K^{\lambda}, z) \\ \exists \vec{x} \in \{0, 1\}^t \quad \forall j : y^j = \sum_i x_i l_i^j \end{array} \wedge \neg \left(\vec{x}' \in D_h \wedge \forall j : y^j \approx^h \sum_{i \in [t]} x'_i l_i^j \right) \right] \leq \text{negl}(k)$$

where j ranges over $\{1, \dots, \lambda\}$, the summations are in the group \mathbb{H} , and the multiplications mean adding an (integer number of) elements of \mathbb{H} .

Compression. If the groups in all the knapsacks in \mathcal{K} are of size $s = s(k)$ then the function ensemble $\mathcal{H}^{t, \lambda, \mathcal{K}}$ compresses (λt) -bit strings to $(\lambda \log s)$ -bit strings.

Discussion: sparseness and amplification. As discussed in Section 6.1, we wish the candidate PECRH (just like a candidate ECRH) to be superpolynomially sparse. Sparseness grows exponentially with the amplification parameter λ : if each knapsack $K \leftarrow \mathcal{K}_k$ is ρ -sparse (i.e., $|\text{Image}(h^{t, K})|/|\mathbb{H}| < \rho$), then with amplification λ we obtain the candidate PECRH $\mathcal{H}^{t, \lambda, \mathcal{K}}$ that is ρ^{λ} -sparse. Thus, for example, as long as ρ is upper-bounded by some nontrivial constant, $\lambda > \omega(\log k)$ suffices to get superpolynomial sparseness. We will indeed use this below, in candidates where the basic knapsacks \mathcal{K} must be just polynomially sparse for the proof of (proximity) collision resistance to go through.

We now proceed to propose instantiations of the Knowledge of Knapsack approach.

8.2.1 Knowledge of Knapsack of Exponents

We first point out that the Knowledge of Knapsack template can be used to express also the Knowledge of Exponent assumptions, by considering subset-sums on pairs of the form (f, f^{α}) . The result is similar to the t -KEA assumption (see Section 8.1), albeit with inferior parameters:

Assumption 8.4 (t -KKE). For $t = t(k) \in \mathbb{N}$, the t -KKE (Knowledge of Knapsack of Exponents) states that there exists an efficiently-samplable ensemble $\mathcal{G} = \{\mathcal{G}_k\}$ where each $(\mathbb{G}, g) \in \mathcal{G}_k$ consists of a multiplicative group of prime order p in $(2^{k-1}, 2^k)$ and a generator $g \in \mathbb{G}$, such that the Knowledge of Knapsack assumption with respect to $(t, 1, \mathcal{K}^{\mathbb{E}}, \equiv_{\mathbb{H}}, \{0, 1\}^t)$ holds for the ensemble $\mathcal{K}^{\mathbb{E}} = \{\mathcal{K}_k^{\mathbb{E}}\}_k$ defined as follows (where $\equiv_{\mathbb{H}}$ is equivalence in the group \mathbb{H} given below):

To sample from $\mathcal{K}_k^{\mathbb{E}}$, draw $(\mathbb{G}, g) \leftarrow \mathcal{G}_k$, let $\mathbb{H} = \mathbb{G} \times \mathbb{G}$ considered as an additive group, draw $\alpha \leftarrow \mathbb{Z}_p$ and $\mathbf{r} \leftarrow \mathbb{Z}_p^t$, let $l_i = (g^{r_i}, g^{\alpha r_i}) \in \mathbb{H}$, and output $(\mathbb{H}, l_1, \dots, l_t)$.

The hash function ensemble $\mathcal{H}^{t, 1, \mathcal{K}^{\mathbb{E}}}$ is readily verified to be $(t(k), 2k)$ -compressing, and collision-resistant assuming the hardness of taking discrete logs. Note that its range is indeed sparse, as prescribed in Section 6.1: for $h \leftarrow \mathcal{H}^{t, 1, \mathcal{K}^{\mathbb{E}}}$, $|\text{Image}(h)|/|\mathbb{H}| = |\mathbb{G}|/|\mathbb{G} \times \mathbb{G}| \approx 1/2^k$. Alas, we lost a factor of k in the compression compared to directly using t -KEA, since we hash t bits as opposed to t group elements as in t -KEA.

8.2.2 Knowledge of Knapsack of Noisy Multiples

Next, we propose a new knowledge assumption based on the following problem: given noisy integer multiples $L = (l_1, \dots, l_t)$ in \mathbb{Z}_N of a secret real number α (of magnitude about \sqrt{N}), find a subset-sum of these multiples.¹⁶ The knowledge assumption says (roughly) that whenever an efficient adversary produces such a subset sum, it knows the corresponding subset. This however requires care, since, taken literally, the assumption is clearly false. To motivate our definition, we describe several attempted attacks, and how the definition avoids them.

- *Perturbation attack.* Any small integer is close to a multiple of α (i.e., 0), and is thus likely to be a sum of *some* subset of L (when L is long enough, as it is in our setting). Thus, the adversary \mathcal{A} could simply output a random small integer and thereby challenge the extractor \mathcal{E} to find a corresponding subset. We let the extractor avoid this difficult task by using the notion of PEGRHs defined above, with the proximity relation $\overset{h}{\approx}$ chosen so that the extractor only needs to output a subset that sums to *approximately* the adversary’s output (in the above example, the extractor can output the empty set).
- *Integer-coefficients attack.* An adversary \mathcal{A} could pick an integer combination of L with coefficients that are small but not all 0 and 1. Even though this is not a valid computation of a sum over a subset of L , the result y is still close to a multiple of the secret real number, and thus, as above, is likely to be a subset sum of for *some* subset, so the extractor \mathcal{E} must “explain” y . We aid \mathcal{E} by enlarging the extended domain D_h to allow small integer coefficients, so that the (non-blackbox) extractor may output the coefficients used by the adversary.
- *Fractional-coefficients attack.* An adversary \mathcal{A} could pick a fractional combination of elements of L . For example, $l_1/2$ will be close to a multiple of α whenever l_1 happens to be close to an even multiple of α (that is, with probability half). However, we amplify our knapsack to consider λ instances concurrently (each consisting of noisy multiples L of some different α), so the extractor is challenged only in the exponentially-unlikely event that all λ instances have l_1 that is close to an even multiple.

Comparison to t -KKE. The above complications arise due to the addition of noise to l_i in the generation of the knapsack instances (otherwise α would be found computing the greatest common divisor on L , easily leading to collisions). Thus the collection of resulting subset-sums constitutes a train of “hills” (each clustered around a multiple of α), which an adversary can traverse by the aforementioned attacks. Conversely, in t -KKE from Section 8.2.1, the underlying discrete log problem does not require injection of noise, hence the subset sums constitute a set of distinct “well-spaced” points in $\mathbb{G} \times \mathbb{G}$ and so (one may hope) the adversary can navigate the structure of the image only by algebraic operations that the extractor can unravel.

Definition. Let $N \in \mathbb{Z}$, $\alpha \in \mathbb{R}$ and $\bar{\sigma} \in (0, 1)$. We define the distribution $\text{NM}_{\alpha, \bar{\sigma}, N}$ of noisy multiples of α in the range $[0, \dots, N - 1]$, with relative noise of standard deviation $\bar{\sigma}$, as follows. Draw an integer $x \xleftarrow{U} \{0, \dots, \lfloor N/\alpha \rfloor\}$ and a noise fraction $y \leftarrow \mathcal{N}_{0, \bar{\sigma}^2}$ (the normal distribution with mean 0 and variance $\bar{\sigma}^2$). Output $\lfloor \alpha(x + y) \bmod N \rfloor$.

Assumption 8.5 ((t, σ) -KKNM). For $t = t(k) > k \in \mathbb{N}$ and noise parameter $\sigma = \sigma(k) \in (0, 1)$, the (t, σ) -KKNM (Knowledge of Knapsack of Noisy Multiples) states that the Knowledge of Knapsack assumption with respect to $(t, \mathcal{K}^{\text{NM}, t, \sigma}, \overset{h}{\approx}, D_h)$ holds for the following distribution of knapsack elements.

¹⁶Our construction is inspired by a cryptosystem of Regev [Reg03, Reg04], where the public key is sampled from a similar distribution, and indeed our analysis of collision resistance and sparsity invokes Regev’s. This is elaborated below.

The ensemble $\mathcal{K}^{\text{NM},t,\sigma} = \left\{ \mathcal{K}_k^{\text{NM},t,\sigma} \right\}_k$ is sampled as follows. To sample from $\mathcal{K}_k^{\text{NM},t,\sigma}$ do the following: let $N = 2^{8k^2}$, draw $h \stackrel{\mathcal{U}}{\leftarrow} \left\{ h \in [\sqrt{N}, 2\sqrt{N}) : |h - \lfloor h \rfloor| < \frac{1}{16t} \right\}$ and draw $\bar{\sigma}$ such that $\bar{\sigma}^2 \stackrel{\mathcal{U}}{\leftarrow} [\sigma^2, 2\sigma^2)$. Let $\alpha = N/h$. Draw t values $l_1, \dots, l_t \leftarrow \text{NM}_{\alpha, \bar{\sigma}, N}$. Output $(\mathbb{Z}_N, l_1, \dots, l_t)$.

For $h \leftarrow \mathcal{K}_k^{\text{NM},t,\sigma}$, let $D_h = \{ \vec{x} \in \mathbb{Z}^t : \|\vec{x}\|_2 < t \log^2 t \}$, and let $\stackrel{h}{\approx}$ be s.t. for $y, y' \in \mathbb{Z}_N$, $y \stackrel{h}{\approx} y'$ if their distance in \mathbb{Z}_N is at most $\sqrt{N}/9$.

Relation to Regev’s cryptosystem [Reg03, Reg04]. The above distributions are essentially the same as in Regev’s cryptosystem, with minor changes for clarity in the present context. Explicitly, the mapping is as follows. The distribution $Q_\beta = (\mathcal{N}_{0,\beta/2\pi} \bmod 1)$ from [Reg04, Section 2.1] is replaced by $\mathcal{N}_{0,\bar{\sigma}^2}$, for $\beta = 2\pi\bar{\sigma}^2$ (the statistical difference between the two is negligible because $\bar{\sigma}$ will be polynomially small). The distribution $\text{NM}_{\alpha, \bar{\sigma}, N}$ is a scaling up by N of $T_{h,\beta}$ as defined in [Reg04, above Definition 4.3], for $h = N/d$ (except for the above deviation, and a deviation due to the event $x + y > h$ which is also negligible in our setting). Thus, the distribution (l_1, \dots, l_t) sampled by \mathcal{K}_{NM} is negligibly close to that of public keys in [Reg04, Section 5] on parameters $n = k$, $m = t$, $\gamma(n) = \sqrt{2/\pi} / \sigma(k)$.

Collision resistance. We show that the hash function ensemble $\mathcal{H}^{\text{KKNM}} = \mathcal{H}^{t,\lambda, \mathcal{K}^{\text{NM},t,\sigma}}$ is proximity-collision-resistant for any $t = O(k^2)$ and suitable λ and σ , assuming on the hardness of the Unique Shortest Vector Problem (uSVP) in lattices. Recall that $f(\mu)$ -uSVP is the computational problem of finding a shortest vector in a lattice of dimension μ given that the shortest vector is at least $f(\mu)$ times shorter than any other (non-parallel) lattice vector (see [Reg04, LM09]).

