

Shielding Probabilistically Checkable Proofs: Zero-Knowledge PCPs from Leakage Resilience

Mor Weiss
Bar-Ilan University
mor.weiss@biu.ac.il

Abstract

Probabilistically Checkable Proofs (PCPs) allow a randomized verifier, with oracle access to a purported proof, to probabilistically verify an input statement of the form “ $x \in \mathcal{L}$ ” by querying only few proof bits. *Zero-Knowledge* PCPs (ZK-PCPs) enhance standard PCPs to additionally guarantee that the view of any (possibly malicious) verifier querying a bounded number of proof bits can be efficiently simulated up to a small statistical distance.

The first ZK-PCP construction of Kilian, Petrank and Tardos (STOC 1997), and following constructions employing similar techniques, necessitate that the *honest* verifier make several rounds of queries to the proof. This undesirable property, which is inherent to their technique, translates into increased round complexity in cryptographic applications of ZK-PCPs.

We survey two recent ZK-PCP constructions – due to Ishai, Yang and Weiss (TCC 2016-A), and Hazay, Venkatasubramanian, and Weiss (ITC 2021) – in which the honest verifier makes a *single* round of queries to the proof. Both constructions use entirely different techniques compared to previous ZK-PCP constructions, by showing connections to the seemingly-unrelated notion of *leakage resilience*. These constructions are incomparable to previous ZK-PCP constructions: while on the one hand the honest verifier only makes a single round of queries to the proof, these ZK-PCPs either obtain a smaller (polynomial) ratio between the query complexity of the honest and malicious verifiers, or obtain a weaker ZK guarantee in which the ZK simulator is not necessarily efficient.

Contents

1	Introduction	3
1.1	Non-Adaptive ZK-PCPs	6
1.1.1	The ZK-PCPs of Hazay et al. [HVW21a]	6
1.1.2	The ZK-PCPs of Ishai et al. [IWY16]	7
1.2	Comparison Between Different ZK-PCP Constructions	9
1.3	Related Notations, Extensions, and Cryptographic Applications	10
2	Preliminaries	12
2.1	PCPs and ZK-PCPs	13
2.1.1	Restrictions, Extensions and Generalizations	15
3	ZK-PCPs from LR Encodings: The Construction of [HVW21a]	18
3.1	Main Building Block: Reconstructable Probabilistic Encodings (RPEs)	20
3.2	The ZK-PCP Construction	21
3.3	ZK-PCPs with Square-Root Gap	23
4	ZK-PCPs from LR Circuits: The Construction of [IWY16]	24
4.1	Leakage-Resilient Circuit Compilers (LRCCs)	26
4.2	A SAT-Respecting LRCC for Arithmetic Circuits	29
4.2.1	Proof of Lemma 39	33
4.3	A SAT-Respecting LRCC Against “Useful” Leakage	42
4.4	The Witness-Indistinguishable PCP	45
5	Discussion	47

1 Introduction

Proofs are a cornerstone of cryptography. They are an essential component of many cryptographic systems, guaranteeing correct execution in the presence of mutually-distrusting parties. Their applications range from mundane tasks such as proving one’s identity when signing into an email account, to general tasks such as proving honest behaviour in distributed systems, i.e., attesting that one had followed its prescribed algorithm.

Due to their centrality, in the past decades a long line of works have studied the notion of proofs, extending it far beyond the traditional notion to also allow for interaction between a *prover* \mathcal{P} and a *verifier* \mathcal{V} , as well as randomization. Various variants have arisen, depending on different properties of the system, such as the amount and type of communication allowed between the prover and verifier, their computational powers, and the type of randomness (for example, whether the verifier’s random coin tosses are private or public). These advances have reshaped theoretical computer science.

Our focus is on *Probabilistically Checkable Proofs (PCPs)* [ALM⁺92, AS92] augmented with a cryptographic *Zero-Knowledge (ZK)* property that is very useful when such proofs are used in cryptographic applications. These zero-knowledge PCPs combine aspects of (interactive) zero-knowledge proofs and PCPs, concepts which we now discuss.

Probabilistically Checkable Proofs (PCPs). A *Probabilistically Checkable Proof (PCP)* for a language \mathcal{L} allows a randomized verifier \mathcal{V} to verify a statement of the form “ $x \in \mathcal{L}$ ”, while querying only few bits of an oracle proof π that was generated by a prover \mathcal{P} .¹ More specifically, \mathcal{V} is given x as input, and \mathcal{P} – who is usually required to be efficient – may be given additional information needed to efficiently generate the proof, e.g., a witness in the case of an NP language. More generally, for an NP-language \mathcal{L} the PCP will usually be designed for some specific NP-relation $\mathcal{R}(x, w)$ associated with \mathcal{L} . In standard PCPs the prover is deterministic, whereas the verifier is randomized. (This is necessary since the verifier only queries few proof bits.) The verifier accepts true claims with probability 1, whereas the *soundness error* – namely the probability that a false claim is accepted – is small, regardless of the purported proof π^* given to the verifier. The celebrated PCP theorem [ALM⁺92, AS92, Din06] asserts that any NP language has a PCP system with soundness error $1/2$ in which the verifier \mathcal{V} reads only a *constant* number of proof bits. Moreover, \mathcal{V} is *non-adaptive*, namely its queries are determined solely by its randomness.

Interactive and Zero-Knowledge Proofs. *Interactive Proofs (IPs)* [GMR85] are a different kind of proof system, in which the efficient probabilistic verifier \mathcal{V} interacts with a prover \mathcal{P} , whose goal is to convince \mathcal{V} that $x \in \mathcal{L}$ for a joint input x . Such proofs are extremely powerful, compared to classical proofs: any PSPACE language has an interactive proof with a polynomial-time verifier [Sha90], whereas classic proofs with a polynomial-time verifier only exist for NP languages. *Zero-Knowledge (ZK) Proofs* [GMR85] are an important and useful generalization of IPs for NP, enhancing them to also guarantee privacy of the NP witness. These proofs carry no extra knowledge other than being convincing, in the sense that any information \mathcal{V} can infer from its interaction with \mathcal{P} , it could have (efficiently) computed given only the input x . This holds even in the presence of a *malicious* verifier \mathcal{V}^* , namely a verifier who arbitrarily deviates from the protocol. It is important to note that while proofs are prominently used to protect *the verifier* (guaranteeing it would reject

¹The prover entity is not traditionally included as part of a PCP system. Explicitly introducing this entity will be useful when PCPs are used for cryptographic applications, which necessitate efficient proof generation (given an appropriate witness, in the case of NP languages).

false claims), ZK is designed to protect the *prover*, guaranteeing that its private information (the NP witness, in this case) remains entirely hidden.

Zero-Knowledge PCPs. *Zero-Knowledge PCP (ZK-PCP)* [KPT97] are proofs systems that combine the advantages of PCPs and ZK proofs. These are PCPs with the additional guarantee that the view – encompassing all the knowledge which the verifier possesses of the interaction – of any (possibly malicious) verifier \mathcal{V}^* who queries an *a-priori bounded* number of proof bits can be efficiently simulated given only the input, up to a small statistical distance.

ZK-PCPs differ from traditional PCPs in several respects. First, whereas the PCP prover is traditionally deterministic, in ZK-PCPs the proof is randomized, and this is inherent to obtaining ZK. Second, ZK-PCPs are used to protect the *prover's* private information (e.g., an NP witness) against *malicious verifiers*. More specifically, the system is associated with an a-priori query bound q^* , where ZK holds only against verifiers \mathcal{V}^* who query at most q^* proof bits. We stress that this is the *only* restriction on \mathcal{V}^* , and no further assumptions or limitations are made on its computational power or the manner in which it operates. It is important to note that bounding the query complexity is inherent in systems with efficient provers – an essential requirement for such systems to be useful for cryptographic applications. Indeed, the proof has polynomial length len , so any (efficient) verifier running in time len could read the *entire* proof, thus necessarily learning some information about the witness. Finally, we note that ZK against query-bounded verifiers is a stronger guarantee than *Honest-Verifier ZK (HVZK)*, namely ZK only against the *honest* verifier. This is because the honest verifier is always query bounded.

The models also differ in the main parameters of interest. Specifically, for PCPs these consist of the randomness and query complexities of the verifier (i.e., the number of coins it tosses and the number of queries it makes to the proof), which also determine the proof length. Moreover, it is standard to consider a constant soundness error – with a verifier that queries a constant number of proof bits – since this setting has strong connections to proving hardness of approximation results. We currently have PCPs whose length is quasi-linear length in the witness length, with a constant soundness error (which can be amplified through repetition), and a non-adaptive honest verifier that queries a constant number of proof bits [BGH⁺04, Din06].

On the other hand, for ZK-PCPs we desire the soundness error to be negligible (in the input length, or some security parameter) as is standard in cryptographic systems, and thus the query complexity is necessarily polylogarithmic. The query bound q^* on malicious verifiers is another important parameter of the system, and would ideally be much larger than the query complexity of the honest verifier, e.g., polynomial in the input length. Thus, the *query gap* between the query complexity needed to verify the proof, and the number of queries a malicious verifier can make without violating ZK, would be exponential. Finally, we would like the *honest* verifier to be non-adaptive, namely to make a single round of queries to the proof. This should be contrasted with adaptive verifiers whose queries might depend on the oracle answers to previous queries, and who therefore necessarily make several rounds of queries to the proof. As we will shortly explain, whereas the verifier in traditional PCPs is non-adaptive, the honest verifier in certain ZK-PCP constructions is (inherently) adaptive. We note that similar to traditional PCPs, the proof length is also of interest.

The focus on these parameters in ZK-PCP constructions stems from their effect on the properties and parameters of cryptographic systems using ZK-PCPs. Specifically, the query complexity and adaptivity of the honest verifier translates into communication and round complexities; the query-bound on malicious verifiers corresponds to the privacy guarantee of the resultant system; and in distributed proof systems (as in, e.g., [IWY16]) the proof length and query bound translate into the total number of parties and the number of corrupted parties, respectively.

ZK-PCP Constructions. The first ZK-PCP for NP, due to Kilian, Petrank and Tardos [KPT97],² obtained a negligible soundness error with an honest verifier that queries $q = \text{polylog}(|x|)$ proof bits, and ZK against verifiers making $q^* = p(|x|)$ queries to the proof for a *fixed* polynomial p that is much smaller than the proof length, but is much larger than q . Later works [IMS12, IW14] simplified the system, making it more modular, and also generalized it to other proof models (specifically, PCPs of proximity with zero knowledge [IW14]). While obtaining a desirable exponential query gap, the honest verifier in all these constructions is *adaptive*, namely it makes several rounds of queries to the proof, a severe limitation when the system is used in cryptographic applications.³ Unfortunately, the honest verifier’s adaptivity is inherent to these constructions, as we now explain.

These ZK-PCP constructions follow the blue-print of [KPT97], who show a 2-step compiler from a standard non-ZK PCP into a ZK-PCP. In the first step, the PCP is transformed into a PCP with HVZK – a weak ZK guarantee that holds only for the *honest verifier*. In the second step, HVZK is “boosted” into full-fledged ZK, against any (possibly malicious) query-bounded verifier \mathcal{V}^* . This is obtained by forcing – through modifications made to the proof – \mathcal{V}^* ’s queries to be distributed similarly to the queries of the honest verifier \mathcal{V} .⁴ (We stress that these works do *not* make any *assumptions* on the query pattern of \mathcal{V}^* , but rather by appropriately constructing the proof they guarantee that *any* query pattern will be “harmless” in the sense that it reveals no information about the NP witness.) This restriction on the verifier’s queries is imposed by combining an information-theoretic analogue of a standard cryptographic commitment, called a “locking scheme” [KPT97], with a modified version of the PCP obtained in the first step. The proof generated in this second step requires adaptive verification, due to the structure of the modified version of the proof, as well as the use of locking schemes.

The second step can be very roughly (and somewhat inaccurately) illustrated through the following example: Alice is trying to locate a particular CD cd in Bob’s CD collection, which Bob has mixed in the following way: (1) the CDs were taken out of their cases and randomly placed back into the cases (where Bob knows which case contains which CD), and then (2) each CD case was locked in a transparent box. To get cd , Alice must first ask Bob in which case he put cd . Once this is known, she can locate the box in which this CD case is, but still needs to ask Bob for the key which unlocks the box. Since Alice cannot predict in which case cd is, she must first wait for Bob’s answer to her first query, before making her second query.⁵ In the ZK-PCP, step (1) of randomly mixing the CDs corresponds to the modifications performed to the proof to guarantee that \mathcal{V}^* ’s queries are “harmless”, specifically that \mathcal{V}^* cannot “cherry pick” specific locations in the HVZK PCP generated in the first step of the compiler. Moreover, step (2) of locking CDs in boxes corresponds to locking proof symbols in locking schemes, and even the process of unlocking (given the key) in itself is adaptive.

The cost of ZK in these ZK-PCPs is high: it incurs adaptive verification, even if the underlying PCP can be verified non-adaptively, which is indeed the case for traditional (non-ZK) PCPs. This naturally gives rise to the following research goal:

²Earlier constructions, e.g. [DFK⁺92], obtained only limited ZK guarantees such as HVZK.

³Another line of works obtain non-adaptive verification by “pushing” adaptivity to the prover side; see Section 1.3. We note that having adaptive proof generation has similar disadvantages to having adaptive verification.

⁴This is in fact a gross over-simplification – for example, \mathcal{V}^* can query many more proof bits compared to \mathcal{V} . Somewhat more accurately, these works effectively force \mathcal{V}^* ’s queries to be distributed similarly to a “repeated” version of \mathcal{V} (obtained by emulating \mathcal{V} multiple times with independent random coins). We refer the interested reader to [IW14, Wei16] for more details.

⁵The PCP-version of this example will have Bob somehow write down the list of pairs (CD, CD case), as well as the keys, in the oracle proof.

Design ZK-PCPs with a non-adaptive honest verifier, and ZK against malicious query-bounded verifiers.

Obtaining non-adaptive verification is motivated by the goal of matching the parameters of non-ZK PCPs, as well as by cryptographic applications of ZK-PCPs, in which adaptive verification translates into increased complexity of the resultant system.

The question of designing non-adaptive ZK-PCPs had remained open for nearly 20 years, until Ishai, Weiss and Yang [IWY16] gave the first construction of a non-adaptive ZK-PCP, which was followed by the non-adaptive ZK-PCP of Hazay, Venkatasubramanian and Weiss [HVW21a]. The focus of this survey is on describing and comparing these constructions.