Claim 8.6. *The samples l_1, \dots, l_t drawn by $\mathcal{K}^{\text{NM},t,\sigma}$ are pseudorandom (i.e., indistinguishable from t random integers in the interval $\{0, \dots, N - 1\}$), assuming hardness of $(\sqrt{2/\pi}\mu/\sigma(\mu))$ -uSVP.*

Proof sketch. It suffices to show pseudorandomness for the distribution obtained by modifying $\mathcal{K}^{\text{NM},t,\sigma}$ to sample $h \stackrel{\mathcal{U}}{\leftarrow} [\sqrt{N}, 2\sqrt{N})$ (for the same reason as in [Reg04, Lemma 5.4]). This pseudorandomness follows from [Reg04, Theorem 4.5] with $g(n) = \sqrt{2\mu/\pi}/\sigma(\mu)$. \square

Claim 8.7. *The function ensemble $\mathcal{H}^{\text{KKNM}}$ is proximity-collision-resistant, with $\stackrel{h}{\approx}$, D_h, \bar{h} defined as in Assumption 8.5, for $t = O(k^2)$, assuming hardness of $\tilde{O}(\max(\mu^{3/2}, \sqrt{\mu}/\sigma(\mu)))$ -uSVP.*

Proof sketch. By Claim 8.6, the hash functions drawn by $\mathcal{H}^{\text{KKNM}}$ are indistinguishable from the ensemble \mathcal{U} of uniformly-random modular subset sums (as defined in [Reg04, Section 6]), assuming $\tilde{O}(\sqrt{\mu}/\sigma(\mu))$ -uSVP. It thus suffices to show that \mathcal{U} is proximity-collision-resistant, since this implies finding collisions in $\mathcal{H}^{\text{KKNM}}$ would distinguish it from \mathcal{U} . The ensemble \mathcal{U} is collision-resistant assuming $\tilde{O}(\mu^{3/2})$ -uSVP, by [Reg04, Theorem 6.5]. Moreover, the proximity relation $\stackrel{h}{\approx}$ is accommodated by noting that the theorem still holds if in its statement, $\sum_{i=1}^m b_i a_i \equiv 0 \pmod{N}$ is generalized to $\text{frc}((\sum_{i=1}^m b_i a_i)/N) < 1/9\sqrt{N}$; inside that theorem’s proof, this implies $\text{frc}((\sum_{i=1}^m b_i z_i)/N) < 1/8\sqrt{N}$ and thus, in the one-but-last displayed equation, $h \cdot \text{frc}((\sum_{i=1}^m b_i z_i)/N) < h/9\sqrt{N} < 1/9$ so the last displayed equation still holds and the proof follows. The extended domain D_h and induced \bar{h} are accommodated by noting that Regev’s bound $\|b\| \leq \sqrt{m}$ (in his notation) generalizes to $\|b\| \leq \tilde{O}(\sqrt{m})$. \square

Sparseness and parameter choice. To make the extractability assumption plausible, we want the function’s image to be superpolynomially sparse within its range, as discussed in Section 6.1. Consider first the distribution $\mathcal{H}^{\text{KKNM}} = \mathcal{H}^{t,1, \mathcal{K}^{\text{NM},t,\sigma}}$ (i.e., $\lambda = 1$, meaning no amplification). The image of h drawn from $\mathcal{H}^{\text{KKNM}}$ becomes “wavy” (hence sparse) when the noise (of magnitude $\sigma\alpha$) added to each multiple of α

is sufficiently small, resulting in distinct peaks, so that any subset sum of t noisy multiples is still a noisy multiple:

Claim 8.8. For $\sigma(k) = 1/16t \log^2 k$, the ensemble $\mathcal{K}^{\text{NM},t,\sigma}$ is $\frac{1}{2}$ -sparse:

$$\Pr_{h \leftarrow \mathcal{H}_k^{t,1,\mathcal{K}^{\text{NM},t,\sigma}}} [|\text{Image}(h)|/N > 1/2] < \text{negl}(k)$$

Proof sketch. In terms of the corresponding Regev public key, this means decryption failure become impossible with all-except-negligible probability over the keys. For this, it clearly suffices that each of the t noisy multiples is at most $\alpha/16t$ away from a multiple of α , so that any sum of them will have accumulated noise at most $\alpha/16$ (plus another $\alpha/16$ term due to modular reductions, as in Regev’s decryption lemma [Reg04, Lemma 5.2]). This indeed holds for $\sigma(k) = 1/16t \log^2 k$, by a tail bound on the noise terms $\alpha \mathcal{N}_{0,\sigma}$ followed by a union bound over the t samples. \square

Thus, the image becomes somewhat sparse when $\sigma = \tilde{o}(1/t)$. However, superpolynomial sparseness would require making σ superpolynomially small (and likewise a tighter distribution over h), in which case Claim 8.7 would require assuming hardness of $\mu^{\omega(1)}$ -uSVP; this assumption is unmerited in light of the excellent heuristic performance of LLL-type lattice reduction algorithms on lattices with large gaps (e.g., [GN08] conjecture, from experimental evidence, that 1.02^μ -uSVP is easy). Instead, we can set $\sigma = \tilde{\Omega}(1/k^2)$ so that Claim 8.7 needs to assume merely hardness of $\tilde{O}(\mu^{3/2})$ -uSVP, and then amplify via repetition, by choosing sufficiently large λ . In particular, by setting $\sigma(k) = \tilde{o}(1/t)$, $\lambda = \omega(\log(k))$ and $t = O(k^2)$, we indeed obtain superpolynomial sparseness.

Regarding the aforementioned integer-coefficient attack, note that the extended domain D_h allows \mathcal{E} to explain y via any vector using a linear combination whose coefficients have ℓ_2 norm at most $t \log^2 t$, since beyond this norm, the linear combination is unlikely to be in the image of h .

Lastly, note that $k = n^2$ (or, indeed, any $k = n^{1+\varepsilon}$) suffices for the SNARK construction.

Relation to other lattice hardness assumptions. The collision resistance is shown assuming hardness of the uSVP lattice problem. This can be generically translated to other (more common) lattice hardness assumptions following Lyubashevsky and Micciancio [LM09].

8.2.3 Knowledge of Knapsack of Noisy Inner Products

Further PECRH candidates can be obtained from Knowledge of Knapsack problems on other lattice-based problems. In particular, the Learning with Error problem [Reg05] problem leads to a natural knapsack ensemble, sampled by drawing a random vector $\vec{s} \in \mathbb{Z}_p^n$ and then outputting a knapsack $K = (\mathbb{Z}_p^{n+1}, l_1, \dots, l_t)$ where each l_i consists of a random vector $\vec{x} \xleftarrow{U} \mathbb{Z}_p^n$ along with the inner product $\vec{s} \cdot \vec{x} + \varepsilon$ where ε is independently-drawn noise of small magnitude in \mathbb{Z}_p . For suitable parameters this ensemble is sparse, and proximity-collision-resistant following an approach similar to KKNM above: first show pseudorandomness assuming hardness of LWE [Reg05], and then rely on the collision resistance of the uniform case (e.g., [Ajt96, GGH96, MR07]).

In this case, amplification can be done more directly, by reusing the same \vec{x} with multiple s_i instead of using the generic amplification of Definition 8.3.

9 Zero-Knowledge SNARKs

In this section we consider the problem of constructing *zero-knowledge SNARKs* (zkSNARKs); that is, we want to ensure that the succinct proof does not leak information about the witness used to generate it.

Recall that SNARKs do not require any set-up assumptions (i.e., are in the plain model). However, now that we seek the additional property of zero-knowledge, we cannot proceed in the plain model because otherwise we would obtain a two-message zero-knowledge protocol in the plain model, which is impossible [GO94]. We thus work in the standard common reference string (CRS) model.

There are two natural candidate zkSNARK constructions that one could consider, both starting with a NIZK system in the CRS model and making it succinct:

- “*SNARK on top of NIZK*”. At high level, the prover first produces a (non-succinct) NIZK argument π_{zk} for the statement y (given a valid witness w), and then produces a non-interactive succinct argument π for the statement y' that the ZK verifier would have accepted the proof π_{zk} for y .
- “*NIZK on top of SNARK*”. At high level, the prover first produces a non-interactive succinct argument π for the statement y (given a valid witness w), and then produces a NIZK argument π_{zk} for the statement y'' that the verifier would have accepted the (succinct) proof π for y .

In both constructions, one needs the ZK system to be adaptively sound. We now describe in further detail each of the above approaches.

9.1 SNARK on top of NIZK

Theorem 9.1. *If there exist adaptively-sound NIZK arguments and SNARKs, then there exist zkSNARKs. If furthermore the NIZK argument has a proof of knowledge then we obtain zkSNARKs.*

Proof sketch. The setup phase consists of generating a common reference string crs_{zk} and publishing it. A verifier then generates $(\text{vgrs}, \text{priv})$ and sends vgrs to the prover, and keeps the private verification state priv for later use. In order to prove membership for an instance y with valid witness w , the prover performs the following steps:

1. Generate, using crs_{zk} , a (non-succinct) NIZK argument of knowledge π_{zk} for the instance y using the valid witness w .
2. Generate, using vgrs , a (succinct) SNARK proof π for the NP statement “there exists a proof π_{zk} that makes the NIZK verifier accept it as a valid proof for the instance y , relative to crs_{zk} ”.
3. Send (y, π) to the verifier.

The verifier can now use $(\text{vgrs}, \text{priv})$ and crs_{zk} to verify (y, π) by running the SNARK verifier on the above NP statement.

By using the SNARK extractor, we can obtain (efficiently) a valid NIZK proof π_{zk} for the claimed theorem y . Invoking the (computational) soundness of the NIZK argument, it must be that y is a true theorem (with all except negligible probability). If the NIZK argument also guarantees an extractor, we could use it to also extract a witness for y .

As for the zero-knowledge property, it follows from the zero-knowledge property of the NIZK argument: the proof π_{zk} is “already” zero knowledge, and thus using it as a witness in the SNARK implies that the resulting succinct proof will also be zero knowledge. More formally, we can first run the simulator of the

NIZK system to obtain a simulated proof π'_{zk} for y , and then honestly generate the proof π using π'_{zk} as the witness to the NP statement. (We note that as long as the simulator for the NIZK is black-box, so will be the simulator of the zkSNARK.) \square

We note that:

- If the NIZK argument extractor requires a trapdoor for the common-reference string crs_{zk} , so will the extractor for the resulting zkSNARK.
- The common-reference string crs_{zk} of the NIZK argument needs to be of size polynomial in the security parameter (and must not depend on the theorem being proved).
- Even if zkSNARKs “live” in the common-reference string model, proofs are still only privately verifiable (if indeed the SNARK that we start with requires a designated verifier): the verifier generates $(\text{vgrs}, \text{priv})$ and sends vgrs to the prover; the prover uses both the vgrs and the common reference string crs_{zk} to produce a proof π for a theorem of his choice; the verifier then uses both $(\text{vgrs}, \text{priv})$ and crs_{zk} to verify the proof. In other words, the crs_{zk} can be used by multiple verifiers, each one of which will generate a $(\text{vgrs}, \text{priv})$ pair “on the fly” whenever they want to contact a prover. (Moreover, if the vgrs of the underlying SNARK can be reused, so can the vgrs of the resulting zkSNARK.)

For example, to obtain zkSNARKs, one may combine any SNARKs with the NIZK arguments of knowledge of Abe and Fehr [AF07] (which are based on an extended Knowledge of Exponent assumption, and happen to have an extractor that does not require a trapdoor).

Proof of knowledge strikes again. We emphasize that, in the “SNARK on top of NIZK” approach, the proof of knowledge of the SNARK was crucial for obtaining Theorem 9.1, even when only aiming for (only-sound) zkSNARKs. (Other instances where proof of knowledge played a crucial role were the results of Section 7, and thus in particular the “converse” of our main technical theorem, as well as some applications discussed in Section 10.)

9.2 NIZK on top of SNARK

Theorem 9.2. *If there exist adaptively-sound NIZK arguments of knowledge, SNARGs, and function-hiding FHE, then there exist zkSNARKs. If furthermore there exist SNARKs then we obtain zkSNARKs.*

Proof sketch. The setup phase again consists of generating a common reference string crs_{zk} and publishing it. A verifier then generates (sk, pk) for the FHE and $(\text{vgrs}, \text{priv})$ for a SNARG; then, it sends vgrs and $e := \text{Enc}_{\text{pk}}(\text{priv})$ to the prover, and keeps e and sk for later use. In order to prove membership for an instance y with valid witness w , the prover performs the following steps:

1. Generate, using vgrs , a (succinct) SNARG proof π for y ,
2. Sample randomness R for the (function-hiding) homomorphic evaluation and compute $\hat{e} = \text{Eval}_R(e, C_{y,\pi})$, where $C_{y,\pi}$ is a circuit that given input priv computes $\mathcal{V}(\text{priv}, y, \pi)$, where \mathcal{V} is the SNARG verifier.
3. Generate using crs_{zk} , a NIZK argument of knowledge π_{zk} for the the NP statement “there exist R, π such that $\hat{e} = \text{Eval}_R(e, C_{y,\pi})$ ”. (Note that verifying this NP statement is a succinct computation!)
4. Send $(y, \pi_{\text{zk}}, \hat{e})$ to the verifier.

The verifier can now use y, e, \hat{e} and crs_{zk} to verify the proof, by running the NIZK verifier and then verifying that \hat{e} decrypts to 1.

By using the NIZK extractor, we can obtain (efficiently) a valid SNARG proof π for the claimed theorem y . Invoking the semantic security of the FHE and the (computational) soundness of the SNARG, it must be that y is a true theorem (with all except negligible probability). If the SNARG also guarantees an extractor (i.e., it's a SNARK), we could use it to also extract a witness for y .

The proof $(\pi_{\text{zk}}, \hat{e})$ is zero-knowledge, because we can simulate \hat{e} (from the evaluation result “1”) by the function-hiding of the FHE and then simulate π_{zk} by running the NIZK simulator. (We note that as long as the simulator for the NIZK is black-box, so will be the simulator of the zkSNARK.) \square

Unlike in the “SNARK on top of NIZK” approach, in the “NIZK on top of SNARG” approach the knowledge property of the SNARG was not needed, but we had to additionally assume the existence of function-hiding FHE. Had the SNARG been publicly verifiable this assumption would not have been needed.

10 Applications of SNARKs and zkSNARKs

In this section we discuss applications of SNARKs and zkSNARKs to delegation of computation (Section 10.1) and secure computation (Section 10.2).

10.1 Delegation of Computation

Recall that, in a *two-message delegation scheme* (in the plain model): to delegate a T -time function F on input x , the delegator sends a message σ to the worker; the worker computes an answer (z, π) to send back to the delegator; the delegator outputs z if π is a convincing proof of the statement “ $z = F(x)$ ”. The delegator and worker time complexity are respectively bounded by $p(|F| + |x| + |F(x)| + \log T)$ and $p(|F| + |x| + |F(x)| + T)$, where p is a universal polynomial (not depending on the specific function being delegated).

(Throughout this section we ignore the requirement for input privacy because it can always be achieved by using a semantically-secure fully-homomorphic encryption scheme.)

10.1.1 Folklore Delegation from Succinct Arguments

There is a natural method to obtain a two-message delegation scheme from a designated-verifier non-interactive succinct argument for NP with *adaptive soundness*: the delegator sends the desired input x and function F to the worker (along with the verifier-generated reference string vgrs , which is independent of the statement being proved), and asks him to prove that he evaluated the claimed output z for the computation $F(x)$ correctly.

In the above paragraph, “for NP” indicates that it is enough for there to exist a protocol specialized for each NP relation (as the delegator, at delegation time, does know which NP relation is relevant for the function being delegated), but the succinctness requirement is still required to be universal, as captured by our Definition 4.1.

We note that, as long as one uses FHE in order to obtain input privacy, designated-verifier SNARGs, as opposed to publicly verifiable SNARGs, suffice, since public verification is lost anyhow upon using FHE. In fact, there is also no need to insist that the verifier-generated reference string is re-usable (beyond the logarithmically-many theorems that it can always support anyways), as a fresh string can be very quickly

generated and “sent out” along with the function and input being delegated. Thus, starting with designated-verifier non-interactive succinct arguments, is usually “enough” for delegation of computation.

Furthermore, the use of succinct arguments, in a sense, provides the “best” properties that one could hope for in a delegation scheme:

- There is *no need for preprocessing* and *no need to assume that the verifier’s answers remain secret*. All existing work providing two-message *generic* delegation schemes are in the preprocessing setting and assume that the verifier’s answers remain secret [KR06, Mie08, GKR08, KR09, GGP10, CKV10]. (Notable exceptions are the works of Benabbas et al. [BGV11] and Papamanthou et al. [PTT11], which however only deal with delegation of specific functionalities, such as polynomial functions or set operations.)
- The delegation scheme can also support inputs of the worker: one can delegate functions $F(x, x')$ where x is supplied in the first message by the delegator, and x' is supplied by the worker. (Indeed, x' acts as a “witness”.) This extension is *delegation with worker input*.

We stress once more that adaptive soundness in the setting of delegation is really needed, unless the delegator is willing to first let the worker claim an output z for the computation $F(x)$ (after communicating to him F and x), and only after that to send out the verifier-generated reference string to the worker. But in such a case the delegation scheme would have more messages, so that we might have well have used the four-message universal arguments of Barak and Goldreich (where the first message is independent of the input being proven, and one is indeed able to show adaptive soundness), and rely on standard assumptions (i.e., the existence of CRHs).

10.1.2 Our Instantiation

Our main technical result, Theorem 1, provides an instantiation, based on a simple and generic knowledge assumption (instantiated by several quite different candidates), of the designated-verifier non-interactive succinct argument for NP with adaptive soundness required for constructing a two-message delegation scheme.

Corollary 10.1. *Assume that there exists extractable collision-resistant hash functions. Then there exists a two-message delegation scheme.*

We note that previous two-message arguments for NP [DCL08, Mie08] did not provide strong enough notions of succinctness or soundness to suffice for constructing delegation schemes.

Our specific instantiation also has additional “bonuses”:

- Not only is the delegation sound, *but also has a proof of knowledge*. Therefore, $F(x, x')$ could involve cryptographic computations, which would still be meaningful because the delegator would know that a “good” input x' can be found in efficient time.

For example, the delegated function $F(x, x')$ could first verify whether the hash of a long x' is x and, if so, proceed to conduct an expensive computation; if the delegation were merely sound, the delegator would not be able to delegate such a computation, for such an x' may always exist!

We discuss more consequences of this point in the paragraph below.

- Even if the construction from our Theorem 1 formally requires the argument to depend on a constant $c \in \mathbb{N}$ bounding the time to verify the theorem, the only real dependence is a simple verification by the verifier (i.e., checking that $t \leq |x|^c$), and thus in the setting of delegation of computation we obtain

a *single* protocol because the delegator gets to choose c . Of course, despite the dependence on c , our construction still delivers the “universal succinctness” (as already remarked in Section 5.2, satisfying our Definition 4.1) required at the beginning of this section.

(Indeed, note that if the polynomial bounding the verifier time complexity is allowed to depend on the function being delegated, then a trivial solution is to just let the verifier compute the function himself, as the function is assumed to be poly-time computable!)

- When our construction is instantiated with the quasilinear-time PCPs of Ben-Sasson et al. [BSCGT13] we get essentially-optimal efficiency (up to polylogarithmic factors):
 - the delegator’s first message requires time complexity $\text{poly}(k, \log T) \tilde{O}(|F| + |x|)$;
 - the worker’s computation requires time complexity $\text{poly}(k) \tilde{O}(|F| + |x| + |F(x)| + T)$; and
 - the delegator’s verification time requires time complexity $\text{poly}(k, \log T) \tilde{O}(|F| + |x| + |F(x)|)$.

Delegating memory, streams, and authenticated data. We would like to point out that the SNARK’s adaptive proof of knowledge enables the delegator to handle “large inputs”.

Indeed, if we are interested in evaluating many functions on a large input x , we could first in an offline stage compute a Merkle hash c_x for x and then communicate x to the worker; later, in order to be convinced of $z = F(x)$, we simply ask the worker to prove us there is some \tilde{x} such that the Merkle hash of \tilde{x} is c_x and, moreover, $z = F(\tilde{x})$. By the collision resistance of the Merkle hash, the delegator is convinced that indeed $z = F(x)$ — the proof of knowledge is crucial to invoke collision resistance!

Note that x , which now “lies in the worker’s untrusted memory”, can be updated by letting the worker compute the updated x and prove that the new Merkle hash is a “good” one; moreover, if new pieces of x become available, c_x can be updated in a streaming fashion. In this way, we are able to give simple *two-message* constructions for both the tasks of *memory delegation* and *streaming delegation*, introduced by [CTY10, CKLR11] — again getting the “best” that can hope for here. The resulting schemes, based on SNARKs, are simpler than previously proposed (but of course, to instantiate the SNARKs, one may have to invoke stronger assumptions).

Of course, special cases of delegating “large datasets” such as [BGV11] and [PTT11] are also implied by our instantiation. (Though, in the case of [PTT11], we only get a designated-verifier variant.) While our construction is definitely not as practically efficient, it provides the only other construction with two messages (i.e., is “non-interactive”).

As yet another example of an application, consider the following scenario. A third party publishes on a public database a large amount of authenticated data (e.g., statistics of public interest) along with his own public key, denoted by x' and pk respectively. A worker comes along and wants to compute a function F over this data, but, because the data is so large, is only interested in learning the result z of the computation but not seeing the data itself. Relying on the proof of knowledge property (and making the inconsequential simplifying assumption that the third party only ever published a single authenticated database), the worker could ask the database to prove that there is some (x'', σ) such that $z = F(x'')$ and each entry of x'' is accompanied by a corresponding signature in σ relative to pk . In other settings, one may prefer to think as the authenticated database as private, and thus a zero-knowledge property would be needed; this can be guaranteed by either Theorem 9.1 or Theorem 9.2 in the CRS model.

In all of the above examples, the delegator is only “paying” polylogarithmically in the size of the data upon verification time.

10.1.3 Are Knowledge Assumptions Needed?

If one insists on two-message delegation of computation for all of NP, then a knowledge assumption is in part justified: the impossibility result of Gentry and Wichs [GW11] implies that there is no proof of security via any black-box reduction to a falsifiable assumption. (Note that, adaptivity in the soundness is indeed needed, because the prover does get to choose the output and thus the theorem that is being proved!)

However, such delegation may still be possible without knowledge assumptions for “natural” NP languages. (Recall that the result of [GW11] involves constructing an “unnatural” NP language.) Moreover, for delegation schemes where no input from the delegator is supported it suffices to capture languages in P, because in such a case no witness is needed. In both of these cases, we are not aware of any evidence suggesting that a knowledge assumption may be needed.