1.1 Non-Adaptive ZK-PCPs

We survey two recent works [IWY16, HVW21a] that construct ZK-PCPs for NP with a *non-adaptive honest verifier*, obtained through a novel connection to the seemingly unrelated field of *leakage-resilient cryptography*. These constructions differ drastically from the ZK-PCP constructions described above. This is not surprising, since adaptive verification is inherent to the latter, so obtaining non-adaptive verification necessitates an entirely new approach.

Malicious Verifiers Through the Leakage-Resilience Lens. Recall that the ZK-PCPs of [KPT97, IMS12, IW14] are obtained from a PCP with weak ZK guarantees (specifically, HVZK) by effectively *restricting the malicious verifier*, i.e., forcing its queries to be distributed similarly to the queries of (multiple independent copies of) the honest verifier. The ZK-PCPs of [IWY16, HVW21a] take a different approach: instead of forcing a certain structure on \mathcal{V}^* 's queries, they classify the *type of information* which an arbitrary query-bounded \mathcal{V}^* obtains by querying the proof, and modify the proof to guarantee this type of information reveals nothing about the underlying NP witness. Their insight is that the partial information obtained by deviating from the honest verifier's query pattern constitutes *leakage* on the proof and consequently, on the underlying witness.⁶ Accordingly, they employ tools from the leakage-resilience literature to protect the witness and proof.

The works of [IWY16, HVW21a] differ in the method they use to protect against leakage. Hazay et al. [HVW21a] chose to protect *the proof* itself, whereas Ishai et al. [IWY16] protect the *process of proof generation* from the witness. Put differently, the latter protect *computation* against leakage, for which they employ leakage-resilient *circuits*, whereas the former protect *information* (the proof, once it has been generated), by using an appropriate leakage-resilient *encoding*. Consequently, these works differ in their requirements from the underlying PCP, and in the parameters and properties of the resultant ZK-PCP. We now elaborate on these differences (see also Table 1).

1.1.1 The ZK-PCPs of Hazay et al. [HVW21a]

Hazay et al. [HVW21a] construct a non-adaptive ZK-PCP in which the ratio between the bound q^* on the query complexity of a malicious verifier, and the query complexity q of the honest verifier, is polynomial.

⁶In a broader cryptographic context, *leakage* roughly refers to the information an adversary obtains by deviating from the assumed adversarial model for which the system was designed. For example, so-called "side channel" attacks – such as measuring the power consumption of an object – exploit adversarial capabilities which were not taken into account when designing the system (adversarial and attack models traditionally disregard such information, that is obtained from the *physical implementation*, and is not part of the more abstract model description). Similarly, a malicious verifier querying a PCP is capable of querying more – and different sets of – proof bits compared to the honest verifier for which the system was designed.

The leakage-resilient primitive they employ is a *Leakage-Resilient Encoding (LRE)*. Roughly, an LRE consists of a randomized efficient encoding procedure Enc mapping a bit string m to an encoding c , and an efficient deterministic decoder algorithm that given an encoding c of m , outputs m . The leakage-resilience guarantee is that for any pair m, m' , if $c \leftarrow \text{Enc}(m)$, $c' \leftarrow \text{Enc}(m')$, then any small subset of bits in c, c' are identically distributed. (They actually need a stronger leakage-resilience guarantee, see Section 3.1.) Non-explicit constructions of such encodings easily follow from the existence of linear error-correcting codes with sufficiently “good” parameters, such as random linear codes (see, e.g., [CCG⁺07]). In Section 3, we also describe an explicit construction of such encodings.

The leakage-resilience guarantee of an LRE is restricted to protecting the codeword *once it has been encoded*, and does not protect the encoding procedure itself (as opposed to leakage-resilient circuits, see below). Therefore to obtain PCPs with full-fledged ZK against arbitrary query-bounded verifiers, the underlying PCP should possess a zero-knowledge property, which is weaker than full-fledged ZK, but stronger than the HVZK property used in [KPT97, IMS12, IW14]. Specifically, they use a ZK-PCP variant *over a large alphabet*, which is much easier to obtain compared to full-fledged ZK for standard PCPs, in which the proof is binary.

A ZK-PCP variant over a large alphabet can be thought of as a (standard) PCP which is divided into “regions”, each corresponding to a single symbol in the large alphabet, where ZK is guaranteed only as long as the verifier queries “full” regions (i.e., all bits in the region). Thus, a malicious verifier does not have to follow the honest verifier’s query pattern, but ZK holds only against verifiers querying at most q^* full regions, for some a-priori bound q^* . Such ZK-PCPs are constructed from general secure multi-party computation protocols in [IKOS07].

Given a ZK-PCP variant over a large alphabet, [HVW21a] design an alphabet reduction that preserves ZK. (Naïve alphabet reduction techniques do not preserve ZK; see Section 3.) This reduction uses the underlying PCP $(\mathcal{P}', \mathcal{V}')$ as a black box, transforming it into a ZK-PCP $(\mathcal{P}, \mathcal{V})$ in which the proof is over bits. The high-level idea is to interpret each symbol of a proof π' generated by \mathcal{P}' as a bit string, and encode it using the LRE. \mathcal{V} then emulates \mathcal{V}' , answering an oracle query i by reading the entire encoding of the i 'th symbol from its proof, decoding it, and then providing \mathcal{V}' with the resultant symbol as the oracle answer.

Hazay et al. [HVW21a] then apply their alphabet reduction to the ZK-PCP variant of [IKOS07] to obtain the following (see Theorem 14 in Section 3 for the formal statement):⁷

Informal Theorem 1. *There exists a constant $\epsilon \in (0, 1)$ such that for any ZK parameter $q^* \in \mathbb{N}$ and any NP-language \mathcal{L} there exists a ZK-PCP for \mathcal{L} with ZK against q^* -query bounded verifiers, and a negligible soundness error with a non-adaptive honest verifier that queries $(q^*)^\epsilon$ proof bits.*

1.1.2 The ZK-PCPs of Ishai et al. [IWY16]

Ishai et al. [IWY16] construct a non-adaptive ZK-PCP in which the ratio between the bound q^* on the query complexity of a malicious verifier, and the query complexity q of the honest verifier, is *exponential*, but ZK holds with an inefficient simulator.

Their starting point is a *standard* PCP with no zero-knowledge guarantees. Consequently, to obtain full-fledged ZK, they rely on a stronger leakage-resilient tool. To describe their construction, it would be easier to consider the NP-relation $\mathcal{R} = \mathcal{R}(x, w)$ associated with the NP-language \mathcal{L} , and its corresponding verification circuit C . We will assume without loss of generality that a witness w for x has canonical form, namely it consists of the entire wire values of C given input x, w' , for

⁷The system described in Informal Theorem 1 requires a tighter analysis of the ZK-PCP variant of [IKOS07], which [HVW21a] provide. See Section 3 and [HVW21a] for further details.

some w' such that $C(x, w') = 1$. (In particular, w' is the information that we wish to keep secret.) The PCP prover generates the proof from this canonical NP witness w . Therefore, each proof bit is an information bit on w , i.e., on the wire values of C when evaluated on x, w' . This can be thought of as *leakage* on the *computation* in C , whose purpose is to reveal information about the secret input w' of C . While in general leakage on the wire values of C might reveal information on the secret input w' , there are tools to compile C into a *leakage-resilient* circuit \widehat{C} that resists such leakage.

However, it is well known that one cannot protect circuits against *general* polynomial-time leakage [BGI⁺01]. Consequently, leakage-resilient circuits are associated with a restricted leakage class from which leakage functions can be chosen, and the circuit only resists leakage computed by a function from the class. The main observation of [IWY16] is that every PCP system naturally has such a restricted class \mathcal{LEAK} of leakage functions associated with it. Indeed, while \mathcal{V}^* can choose which proof bits to query, it has no control over the *type of functions* applied to w to generate the proof bits – these functions are determined solely by the prover algorithm. Specifically, for every subset \mathcal{I} of at most q^* indices in the proof, the corresponding leakage function $\ell_{\mathcal{I}} \in \mathcal{LEAK}$ applies the PCP prover function to w , then outputs the restriction of the proof to the indices in \mathcal{I} .

The main building block is therefore a *Leakage-Resilient Circuit compiler (LRCC)* – a compiler that transforms a given circuit C into a leakage-resilient circuit \widehat{C} – that resists leakage from the class \mathcal{LEAK} associated with the underlying PCP system. Informally, an LRCC is associated with a function class \mathcal{LEAK} (the *leakage class*) and a (randomized) input encoding scheme $E = (\text{Enc}, \text{Dec})$, and compiles a deterministic circuit C into a deterministic circuit \widehat{C} that emulates C 's operation over encoded inputs. \widehat{C} is guaranteed to emulate C on *properly encoded inputs*, and is leakage resilient in the sense that for any pair of inputs z, z' for C such that $C(z) = C(z')$, and any $\ell \in \mathcal{LEAK}$, the output of ℓ on the wire values of \widehat{C} when evaluated on a random encoding $\widehat{z} \leftarrow \text{Enc}(z)$ is statistically close to its output when \widehat{C} is evaluated on a random encoding $\widehat{z}' \leftarrow \text{Enc}(z')$.⁸ In particular, if the prover generated the proof from the wire values \widehat{w} of \widehat{C} (instead of the wire values w of C), then \mathcal{V}^* 's queries to the proof – which constitute leakage from \mathcal{LEAK} on \widehat{w} – would reveal no information about w' .

This observation gives a general blue-print for compiling traditional PCPs into ZK-PCPs: given a PCP system $(\mathcal{P}', \mathcal{V}')$ with an associated leakage class \mathcal{LEAK} as described above, and a compiler that transforms a given circuit into one that resists leakage from \mathcal{LEAK} , the ZK-PCP system $(\mathcal{P}, \mathcal{V})$ operates as follows. \mathcal{P} on input x, w generates the leakage-resilient version \widehat{C}_x of the circuit $C_x = C(x, \cdot)$ (i.e., C with x hard-wired into it), uses w to generate the entire wire values \widehat{w} of \widehat{C}_x , then emulates \mathcal{P}' on \widehat{w} to generate a PCP π . \mathcal{V} given input x and oracle access to π generates \widehat{C}_x similarly to \mathcal{P} , and emulates \mathcal{V}' on π to check that \widehat{C}_x is satisfiable. In particular, whereas the claim which \mathcal{V} set out to verify can be phrased as “ C_x is satisfiable”, the underlying PCP system $(\mathcal{P}', \mathcal{V}')$ is used to verify a different claim, namely that *the leakage-resilient circuit \widehat{C}_x is satisfiable*.

Unfortunately, this blue-print does not actually work. The reason is that the internal system $(\mathcal{P}', \mathcal{V}')$ is used to prove a *different* statement, namely that \widehat{C}_x is satisfiable. For $(\mathcal{P}, \mathcal{V})$ to be sound, it should be the case that \widehat{C}_x is satisfiable only if C_x is. This, however, is not generally guaranteed by LRCCs, as we now explain. Recall that the LRCC is correct for properly encoded inputs, in

⁸We note that the leakage-resilience literature often considers stateful circuits C whose internal state is secret, and should be protected against leakage, but we will only need the simpler form of leakage resilience for *stateless* circuits, in which we wish to hide the circuit's input. Additionally, the leakage-resilient circuit \widehat{C} is often allowed to have *randomized, leak-free* components. Instead, we require that \widehat{C} be deterministic, where the needed randomness will be provided as part of its input encoding. As will become evident shortly, this is necessary both for correctness – to allow the verifier to generate \widehat{C} on its own – as well as for soundness.

the sense that on such inputs, \widehat{C} emulates C . However, LRCCs in general have no guarantee for inputs which are *not* properly encoded. This is not just an artifact of the definition, but is rather essential for leakage resilience to hold in existing constructions. A main technical contribution of [IWY16] is in defining and constructing an LRCC that also guarantees soundness, in the sense that the leakage-resilient circuit \widehat{C} is satisfiable (even by using invalid encodings as inputs) only if the original circuit C is satisfiable.

To turn this into an actual construction, one needs to design a sound LRCC for a leakage class \mathcal{LEAK} associated with some standard PCP system. The most common leakage classes considered in the literature are either the “Only Computation Leaks” (OCL) model that assumes leakage is “local” in the sense that different “regions” of the circuit leak independently [MR04], or classes of functions that are “computationally simple” [ISW03], i.e., from a low complexity class such as \mathcal{AC}^0 . It is known that the leakage classes associated with a PCP system cannot be of the former type [ARW17], so Ishai et al. [IWY16] focus on the latter. Specifically, they show that the PCP system of Arora and Safra [AS92] has the property that “small” subsets of proof bits can be generated using the class \mathcal{LEAK} of \mathcal{AC}^0 circuits (i.e., constant-depth polynomial-sized circuits with \wedge, \vee, \neg gates of unbounded fan-in and fan-out) augmented with a small number of \oplus gates. (See section 4.3 for a formal definition of this leakage class.) Then, they use correlation bounds of Lovett and Srivinasan [LS11] to show that their sound LRCC resists leakage from \mathcal{LEAK} . This yields the following result, where a witness-indistinguishable PCP is a ZK-PCP in which the ZK property holds with an *inefficient* simulator (see Theorem 23 in Section 4 for the formal statement).

Informal Theorem 2. *For any ZK parameter $q^* \in \mathbb{N}$ and any NP-language \mathcal{L} there exists a witness-indistinguishable PCP for \mathcal{L} with witness-indistinguishability against q^* -query bounded verifiers, and a negligible soundness error with a non-adaptive honest verifier that queries $\text{polylog}(q^*)$ proof bits.*

Assuming the existence of one-way functions (a minimal assumption in cryptography), as well as a Common Random String (CRS) that is available to both parties, and using a standard cryptographic technique – the so-called “FLS technique” [FLS90] – the witness-indistinguishable PCP of Informal Theorem 2 can be transformed into a ZK-PCP system where ZK holds against *computationally-bounded* query-bounded verifiers in the CRS model. (We refer the interested reader to [IWY16] for further details, including a formal definition of the model.)

1.2 Comparison Between Different ZK-PCP Constructions

The ZK-PCPs of [HVW21a] (Informal Theorem 1) obtain ZK with efficient simulation as in the ZK-PCPs of [KPT97, IW14], with a query gap of $q^*/(q^*)^\epsilon$. The query gap is inherited directly from the underlying ZK-PCP variant of [IKOS07], which requires the honest verifier to query $\Omega((q^*)^\epsilon)$ proof symbols to obtain a negligible soundness error. Therefore, the query gap could potentially be improved by replacing the underlying building block with a ZK-PCP variant with a larger query gap.