Unfortunately, at present, a two-message delegation scheme is only known to be achieved via designated-verifier non-interactive succinct arguments for NP, which fall under the negative result of Gentry and Wichs [GW11]. It is an interesting open question to avoid this negative result.

10.2 Succinct Non-Interactive Secure Computation

Non-interactive secure computation (NISC) [IKO⁺11] allows a receiver R to publish a string containing an encryption of his secret input x , so that a sender S , holding a possibly very long input y , can reveal $f(x, y)$ to R by sending him a single message. This should be done while simultaneously protecting the secrecy of y against a malicious R and preventing S from any malicious influence on the output of R . In *succinct NISC* (SNISC), we also require that the amount of work performed by R (and thus the communication complexity) is polynomial in the security parameter k and the input/output length of f (and is in particular independent of the complexity of f).

When the parties are semi-honest there are known solutions (e.g., based on fully homomorphic encryption with function privacy [Gen09]). Naor and Nissim [NN01] observe that using the succinct zero-knowledge arguments of Kilian [Kil92] one can enhance the GMW semi-honest-to-malicious compiler to be communication preserving. However, the resulting protocol is not round preserving and hence cannot be used to achieve SNISC.

Using the zkSNARKs, however, we can obtain SNISC against malicious parties in the CRS model. (Though the string published by the receiver R can only be used logarithmically many times; if the receiver wishes to receive more messages, he should publish a new string.) We achieve UC security or non-concurrent security, depending on whether the sender’s input is long or short.

Corollary 10.2. *In the CRS model, if zkSNARKs exist, so does SNISC with malicious parties, with non-concurrent security in the case of long sender input, and UC security in the case of short sender input.*

Proof sketch. We begin by discussing the long sender input case, and describe a protocol that naturally extends the semi-honest FHE-based SNISC protocol to the malicious setting.

The protocol. To jointly compute a function f :

1. The receiver R sends the verifier-generated reference string v_{grs} , an encryption c of its input x , and π_{zk}^R , a NIZK proof of knowledge of input x (and randomness for the encryption) attesting that c is a valid encryption of x .
2. The sender S verifies that π_{zk}^R is valid and aborts if not. The sender S then homomorphically evaluates $f(\cdot, y)$ on the cipher c (the evaluation is randomized to keep y private), and sends the resulting

evaluated cipher \hat{c} , together with π_{zk}^S , a zkSNARK proving knowledge of y (and randomness for the evaluation algorithm) attesting that \hat{c} is a valid (homomorphic) evaluation of $f(\cdot, y)$ on c .

3. The receiver R verifies that π_{zk}^S is a valid zkSNARK and, if so, outputs the decryption of \hat{c} or else \perp .

We stress that the amount of work done by R (including his NIZK π_{zk}^R) is independent of f 's complexity. We next briefly describe how each party is simulated.

Simulating a malicious R^* . To simulate R^* , we proceed as follows. First, generate the CRS together with a trapdoor for the NIZK of knowledge and then provide R^* with the CRS to obtain $(\text{vgrs}, c, \pi_{\text{zk}}^R)$. In case the NIZK π_{zk}^R doesn't verify, abort; otherwise, use the trapdoor to extract the input x , hand it to the trusted party, and receive $f(x, y)$. To simulate the message sent by S , simulate the evaluation result \hat{c} using the underlying plaintext $f(x, y)$; this can be done by the function privacy guarantee. Next, invoke the simulator for the zkSNARK with respect to the statement given by the simulated evaluation \hat{c} .

The validity of the simulation follows from the function privacy guarantee of the FHE and the validity zkSNARK simulator, as well as from the fact that the NIZK in use is a proof of knowledge.

Simulating a malicious S^* . To simulate S^* , we proceed as follows. Generate the CRS together with a trapdoor for the NIZK of knowledge. Simulate R^* 's message by encrypting an arbitrary string (of the proper length) to create a simulated encryption c and running the NIZK simulator with respect to the statement given by c . Also, simulate the vgrs by employing the generator of the zkSNARK. Feed S^* with the generated message to obtain a proof π_{zk}^R . Check the validity of π_{zk}^R using the CRS and the private verification state generated with vgrs. If the proof is invalid abort; otherwise, use the zkSNARK extractor to obtain the input y and hand it to the trusted party. (Note that here is where simulation makes non-black-box use of the adversary S^* , because the extractor for S^* guaranteed by the knowledge property of the zkSNARK might need to use the code of S^* .)

The validity of the simulation follows from the semantic security of the encryption and the validity of the NIZK simulator, as well as from the fact that the zkSNARK is a proof of knowledge.

We note that the FHE-based protocol above can be naturally generalized to yield a simple compiler that transforms any SNISC protocol in a slightly strengthened semi-honest model to a (still non-interactive) protocol that is secure against malicious parties in the CRS model. The strengthening of the semi-honest model is to require security with respect to parties that may choose arbitrary randomness. (In the interactive setting, the GMW compiler itself deals with this issue using "coin tossing into the well", which can not be done in the non-interactive setting.)

In the case where the input of the sender S is short, S acts instead as follows: he gives a NIZK of knowledge π' attesting that he knows his own input y and then uses a zkSNARK to prove that "there exists y (as well as randomness for the evaluation algorithm and randomness for the NIZK) for which \hat{c} is a valid evaluation of $f(\cdot, y)$ on c , and y is the witness for the NIZK π' ". Now the simulation of a malicious S^* can be performed in a black-box manner, by simply invoking the black-box extractor of the NIZK proof of knowledge; a non-black-box use of S^* is only made within the proof when the (computational) soundness of the zkSNARK is invoked. (In particular, the proof of knowledge property of the zkSNARK is not essential for the short sender input case.) \square

11 Extractable One-Way Functions and Their Applications

In this section, we consider a notion that is closely related to ECRHs: *extractable one-way functions* (EOWFs). Specifically, we formalize two strong variants of EOWFs and show that, assuming the exis-

tence of EOWFs and enhanced trapdoor permutations, there exist a *non-interactive* (two-message) selective-opening-attack-secure (SOA-secure) commitment scheme and a *3-round* concurrent zero-knowledge (ZK) protocol. Previous works showed that it is impossible to construct the former primitive from standard assumptions using black-box security reductions [Pas11] (provided one-to-one one-way functions exist), and that it is impossible to have sublogarithmic-round concurrent zero-knowledge protocols with black-box simulation [CKPR01]. Our constructions circumvent previous impossibility results by relying on the (non-black-box) extractability property of EOWFs.

The rest of this section is organized as follows. In Section 11.1 we formalize two variants of EOWFs: *strong extractable OWFs* (sEOWFs) and *strong concurrently-extractable OWFs* (scEOWFs). Next, in Section 11.4, we use sEOWFs to construct a non-interactive SOA-secure commitment scheme; the technical core of our construction is obtaining a new 3-round zero-knowledge argument-of-knowledge (ZKAOK) protocol having the special property that only the last message of the protocol depends on the statement to be proved; we believe this construction is of independent interest, and we present it first separately in Section 11.2. Next, in Section 11.3, we show that when the sEOWF used in our 3-round ZKAOK protocol is replaced with a scEOWF, the protocol is also concurrent ZK, yielding a 3-round concurrent ZK protocol. Finally, in Section 11.5 we give candidate constructions of sEOWFs and scEOWFs, based on the (original) Knowledge-of-Exponent Assumption (KEA) of [Dam92].

11.1 Definitions of sEOWFs and scEOWFs

A *strong extractable OWF* (sEOWF) is an ensemble of extractable functions that are one-to-one and everywhere one-way (namely, for every sufficiently large security parameter, *every* function in the family is hard to invert); furthermore, given a function, it is possible to efficiently verify whether the function belongs to the ensemble or not.

Definition 11.1. *A sEOWF is a function ensemble $\mathcal{F} = \{\mathcal{F}_k\}_k$ mapping $\{0, 1\}^k$ to $\{0, 1\}^{\ell(k)}$ that is extractable (see Definition 1) and additionally satisfies the following properties:*

1. **One-to-one:** *every function in \mathcal{F} is one to one (so that $\ell(k) \geq k$).*
2. **Verifiability:** *there is a polynomial-time algorithm that, given $k \in \mathbb{N}$ and f , decides whether f is in \mathcal{F}_k .*
3. **Everywhere one-wayness:** *For every polynomial-size adversary A , every sufficiently large security parameter $k \in \mathbb{N}$, every function $f \in \mathcal{F}_k$, and every auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$, the following holds:*

$$\Pr_{x \leftarrow \{0,1\}^k} \left[f(x') = f(x) \mid \begin{array}{l} y \leftarrow f(x) \\ x' \leftarrow A(1^k, z, y) \end{array} \right] \leq \text{negl}(k) .$$

Remark 11.2 (sEOWF vs. EPOW). Our notion of sEOWF is similar to the notion of extractable perfectly one-way (EPOW) function defined by Canetti and Dakdouk [CD09]. While both notions seek to formalize the notion of extractability for a one-way function, the two notions differ in their concrete hardness requirements: a EPOW requires every function in the ensemble to be a perfectly one-way function [Can97] (which is a probabilistic function whose images hides all partial information of the preimage) whereas a sEOWF ensemble only requires every function in the ensemble to be hard to invert.

We now prove a useful lemma that says that if a function ensemble \mathcal{F} is extractable then it is also *parallel extractable*, in the sense that if the adversary outputs, in parallel, many values for different functions, then, for every value that is in the image of the corresponding function, the extractor succeeds in extracting a corresponding preimage.

Lemma 11.3. *If a family ensemble $\mathcal{F} = \{\mathcal{F}_k\}_k$ is extractable then it is also parallel extractable in the following sense: for any polynomial-size adversary \mathcal{A} , there exists a polynomial-size extractor \mathcal{E} such that for all large enough $k \in \mathbb{N}$, any polynomial m , and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$:*

$$\Pr_{\mathbf{f} \leftarrow \mathcal{F}_k} \left[\begin{array}{c} \exists i \in [m(k)] \text{ s.t.} \\ y_i \in \text{Image}(f_i) \text{ and } y_i \neq f(x'_i) \end{array} \middle| \begin{array}{l} \mathbf{y} \leftarrow \mathcal{A}(\mathbf{f}, z) \\ \mathbf{x}' \leftarrow \mathcal{E}(\mathbf{f}, z) \end{array} \right] \leq \text{negl}(k) .$$

Proof. Consider the adversary $\tilde{\mathcal{A}}$ that, on input $f_i, z'_i = (z, i, f_1, \dots, f_{i-1}, f_{i+1}, \dots, f_{m(k)})$, runs $\mathcal{A}(\mathbf{f}, z)$, except that it only outputs the i -th value x_i that \mathcal{A} outputs. Let $\tilde{\mathcal{E}}$ be the extractor corresponding to $\tilde{\mathcal{A}}$. Then the parallel extractor \mathcal{E} for \mathcal{A} simply internally runs $\tilde{\mathcal{E}}$ in parallel with inputs f_i, z'_i for every i and outputs the values extracted by the parallel executions of $\tilde{\mathcal{E}}$. By definition of $\tilde{\mathcal{E}}$, except with negligible probability the i -th invocation of $\tilde{\mathcal{E}}$ returns a valid preimage for the i -th value the adversary \mathcal{A} outputs. Thus, overall, \mathcal{E} is a valid parallel extractor. \square

Next, a *strong concurrently-extractable OWF* (scEOWF) is a sEOWF where the extractability property is strengthened to *concurrent extractability*:

Definition 11.4. *A scEOWF is a sEOWF that is concurrently extractable (see below).*

We now describe the notion of concurrent extractability. The adversary and extractor participate in an *interactive* game, in which the adversary adaptively outputs multiple values while receiving the preimages the extractor obtains for each output value. Concurrent extractability requires that for every adversary there is an extractor such that, in the interactive game, the extractor succeeds (with overwhelming probability) in extracting the correct preimage for every value the adversary outputs.