On the other hand, the witness-indistinguishable PCPs of [IWY16] (Informal Theorem 2) obtain an exponential query gap, similar to the ZK-PCPs of [KPT97, IW14]. This is possible because the underlying (non-ZK) PCP has a negligible soundness error with an honest verifier that queries a poly-logarithmic number of proof bits. However, the construction is only witness-indistinguishable, which is weaker than ZK (see Section 2.1.1). [IWY16] show that unless $\text{NP} = \text{BPP}$, obtaining ZK-PCPs (i.e., with efficient simulation) using their technique would require a new and entirely different approach towards designing and analyzing security of leakage-resilient circuits. We further note that while the techniques of [KPT97, HVW21a] extend also to PCPs of

Proximity (see Section 2.1.1 for a description of this model, and [IW14, HVW21a] for the constructions), the technique of [IWY16] does not seem to readily lend itself to designing zero-knowledge (or even witness-indistinguishable) PCPs of proximity.

Finally, all the aforementioned ZK-PCP constructions [KPT97, IW14, IWY16, HVW21a] have polynomial-length proofs (whereas traditional PCPs can have quasi-linear length). However, while a polynomial blowup in proof length is inherent to the constructions of [KPT97, IW14] due to their use of locking schemes, this is not the case for the leakage resilience based constructions [IWY16, HVW21a], which could potentially have shorter proofs. This is an additional advantage of taking the leakage resilience based approach. The discussion is summarized in Table 1.

	Underlying LR Primitive	ZK Quality	Query Ratio	Underlying PCP	Honest Verification
[IWY16]	Circuits	WI	Exponential	Any standard PCP	Non adaptive
[HVW21a]	Encoding	ZK	Square-Root	ZK-PCP, large alphabet	Non adaptive
[KPT97, IW14]	—	ZK	Exponential	Any standard PCP	Adaptive

Table 1: Comparison of Existing ZK-PCP constructions.

Here, “Underlying LR primitive” refers to the type of building block used in the leakage resilience based constructions; “WI” stands for witness-indistinguishable; “ZK quality” refers to the efficiency of the ZK simulator (which is efficient in ZK systems and inefficient in WI systems); “Query ratio” is q^*/q , where q^* is the bound on the query complexity of a malicious verifier (ZK holds against any verifier querying at most q^* proof bits), and q is the query complexity of the honest verifier (needed to achieve soundness); “Underlying PCP” describes the properties which the transformations needs the underlying PCP system to have; and “Honest verification” refers to the adaptivity of the honest verifier.

1.3 Related Notations, Extensions, and Cryptographic Applications

Many different variants of proof systems, and ZK proof systems, have been considered in the literature (see, e.g., Thaler’s survey [Tha22], and references therein). We briefly mention two notable notions that are closely related to ZK-PCPs; see also Table 2. The first is *Interactive Oracle Proofs (IOPs)* [BCS16, RRR16] (a special case of IOPs appeared earlier in [BCGV16]) which combine aspects of IPs and PCPs. Specifically, in an IOP the verifier and prover interact as in an IP, but the verifier has oracle access to prover messages as in PCPs. These proof systems have received increasing attention, partly due to their uses in blockchain applications. Unlike ZK-PCPs, which can be obtained from standard PCPs through generic compilers, existing ZK-IOPs (see, e.g., [BCGV16, BCL22] and references therein) are constructed in an ad-hoc manner, in which ZK is “tailored” to a specific non-ZK IOP. (It is of course preferable to have a generic compiler, since it can be used to enhance any new IOP construction to also guarantee ZK.) The second notion is PCPs of *Proximity (PCPPs)* [DR04, BGH⁺04, Din06] – a generalization of PCPs in which the verifier does not read its entire input. Instead, \mathcal{V} has oracle access to x, π , and wishes to check whether x is close to \mathcal{L} in relative Hamming distance. *Zero-Knowledge PCPPs (ZK-PCPPs)* [IW14] extend ZK-PCPs to the PCPP realm. They guarantee that the view of any verifier \mathcal{V}^* making q^* queries to the input and the proof can be efficiently simulated, up to a small statistical distance, by making only q^* queries to the input. Ishai and Weiss [IW14] construct ZK-PCPPs for NP with comparable parameters to the ZK-PCPs of [KPT97, IMS12], where soundness holds for inputs which are δ -far from the language, for δ which is constant or inverse polylogarithmic. The honest verifier in their construction is adaptive. Hazay et al. [HVW21a] show how to extend their techniques to the set-

ting of PCPPs, constructing ZK-PCPPs for NP with a polynomial query gap with a *non-adaptive* honest verifier.

	$\mathcal{P} \leftrightarrow \mathcal{V}$ Comm.	Access to Prover Messages	# Prover Messages	Access to Input	Soundness Guarantee	ZK Variant Hides
IP	Yes	Full	Multiple	Full	Full	Witness
PCP	No	Oracle	Single	Full	Full	Witness
IOP	Yes	Oracle	Multiple	Full	Full	Witness
PCPP	No	Oracle	Single	Oracle	Promise (input)	Witness, input (partial)
IOPP	Yes	Oracle	Multiple	Full	Promise (witness)	Witness (partial)

Table 2: Comparison Between Different Probabilistic Proof Systems.

IOPPs are IOPs of *Proximity*, considered in, e.g., [BCF⁺16], in which the verifier has full access to the input, and oracle access to the witness. Here, “ $\mathcal{P} \leftrightarrow \mathcal{V}$ Comm.” refers to whether there is direct communication between the prover and verifier (in particular, whether the verifier can send messages to the prover); “Access to Prover Messages” states whether the verifier reads prover messages in full, or has oracle access to them; similarly, “Access to Input” indicates whether the verifier reads the input in full, or only has oracle access to it; “Soundness Guarantee” refers to the type of inputs which are guaranteed to be rejected, where “Full” means that all inputs not in the languages are rejected, whereas “Promise (input)” means that only inputs that are *far* (in relative Hamming distance) from the language are guaranteed to be rejected, and “Promise (witness)” only guarantees that \mathcal{V} rejects when given oracle access to a witness which is far (in relative Hamming distance) from all valid witnesses for the input; “ZK Variant Hides” of a system \mathcal{X} states, for the ZK variant ZK- \mathcal{X} of system \mathcal{X} , which input of the prover remains hidden from the verifier, where “Witness” means verification reveals no information about the underlying NP witness, and “Witness, input (partial)” (resp. “Witness (partial)”) roughly means that a verifier making q queries to the input and proof(s) (resp. witness) learns only q physical bits of the input (resp., witness).

ZK-PCPs (and ZK-PCPPs) are motivated not only from a purely-theoretical perspective as a natural model of a proof system, but also from their usefulness for cryptographic applications. Specifically, they enable modular design of cryptographic proofs systems, by separating a “clean” information-theoretic proof system component, from the cryptographic assumptions (such as hardness assumptions or an augmented model of computation) which can then be used to transform – i.e., “compile” through a *cryptographic compiler* – the information-theoretic system into a computationally-secure system than can be implemented. Thus, one can design, analyze and optimize the information-theoretic proof system, then apply different cryptographic compilers to obtain different properties of the resultant system. This paradigm has been extremely successful, and is widely used. For example, ZK-PCPs are the underlying combinatorial building blocks in constructions of succinct zero-knowledge arguments [BCG⁺13].

ZK-PCPs and ZK-PCPPs also have more direct cryptographic applications both to 2-party and multiparty scenarios which require highly efficient verification methods on secret data. In the 2-party setting, these include constructions of constant-round (or noninteractive in the random-oracle model) black-box arguments for NP with statistical ZK [IMS12] (whereas constructions based on *non*-ZK PCPs are not black-box [Kil92, Mic94]). ZK-PCPPs can additionally be used to design 2-party and multiparty black-box commit-and-prove protocols for NP [IW14]. While these applications only require a weaker ZK guarantee (specifically, ZK against the *honest* verifier), applications in multiparty settings require full-fledged ZK against *malicious* verifiers. These include constructions of certifiable versions of verifiable secret sharing from ZK-PCPPs, as well as sublinear ZK proofs in a distributed setting from ZK-PCPs. (Various other notions of ZK proofs in a distributed setting have been considered recently.) In the latter application, the prover and

verifier are aided by multiple (potentially corrupted) servers. The motivation for this distributed setting is to minimize the round complexity, and underlying assumptions, of sublinear ZK proofs. Specifically, using PCPs with ZK against *malicious* verifiers, [IWY16] construct distributed 3-round witness-indistinguishable proofs (respectively, ZK proofs in the computational setting with a common random string) for NP, which are unconditionally secure (respectively, based on the existence of one-way functions) in which the total communication involving the verifier is sublinear in the input length. This should be contrasted with standard sublinear ZK arguments, that require at least 4 rounds of interaction, and require the existence of collision resistant hash functions [Kil92, IMS12].

2 Preliminaries

Basic notations. We denote the security parameter by κ . A function $\mu : \mathbb{N} \rightarrow \mathbb{N}$ is *negligible* if for every positive polynomial $p(\cdot)$ and all sufficiently large κ 's it holds that $\mu(\kappa) < \frac{1}{p(\kappa)}$, and $\text{negl}(\kappa)$ denotes the set of all negligible functions. We use the asymptotic notations $O(\cdot)$ and $\Omega(\cdot)$, where $\tilde{O}(n)$ and $\tilde{\Omega}(n)$ denote $n \cdot \text{poly}(\log n)$ and $n/\text{poly}(\log n)$, respectively. We use the abbreviation PPT to denote Probabilistic Polynomial-Time, and denote by $[n]$ the set of elements $\{1, \dots, n\}$. For a string s of length n , and a subset $I \subseteq [n]$, we denote by $s|_I$ the restriction of s to the coordinates in I .

We usually denote vectors using boldface letters (e.g., \mathbf{a}), or as \vec{a} . For a pair \mathbf{x}, \mathbf{y} of vectors, $\langle \mathbf{x}, \mathbf{y} \rangle$ denotes their inner product, and $\text{Ham}(\mathbf{x}, \mathbf{y})$ to denote their Hamming distance, i.e., $\text{Ham}(\mathbf{x}, \mathbf{y}) = |\{i : x_i \neq y_i\}|$. For functions f, g , we denote their composition as $(f \circ g)(x) := f(g(x))$. The composition of families F, G of functions is defined as $F \circ G := \{f \circ g : f \in F, g \in G\}$.

For a distribution \mathcal{D} , sampling according to \mathcal{D} is denote by $X \leftarrow \mathcal{D}$, or $X \in_R \mathcal{D}$. For a pair of random variables X, Y , we use $X \equiv Y$ to denote that X, Y are identically distributed. For random variables X_κ and Y_κ over a finite domain Ω , the *statistical distance* between them is defined as

$$\text{SD}(X_\kappa, Y_\kappa) = \frac{1}{2} \sum_{w \in \Omega} |\Pr[X_\kappa = w] - \Pr[Y_\kappa = w]|.$$

X_κ and Y_κ are ϵ -*statistically close* if their statistical distance is at most $\epsilon(\kappa)$. Ensembles $\{X_\kappa\}_\kappa, \{Y_\kappa\}_\kappa$ are *statistically close*, denoted $X_\kappa \approx Y_\kappa$, if there exists an $\epsilon(\kappa) = \text{negl}(\kappa)$ such that X_κ, Y_κ are $\epsilon(\kappa)$ -close for every κ . We say that $\{X_\kappa\}_\kappa, \{Y_\kappa\}_\kappa$ are *computationally indistinguishable* if they have a $\text{negl}(\kappa)$ *computational distance*, i.e., for every PPT distinguisher \mathcal{D} there exists an $\epsilon(\kappa) = \text{negl}(\kappa)$ such that for every κ :

$$|\Pr[\mathcal{D}(X_\kappa) = 1] - \Pr[\mathcal{D}(Y_\kappa) = 1]| \leq \epsilon(\kappa).$$

Languages and Relations. We will consider NP-relations $\mathcal{R} = \mathcal{R}(x, w)$, and the corresponding NP-languages $\mathcal{L} = \{x : \exists w \text{ s.t. } (x, w) \in \mathcal{R}\}$. We sometimes write $\mathcal{R}_\mathcal{L}$ to refer to the NP-relation whose corresponding language is \mathcal{L} .

Encoding Schemes and Leakage-Resilient Encoding Schemes. Our constructions will rely on encodings schemes with leakage-resilience properties. We now provide a simple definition of an encoding scheme, and refer the reader to Section 3.1 for a more detailed discussion of the notion.

Definition 3. Let $k, n \in \mathbb{N}$. An Encoding Scheme over an alphabet Σ is a pair (Enc, Dec) where Enc is a PPT algorithm, and Dec is a (deterministic) polynomial-time algorithm, that satisfy the following.

- **Syntax.** Enc on input a secret $x \in \Sigma^k$ outputs a codeword $c \in \Sigma^n$. Dec on input $c \in \Sigma^n$ outputs $x \in \Sigma^k$ or a special error symbol \perp .

- **Correctness.** *There exists a $t \geq 0$ such that the following holds for every $x \in \Sigma^k$, and every $c \in \Sigma^n$: if there exists $c_x \in \text{Supp}(\text{Enc}(x))$ such that $\text{Ham}(c, c_x) \leq t$ then $\text{Dec}(c) = x$, otherwise Dec outputs \perp .*

A note on terminology. In this section and in Section 3, we use k, n respectively to denote the input and output lengths of Enc , as is customary in the context of error-correcting codes. In Section 4, the input and output lengths are denoted by n, \hat{n} , respectively, which is sometimes used in the context of leakage resilience.

We now define two useful properties of encoding schemes. The first is that the scheme is *linear*. The second is that the scheme is *onto*, meaning any $c \in \Sigma^n$ can be interpreted as the encodings of some $x \in \Sigma^k$, in the sense that c would be decoded to x .

Definition 4. *An encoding scheme (Enc, Dec) over a field \mathbb{F} is linear if for every k : (1) k divides n ; and (2) there exists a decoding vector \mathbf{d} such that for every $x \in \mathbb{F}^k$: (1) every $\mathbf{x} \in \text{Supp}(\text{Enc}(x))$ can be partitioned into k sub-vectors $\mathbf{x} = (\mathbf{x}^1, \dots, \mathbf{x}^k)$, such that $\text{Dec}(\mathbf{x}) = (\langle \mathbf{d}, \mathbf{x}^1 \rangle, \dots, \langle \mathbf{d}, \mathbf{x}^k \rangle)$.*

(Enc, Dec) is onto, if Dec is defined (i.e., does not output \perp) for every $c \in \Sigma^n$.

Section 4 will use a *parameterized* notion of encoding schemes. In such encoding schemes, the encoding and decoding algorithms are given an additional input 1^σ , which is used as a security parameter. The encoding length can then depend also on σ (and not only on k), and we require that for every σ the resultant scheme is an encoding scheme. A parameterized encoding scheme is onto (linear, respectively) if it is onto (linear, respectively) for every σ .

We now define leakage resilience of distributions and encodings.

Definition 5 (Leakage Resilience – Distributions and Encoding Schemes). *Let \mathcal{LEAK} be a family of functions, and $\epsilon > 0$. For a finite set D , a pair of distributions X, Y over D are $(\mathcal{LEAK}, \epsilon)$ -leakage resilient if for any function $\ell \in \mathcal{LEAK}$ with domain D it holds that $\text{SD}(\ell(X), \ell(Y)) \leq \epsilon$.*

A randomized function $f : \Sigma^n \rightarrow \Sigma^m$ is $(\mathcal{LEAK}, \epsilon)$ -leakage resilient if for every $x, y \in \Sigma^n$, the distributions $f(x), f(y)$ are $(\mathcal{LEAK}, \epsilon)$ -leakage resilient.

An encoding scheme (Enc, Dec) is $(\mathcal{LEAK}, \epsilon)$ -leakage resilient if for every large enough $\sigma \in \mathbb{N}$, $\text{Enc}(\cdot, 1^\sigma)$ is $(\mathcal{LEAK}, \epsilon)$ -leakage resilient.

For the construction of Section 3 we will be particularly interested in *probing-resilient* encoding schemes, namely ones that are leakage resilient against leakage functions that probe bits of the codeword.

For a function family \mathcal{LEAK} , we sometimes use the term “leakage family \mathcal{LEAK} ”, or “leakage class \mathcal{LEAK} ”, and refer to functions in \mathcal{LEAK} as “leakage functions”.

2.1 PCPs and ZK-PCPs

The probabilistic proof system we focus on in this work is *Probabilistically Checkable Proofs (PCPs)* with zero-knowledge guarantees. We first describe the (standard, non-zero-knowledge) notion of PCPs.

At a high level, PCPs allow a randomized verifier to probabilistically verify the validity of some input statement by querying few bits of a purported proof to which it has oracle access. PCPs can be defined (and have been studied) in relation to various complexity classes (e.g., $\text{DTIME}(n)$). In this work, we focus on PCPs for NP since cryptographic applications usually require proof systems for NP (and this also simplified the presentation). In the context of complexity theory, the system consists solely of the verifier and the proof oracle, where the proof generation process is implicit. However, cryptographic applications necessitate that proof generation be efficient, given the NP

witness. Thus, we define PCPs as a system consisting of efficient prover and verifier. (This is by now standard in the literature of PCPs for cryptographic applications, e.g., [KPT97], as well as other proof systems such as interactive oracle proofs [BCS16, RRR16].)

More specifically, a PCP system for a language $\mathcal{L} \in \text{NP}$ consists of a polynomial-time prover \mathcal{P} that given $x \in \mathcal{L}$ and a corresponding witness generates a proof π for x , and a PPT verifier \mathcal{V} having direct access (“oracle access”) to individual symbols of π . \mathcal{V} will read only part of its proof string π (called oracle), where the queries to π are determined by \mathcal{V} ’s input and coin tosses. Formally,

Definition 6 (PCP). *A Probabilistically Checkable Proof (PCP) for a language $\mathcal{L} \in \text{NP}$ consists of a polynomial-time prover \mathcal{P} and a PPT verifier \mathcal{V} such that there exists a negligible function $\text{ngl} = \text{negl}(\kappa)$ for which the following holds:*

- **Syntax:** *The prover \mathcal{P} has input $1^\kappa, x, w$ (κ is the security parameter), and outputs a proof $\pi \in \{0, 1\}^*$ for x . The verifier \mathcal{V} has input $1^\kappa, x$, and oracle access to π . It makes q queries to π , and outputs either 0 or 1 (representing reject or accept, respectively). q is called the query complexity of the system, and the system is called a q -query PCP.*

- **Semantics:** *The system satisfies the following semantic properties:*

- **Completeness:** *For every $(x, w) \in \mathcal{R}_{\mathcal{L}}$, and every proof $\pi \in \mathcal{P}(1^\kappa, x, w)$,*

$$\Pr[\mathcal{V}^\pi(1^\kappa, x) = 1] = 1$$

where the probability is over the randomness of \mathcal{V} .

- **Soundness:** *For every $x \notin \mathcal{L}$ and every oracle π^* ,*

$$\Pr[\mathcal{V}^{\pi^*}(1^\kappa, x) = 1] \leq \text{negl}(\kappa)$$

where the probability is over the coin tosses of the verifier. $\text{negl}(\kappa)$ is called the soundness error of the system.

Zero-Knowledge PCPs (ZK-PCPs). Intuitively, ZK-PCPs are PCPs in which the witness remains entirely hidden throughout the verification procedure, in the sense that even a *malicious* verifier which deviates from the specified verification procedure learns nothing but the validity of the claim. Achieving this property requires some modifications to the standard notion of a PCP. First, the prover must now be *probabilistic*. Indeed, completeness of the PCP requires (intuitively, at least) that the proof encode the witness. That is, certain proof bits will carry *some* information about the witness, and by querying them a malicious verifier \mathcal{V}^* may learn information about the witness. Allowing randomized proof generation is akin to having a probabilistic prover in zero-knowledge *interactive* proofs (whereas in standard interactive proofs the prover may be deterministic).

Second, while we do not impose any restrictions on the query *pattern* (i.e., the strategy) of a malicious verifier \mathcal{V}^* , we impose a bound on its query *complexity*, namely the number of queries it makes to the proof. To see why imposing *some* restriction is needed, recall that cryptographic applications necessitate an efficient honest prover, meaning the proof will have polynomial length. Therefore, a polynomial-time verifier \mathcal{V}^* that is unrestricted in its access to the proof could potentially read the entire proof and thus necessarily learn information about the witness (since, as noted above, the proof carries information about the witness). There are several possible methods of restricting \mathcal{V}^* to reading only *part* of the proof (e.g., requiring its runtime to be smaller than the proof length), where the one used in the literature is to restrict its *query complexity* to some a-priori

bound q^* which is known prior to proof generation (in particular, the proof length may depend on q^*). One advantage of restricting the query complexity is that it allows us to obtain *information theoretic* ZK, namely that ZK holds against malicious verifiers with *unbounded computational power*, so long as the verifier makes at most q^* queries to the proof. It is also more inline with the main efficiency measures of standard PCPs, which focus on the query (and randomness) complexity.

To define ZK-PCPs, we first formalize the restriction on the query complexity of the verifier.

Definition 7 (Query-bounded verifier). *We say that a (possibly malicious) verifier \mathcal{V}^* with oracle access to a proof π is q^* -query-bounded if it makes at most q^* queries to π .*

As noted above, we will allow a malicious verifier to be computationally unbounded. Moreover, we will allow its query bound q^* to be much larger than that of the honest verifier. Ideally, the honest verifier will only make $\text{polylog}(q^*)$ queries,⁹ and the proof will have length $\text{poly}(q^*)$. Another aspect in which we allow a malicious verifier to be more powerful than the honest verifier is *adaptivity*: whereas we would like the honest verifier to be non-adaptive – namely, make a single round of queries to the proof (as in standard PCPs), we allow a malicious verifier to be adaptive – i.e., make several rounds of queries to the proof. This is formalized in the following definition.

Definition 8 (Adaptive and-non adaptive verifiers). *We say that a (possibly malicious) verifier \mathcal{V}^* is non adaptive if its queries are determined solely by its input x and randomness (in particular, \mathcal{V}^* can make all its queries to the proof in a single round). Otherwise, we say that \mathcal{V}^* is adaptive (in particular, the queries of an adaptive verifier may depend on the answers to previous queries).*

We are now ready to define ZK-PCPs. Similar to zero-knowledge *interactive* proofs (and, more generally, cryptographic protocols), we formalize zero knowledge by requiring the existence of an efficient *simulator* algorithm that, given the input $1^\kappa, x$, can simulate the *view* of the verifier, namely the verifier’s input, random coins, and the oracle answers (the queries of the verifier can be computed from these values). Intuitively, this guarantees zero knowledge because the view of the verifier captures the entire “knowledge” it obtained through verification, and in particular by querying the proof oracle. Since the view can be simulated from the input alone, without any access to the NP witness, this implies that the view contains no information about the witness. We now formalize this intuition.

Notation 9. *For a PCP system $(\mathcal{P}, \mathcal{V})$ and a (possibly malicious) verifier \mathcal{V}^* , we use $\mathbf{View}_{\mathcal{V}^*, \mathcal{P}}(\kappa, x, w)$ to denote the view of \mathcal{V}^* when it has input $1^\kappa, x$ and oracle access to a proof that was randomly generated by \mathcal{P} on input $(1^\kappa, x, w)$.*

Definition 10 (ZK-PCP). *We say that a PCP system $(\mathcal{P}, \mathcal{V})$ for \mathcal{L} is a (q^*, ε) -Zero-Knowledge PCP (ZK-PCP) if for every (possibly malicious and adaptive) q^* -query-bounded verifier \mathcal{V}^* there exists a PPT simulator Sim , such that for every $(x, w) \in \mathcal{R}$, $\text{Sim}(1^\kappa, x)$ is distributed ε -statistically close to $\mathbf{View}_{\mathcal{V}^*, \mathcal{P}}(x, w)$.*

One prevalent approach for designing simulators (which we will use) is to have the simulator emulate the verifier \mathcal{V}^* , simulating the answers to \mathcal{V}^* ’s oracle queries.

2.1.1 Restrictions, Extensions and Generalizations

Definition 10 can be restricted, or alternatively, generalized, in several ways, as we now discuss.

⁹Traditional PCPs can be verified with a constant number of queries, achieving a constant soundness error. However, since cryptographic applications usually necessitate a *negligible* soundness error, the query complexity of the honest verifier is necessarily polylogarithmic in the security parameter κ , or in q^* .

A taxonomy based on ZK quality. While we define a statistical notion of ZK as the default for ZK-PCPs, one can also consider stronger or weaker forms of ZK. Specifically, we can require a stronger *perfect* ZK guarantee in which the simulated view is distributed identically to the real view of the verifier.

Notation 11. We say that a ZK-PCP has q^* -ZK if it has perfect ZK against q^* -query bounded verifiers.

Alternatively, we can settle for a weaker *computational* ZK guarantee which only holds against *computationally-bounded* verifiers \mathcal{V}^* , where the simulated view is computationally indistinguishable from the real view. One particular relaxation of ZK which we consider in this work is *Witness-Indistinguishability* which does not require that the ZK simulator be efficient. This can be obtained by removing the requirement that Sim be PPT in Definition 10, but the following alternative formulation would be more useful.

Definition 12 (WI-PCP). We say that a PCP system $(\mathcal{P}, \mathcal{V})$ for \mathcal{L} is a (q^*, ε) -Witness-Indistinguishable PCP (WI-PCP) if for every (possibly malicious and adaptive) q^* -query-bounded verifier \mathcal{V}^* , for any $x \in \mathcal{L}$, and any pair of corresponding witnesses w_1, w_2 such that $(x, w_1), (x, w_2) \in \mathcal{R}_{\mathcal{L}}$:

$$\text{SD}(\mathbf{View}_{\mathcal{V}^*, \mathcal{P}}(x, w_1), \mathbf{View}_{\mathcal{V}^*, \mathcal{P}}(x, w_2)) \leq \varepsilon.$$

There are also various flavors of ZK based on different qualities of the ZK simulator. These include, for example, whether the simulator is straight-line – i.e., it emulates the verifier without having to rewind it, and whether the simulator interacts with the verifier as a black-box. The ZK-PCP constructions described in this work have straight-line, black-box simulators.

Honest-Verifier ZK. Another natural restriction of the definition is by considering zero knowledge only against the *honest verifier* \mathcal{V} . Such systems are called *Honest-Verifier ZK (HVZK)*. As we explain below, it is fairly simple to obtain an HVZK PCP system from a standard PCP system (e.g., such a system was presented already in the paper of Kilian, Petrank and Tardos on ZK-PCPs [KPT97], and a weaker construction with a large soundness error was presented in [DFK⁺92]). However, such a restriction is generally too weak to be used in cryptographic applications, for example, to prove honest behaviour in cryptographic protocols, in which case the verifying party might be maliciously corrupted. To see why settling for HVZK simplifies the problem considerably, notice that HVZK is preserved under standard soundness amplification techniques. More specifically, assume we have an HVZK PCP system with a large soundness error (obtaining such systems is relatively easy given known techniques such as “MPC-in-the-head” [IKOS07]). Soundness can then be amplified by having the prover generate many fresh, independent copies of the proof, and having the honest verifier repeat the verification procedure several times, each time using a fresh proof copy. HVZK of this verification procedure easily reduces to the HVZK of the original system. However, this amplification does not preserve *full-fledged* ZK (i.e., against malicious verifiers) *even if* the original system is ZK against malicious verifiers. Indeed, the reason is that the query complexity – even of the honest verifier – increases through this transformation, and in particular would exceed the ZK query bound of the original system. Thus, a malicious verifier that “concentrates” all its queries to a *single* proof copy might be able to violate the ZK of the underlying system.

Verifier Adaptivity. In *non-ZK* PCP constructions, the honest verifier is non-adaptive. In contrast, the classic ZK-PCP of [KPT97], and all consequent ZK-PCPs – *except the ones based on leakage resilience* [IWY16, HVW21a] – have an *adaptive* honest verifier. As discussed in Section 1, this is because, very roughly, they enhance an HVZK PCP to have full-fledged ZK, by modifying the proof

such that even honest verification requires multiple rounds of queries to the proof. Moreover, this is inherent to their technique of using “locking schemes” [KPT97]. Having a non-adaptive honest verifier is a major advantage of leakage resilience based ZK-PCPs (see Table 1 for a comparison of existing ZK-PCP systems), since having non-adaptive *honest* verification is a desirable feature of the system. Indeed, an adaptive honest verifier translates into multiple interaction rounds in cryptographic applications of ZK-PCPs.

Notation 13. We say that a ZK-PCP system $(\mathcal{P}, \mathcal{V})$ is a non-adaptive ZK-PCP if the honest verifier \mathcal{V} is non-adaptive.