Definition 11.5 (Concurrent Extractability). *A function ensemble $\mathcal{F} = \{\mathcal{F}_k\}_k$ mapping $\{0, 1\}^k$ to $\{0, 1\}^{\ell(k)}$ is concurrently extractable if for any polynomial-size adversary \mathcal{A} there exists a polynomial-size extractor \mathcal{E} such that for every large enough security parameter $k \in \mathbb{N}$, any polynomial m , and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$, the advantage of \mathcal{A} against \mathcal{E} in the experiment $\text{EXP}_{\mathcal{A}, \mathcal{E}}^{\mathcal{F}}(1^k)$ (described below) is $\text{negl}(k)$.*

The definition of $\text{EXP}_{\mathcal{A}, \mathcal{E}}^{\mathcal{F}}(1^k)$ is as follows:

$$\text{EXP}_{\mathcal{A}, \mathcal{E}}^{\mathcal{F}}(1^k) \equiv$$

1. $f_1 \cdots f_{m(k)} \leftarrow \mathcal{F}_k$;
2. $\text{st}_0 \leftarrow (f_1, \dots, f_{m(k)}, z)$;
3. run $\mathcal{A}^{\mathcal{O}}(\text{st}_0)$ until it halts, replying to the i -th oracle query (f_{j_i}, y_i) as follows:
 - (a) $(x_i, \text{st}_i) \leftarrow \mathcal{E}(f_{j_i}, y_i, \text{st}_{i-1})$;
 - (b) if $f_{j_i} \in \{f_1, \dots, f_{m(k)}\}$, $y_i \in \text{Image}(f_{j_i})$, and $y_i \neq f_{j_i}(x_i)$, abort and output fail;
 - (c) return x_i to \mathcal{A} ;
4. output 1.

The advantage of \mathcal{A} against \mathcal{E} in the above game, denoted as $\text{ADV}_{\mathcal{A}, \mathcal{E}}^{\mathcal{F}}(1^k)$, is the probability that the game's output is fail; that is,

$$\text{ADV}_{\mathcal{A}, \mathcal{E}}^{\mathcal{F}}(1^k) := \Pr \left[\text{EXP}_{\mathcal{A}, \mathcal{E}}^{\mathcal{F}}(1^k) = \text{fail} \right] .$$

Remark 11.6. Concurrent extractability is very similar to the notion of extractability required of extractable hash functions (Definition 2.8) in [DFH11], except that, in their security game, the adversary attacks only a single randomly chosen hash function, whereas in the above definition, the adversary is allowed to choose to attack any of an arbitrary polynomial number of functions adaptively.

Remark 11.7 (Auxiliary input). As noted after Definition 6.1, EOWF (and surely sEOWF or scEOWF) cannot be achieved with respect to arbitrary auxiliary-input distributions of a priori unbounded polynomial size, assuming indistinguishability obfuscation. In the applications that follow, we use EOWFs within a larger protocol, and cannot necessarily restrict the distribution on the auxiliary-input. (For example, in a zero-knowledge protocol, the instance itself should can be seen as auxiliary-input that is adversarially chosen in a worst-case manner. Yet, using similar techniques to those used in [BCPR14], we can make sure that any auxiliary-input is a priori-bounded by a fixed polynomial, a setting in which no impossibility results for EOWFs are known. Thus, all of our results can be scaled down to consider only verifiers with bounded auxiliary information.

In Section 11.5, we discuss candidate constructions of sEOWFs and scEOWFs.

11.2 A Special 3-round ZKAOK protocol

In this section, we construct a 3-round ZKAOK protocol for NP that has the special property that only the last message of the protocol depends on the statement to be proved, whereas the first two messages only depend on the size of the statement. Therefore, the statement only needs to be specified before the generation of the last message. Our construction relies on three building blocks:

1. A sEOWF family ensemble \mathcal{F} .
2. A ZAP protocol (P_Z, V_Z) for NP, that is, a two-round public-coin witness-indistinguishable (WI) argument. Such protocols exist assuming the existence of trapdoor permutations [DN00].
3. A 3-round WIAOK protocol (P_W, V_W) for NP that has two special properties. First, it satisfies that only the last message in the protocol depends on the statement to be proved (and the first two messages depend only on the size of the statement). Second, it has a strong argument of knowledge property. Namely, given two accepting transcripts (m_1, m_2^1, m_3^1) and (m_1, m_2^2, m_3^2) for two different statements x_1 and x_2 that have the same first message but different second messages $m_2^1 \neq m_2^2$, a witness of either x_1 or x_2 can be deterministically computed—we refer to this property as the *special soundness* property. It has been shown in [LS90] that such a protocol can be constructed from one-to-one one-way functions.

Given these building blocks, the 3-round ZKAOK protocol (P, V) proceeds as follows: To prove a NP statement x , the prover P and the verifier V on common input 1^k and private input a witness w of x to P exchanges the following three messages.

First message: The prover sends the following:

- A randomly sampled function $f \leftarrow \mathcal{F}_k$,
- the first message α of a (P_Z, V_Z) proof where the prover acts as the receiver, and
- the first message m_1 of a (P_W, V_W) proof where the prover acts as the prover.

Second message: The verifier sends the following:

- The images y_1, y_2 of two randomly sampled k -bit strings r_1, r_2 through the function f ,
- the second message β of (P_Z, V_Z) in response to α , proving that either y_1 or y_2 is in the range of f (the honest verifier uses the preimage of y_b chosen at random as the witness), and
- the second message m_2 of (P_W, V_W) in response to m_1 .

Third Message: The committer sends the following:

- The third message m_3 of (P_W, V_W) in response to m_1, m_2 , proving that either x is true, or the value committed to in (c_1, c_2) is a preimage of either y_1 or y_2 through f . (The honest prover uses the witness w of x as the witness.)

By construction, it is easy to see that the first two messages of the protocol do not depend on the statement to be proved. Next we first show in Lemma 11.8 that indeed this protocol is a ZKAOK for NP. Then we observe in Lemma 11.9 that by slightly extending the proof of Lemma 11.8, we can in fact show that (P, V) satisfies two stronger properties, namely, parallel ZK and *adaptive soundness*, where the latter guarantees that no efficient prover can prove a false statement with non-negligible probability, even if it can adaptively choose the false statement to prove about adaptively depending on the verifier’s message (before the generation of the last message).

Lemma 11.8. *The protocol (P, V) is a ZKAOK for NP.*

Proof. We first show that (P, V) is ZK. Fix a malicious verifier V^* , a security parameter k and an auxiliary input z to V^* . We construct the simulator S for V^* . S internally runs V^* with input $(1^k, z)$ and a uniformly sampled random string r , and emulates the prover’s messages as follows:

- S emulates the first message (f, α, m_1) to V^* honestly;
- Upon receiving the second message (y_1, y_2, β, m_2) from V^* , it tries to extract a preimage of y_1 or y_2 by relying on the extractability property of sEOWF. More precisely, consider a wrapper machine \mathcal{A} that on input f and auxiliary input $z' = (1^k, z, r, \alpha, m_1)$ runs V^* internally with input $(1^k, z)$, random tape r and first message (f, α, m_1) , and outputs the two images y_1, y_2 output by V^* . It follows from the (parallel) extractability of the sEOWF \mathcal{F} that there is an extractor \mathcal{E} that on the same input f, z' outputs x_1, x_2 such that x_i is a valid preimage of y_i as long as y_i is in the range of y with overwhelming probability. The simulator internally incorporates \mathcal{E} and runs it with input f, z' to obtain x_1 and x_2 . If neither x_1 nor x_2 is a valid preimage, it aborts and outputs fail. Otherwise, it records a valid preimage x_i .
- S simulates the third message m_3 by proving that one of y_1, y_2 is in the range of f using x_i as the witness.

We show that S emulates the view of V^* correctly. First, it follows from the soundness of the ZAP protocol (P_Z, V_Z) that at least one of the images y_1 and y_2 output by V^* is in the range of f . Then by the extractability of \mathcal{F} , \mathcal{E} succeeds in extracting at least one valid preimage with overwhelming probability, and thus S outputs fail with only negligible probability. Whenever \mathcal{E} succeeds in extracting a valid preimage, S uses it as a “trapdoor” to cheat in the last message. It then follows from the witness indistinguishability property of the protocol (P_W, V_W) that the simulated view is indistinguishable from the real view.

Next we show that (P, V) is an AOK for NP. Consider an efficient prover P^* (w.l.o.g. deterministic) that for infinitely many $k \in \mathbb{N}$ and auxiliary input z , proves a statement x with probability $1/p(k)$. We

show that there is an efficient extractor \mathcal{E} , that on input $(1^k, x, z)$ and with black-box access to P^* extracts a valid witness w of x with probability at least $1/q(k) = 1/p(k)^2 - \text{negl}(k)$. The extractor \mathcal{E} simply emulates two executions of P^* with an honest verifier V , obtaining two transcripts T_1, T_2 , which share the same first message (f, α, m_1) , and have different second and third messages $(y_1^1, y_2^1, \beta^1, m_2^1), m_3^1$ and $(y_1^2, y_2^2, \beta^2, m_2^2), m_3^2$; if both T_1 and T_2 are accepting, then it extracts a witness w from the two transcripts (m_1, m_2^1, m_3^1) and (m_1, m_2^2, m_3^2) of (P_W, V_W) ; it outputs w if it is indeed a valid witness of x ; otherwise, it outputs fail.

We argue that \mathcal{E} extracts a valid witness with probability at least $1/q(k) = 1/p(k)^2 - \text{negl}(k)$. To see this, first note that since \mathcal{E} emulates two executions between P^* and the honest verifier V honestly, it happens with probability $1/p(k)^2$ that the two transcripts T_1 and T_2 collected are both accepting proofs for statement x . We show below that conditioned on this happening, except with negligible probability, the value w extracted from the two proofs of (P_W, V_W) must be a valid witness of x . If so, the probability that \mathcal{E} extracts a valid witness successfully is at least $1/p(k)^2 - \text{negl}(k)$.