We note that an orthogonal measure of adaptivity is whether a *malicious* verifier is restricted to being non-adaptive. Unlike having a non-adaptive honest verifier, guaranteeing ZK only against non-adaptive malicious verifiers is an *undesirable restriction* of the system, since it means there are no guarantees against adaptive verifiers.

Universal vs. Non-Universal Simulation. Definition 10 requires ZK to hold with a non-universal simulator, requiring, for every malicious verifier \mathcal{V}^* , the existence of a simulator $\text{Sim}_{\mathcal{V}^*}$. A stronger possible definition would require the existence of a *universal* simulator Sim that can simulate the view of *any* query-bounded verifier \mathcal{V}^* . We note that all the ZK-PCPs described in this work (in fact, to the best of our knowledge, all existing ZK-PCP constructions) satisfy this stronger definition.

The Alphabet. Similar to PCPs, ZK-PCPs are defined as bit-strings. One could also consider a relaxed notion in which the proof is over some larger alphabet Σ (and this indeed has been done in the PCP context). We note, however, that while for standard PCPs the choice of alphabet affects only the parameters of the scheme (but not its semantic properties), this is not the case for *zero-knowledge* PCPs. Indeed, any PCP over Σ can be transformed to a PCP over $\{0, 1\}$ by replacing each symbol with a bit-string representation of it, without violating completeness or soundness.¹⁰ However, as we explain in Section 3, doing so for a ZK-PCP does not preserve ZK against *malicious* verifiers. ZK-PCPs over a large alphabet Σ can still be useful as a building block for obtaining ZK-PCPs (with proofs over $\{0, 1\}$), see Section 3.

PCPs of Proximity. A useful generalization of PCPs are PCPs of *Proximity* (PCPPs), that allow verification of an input claim while reading only a small portion of it. This is formalized by giving the verifier oracle access to the input, similar to how it accesses the proof. Of course, in this case the verifier cannot be expected to distinguish a true claim from a claim that is false, but very close to being true (e.g., a 3-CNF for which there exists an assignment that satisfies all but a tiny fraction of the clauses). Instead, soundness is defined similarly to correctness of promise problems: any input which is sufficiently far from the corresponding NP language will be rejected with high probability. PCPPs are an important building block in PCP constructions, and a useful notion in its own right. There are PCPP constructions matching the properties of the best-known standard PCPs [BS08, Mie09].

Zero-Knowledge PCPPs (ZK-PCPPs) [IW14] have a stronger ZK guarantee than ZK-PCPs: while ZK-PCPs guarantee that the *witness* remains entirely hidden through verification, ZK-PCPPs additionally guarantee that the *input* itself remains mostly hidden, in the sense that the verifier

¹⁰More elaborate alphabet reduction techniques have been used also in the context of traditional PCP, but with the goal of improving the system’s parameters, e.g., in [Din06].

(even a malicious one) learns only few *physical* input bits. This is formalized using the simulation paradigm as in ZK-PCPs, where instead of giving the entire input to the simulator, it has oracle access to it, and is restricted to making q^* queries (where q^* is the query complexity of the verifier).

Certain techniques for constructing ZK-PCPs extend also to PCPPs, while others do not (or, at least, it is not clear how to extend them). In particular, the original ZK-PCP construction of [KPT97] can be extended to also apply to ZK-PCPPs [IW14], and the ZK-PCPs based on leakage-resilient encodings described in Section 3 also applies to PCPPs (see [HVW21a] for a full description of the construction). On the other hand, the construction of ZK-PCPs from LR circuits (Section 4) does not seem to easily extend to the PCPP realm.

3 ZK-PCPs from LR Encodings: The Construction of [HVW21a]

The main result of this section is a construction of ZK-PCPs for NP with a non-adaptive honest verifier and a polynomial query gap (between the query complexity of the honest and malicious verifiers) due to [HVW21a]:

Theorem 14 (ZK-PCPs for NP, Formal statement of Informal Theorem 1). *There exists a constant $\epsilon \in (0, 1)$ such that for any ZK parameter $q^* \in \mathbb{N}$ there exists a non-adaptive $(q^*)^\epsilon$ -query $\Omega(q^*)$ -ZK-PCP for NP.*

We describe a simplified version of the construction of [HVW21a] which nonetheless suffices for designing ZK-PCPs. The interested reader is referred to [HVW21a] for a description of the more general paradigm which employs an equivocal notion of secret sharing instead of the weaker leakage-resilient encodings used here.

The construction employs Leakage-Resilient (LR) Encodings. The starting point is a ZK-PCP variant $(\mathcal{P}', \mathcal{V}')$ over a *large alphabet* Σ , namely where the proof π' is over Σ . To obtain a standard ZK-PCP $(\mathcal{P}, \mathcal{V})$ – i.e., one in which the proof π is over *bits* – we need an *alphabet reduction*. That is, we are looking for a transformation that replaces each symbol $\pi'_i \in \Sigma$ with a bit-string “segment” $\text{segm}_i \in \{0, 1\}^*$. Then, given a proof $\pi' = (\pi'_1, \dots, \pi'_N)$, the resultant proof would be $\pi = (\text{segm}_1, \dots, \text{segm}_N)$.

As mentioned in Section 2.1.1, the naive alphabet reduction which replaces each symbol of Σ with a bit string representing it does not preserve ZK. Indeed, this alphabet reduction necessarily increases the query complexity of the honest verifier \mathcal{V} , who will need to query $q \cdot |\text{segm}|$ proof bits (where q is the query complexity of \mathcal{V}'). Thus, a malicious verifier – whose query complexity is at least as that of the honest verifier – with oracle access to the resultant proof π may query subsets of bits in *many* segments segm_i , effectively learning partial information about many proof symbols (in particular, more than the ZK guarantee of π'), and violating ZK. Therefore, we need an alphabet reduction which *preserves ZK*.

Viewed through the leakage-resilience lens, the information which a malicious verifier obtains on a symbol $\pi'_i \in \Sigma$ by querying bits of segm_i constitutes *probing leakage* on segm_i , and consequently also on π'_i . Thus, intuitively, ZK can be guaranteed by protecting the segments segm_i from probing leakage. This gives a general blueprint for a ZK alphabet reduction: replace each symbol σ of Σ with its binary representation s_σ , then encode s_σ using a probing-resilient encoding. While this roughly describes the alphabet reduction of [HVW21a], there are a few subtleties, as we now describe.

Simulation Strategy for Malicious Verifiers. A probing-resilient encoding can only protect against probing of a sufficiently small subset of bits of the encoding, namely against probing of some a-priori fixed fraction τ of bits. (Indeed, since the message can be decoded from the encoding, an adversary that probes the entire encoding necessarily learns the underlying message.) However, a malicious verifier \mathcal{V}^* may query an entire segment segm_i ¹¹ – or, more generally, more than a τ -fraction of it – in which case the probing-resilience of the encoding cannot be used. We solve this issue in the simulation by dividing the segments segm_i into two types: “heavy” and “light” segments. Intuitively, heavy segments are ones from which \mathcal{V}^* queried many bits, in particular, more than a τ -fraction. Light segments are segments that are not heavy. We use the probing-resilience of the underlying encoding to claim that \mathcal{V}^* learns no information about the symbols encoded in the light segments, and use the ZK guarantee of the underlying ZK-PCP system $(\mathcal{P}', \mathcal{V}')$ to simulate the symbols encoded in heavy segments. This gives us a simulation strategy for $(\mathcal{P}, \mathcal{V})$: simulate heavy symbols using the simulator of the underlying system $(\mathcal{P}', \mathcal{V}')$, and simulate light symbols using random and independent encoding of an arbitrary value (e.g., the all-zeros string).

Simulating Partially-Leaked Symbols. The simulation strategy defined in the previous paragraph necessitates that the simulator Sim knows in advance which segments are heavy and which are light, since this determines how to generate the answer to a query. Whether or not a segment is heavy is a function of *the entire query pattern* of \mathcal{V}^* . Thus, this proof strategy only works against *non-adaptive* malicious verifiers, namely ones which make a single round of queries to the proof. (Indeed, in this case Sim learns all of \mathcal{V}^* 's queries before it needs to simulate the oracle answers.) ZK against *adaptive* malicious verifiers, namely ones which make several rounds of queries to the proof, where each query may depend on the answers to the previous queries, requires a somewhat different simulation strategy, and a stronger LR guarantee from the probing-resilient encoding, as we now explain.

The high-level idea is as follows. At the onset of the simulation, Sim treats all segments as light, answering queries using random and independent encodings of $\vec{0}$. At certain points in the simulation, a certain segment i may become heavy – namely, the number of queries \mathcal{V}^* made to it exceeds the probing threshold τ . At this point, Sim uses the simulator Sim' of the underlying ZK-PCP system $(\mathcal{P}', \mathcal{V}')$ to simulate the symbol π'_i . To continue with the simulation, Sim must now generate an encoding of π'_i which is (1) consistent with the bits already probed from the i 'th segment; and (2) is distributed as a random encoding of π'_i subject to (1). For this, we need the underlying probing-resilient encoding to be *equivocal* – allowing one to efficiently sample from this distribution. Such encodings are called *Reconstructable Probabilistic Encodings (RPEs)* [CDMW08, CDMW18].

Putting It Together. We are now ready to describe the full alphabet reduction (see Section 3.2, and Figure 1 in particular) that transforms a ZK-PCP variant $(\mathcal{P}', \mathcal{V}')$ over alphabet Σ into a ZK-PCP $(\mathcal{P}, \mathcal{V})$ (over bits). The reduction employs an RPE (see Section 3.1). For proof generation, the prover \mathcal{P} first runs \mathcal{P}' to generate a proof $\pi' = (\pi'_1, \dots, \pi'_N)$ over Σ . Then, it replaces each proof symbol π'_i with its binary representation s_i , and uses the RPE to encode s_i into a segment segm_i . \mathcal{P} outputs the proof $\pi = (\text{segm}_1, \dots, \text{segm}_N)$. The verifier \mathcal{V} , given oracle access to π , verifies the proof by emulating \mathcal{V}' . Whenever \mathcal{V}' queries a symbol π'_q of π' , \mathcal{V} queries segm_q from π , RPE-decodes it to obtain the binary representation s_q of a symbol $\sigma_q \in \Sigma$, and provides σ_q to

¹¹This is because the query bound q^* imposed on the malicious verifier is expected to be much larger than the length of the encoding. Indeed, q^* should be at least as large as the query complexity of the *honest* verifier, which would need to read at least a few symbols of the original PCP, i.e., a few *full* encodings of symbols of the original PCP.

\mathcal{V}' as the answer of the oracle. When the emulation ends, \mathcal{V} outputs whatever \mathcal{V}' outputs. In the following sections, we describe the RPE building block (Section 3.1), and analyze the resultant ZK-PCP scheme (Section 3.2).

3.1 Main Building Block: Reconstructable Probabilistic Encodings (RPEs)

The main building block of the ZK-preserving alphabet reduction is an encoding scheme with equivocation properties called *Reconstructable Probabilistic Encoding (RPE)* [CDMW08, BDKM16, BDG⁺18, CDMW18]. In this section, we formally define these objects.

Codes, or encoding schemes, are extensively used in computer science, the most notable example being Error-Correcting Codes (ECCs), which are used to guarantee that the data can still be recovered even if faults occur (i.e., some of the symbols of the data are erased or corrupted). A code consists of an encoding procedure Enc which maps a message to a codeword, and a decoding procedure Dec which decodes the message from a (possibly corrupted) codeword. In the context of error-correction, Enc , Dec are usually deterministic. Most ECCs are linear codes, where the code is defined by a *generator matrix*, and encoding simply multiplies the generator matrix with the message. The main parameters of interest for such codes are: (1) the *rate* of the code – the ratio between the length of the encoding (also known as a *codeword*) and the length of the original message; (2) its *distance* – namely the minimal distance between a pair of codewords¹²; and (3) the alphabet size (where most ECCs are binary). There are numerous extensions and generalizations of ECCs that guarantee additional properties beyond error correction.

We will be interested in a generalized notion of an ECC which also guarantees *probing-resilience* in the sense that few codeword symbols reveal no information (in an information-theoretic sense) on the encoded message. Of course, such a guarantee cannot be satisfied if encoding is deterministic. Probing-resilient encodings therefore allow for a *randomized* encoding procedure, where each message has a subset of codewords to which it can be mapped, and encoding chooses one of them at random. It is fairly simple to obtain such a leakage-resilient encoding from a linear code, as long as its generator matrix has a “good” structure. (See, e.g., [ISVW13] for a description of the needed properties and how encoding works.)

As explained above, probing-resilience alone is insufficient to guarantee ZK against *adaptive* malicious verifiers. Instead, we rely on a stronger *equivocation* property which guarantees that as long as the probing threshold τ had not been violated, the probed bits can be efficiently “explained” as the bits in an encoding of *any* arbitrary message msg . Intuitively, an RPE is an encoding scheme (see defined in Section 2) which is probing-resilient, and is additionally associated with a *resampling/reconstruction* algorithm Rec that can “explain” the probed bits. Formally:

Definition 15 (Reconstructable Probabilistic Encoding (RPE)). *Let $k, n, \ell \in \mathbb{N}$. A (k, n, ℓ) -Reconstructable Probabilistic Encoding (RPE) is a triple $(\text{Enc}, \text{Dec}, \text{Rec})$ where Enc , Rec are PPT algorithms, and Dec is a (deterministic) polynomial-time algorithm, that satisfy the following.*

- **Syntax.** Enc on input a secret $x \in \{0, 1\}^k$ outputs a codeword $c \in \{0, 1\}^n$. Dec on input $c \in \{0, 1\}^n$ outputs $x \in \{0, 1\}^k$ or a special error symbol \perp . Rec on input a secret x , a set $\mathcal{I} \subset [n]$ of size $|\mathcal{I}| \leq \ell$, and ℓ bits $(c_i)_{i \in \mathcal{I}}$, outputs $c' \in \{0, 1\}^n$.
- **Correctness.** There exists a $t \geq 0$ such that the following holds for every $x \in \{0, 1\}^k$, and every $c \in \{0, 1\}^n$: if there exists $c_x \in \text{Supp}(\text{Enc}(x))$ such that $\text{Ham}(c, c_x) \leq t$ then $\text{Dec}(c) = x$, otherwise Dec outputs \perp .

¹²The distance between a pair of codewords is the number of coordinates in which they differ.

Construction 16 (Alphabet reduction for ZK-PCPs). *Let κ be a security parameter.*

Building blocks:

- A PCP system $(\mathcal{P}', \mathcal{V}')$ over alphabet Σ of size $|\Sigma| = 2^k$.
- An RPE $(\text{Enc}, \text{Dec}, \text{Rec})$ for secrets in $\{0, 1\}^k$.

Prover algorithm. \mathcal{P} has input $1^\kappa, x, w$. It runs \mathcal{P}' with input $1^\kappa, x, w$ to obtain a proof $\pi' = (\pi'_1, \dots, \pi'_N)$ over Σ . For every proof symbol π'_i , let s_i denote its bit-string representation, then \mathcal{P} encodes $\text{segm}_i \leftarrow \text{Enc}(s_i)$. Finally, \mathcal{P} outputs the proof $\pi = (\text{segm}_1, \dots, \text{segm}_N)$.

Verifier algorithm. \mathcal{V} is given input $1^\kappa, x$ and oracle access to π . It runs \mathcal{V}' with input $1^\kappa, x$, and emulates the oracle π' for \mathcal{V}' as follows. Whenever \mathcal{V}' reads a symbol σ from π' , \mathcal{V} reads the entire encoding of σ from π . Then, it uses Dec to recover the symbol σ , which it gives to \mathcal{V}' as the answer of the oracle. When the emulation terminates, \mathcal{V} outputs whatever \mathcal{V}' outputs.

Figure 1: Alphabet Reduction For ZK-PCPs [HVW21a]

- **ℓ -Secrecy (of partial views).** For every pair of secrets x, x' , and any subset $\mathcal{I} \subseteq [n]$ such that $|\mathcal{I}| \leq \ell$, $\text{Enc}(x)|_{\mathcal{I}} \equiv \text{Enc}(x')|_{\mathcal{I}}$.
- **ℓ -Reconstruction (from partial views).** For any secret x , any subset $\mathcal{I} \subseteq [n]$ of size $|\mathcal{I}| \leq \ell$, and any set $(c'_i)_{i \in \mathcal{I}}$ of bits, $\text{Rec}(x, \mathcal{I}, (c'_i)_{i \in \mathcal{I}})$ is distributed identically to an encoding $c \in \text{Supp}(\text{Enc}(x))$ that is random subject to being consistent with $(c'_i)_{i \in \mathcal{I}}$.

A few remarks are in order. First, our constructions can make due with a relaxed RPE notion in which the secrecy and reconstruction properties hold statistically with statistical distance ϵ . (In this case, the resultant ZK-PCP will have statistical ZK with a statistical error of roughly $N \cdot \epsilon$, where N denotes the proof length, see [HVW21a] for details.) Second, while we define RPEs for a single message length k , the notion naturally generalizes to *families* of codes such that for every $k \in \mathbb{N}$ there exists a code whose codewords have length $n = n(k)$; and there exists a uniform algorithm that given 1^k as input, generates the encoding, decoding and resampling procedures for message length k . We will only consider (efficiently encodable and decodable) *families* of codes in this work. For simplicity and clarity of the definitions, we do not explicitly refer to a family of codes, but k should be understood as a general input length parameter. (This is standard in the literature.) Finally, we note that non-explicit constructions of RPEs follow easily from the existence of linear error-correcting codes with sufficiently “good” parameters, which are satisfied by random linear codes (see, e.g., [CCG⁺07]). Indeed, such codes possess the secrecy property of RPEs, which guarantees that for every subset of ℓ codeword symbols, the resultant system of linear equations has a solution for *any* possible secret. This gives an efficient reconstructor Rec . We note that the final ZK-PCP construction will use an *explicit* RPE construction due to [GM10, BDG⁺18].

3.2 The ZK-PCP Construction

We now describe the alphabet reduction for ZK-PCPs, which uses RPEs to transform a ZK-PCP variant over a large alphabet to a ZK-PCP (over bits).

Construction 16 transforms a PCP system over a large alphabet into a PCP (over bits). The following theorem of [HVW21b, Theorem 9] states that if the underlying PCP system is ZK, and the RPE is secure, then the resultant scheme is also ZK.

Theorem 17 (ZK-PCPs from LR encodings [HVW21a]). *Assume Construction 16 is instantiated with:*

- A (q^*, ϵ) -ZK-PCP $(\mathcal{P}', \mathcal{V}')$ over alphabet Σ for a language \mathcal{L} .
- A (k, n, ℓ) -RPE $(\text{Enc}, \text{Dec}, \text{Rec})$.

Then Construction 16 is a $((q^* + 1) \cdot (\ell + 1) - 1, \epsilon)$ -ZK-PCP for \mathcal{L} .

Moreover, the transformation preserves the soundness and completeness of $(\mathcal{P}', \mathcal{V}')$. Furthermore, if $(\mathcal{P}', \mathcal{V}')$ has proofs of length N that can be verified non-adaptively with q' queries, then Construction 16 has proofs of length $N \cdot n$ that can be verified non-adaptively with $q = q' \cdot n$ queries.

Proof. **Completeness** follows directly from a combination of the completeness of $(\mathcal{P}', \mathcal{V}')$ and the correctness of the RPE (which guarantees that \mathcal{V} perfectly emulates the proof oracle for \mathcal{V}'). The claim regarding q follows directly from the construction.

Soundness. Let $x^* \notin \mathcal{L}$, and let π^* be a purported proof oracle for \mathcal{V} . We show that \mathcal{V} rejects x^* with the same probability as \mathcal{V}' . We partition π^* into N length- n segments $\pi^{*,1} \dots \pi^{*,N}$, where the i 'th segment $\pi^{*,i}$ contains the bits in locations $(i-1)n + 1, \dots, i \cdot n$. (Notice that the i 'th segment in an *honestly-generated proof* would contain the RPE-encoding of the i 'th symbol in a proof over Σ .) Then the correctness of the RPE implies that for every $1 \leq i \leq N$ there exists a $\sigma_i \in \Sigma$ such that $\text{Dec}(\pi^{*,i}) = \sigma_i$ (indeed, if $\text{Dec}(\pi^{*,i}) = \perp$ then we set σ_i to some arbitrary symbol in Σ). Let $\pi^{*'} = (\sigma_1, \dots, \sigma_N) \in \Sigma^N$, and notice that when \mathcal{V} has oracle access to π^* , it emulates \mathcal{V}' with oracle access to $\pi^{*'}$. The soundness of $(\mathcal{P}', \mathcal{V}')$ therefore guarantees that \mathcal{V}' (and consequently also \mathcal{V}) accepts with probability ϵ .

Zero Knowledge. Let $q^{**} = (q^* + 1) \cdot (\ell + 1) - 1$, and let \mathcal{V}^* be a (possibly malicious and adaptive) q^{**} -query bounded verifier. For simplicity of the description, we assume \mathcal{V}^* makes its queries one at a time, and never repeats queries (this is without loss of generality). We describe a simulator Sim for \mathcal{V}^* , which relies on the simulator Sim' of the underlying ZK-PCP system $(\mathcal{P}', \mathcal{V}')$, and emulates a proof oracle π^* for \mathcal{V}^* as follows:

1. $\pi^* \in \{0, 1\}^{N \cdot n}$ is the concatenation of N segments $\text{segm}_1, \dots, \text{segm}_N \in \{0, 1\}^n$, which Sim initializes as random and independent RPE-encoding of 0^k , by computing $\text{segm}_i \leftarrow \text{Enc}(0^k)$.
 Sim additionally maintains N counters $\text{Cnt}_1, \dots, \text{Cnt}_N$, initialize to 0, and N sets $\mathcal{I}_1, \dots, \mathcal{I}_N$, initialized to \emptyset .
2. Sim answers each oracle query Q of \mathcal{V}^* as follows. Assume that Q is a query to the j 'th bit of the i 'th segment (meaning Q queries the $(i-1)n + j$ 'th bit of the proof).
 - (a) If $\text{Cnt}_i < \ell$, or $\text{Cnt}_i \geq \ell + 1$, then Sim answers with the j 'th bit of segm_i , increases Cnt_i by 1, and adds j to \mathcal{I}_i .
 - (b) Otherwise, $\text{Cnt}_i = \ell$, meaning \mathcal{V}^* has already queried ℓ bits from the encoding in the i 'th segment. In this case, Sim uses Sim' to simulate the i 'th symbol σ_i of a proof of the underlying ZK-PCP system. Then, Sim resamples an encoding of σ_i by computing $\text{segm}'_i \leftarrow \text{Rec}(\sigma_i, \mathcal{I}_i, (\text{segm}_i)_{|\mathcal{I}_i})$ (this resamples a fresh encoding of σ_i consistently with the oracle answers already simulated), sets $\text{segm}_i := \text{segm}'_i$, and provides the j 'th bit of segm_i as the answer of the oracle. Finally, it increases Cnt_i by 1.

We now prove that the simulated and real views of \mathcal{V}^* are ϵ -statistically close, using a hybrid argument.

\mathcal{H}_0 : This is the view of \mathcal{V}^* in the simulation described above.

\mathcal{H}_1^0 : \mathcal{H}_1^0 is obtained from \mathcal{H}_0 by replacing the simulated answers of Sim' with the actual proof symbols of a proof π' honestly generated by \mathcal{P}' .

Then \mathcal{H}_0 and \mathcal{H}_1^0 are ϵ -statistically close by the (q^*, ϵ) -ZK of $(\mathcal{P}', \mathcal{V}')$.

Indeed, since Sim' is used to simulate the i' th proof symbol only when $\text{Cnt}_i = \ell$, i.e., only on the $\ell + 1$ query to the i' th segment, and since $q^{**} = (q^* + 1)(\ell + 1) - 1$, Sim' is only used to simulate at most q^* symbols, and so the (adaptive) ZK of $(\mathcal{P}, \mathcal{V})$ implies that the simulated answers are ϵ -statistically close to the corresponding symbols in a real proof π' .

$\mathcal{H}_1^i, 1 \leq i \leq N$: \mathcal{H}_1^i is obtained from \mathcal{H}_1^{i-1} by replacing the simulated answers of Sim with the actual bits in the i' th section of the proof π . (In particular, these are bits in random RPE-encodings of π'_1, \dots, π'_N .)

Then \mathcal{H}_1^{i-1} and \mathcal{H}_1^i are identically distributed by the ℓ -secrecy or ℓ -reconstruction of the RPE.

To see why this holds, we consider two cases depending on whether or not \mathcal{V}^* made more than ℓ queries to the i' th segment. If \mathcal{V}^* made at most ℓ queries to the i' th segment, then all queries in \mathcal{H}_1^{i-1} were answered according to an RPE-encoding of 0^k , whereas all queries in \mathcal{H}_1^i were answered according to an RPE-encoding of (the binary representation of) π'_i . The bits queried in the i' th segment (and consequently, also the entire hybrids) are therefore identically distributed by the ℓ -secrecy of the RPE.

If, on the other hand, \mathcal{V}^* made *more than* ℓ queries to the i' th segment, then the first ℓ queries in \mathcal{H}_1^{i-1} were answered according to an RPE-encoding of 0^k , and the remaining queries were answered using a resampled encoding of (the binary representation of) π'_i , resampled consistently with the answers to the first ℓ queries; whereas in \mathcal{H}_1^i all queries to the i' th segment were answered according to a random encoding of (the binary representation of) π'_i . In this case, the distributions are identically distributed by the ℓ -reconstruction of the RPE.

We conclude the proof by noting that \mathcal{H}_1^N is distributed identically to the real view of \mathcal{V}^* . ■

Remark 18 (ZK-PCPs from weaker primitives). *We note that if one only requires that ZK hold against malicious non-adaptive verifiers, then it suffices for the underlying ZK-PCP system over Σ to have ZK against non-adaptive verifiers. Additionally, the reduction can be instantiated with an RPE in which secrecy and reconstruction hold statistically with some error ϵ' , in which case the overall simulation error will be $\epsilon + \epsilon' (N - q^*)$. We refer the interested reader to [HVW21b] for a proof of these claims.*

3.3 ZK-PCPs with Square-Root Gap

In this section we describe the $\sqrt{q^*}$ -query q^* -ZK-PCPs of [HVW21a] (i.e., the system has q^* -ZK with an honest verifier that makes $\sqrt{q^*}$ queries), which are obtained by appropriately instantiating the building blocks of Construction 16. These are the only known PCPs to date that have full-fledged ZK with a non-adaptive honest verifier. We first describe how we instantiate the building blocks.

The Building Blocks. Hazay et al. [HVW21a] instantiate Construction 16 with a ZK-PCP of [IKOS07] (using an improved soundness analysis given in [HVW21a]), and an RPE based on linear codes.

More specifically, the ZK-PCP over Σ is obtained using the “MPC-in-the-head” technique. It has perfect ZK against malicious verifiers querying at most a constant fraction of proof symbols (for an a-priori bounded constant), where the honest verifier obtains a negligible soundness error by non-adaptively querying only a square-root of the proof symbols. We note that the original

soundness analysis of [IKOS07] required the honest verifier to make as many queries as a malicious verifier, but this analysis was recently improved by [HVW21a].

Theorem 19 (Non-adaptive ZK-PCPs over large alphabets with \sqrt{Q} -gap, implicit in [IKOS07]). *For any $\mathcal{L} \in \text{NP}$, any $Q \geq 3$, and any input length n , there exists an alphabet Σ of size $|\Sigma| = 2^{\text{poly}(n, \log Q)}$ for which there exists a ZK-PCP for \mathcal{L} over Σ , with $\text{negl}(Q)$ soundness error with a non-adaptive honest verifier that makes $\log Q \cdot \sqrt{Q}$ queries, proofs of length Q , and perfect $\Omega(Q)$ -ZK.*

The RPE which [HVW21a] use is obtained by applying a general observation of [GM10, BDG⁺18] – that the existence of linear codes implies the existence of RPEs – to the linear codes of Decatur et al. [DGR99]. Specifically, Ball et al. [BDG⁺18, Lemma 2] prove the following, where a code $\mathcal{C} \subseteq \{0, 1\}^n$ is linear if its encoding procedure simply multiplies the input with a public generator matrix, and the distance of the code is $\min_{c \in \mathcal{C}} \text{Ham}(c, 0^n)$ (i.e., the minimal weight of a non-zero codeword):

Lemma 20 (RPEs from linear error-correcting codes [BDG⁺18]). *If there exists a linear error-correcting code $\mathcal{C} \subseteq \{0, 1\}^n$ with messages in $\{0, 1\}^k$ and distance d , then there exists a $(k, n, d - 1)$ -RPE.*

To obtain an RPE with good parameters, we apply Lemma 20 to any explicit family of linear codes with constant rate and constant relative distance (e.g., the codes of [DGR99, Theorem 2.1], which already possess secrecy from partial views)¹³. In particular, we have

Corollary 21 (RPEs). *For every message length $k \in \mathbb{N}$, there exists a $(k, O(k), \Omega(k))$ -RPE.*

A ZK-PCP with Square-Root Query Gap. With these building blocks in place, we are ready to prove Theorem 14.