Assume for contradiction that conditioned on that T_1 and T_2 are both accepting proofs of x , the value w extracted from (m_1, m_2^1, m_3^1) and (m_1, m_2^2, m_3^2) in T_1 and T_2 is not a valid witness of x with non-negligible probability $1/p'(k)$. Then we can construct a machine B that violates the everywhere one-wayness of \mathcal{F} . The machine B on input $(1^k, x, z)$ internally proceeds as \mathcal{E} does, except that, when emulating the two executions between P^* and V , it first forwards the function f from P^* externally; then upon receiving an image y of a random value externally, it assigns y to one of $y_1^1, y_2^1, y_1^2, y_2^2$ at random—let it be y_b^d —and generate the other three values honestly; furthermore, in the d^{th} execution, it emulates the second message of ZAP by proving that y_{1-b}^d is in the range of f . Finally, after extracting the value w , it checks if w is a preimage of y ; if so, it outputs w ; otherwise, it outputs fail. Since B emulates the view of P^* perfectly and extracts w as \mathcal{E} does, by our hypothesis, it occurs with probability $1/p'(k)$ that T_1, T_2 are accepting but the extracted value w is not a valid witness of x . Then it follows from the special soundness of $\langle P_W, V_W \rangle$ that except with negligible probability w must be a valid witness of either the statement of (m_1, m_2^1, m_3^1) or that of (m_1, m_2^2, m_3^2) , in other words, w must be a preimage of one of $y_1^1, y_2^1, y_1^2, y_2^2$. Furthermore, it follows from the fact that the distribution of the verifier's message are statistically close in the two executions and the WI property of ZAP that which of the four values y is assigned to is computationally hidden. Therefore, with probability $1/4 - \text{negl}(k)$, w is a preimage of y . Therefore B violates the everywhere one-wayness of \mathcal{F} with probability at least $1/5p'(n)$, which gives a contradiction. \square

Lemma 11.9. *The protocol (P, V) is parallel ZK and has adaptive soundness.*

Proof Sketch. The proof of the parallel ZK property can be easily extended from that of the ZK property. Given a malicious verifier V^* , to simulate many parallel proofs of (P, V) to it, simply consider a simulator that simulates the prover's message in each parallel proof as how the simulator in the proof of Lemma 11.8 simulates a single proof, except that it extracts the “trapdoors” in all parallel proofs by relying on the parallel extractability property of the sEOWF as shown in Lemma 11.3.

Next we show that (P, V) has adaptive soundness. Assume for contradiction that there is a malicious prover P^* , w.l.o.g. deterministic, such that for infinitely many $k \in \mathbb{N}$ and auxiliary input z , it can prove a false statement of its choice with a non-negligible probability $1/p(k)$. Then we show that we can construct a machine B' that can violate the everywhere one-wayness of \mathcal{F} . The machine B' proceeds identically to the machine B constructed in the proof of the AOK property of Lemma 11.8. As argued in the proof of Lemma 11.8, by our hypothesis, with probability $1/p(k)^2$, B' obtains two accepting transcripts T_1 and T_2 for two (potentially different) false statements x_1 and x_2 . Then by the special-soundness of (P_W, V_W) , B' must extract a preimage w of one of the four values $y_1^1, y_2^1, y_1^2, y_2^2$ with probability $1/q(k) = 1/p(k)^2 - \text{negl}(k)$.

Then by the same argument as in the proof of Lemma 11.8, w must be a preimage of the value y that B' receives externally with probability at least $1/5q(k)$. Therefore B' violates the everywhere one-wayness of \mathcal{F} and this gives a contradiction. □

Remark 11.10. Bitansky, Canetti, Paneth, and Rosen [BCPR14] show that when considering a restricted class of adversaries with *bounded* polynomial advice and unbounded polynomial running time, a weaker variant of extractable one-way functions can be constructed from standard assumptions (e.g., subexponential security of Decision Diffie-Hellman or Quadratic Residuosity). They then show how to use this weaker variant of extractable one-way functions to construct 2-message zero-knowledge arguments and 3-message zero-knowledge arguments of knowledge against adversaries of the same class. Their protocols follow the same structure as our construction above, with modifications tailored to their weaker extractable one-way functions.

11.3 A 3-round Concurrent ZK Protocol

In this section, we first recall the definition of concurrent ZK and then show that assuming that the underlying strong extractable OWF's are concurrently extractable, that is, a scEOWF, then the 3-round ZKAOK protocol (P, V) described in Section 11.2 is a concurrent ZK protocol for NP.

Definition of concurrent zero-knowledge protocols. Let (P, V) be an interactive argument for a language L . Consider a concurrent adversarial verifier V^* that on common input $1^k, x$ and auxiliary input z , interacts with *any polynomial number* of independent copies of P concurrently, without any restrictions over the scheduling of the messages in the different interactions with P . Let $\text{View}_{V^*}^P(1^k, x, z)$ denote the random variable describing the view of the adversary V^* in an interaction with P .

Definition 11.11. Let (P, V) be an interactive argument system for a language L . We say that (P, V) is *concurrent ZK* if for every PPT concurrent adversary V^* , there exists a PPT simulator \mathcal{S} , such that, it holds that the ensembles $\{\text{View}_{V^*}^P(1^k, x)\}_{k \in \mathbb{N}, x \in \{0,1\}^k \cap L}$ and $\{\mathcal{S}(1^k, x)\}_{k \in \mathbb{N}, x \in \{0,1\}^k \cap L}$ are computationally indistinguishable over $k \in \mathbb{N}$.

It was shown in [CKPR01] that it is impossible to construct a slightly sub-logarithmic round concurrent ZK protocol with black-box simulation, that is the simulator \mathcal{S} only uses black-box access to the malicious verifier V^* . Next, by relying on the existence of a scEOWF and enhanced trapdoor permutations, we show that there is a 3-round concurrent ZK protocol, which circumvents the impossibility result by using non-black-box simulation. We note that in a recent work [GS12], Gupta and Sahai formulated a new knowledge assumption and showed that it implies constant-round concurrent ZK arguments. Although their knowledge assumption is formulated in a “stand-alone” fashion, it essentially implies concurrent extractability; moreover, their protocol has five rounds.

The 3-round protocol (P, V) is concurrent ZK. We show that assuming that the underlying family ensemble \mathcal{F} used in the protocol (P, V) in section 11.2 is a scEOWF, then the protocol (P, V) is concurrent ZK.

Theorem 11.12. Let \mathcal{F} be a scEOWF. Then (P, V) is concurrent ZK.

Proof. Fix a PPT concurrent adversary V^* , a security parameter $k \in \mathbb{N}$ a statement x , and an auxiliary input z . We first construct a simulator \mathcal{S}' that, with access to an oracle \mathcal{O} that inverts the one-way functions in \mathcal{F} , simulates the view of V^* . More precisely, let \mathcal{O} be an oracle satisfying that, if it is fed with a pair

(f, y) with $f \in \mathcal{F}$ and y in the range of f , it returns a valid preimage through f ; (otherwise, it can return any value).

The machine \mathcal{S}' on auxiliary input $z' = (1^k, x, z)$ and a vector of randomly chosen functions $\mathbf{f} \leftarrow \mathcal{F}$, internally simulates a concurrent execution with $V^*(1^k, x, z)$ as follows: it emulates the first messages for V^* honestly by forwarding f_i as the randomly chosen function the i^{th} interaction; it simulates the third messages by committing to a “fake” witness—a preimage of one of the two values y_1 or y_2 from V^* sent in the second message—and cheating in the ZAP proof that it has committed to such a “fake” witness. For every interaction of (P, V) , \mathcal{S}' obtains a “fake” witness by querying the oracle on the two values y_1 and y_2 from V^* , obtaining x_1 and x_2 ; if neither x_1 nor x_2 is a valid preimage of y_1 or y_2 , \mathcal{S}' aborts and outputs fail; otherwise, \mathcal{S}' records a valid preimage and later simulates the ZAP proof of this interaction using it as a “fake” witness. Finally, \mathcal{S}' outputs the simulated view of V^* . By the soundness of ZAP proof from the malicious verifier, except with negligible probability, for every interaction of (P, V) , one of the two values y_1 and y_2 has a valid preimage; then, the oracle must return one valid preimage, and thus the probability that \mathcal{S}' outputs fail is negligible. In this case, it follows directly from the witness indistinguishability of the ZAP proof and the hiding property of the commitment scheme (C, R) that the simulated view of V^* by \mathcal{S}' is indistinguishable from the real view of V^* when interacting with an honest prover.

Next, we construct the actual simulator \mathcal{S} that emulates the oracle \mathcal{O} for \mathcal{S}' by relying on the concurrent extractability of \mathcal{F} . More precisely, by the concurrent extractability of \mathcal{F} , there is an extractor \mathcal{E} such that when replacing the oracle answers to \mathcal{S}' with the preimages extracted by $\mathcal{E}(\mathbf{f}, z')$, the probability that \mathcal{E} fails to extract a valid preimage for a value output by \mathcal{S}' that has a preimage is negligible. Furthermore, it follows from the one-to-one property of \mathcal{F} that except with negligible probability, \mathcal{E} emulates the oracle \mathcal{O} perfectly. Therefore, the output view of \mathcal{S} is indistinguishable from the real view of V^* , and we conclude the theorem. \square

Beyond Concurrent ZK. By applying the transformation of [BGGL01] to our protocol, we obtain a three-round resettably-sound concurrent ZK protocol. By additionally applying the transformation of [DGS09] to the resulting resettably-sound concurrent ZK protocol, we obtain a three-round *simultaneously-resettably* ZK protocol.

Theorem 11.13. *Assuming the existence of a scEOWF and enhanced trapdoor permutations, there is a 3-round simultaneously-resettably ZK protocol.*

11.4 Two-message Selective Opening Attack Secure Commitments

In this section, we first provide a formal definition of a SOA-secure commitment scheme and then provide a non-interactive construction of it using the 3-round adaptively-sound parallel ZK protocol constructed in Section 11.2.

Definition of SOA-secure commitments. A commitment scheme secure against selective opening attack has a strong hiding property that holds even if the adversary gets to selectively ask for the decommitment of some of the commitments it receives. We consider a strong notion of SOA-security, which is essentially the same as the simulation based definition of a SOA-secure commitment scheme in [DNRS03], but strengthens it to require indistinguishability of the simulation rather than a relation-based security guarantee. Formally, let (C, R) be a commitment scheme. We compare between a real and an ideal execution. In the real experiment, the adversary gets to interact with the honest committer C in the commit stages of $m = m(k)$ commitments to values v_1, \dots, v_m sampled from some distribution D , and may adaptively ask for decommitments of any commitment c_i , where i is a part of a “legal” subset $I \subseteq \{1, \dots, m(k)\}$. In the

ideal experiment, the adversary simply gets to ask for the values v_i in the legal subset $i \in I$. Let $\text{real}((C, R), D, I, \mathcal{A}, 1^k)$ denote the view of an adversary \mathcal{A} in the following experiment:

- Sample (\mathbf{x}, z) from D , and for each $i \in |\mathbf{x}|$, engage with the adversary $\mathcal{A}(z)$ in the commit stage of (C, R) committing to value x_i , where \mathbf{x} is a vector of length $m(k)$ and x_i is the i^{th} component of \mathbf{x} .
- \mathcal{A} chooses a subset $J \subseteq I$. For every $j \in J$, decommit the j^{th} commitment to value x_j by sending the corresponding decommitment string d_j .

Let $\text{ideal}(D, I, \mathcal{S}, 1^k)$ denote the output of the machine \mathcal{S} in the following experiment:

- Sample (\mathbf{x}, z) from D . Feed $(1^k, z)$ to \mathcal{S} .
- \mathcal{S} chooses a subset $J \subseteq I$. For every $j \in J$, feed x_j to \mathcal{S} .

Definition 11.14 (Hiding under Selective Decommitment). *Let (C, R) be a commitment scheme. We say that (C, R) is secure under selective decommitment w.r.t. the legal set $\mathcal{I} = \{I_k\}_{k \in \mathbb{N}}$, where $I_k \in [m(k)]$ and the ensemble of distributions $\mathcal{D} = \{D_k\}_{k \in \mathbb{N}}$, where D_k is a distribution over $(\{0, 1\}^{\text{poly}(k)})^{m(k)+1}$; if for every non-uniform PPT adversary \mathcal{A} , there exists a non-uniform PPT \mathcal{S} such that the following two ensembles are indistinguishable.*

- $\{\text{real}((C, R), D_k, I_k, \mathcal{A}, 1^k)\}_{k \in \mathbb{N}}$
- $\{\text{ideal}(D_k, I_k, \mathcal{S}, 1^k)\}_{k \in \mathbb{N}}$

It is implied by the impossibility result in [Pas11] that assuming the existence of one-to-one one-way functions, then there exist legal set \mathcal{I} and distribution \mathcal{D} , such that, it is impossible to construct a non-interactive SOA-secure commitment scheme w.r.t. \mathcal{I} and \mathcal{D} based on any *bounded-round* assumptions¹⁷ via a black-box security reduction.

In contrast to the impossibility result, next we show that assuming the existence of a sEOWF and enhanced trapdoor permutations, we can construct a non-interactive SOA-secure commitment scheme. Our construction circumvents the impossibility result of [Pas11] at two aspects: First the assumption of the existence of sEOWF is not a bounded-round assumption, and second, our security reduction is non-black-box.

A non-interactive SOA-secure commitment scheme (\tilde{C}, \tilde{R}) . The construction makes use of the above 3-round adaptively-sound parallel ZK protocol (P, V) and a basic one-message statistically-binding commitment scheme Com , which exists assuming one-to-one one-way functions. To commit to a value v , the committer \tilde{C} and the receiver \tilde{R} on common input a security parameter 1^k proceeds as follow:

First message: The committer sends a commitment c to v using Com , together with the first message m_1 of (P, V) .

Second message: The receiver replies with the second message m_2 of (P, V) .

Decommitment message: The committer sends v and the third message m_3 of (P, V) proving that v is indeed the value committed to in c . The receiver accepts if the proof (m_1, m_2, m_3) is accepting.

¹⁷A bounded-round assumption is an intractability assumption formalized using a bounded-round interactive game between an adversary and a challenger. It is similar to the notion of falsifiable assumption proposed by [Nao03] but does not require the challenger to be efficient [Pas11].

Theorem 11.15. (\tilde{C}, \tilde{R}) is computationally binding and computationally hiding under selective decommitment w.r.t. any distribution ensemble \mathcal{D} and legal set ensemble \mathcal{I} .

Proof Sketch. It follows from the adaptive soundness of (P, V) and the statistically binding property of Com that the commitment scheme (\tilde{C}, \tilde{R}) is computationally binding. To show that it is also hiding under selective decommitment, consider a distribution ensemble $\mathcal{D} = \{D_k\}_k$, a legal set ensemble $\mathcal{I} = \{I_k\}_k$ and an adversary \mathcal{A} , and fix a security parameter $k \in \mathbb{N}$. We construct a simulator \mathcal{S} that in an ideal experiment $\text{ideal}(D_k, I_k, \mathcal{S}, 1^k)$ proceeds as follows:

- Upon receiving $1^k, z$, \mathcal{S} internally runs $\mathcal{A}(1^k, z)$, and sends the first messages of $m = m(k)$ commitments of 0^n to \mathcal{A} .
- Upon receiving the request from \mathcal{A} for the decommitment of commitments in subset $J \subseteq I_k$, \mathcal{S} forwards J externally and obtains values x_j for every $j \in J$.
- \mathcal{S} sends \mathcal{A} x_j 's and simulates the decommitment messages to x_j 's, which are simply the third messages of (P, V) , by relying on the parallel ZK property of (P, V) .

It follows from the parallel ZK property of (P, V) and the hiding property of Com that the simulation is indistinguishable. \square

11.5 Candidate Constructions of sEOWF and scEOWF

A candidate construction of sEOWF. We show that the construction of ECRH in Section 8.1 when instantiated with the 1-KEA assumption (Assumption 8.1 in Section 8.1) is essentially already a sEOWF (up to a slight modification). For completeness, we provide the construction below.

Let $\mathcal{G} = \{\mathcal{G}_k\}$ be an efficiently-samplable ensemble for which the 1-KEA assumption holds. That is, each \mathcal{G}_k consists of a group \mathbb{G} of prime order $p \in (2^{k-1}, 2^k)$ and a generator $g \in \mathbb{G}$, and it holds that for any polynomial-size adversary \mathcal{A} there exists a polynomial-size extractor \mathcal{E} , such that, whenever \mathcal{A} , given $(g^r, g^{\alpha r}, z)$ for r and α chosen randomly from \mathbb{Z}_p , outputs a valid tuple (c, c') such that $c' = c^\alpha$, the extractor $\mathcal{E}(g^r, g^{\alpha r}, z)$ finds a discrete logarithm x such that $g^{xr} = c$. Now a sEOWF can be constructed from the KEA assumption as follows:

- To sample from \mathcal{F}_k : set $(\mathbb{G}, g) \leftarrow \mathcal{G}_k$ and sample $(r, \alpha) \xleftarrow{U} \mathbb{Z}_p^* \times \mathbb{Z}_p^*$ where $p = |\mathbb{G}|$, and output $f := (\mathbb{G}, g^r, g^{\alpha r})$. (Note that, different from the construction of ECRH in Section 8.1, α and r are sampled from \mathbb{Z}_p^* instead of \mathbb{Z}_p .)
- To compute $f(x)$: output the pair $(g^{xr}, g^{x\alpha r})$.

Assuming the 1-KEA assumption and that the hardness of computing discrete logarithms in \mathcal{G} , the above construction is a sEOWF. First, it follows directly from the 1-KEA assumption and the almost the same argument for the construction of ECRH from t -KEA in Section 8.1 that the above construction is extractable. Furthermore, by construction, g^{xr} for every $r \in \mathbb{Z}_p^*$ uniquely determines x and thus every function in the ensemble is one-to-one. Finally, it follows from the hardness of discrete logs that for sufficiently large security parameter k , every function in the family \mathcal{F}_k is hard to invert.

A candidate construction of scEOWF. To obtain a scEOWF, we use the same construction as above, but assuming a stronger *concurrent* 1-KEA assumption, which simply extends the 1-KEA assumption to allow concurrent extraction in a similar way as in the definition of concurrent extractability property (Definition 11.5).

Assumption 11.16 (Concurrent 1-KEA). *There exists a family $\mathcal{G} = \{\mathcal{G}_k\}$ where each \mathcal{G}_k consists of a group \mathbb{G} of prime order $p \in (2^{k-1}, 2^k)$ and a generator $g \in \mathbb{G}$, such that the following holds. For any polynomial-size adversary \mathcal{A} there exists a polynomial-size extractor \mathcal{E} such that for all large enough $k \in \mathbb{N}$, any polynomial m , and any auxiliary input $z \in \{0, 1\}^{\text{poly}(k)}$, the advantage of the adversary in the following experiment is bounded by $\text{negl}(k)$.*

$\text{EXP}_{\mathcal{A}, \mathcal{E}}^{\mathcal{G}}(1^k)$:

- Let $(\mathbb{G}, g) \leftarrow \mathcal{G}_k$. Sample $\alpha_1 \cdots \alpha_{m(k)} \xleftarrow{U} \mathbb{Z}_p$ and $r_1 \cdots r_{m(k)} \xleftarrow{U} \mathbb{Z}_p$, where $p = |\mathbb{G}|$. Set $St_0 = (g^{r_1}, \dots, g^{r_{m(k)}}, g^{\alpha_1 r_1}, \dots, g^{\alpha_{m(k)} r_{m(k)}}, z)$.
- Run $\mathcal{A}^{\mathcal{O}}(St_0)$ until it halts, replying to the i -th oracle query (j_i, c_i, c'_i) as follows:
 - $(x_i, St_i) = \mathcal{E}(j_i, c_i, c'_i, St_{i-1})$
 - If $j_i \in [m(k)]$, $c'_i = c_i^{\alpha_{j_i}}$, and $c_i \neq g^{r_{j_i} x_i}$, abort and output fail.
 - Return x_i to \mathcal{A} .
- Output 1.

The advantage of the adversary in the above game, denoted as $\text{ADV}_{\mathcal{A}, \mathcal{E}}^{\mathcal{G}}(1^k)$, is the probability that the game outputs fail, that is,

$$\text{ADV}_{\mathcal{A}, \mathcal{E}}^{\mathcal{G}}(1^k) = \Pr[\text{EXP}_{\mathcal{A}, \mathcal{E}}^{\mathcal{G}}(1^k) = \text{fail}]$$

It is easy to see that assuming the concurrent 1-KEA assumption and that the discrete logs problem is hard in \mathcal{G} , the above mentioned construction for sEOWF is in fact a scEOWF.

Acknowledgements

Nir and Ran wish to thank Ben Riva and Omer Paneth for enlightening discussions in the early stages of this research. Eran wishes to thank Shai Halevi for early discussions about using extractable collision resistance as a solution approach, and Daniele Micciancio for a discussion of lattice-based Knowledge of Knapsacks assumptions.

References

- [ABOR00] William Aiello, Sandeep N. Bhatt, Rafail Ostrovsky, and Sivaramakrishnan Rajagopalan. Fast verification of any remote procedure call: Short witness-indistinguishable one-round proofs for NP. In *Proceedings of the 27th International Colloquium on Automata, Languages and Programming*, pages 463–474, 2000.
- [AF07] Masayuki Abe and Serge Fehr. Perfect NIZK with adaptive soundness. In *Proceedings of the 4th Theory of Cryptography Conference*, pages 118–136, 2007.
- [AIK10] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. From secrecy to soundness: efficient verification via secure computation. In *Proceedings of the 37th International Colloquium on Automata, Languages, and Programming*, pages 152–163, 2010.
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems. In *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing*, pages 99–108, 1996.
- [BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences*, 37(2):156–189, 1988.
- [BCCT11] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. Cryptology ePrint Archive, Report 2011/443, 2011. <http://eprint.iacr.org/>.
- [BCCT12] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS '12*, pages 326–349, 2012.
- [BCCT13] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. Recursive composition and bootstrapping for SNARKs and proof-carrying data. In *Proceedings of the 45th ACM Symposium on the Theory of Computing, STOC '13*, pages 111–120, 2013.
- [BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In *STOC*, 2014.
- [BG08] Boaz Barak and Oded Goldreich. Universal arguments and their applications. *SIAM Journal on Computing*, 38(5):1661–1694, 2008. Preliminary version appeared in CCC '02.
- [BGGL01] Boaz Barak, Oded Goldreich, Shafi Goldwasser, and Yehuda Lindell. Resettably-sound zero-knowledge and its applications. In *FOCS*, pages 116–125, 2001.
- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *SIAM Journal on Computing*, 59(2), 2012.
- [BGV11] Siavosh Benabbas, Rosario Gennaro, and Yevgeniy Vahlis. Verifiable delegation of computation over large datasets. In *Proceedings of the 31st Annual International Cryptology Conference*, pages 111–131, 2011.
- [BHK11] Mark Braverman, Avinatan Hassidim, and Yael Tauman Kalai. Leaky pseudo-entropy functions. In *Proceedings of the 2nd Symposium on Innovations in Computer Science*, pages 353–366, 2011.
- [BHZ87] Ravi B. Boppana, Johan Håstad, and Stathis Zachos. Does co-NP have short interactive proofs? *Information Processing Letters*, 25(2):127–132, 1987.
- [Blu81] Manuel Blum. Coin flipping by telephone. In *Proceedings of the 18th Annual International Cryptology Conference*, pages 11–15, 1981.
- [BP04] Mihir Bellare and Adriana Palacio. The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In *Proceedings of the 24th Annual International Cryptology Conference*, pages 273–289, 2004.