Proof of Theorem 14: We instantiate Theorem 17 with the ZK-PCP system of Theorem 19 and the RPE of Corollary 21. We assume without loss of generality that $q^* \geq n$ (where n is the input length). Since we set Q below to be polynomially-related to q^* (and consequently also to n), there exists a constant c such that the PCP of theorem 19 is over an alphabet of size 2^{Q^c} . Let $\alpha = 1/(c + 1)$, then we instantiate Theorem 19 with $Q := (q^*)^\alpha$, and set $k = Q^c$ in Corollary 21. Then Theorem 17 guarantees that the resultant ZK-PCP has proofs of length $Q \cdot O(Q^c) = O(Q^{c+1}) = (q^*)^{\alpha \cdot (c+1)} = O(q^*)$ with perfect ZK against (possibly malicious and adaptive) verifiers making $\Omega(Q) \cdot \Omega(Q^c) = \Omega(q^*)$ queries, and a $\text{negl}(Q) = \text{negl}(q^*)$ soundness error with a non-adaptive honest verifier whose query complexity is $\log Q \cdot \sqrt{Q} \cdot O(Q^c) = \tilde{O}(Q^{c+1/2}) = \tilde{O}\left((q^*)^{(c+1/2)/(c+1)}\right)$. The theorem now follows for any ϵ which is larger than $(c + 1/2)/(c + 1)$ (and for a sufficiently large q^*). ■

Theorem 14 allows one to choose the ZK query bound q^* . Hazay et al. [HVW21a] also give an alternative formulation of Theorem 14, in which the square-root query gap obtained by the ZK-PCP system is more clearly manifested.

Corollary 22 (ZK-PCP with \sqrt{n} query gap, Corollary 10 of [HVW21b]). *There exists a constant $c > 0$ such that there exists ZK-PCP with perfect $\Omega(n^{c+1})$ -ZK, and $\text{negl}(n)$ soundness error with an honest verifier that non-adaptively queries $\tilde{O}(n^{c+1/2})$ proof bits, where n denotes the input length.*

Proof. We instantiate Theorem 17 with the ZK-PCP system of Theorem 19 and the RPE of Corollary 21. Setting $Q = n$ in Theorem 19, let c be a constant such that the PCP of theorem 17 is over an

¹³We note that the construction of [DGR99] relies on Toeplitz matrices of logarithmic size which generate codes with good parameters. Such a matrix can be efficiently found by traversing all these matrices by some pre-defined order, and using the first matrix satisfying the desired properties.

alphabet of size n^c . We set $k = n^c$ in Corollary 21. Then Theorem 17 guarantees that the resultant ZK-PCP has proofs of length $n \cdot O(n^c) = O(n^{c+1})$ with perfect ZK against (possibly malicious and adaptive) verifiers making $\Omega(n) \cdot \Omega(n^c) = \Omega(n^{c+1})$ queries, and a $\text{negl}(n)$ soundness error with a non-adaptive honest verifier whose query complexity is $\log n \cdot \sqrt{n} \cdot O(n^c) = \tilde{O}(n^{c+1/2})$. ■

4 ZK-PCPs from LR Circuits: The Construction of [IWY16]

The main result of this section is a construction of witness-indistinguishable PCPs for NP with a non-adaptive honest verifier and an exponential query gap (between the query complexity of the honest and malicious verifiers), which was given in [IWY16]:

Theorem 23 (WI-PCPs for NP, formal statement of Informal Theorem 2). *Let $n \in \mathbb{N}$ be an input length parameter. For any query bound $q^* = \text{poly}(n)$ there exists a non-adaptive $\text{poly}(\log q^*, \kappa)$ -query $(q^*, \text{negl}(q^*))$ -WI-PCP for NP with $\text{negl}(\kappa)$ soundness error, where κ is a statistical security parameter.*

The WI-PCP system is constructed from leakage-resilient *circuits*. Historically, this construction was presented before the ZK-PCPs of Section 3, and was the first to show a connection between ZK-PCPs and leakage resilience. Moreover, it was the first construction of PCPs with (relaxed) ZK against *malicious* verifiers with *non-adaptive honest* verification. Compared to the ZK-PCPs of Section 3, the scheme we describe in this section has the advantage of obtaining an *exponential* query gap: the scheme is ZK against q^* -query bounded verifiers, but the honest verifier only needs to query $\text{polylog}(q^*)$ proof bits to verify the proof (setting $\kappa = \text{polylog}(q^*)$). However, the PCP system described in this section obtains a weaker form of ZK called *Witness Indistinguishability (WI)* in which the ZK simulator is not guaranteed to be efficient. The construction of Section 3 has the added feature of being simpler.

The main building-block in the transformation of [IWY16] are leakage-resilient *circuits* – a stronger primitive than the leakage-resilient *encodings* used by [HVW21a]. Indeed, leakage-resilient encodings only protect *information*, whereas leakage-resilient *circuits* protect *computations*. Thus, while Hazay et al. [HVW21a] could only use leakage-resilient encodings to protect the proof *once it was already generated*, Ishai et al. [IWY16] employ leakage-resilient circuits to protect *proof generation itself*. When used in the context of proof generation, these leakage-resilient circuits in effect amplify leakage resilience: from a relatively low leakage bound on the witness, to a much larger leakage bound on the entire computation. This amplification results in an exponential query gap, which eluded the ZK-PCPs of Section 3.

High-Level Idea: Leakage-Resilient Proof Generation. The goal of the prover \mathcal{P} is to convince the verifier \mathcal{V} that $x \in \mathcal{L}$, where \mathcal{P} has a corresponding witness w . Recall from Section 1.1.2 that one way of doing so is to emulate the verification circuit C of the corresponding NP-relation $\mathcal{R}_{\mathcal{L}}$ on (x, w) , where \mathcal{V} then checks that this computation was executed correctly. Indeed, without loss of generality we can assume that $\mathcal{R}_{\mathcal{L}}$ has a canonical form in which the witness consists of the entire wire values of C . In particular, in a PCP system the prover would generate the PCP from the entire wire values of C , which we denote by $[C, x, w]$. While this might blatantly violate ZK (e.g., w itself is part of $[C, x, w]$, and the PCP might explicitly contain these wire values), the main observation of [IWY16] is that by querying few proof bits, \mathcal{V} is leaking on a *computation* – namely, the evaluation of C on (x, w) . Therefore, one can use leakage-resilient circuits to guarantee that this leakage gives \mathcal{V} no information on the witness w (i.e., C 's input).

Concretely, one can first replace C with a *leakage-resilient version* \hat{C} , then have \mathcal{V} check the computation executed in \hat{C} . If \mathcal{V} were reading directly from the wire values of \hat{C} , then we would need

\widehat{C} to resist probing leakage (similar to Section 3). However, \mathcal{V} is actually probing bits in the PCP π , which was generated by applying the prover algorithm to the wire values of \widehat{C} . Therefore, \mathcal{V} obtains more evolved forms of leakage on these wire values, and we thus need leakage resilience against a wider class of potential leakage functions. Still, \mathcal{V} cannot obtain any leakage that it wants on the wire values of \widehat{C} , but rather it is restricted to whatever the prover computes on these wire values during proof generation. In particular, by restricting the types of functions the prover applies to the wire values of \widehat{C} , we can “control” the type of information which *even a malicious* verifier \mathcal{V}^* obtains on the witness. (Indeed, \mathcal{V}^* can deviate from an honest verification strategy by choosing to read different – and a larger number of – proof bits than the honest verifier, but cannot affect the proofs generation itself.) This gives a general blue-print for transforming standard PCPs into ZK-PCPs: the prover and verifier both hard-wire x into the verification circuit C , to obtain a circuit C_x . Then, they generate the LR-version of it \widehat{C}_x . Then, \mathcal{P} randomly encodes the witness w as \widehat{w} ¹⁴, and generates the PCP from the wire values $[\widehat{C}, \widehat{w}]$. Finally, the verifier \mathcal{V} runs the PCP verifier to verify the proof.

Achieving Soundness. While (for an appropriate choice of the family of leakage functions) the blue-print described above would indeed yield ZK proofs, they would not be sound. The reason is that soundness of the original PCP relied on the fact that $C(x, \cdot)$ is satisfiable only if $x \in \mathcal{L}$, i.e., there exists a corresponding witness w such that $(x, w) \in \mathcal{R}_{\mathcal{L}}$. This, however, is not preserved in the *leakage-resilient* version \widehat{C}_x of $C(x, \cdot)$. In fact, in many (if not all) existing constructions, \widehat{C}_x is *necessarily* satisfiable – leakage resilience relies on this fact! To understand why this is the case, and how we can still achieve soundness, we need to first take a closer look at how these LR constructions work.

4.1 Leakage-Resilient Circuit Compilers (LRCCs)

In this section we describe the main building block used in the ZK-PCPs of this section: Leakage-Resilient Circuits Compilers (LRCCs). Compared to the LR encodings used in Section 3, LRCCs offer a stronger LR guarantee: they protect *computation* rather than *information*. More specifically, LRCCs are compilers which transform a given circuit C into a functionally-equivalent circuit \widehat{C} which is leakage resilient in the sense that leakage on the wire values of \widehat{C} reveals no information on its input. Of course, this cannot be obtained if the input of C is given to \widehat{C} in the clear, since then leakage on the wire values reveals information about the input. Instead, the LRCC is associated with a (LR) encoding scheme, which is used to encode the inputs to (and intermediate values in the computation of) the LR circuit \widehat{C} .

We first define the circuit model which we use, then define circuit compilers, and finally define leakage-resilient circuit compilers.

Circuit Model. An (arithmetic) circuit C over the field \mathbb{F} and the set $X = \{x_1, \dots, x_n\}$ of variables is a directed acyclic graph whose vertices are called *gates* and whose edges are called *wires*, and are labeled with functions over X . Every gate in C of in-degree 0 has out-degree 1 and is either an *input* gate labeled by a variable from X ; or a *constant* gate const_α labeled by a constant $\alpha \in \mathbb{F}$. Gates of in-degree 2 and out-degree 1 are labeled by one of the operations $+$, $-$, \times (i.e., addition, subtraction,

¹⁴Recall that a LR circuit operates on encoded inputs, see also Section 4.1 below.

and multiplication over \mathbb{F}).¹⁵ The *size* $|C|$ of a circuit C is the sum of the number of wires, input gates, and output gates, in C . The *depth* of C is the number of gates on the longest path from inputs to outputs. We also consider *Boolean* circuits, with \wedge, \vee gates (replacing the $+, -, \times$ gates of arithmetic circuits), const_0 and const_1 gates, and \neg gates with in- and out-degree 1.

We define several circuit complexity classes, which restrict the size and depth of boolean and arithmetic circuits. Specifically,

Notation 24. *SHALLOW* (d, s) denotes the class of all depth- d , size- s arithmetic circuits over \mathbb{F} . Similarly, *BOOL* (d, s) denotes the class of all depth- d , size- s boolean circuits. Somewhat abusing notation, we use the same notations to denote the families of functions computable by circuits in the respective class of circuits. \mathcal{AC}^0 denotes all constant-depth and polynomial-sized boolean circuits over unbounded fan-in and fan-out (i.e., in-degree and out-degree) $\wedge, \vee, \neg, \text{const}_0$ and const_1 gates.

Definition 25 (Satisfiable Circuits). A circuit $C : \mathbb{F}^n \rightarrow \mathbb{F}$ is satisfiable if there exists an $x \in \mathbb{F}^n$ such that

$$C(x) = \begin{cases} 1 & C \text{ is boolean} \\ 0 & C \text{ is arithmetic} \end{cases}$$

Circuit Compilers. We define the notion of a circuit compiler. Informally, it consists of an encoding scheme and a compiler algorithm, that compiles a given circuit C into a circuit \widehat{C} that emulates the operation of C over encoded inputs. Formally,

Definition 26 (Circuit compiler over \mathbb{F}). A circuit compiler over \mathbb{F} is a pair $(\text{Comp}, (\text{Enc}, \text{Dec}))$ such that the following holds:

- **Syntax:**
 - (Enc, Dec) is an encoding scheme, where Enc is a PPT algorithm that on input a vector $x \in \mathbb{F}^n$, and 1^{len} ,¹⁶ outputs a vector \widehat{x} , and Dec is a polynomial-time algorithm. We assume that $\widehat{x} \in \mathbb{F}^{\widehat{n}}$ for some $\widehat{n} = \widehat{n}(n, \text{len})$.
 - Comp is a polynomial-time algorithm that on input an arithmetic circuit C over \mathbb{F} outputs an arithmetic circuit \widehat{C} .
- **correctness:** For any arithmetic circuit C , and any input x for C , we have $\Pr \left[\widehat{C}(\widehat{x}) = C(x) \right] = 1$, where $\widehat{x} \leftarrow \text{Enc}(x, 1^{|C|})$.

As discussed above, we will need circuit compilers that are also “sound” in the sense that the compiled circuit \widehat{C} is satisfiable only if the original circuit C is satisfiable. We stress that this property should hold even when the inputs of \widehat{C} are *not* valid encodings according to Enc .

Definition 27 (SAT-respecting circuit compiler). A circuit compiler $(\text{Comp}, (\text{Enc}, \text{Dec}))$ is SAT-respecting if for every circuit $C : \mathbb{F}^n \rightarrow \mathbb{F}$, if $\widehat{C} = \text{Comp}(C)$ is satisfiable then C is satisfiable. That is,

- For arithmetic C : if there exists an $\widehat{x}^* \in \mathbb{F}^{\widehat{n}}$ such that $\widehat{C}(\widehat{x}^*) = 0$, then there exists an $x \in \mathbb{F}^n$ such that $C(x) = 0$.
- For Boolean C : if there exists an $\widehat{x}^* \in \{0, 1\}^{\widehat{n}}$ such that $\widehat{C}(\widehat{x}^*) = 1$, then there exists an $x \in \{0, 1\}^n$ such that $C(x) = 1$.

¹⁵Jumping ahead, we will use an LRCC of $[\text{FRR}^+10]$, which employs additional gates (e.g., to duplicate values in the circuit). Since we do not explicitly use these additional gates, we omit their description to simplify the presentation.

¹⁶This additional parameter is a length parameter, which is used to determine the amount of random masks needed to protect the computation from leakage, see Definition 29 below.