- [BP13] Elette Boyle and Rafael Pass. Limits of extractability assumptions with distributional auxiliary input. Cryptology ePrint Archive, Report 2013/703, 2013. <http://eprint.iacr.org/>.
- [BSCGT13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. On the concrete efficiency of probabilistically-checkable proofs. In *Proceedings of the 45th ACM Symposium on the Theory of Computing*, STOC '13, 2013.
- [BSGH⁺05] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. Short PCPs verifiable in polylogarithmic time. In *Proceedings of the 20th Annual IEEE Conference on Computational Complexity*, pages 120–134, 2005.
- [BSS08] Eli Ben-Sasson and Madhu Sudan. Short PCPs with polylog query complexity. *SIAM Journal on Computing*, 38(2):551–607, 2008.
- [BSW11] Dan Boneh, Gil Segev, and Brent Waters. Targeted malleability: Homomorphic encryption for restricted computations. Cryptology ePrint Archive, Report 2011/311, 2011.
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science*, 2011.
- [Can97] Ran Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In *Proceedings of the 17th Annual International Cryptology Conference*, pages 455–469, 1997.
- [CD08] Ran Canetti and Ronny Ramzi Dakdouk. Extractable perfectly one-way functions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, pages 449–460, 2008.
- [CD09] Ran Canetti and Ronny Ramzi Dakdouk. Towards a theory of extractable functions. In *Proceedings of the 6th Theory of Cryptography Conference*, pages 595–613, 2009.
- [CKLR11] Kai-Min Chung, Yael Kalai, Feng-Hao Liu, and Ran Raz. Memory delegation. In *Proceeding of the 31st Annual Cryptology Conference*, pages 151–168, 2011.
- [CKPR01] Ran Canetti, Joe Kilian, Erez Petrank, and Alon Rosen. Black-box concurrent zero-knowledge requires $\tilde{\omega}(\log n)$ rounds. In *STOC '01*, pages 570–579, 2001.
- [CKV10] Kai-Min Chung, Yael Kalai, and Salil Vadhan. Improved delegation of computation using fully homomorphic encryption. In *Proceedings of the 30th Annual International Cryptology Conference*, pages 483–501, 2010.
- [CMS99] Christian Cachin, Silvio Micali, and Markus Stadler. Computationally private information retrieval with polylogarithmic communication. In *Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques*, pages 402–414, 1999.
- [CRR11] Ran Canetti, Ben Riva, and Guy N. Rothblum. Two 1-round protocols for delegation of computation. Cryptology ePrint Archive, Report 2011/518, 2011.
- [CT10] Alessandro Chiesa and Eran Tromer. Proof-carrying data and hearsay arguments from signature cards. In *Proceedings of the 1st Symposium on Innovations in Computer Science*, pages 310–331, 2010.
- [CTY10] Graham Cormode, Justin Thaler, and Ke Yi. Verifying computations with streaming interactive proofs. Technical report, 2010. ECCS TR10-159.
- [Dak09] Ronny Ramzi Dakdouk. *Theory and Application of Extractable Functions*. PhD thesis, Yale University, Computer Science Department, December 2009.
- [Dam92] Ivan Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In *Proceedings of the 11th Annual International Cryptology Conference*, pages 445–456, 1992.
- [DCL08] Giovanni Di Crescenzo and Helger Lipmaa. Succinct NP proofs from an extractability assumption. In *Proceedings of the 4th Conference on Computability in Europe*, pages 175–185, 2008.

- [Den06] Alexander W. Dent. The hardness of the DHK problem in the generic group model. Cryptology ePrint Archive, Report 2006/156, 2006.
- [DFH11] Ivan Damgård, Sebastian Faust, and Carmit Hazay. Secure two-party computation with low communication. Cryptology ePrint Archive, Report 2011/508, 2011.
- [DG06] Alexander Dent and Steven Galbraith. Hidden pairings and trapdoor DDH groups. In Florian Hess, Sebastian Pauli, and Michael Pohst, editors, *Algorithmic Number Theory*, volume 4076 of *Lecture Notes in Computer Science*, pages 436–451. 2006.
- [DGS09] Yi Deng, Vipul Goyal, and Amit Sahai. Resolving the simultaneous resettability conjecture and a new non-black-box simulation strategy. In *FOCS*, pages 251–260, 2009.
- [DK08] Stefan Dziembowski and Pietrzak Krzysztof. Leakage-resilient cryptography. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 293–302, 2008.
- [DLN⁺04] Cynthia Dwork, Michael Langberg, Moni Naor, Kobbi Nissim, and Omer Reingold. Succinct NP proofs and spooky interactions, December 2004. Available at www.openu.ac.il/home/mikel/papers/spooky.ps.
- [DN00] Cynthia Dwork and Moni Naor. Zaps and their applications. In *FOCS*, pages 283–293, 2000.
- [DNRS03] Cynthia Dwork, Moni Naor, Omer Reingold, and Larry J. Stockmeyer. Magic functions. *J. ACM*, 50(6):852–921, 2003.
- [DRS09] Yevgeniy Dodis, Thomas Ristenpart, and Thomas Shrimpton. Salvaging merkle-damgård for practical applications. In *Proceedings of the 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 371–388, 2009.
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: practical solutions to identification and signature problems. In *Proceedings of the 6th Annual International Cryptology Conference*, pages 186–194, 1987.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pages 169–178, 2009.
- [GGH96] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Collision-free hashing from lattice problems. Technical report, 1996. ECCC TR95-042.
- [GGP10] Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: outsourcing computation to untrusted workers. In *Proceedings of the 30th Annual International Cryptology Conference*, pages 465–482, 2010.
- [GH98] Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Information Processing Letters*, 67(4):205–214, 1998.
- [GK96] Oded Goldreich and Hugo Krawczyk. On the composition of zero-knowledge proof systems. *SIAM Journal on Computing*, 25(1):169–192, 1996.
- [GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for Muggles. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 113–122, 2008.
- [GKR10] Rosario Gennaro, Hugo Krawczyk, and Tal Rabin. Okamoto-Tanaka revisited: Fully authenticated Diffie-Hellman with minimal overhead. In *Proceedings of the 8th International Conference on Applied Cryptography and Network Security*, pages 309–328, 2010.
- [GLR11] Shafi Goldwasser, Huijia Lin, and Aviad Rubinfeld. Delegation of computation without rejection problem from designated verifier CS-proofs. Cryptology ePrint Archive, Report 2011/456, 2011.

- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989. Preliminary version appeared in STOC '85.
- [GN08] Nicolas Gama and Phong Q. Nguyen. Predicting lattice reduction. In *Proceedings of the 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 31–51, 2008.
- [GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, December 1994.
- [Gol01] Oded Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, New York, NY, USA, 2001.
- [Gol04] Oded Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, New York, NY, USA, 2004.
- [GR05] Craig Gentry and Zulfikar Ramzan. Single-database private information retrieval with constant communication rate. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming*, pages 803–815, 2005.
- [Gro10] Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In *Proceedings of the 16th International Conference on the Theory and Application of Cryptology and Information Security*, pages 321–340, 2010.
- [GS12] Divya Gupta and Amit Sahai. On constant-round concurrent zero-knowledge from a knowledge assumption. *Cryptology ePrint Archive*, Report 2012/572, 2012. <http://eprint.iacr.org/>.
- [GVW02] Oded Goldreich, Salil Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1/2):1–53, 2002.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 99–108, 2011.
- [HHR06] Iftach Haitner, Danny Harnik, and Omer Reingold. Efficient pseudorandom generators from exponentially hard one-way functions. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming*, pages 228–239, 2006.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- [HT98] Satoshi Hada and Toshiaki Tanaka. On the existence of 3-round zero-knowledge protocols. In *Proceedings of the 18th Annual International Cryptology Conference*, pages 408–423, 1998.
- [IKO⁺11] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, Manoj Prabhakaran, and Amit Sahai. Efficient non-interactive secure computation. In *Proceedings of the 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 406–425, 2011.
- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 723–732, 1992.
- [KR06] Yael Tauman Kalai and Ran Raz. Succinct non-interactive zero-knowledge proofs with preprocessing for LOGSNP. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 355–366, 2006.
- [KR09] Yael Tauman Kalai and Ran Raz. Probabilistically checkable arguments. In *Proceedings of the 29th Annual International Cryptology Conference*, pages 143–159, 2009.

- [LM09] Vadim Lyubashevsky and Daniele Micciancio. On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In *Proceedings of the 29th Annual International Cryptology Conference*, pages 577–594, 2009.
- [LS90] Dror Lapidot and Adi Shamir. Publicly verifiable non-interactive zero-knowledge proofs. In *CRYPTO*, pages 353–365, 1990.
- [Mer89] Ralph C. Merkle. A certified digital signature. In *Proceedings of the 9th Annual International Cryptology Conference*, pages 218–238, 1989.
- [Mic00] Silvio Micali. Computationally sound proofs. *SIAM Journal on Computing*, 30(4):1253–1298, 2000. Preliminary version appeared in FOCS '94.
- [Mie08] Thilo Mie. Polylogarithmic two-round argument systems. *Journal of Mathematical Cryptology*, 2(4):343–363, 2008.
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 37:267–302, April 2007.
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4:151–158, 1991.
- [Nao03] Moni Naor. On cryptographic assumptions and challenges. In *Proceedings of the 23rd Annual International Cryptology Conference*, pages 96–109, 2003.
- [Nec94] Vassiliy Ilyich Nechaev. Complexity of a determinate algorithm for the discrete logarithm. *Mathematical Notes*, 55:165–172, 1994.
- [NN01] Moni Naor and Kobbi Nissim. Communication preserving protocols for secure function evaluation. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 590–599, 2001.
- [OT89] Eiji Okamoto and Kazue Tanaka. Key distribution system based on identification information. *Selected Areas in Communications, IEEE Journal on*, 7(4):481–485, May 1989.
- [Pas11] Rafael Pass. Limits of provable security from standard assumptions. In *STOC*, pages 109–118, 2011.
- [PTT11] Charalampos Papamanthou, Roberto Tamassia, and Nikos Triandopoulos. Optimal verification of operations on dynamic sets. In *Proceeding of the 31st Annual Cryptology Conference*, pages 91–110, 2011.
- [PX09] Manoj Prabhakaran and Rui Xue. Statistically hiding sets. In *Proceedings of the The Cryptographers' Track at the RSA Conference 2009*, pages 100–116, 2009.
- [Reg03] Oded Regev. New lattice based cryptographic constructions. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 407–416, 2003.
- [Reg04] Oded Regev. New lattice-based cryptographic constructions. *Journal of the ACM*, 51(6):899–942, 2004.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 84–93, 2005.
- [RR94] Alexander A. Razborov and Steven Rudich. Natural proofs. *Journal of Computer and System Sciences*, 55:204–213, 1994.
- [RTTV08] Omer Reingold, Luca Trevisan, Madhur Tulsiani, and Salil P. Vadhan. Dense subsets of pseudorandom sets. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 76–85, 2008.
- [Sha92] Adi Shamir. $IP = PSPACE$. *Journal of the ACM*, 39(4):869–877, 1992.
- [Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In *Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques*, pages 256–266, 1997.

- [Val08] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In *Proceedings of the 5th Theory of Cryptography Conference*, pages 1–18, 2008.
- [Wee05] Hoeteck Wee. On round-efficient argument systems. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming*, pages 140–152, 2005.