Leakage-Resilient Circuit Compilers. We now define *Leakage-Resilient* circuit compilers. An LRCC is associated with a class \mathcal{LEAK} of leakage functions, and guarantees that when the input to \widehat{C} is properly encoded, then leakage from \mathcal{LEAK} on the wire values of \widehat{C} reveals no information about the input and internal computations, except for the output of \widehat{C} . This is formalized by requiring that the wire values are distributed statistically close for every pair of inputs x, x' such that $C(x) = C(x')$. We first set some notation.

Notation 28. For a Circuit C , a leakage function $\ell : \mathbb{F}^{|C|} \rightarrow \mathbb{F}^m$ for some $m \in \mathbb{N}$, and an input x for C , $[C, x]$ denotes the wire values of C when evaluated on x , and $\ell[C, x]$ denotes the output of ℓ on $[C, x]$.

Definition 29 (LRCC). Let \mathbb{F} be a finite field, \mathcal{LEAK} be a function class, $S(n) : \mathbb{N} \rightarrow \mathbb{N}$ be a size function, and $\epsilon(n) : \mathbb{N} \rightarrow \mathbb{R}^+$. A circuit compiler $(\text{Comp}, (\text{Enc}, \text{Dec}))$ is $(\mathcal{LEAK}, \epsilon(n), S(n))$ -leakage resilient if for all sufficiently large n 's, every arithmetic circuit $C : \mathbb{F}^n \rightarrow \mathbb{F}^m$ (for some m) of size $|C| \leq S(n)$, every $\ell \in \mathcal{LEAK}$ of input length $|\widehat{C}|$, and every $x, x' \in \mathbb{F}^n$ such that $C(x) = C(x')$, we have

$$\text{SD} \left(\ell \left[\widehat{C}, \widehat{x} \right], \ell \left[\widehat{C}, \widehat{x}' \right] \right) \leq \epsilon(n)$$

where $\widehat{x} \leftarrow \text{Enc}(x, 1^{|C|})$ and $\widehat{x}' \leftarrow \text{Enc}(x', 1^{|C|})$.

LRCCs for boolean circuits are defined similarly.

A few remarks are in order. First, we note that the LR guarantee of Definition 29 is a relaxed version of the standard notion. Specifically, while Definition 29 only guarantees *indistinguishability* – namely, leakage functions cannot distinguish between the wire values of \widehat{C} when evaluated on two different inputs (so long as C has the same output on both), the standard definition (e.g., in [ISW03, FRR⁺10, DF12, MV13]) is *simulation-based*. That is, the standard definition requires the existence of a *PPT* simulator which, for every circuit C , every input x for C , and any leakage function $\ell \in \mathcal{LEAK}$, can simulate the leakage on the wire values of \widehat{C} given only the description of C and its output $C(x)$. To see why this guarantee is stronger, notice that if for some input x there exists *no* x' such that $C(x) = C(x')$ then Definition 29 provides no secrecy guarantee for x . Moreover, Definition 29 is equivalent to a slightly modified version of the standard definition – specifically, in which there exists a simulator as specified above, but it is *not* guaranteed to be efficient. Indeed, given $C, C(x)$ the simulator could find on its own an x' such that $C(x) = C(x')$, then simulate the leakage by computing $\ell \left[\widehat{C}, \widehat{x}' \right]$. We focus on the relaxed version because it captures the security guarantee which we achieve. (In fact, Ishai et al. [IWY16] give strong indications that SAT-respecting LRCCs for “useful” leakage, with the stronger LR guarantee with *efficient* simulation, do not exist – see Section 1.2.)

Second, we note that for simplicity, the error in Definitions 27 and 29 depends (only) on the input length n . This can be naturally extended such that the error depends also on a security parameter κ , which is given as input to Comp .

Gadget-Based LRCCs. A leading technique in constructing LRCCs – which is the one employed in all LRCCs described in this work – is *gadget-based*. Such constructions employ a double-layer of encoding,¹⁷ where the LRCC (Comp, E) is associated with an internal encoding scheme $(\text{Enc}^{\text{in}}, \text{Dec}^{\text{in}})$,

¹⁷Recall from Section 1.1.2 that standard LRCCs in the literature are described as *randomized* circuits that generate the needed randomness internally (sometimes, using “opaque” gates which are assumed to be leak-free), whereas we describe the LR circuit as deterministic, and provide the needed randomness as part of the input to the circuit. For this reason, we need to use a double layer of encoding, which is usually implicit in the literature.

and Comp outputs a circuit \widehat{C} in which the gates and wires of the circuit C over \mathbb{F} are replaced with *gadgets* and *bundles*. A *bundle* \mathbf{b} is an encoding (according to Enc^{in}) of some $b \in \mathbb{F}$, representing the value of a wire in C ; and a *gadget* \mathcal{G}_g is a (boolean or arithmetic) circuit over \mathbb{F} which operates over bundles and emulates the operation of the corresponding gate g . More specifically, in addition to taking as input bundle-encodings of g 's inputs, \mathcal{G}_g has additional *masking inputs*. These masking inputs are encodings (according to Enc^{in}) of some *masking values* in \mathbb{F}^* with a *specific, pre-determined* structure (for example, in [FRR⁺10] a masking value is the all-0 string), which are used to obtain LR. We associate with \mathcal{G}_g a set $\mathcal{WF}_g \subset \mathbb{F}^*$ (for *Well-Formed*) which consists of all encodings of masking values with the “correct” structure. We say that the masking inputs of \mathcal{G}_g are *well formed* if they are in \mathcal{WF}_g (otherwise we say they are *ill-formed*). A gadget \mathcal{G}_g is guaranteed to emulate g when its inputs are valid encodings of inputs for g , and its masking inputs are well-formed. For example, if $g = \times$, then for every $x_1, x_2 \in \mathbb{F}$, for every bundle encodings $\mathbf{x}_i \leftarrow \text{Enc}^{\text{in}}(x_i)$, $i = 1, 2$, and for every well-formed masking input bundles \mathbf{m} , $\mathcal{G}_g(\mathbf{x}_1, \mathbf{x}_2, \mathbf{m})$ encodes $x_1 \times x_2$.

Leakage resilience of \widehat{C} will follow from a combination of the leakage resilience of the internal encoding scheme Enc^{in} , and the following leakage-resilience guarantee of the gadgets. If the masking inputs are uniformly distributed over $\text{Enc}^{\text{in}}(\mathcal{WF}_g)$, then given any valid encodings as standard inputs to the gadget: (1) the outputs are random subject to encoding the correct output (as determined by the standard inputs and the gate operation); and (2) the internal wire values of \mathcal{G} can be reconstructed (i.e., regenerated) in a low complexity class, where the reconstructed and actual wire values are statistically close.

The masking inputs for all gadgets of \widehat{C} are provided as part of the inputs to \widehat{C} , using the double-layered encoding, as we now explain. To simplify the description, we focus on describing how this is done for the LRCC of [FRR⁺10] (see Remark 30 below on how the double-layered encoding extends to other LRCCs as well). In the LRCC of [FRR⁺10], the masking values for *all gates* g are simply the all-0 string, and different gates only differ in the *length* of the string (i.e., the number of 0's). Let M denote the upper bound on the number of masking values used by any gate g . Then $E = (\text{Enc}, \text{Dec})$ where Enc takes as input an $x \in \mathbb{F}^*$, and a size bound S , and outputs an encoding $(\mathbf{x}, \mathbf{m}) = \text{Enc}^{\text{in}}(x, 0^{M \cdot S})$. For a given a circuit C of size $|C| \leq S$, its leakage-resilient version $\widehat{C} = \text{Comp}(C)$ takes as input encodings of the form (\mathbf{x}, \mathbf{m}) , where \mathbf{x} are the input bundles of \widehat{C} , and \mathbf{m} are used as the masking inputs of the gadgets of \widehat{C} . Since each gadget uses at most M encodings from \mathbf{m} , and \widehat{C} has at most S gadgets, \mathbf{m} contains sufficiently-many encodings (according to Enc^{in}) of 0 to enable the evaluation of all gadgets of \widehat{C} .

The computation in \widehat{C} is executed over encodings, so the outcome of this computation – at least as it was described above – results in an *encoding* of the output of C . Since \widehat{C} should have the same output as C , the final step in \widehat{C} consists of a decoder which decodes the output.

Remark 30. *We note that the double-layer encoding idea described above naturally extends to other LRCCs as well. Indeed, all that is needed is that all gadgets use the same masking values (differing only in the number of masking values they use). In fact, this idea can also be used if different gadgets use different masking values, in which case \mathbf{m} can include S sets of masking inputs, each set containing masking inputs for every possible gate g , from which \widehat{C} can choose the appropriate masking inputs for each gate. This will result in a correct and leakage resilient (albeit less efficient) construction.*

4.2 A SAT-Respecting LRCC for Arithmetic Circuits

In this section we describe a *SAT-respecting* LRCC of [IWY16] for arithmetic circuits, which is based on the LRCC of [FRR⁺10]. We first explain why the LRCC of [FRR⁺10] (and similar LRCCs) are *not* SAT-respecting.

Why are standard LRCCs not SAT-Respecting? Recall that an LRCC is SAT-respecting if for every circuit C , the following holds: if its leakage-resilient variant \widehat{C} is satisfiable, then so is C . This is not the case for many LRCCs, and specifically for gadget-based ones such as [FRR⁺10, DF12, MV13].¹⁸ Indeed, in such LRCCs correctness of the computation relies on the assumption that the masking inputs of the gadgets are well-formed.¹⁹ However, \widehat{C} does not necessarily emulate C if its masking inputs are *ill-formed*. In fact, in these LRCCs one can force the output of \widehat{C} to be any desired value by appropriately choosing the masking values. This property is *crucial* for leakage resilience, since the security proof uses such ill-formed masking inputs to switch the input from x to a different x' . Therefore, the main challenge in obtaining the SAT-respecting property is to guarantee that the masking inputs are well-formed while still allowing the security reduction to use ill-formed masking inputs.

Checking Validity of Encodings While Preserving Leakage Resilience. The main technical ingredient in the SAT-respecting LRCC of [IWY16] is their component which checks well-formedness of encodings in a leakage-resilient manner. The technique is reminiscent of the “2-key trick” of [NY90] (used to convert a CPA-secure encryption scheme into a CCA-secure one) where they hold two copies of C , and in the security reduction one of the copies is used to achieve the SAT-respecting property, whereas the other is used to obtain leakage resilience. This component is tailored to the LRCC of [FRR⁺10], exploiting the fact that well-formed masking inputs of [FRR⁺10] are simply encodings of the all-0 string.

More specifically, the SAT-respecting LRCC of [IWY16] consists of 3 parts. The first is two copies $\widehat{C}_1, \widehat{C}_2$ of the circuit \widehat{C} , obtained using the LRCC of [FRR⁺10]. The second is a mask-checking component $\mathcal{C}_{0\text{-check}}$, which checks that at least one of the copies $\widehat{C}_1, \widehat{C}_2$ uses well-formed masking inputs. The important point is that, for the security proof to go through, $\mathcal{C}_{0\text{-check}}$ must hide *which* of the copies uses well-formed masking inputs. (This is because in the security proof first \widehat{C}_1 and then \widehat{C}_2 use ill-formed masking inputs; for the hybrids to be indistinguishable it must be the case that one cannot distinguish between the two cases.) As it turns out, this can be achieved by replacing $\mathcal{C}_{0\text{-check}}$ with its leakage-resilient version $\widehat{\mathcal{C}}_{0\text{-check}}$ (obtained using the LRCC of [FRR⁺10]). This now introduces a further complication since $\widehat{\mathcal{C}}_{0\text{-check}}$ has masking inputs of its own, whose well-formedness must be verified. Indeed, if $\widehat{\mathcal{C}}_{0\text{-check}}$ is allowed to use any masking inputs, then it is no longer guaranteed to emulate $\mathcal{C}_{0\text{-check}}$. In particular, by using ill-formed masking inputs in $\widehat{\mathcal{C}}_{0\text{-check}}$ one can flip its output, thus causing it to accept even when the masking inputs of *both* $\widehat{C}_1, \widehat{C}_2$ are ill-formed, rendering the $\widehat{\mathcal{C}}_{0\text{-check}}$ component completely useless. The third component – a mask decoder \mathcal{C}_{dec} – is used to guarantee that $\widehat{\mathcal{C}}_{0\text{-check}}$ uses well-formed masking inputs. Using a double-layer of mask checking (i.e., checking the masks of $\widehat{C}_1, \widehat{C}_2$, and then checking the masks of the mask-checker) is helpful because the computations in $\widehat{\mathcal{C}}_{0\text{-check}}$ are *not directly related* to the inputs of $\widehat{C}_1, \widehat{C}_2$. As such, the masking inputs of $\widehat{\mathcal{C}}_{0\text{-check}}$ can be *checked directly* by simply decoding them (and checking that all decoded values are 0) without violating LR. We now formally describe each of the components, and the resultant SAT-respecting LRCC.

¹⁸One exception is the LRCC of [ISW03] which resists \mathcal{AC}^0 leakage [BIS19], but this leakage class does not seem sufficiently strong to contain the function one needs to apply to an NP witness to obtain a PCP.

¹⁹As noted above, in the leakage-resilience literature, such encodings are usually assumed to be generated by leak-free components, or opaque gates, whose outputs are guaranteed to be distributed according to the correct distribution (i.e., be well-formed), and whose internal wires are unavailable to the leakage function. To eliminate this trust assumption, we instead have these encodings provided as part of the input, and check their validity to achieve the SAT-respecting property.

The Mask Checker $\widehat{\mathcal{C}}_{0\text{-check}}$. The mask checker verifies that at least one of the copies $\widehat{\mathcal{C}}_1, \widehat{\mathcal{C}}_2$ uses well-formed masking inputs, while hiding which one. Recall that in the LRCC of [FRR⁺10] well-formed masking inputs are encodings of the all-0 string. Thus, if \vec{m}^1, \vec{m}^2 denote the masking values whose encodings are used in $\widehat{\mathcal{C}}_1, \widehat{\mathcal{C}}_2$, it suffices to check that $m_j^1 \times m_l^2 = 0$ for every j, l .²⁰ We construct $\mathcal{C}_{0\text{-check}}$ in two stages. We first describe a “binarization” component \mathcal{T} which checks that a given field element is 0, then use it to construct $\mathcal{C}_{0\text{-check}}$.

Notation 31 (“Binarization” component \mathcal{T}). $\mathcal{T} : \mathbb{F} \rightarrow \mathbb{F}$ is defined as $\mathcal{T}(z) = -\prod_{0 \neq a \in \mathbb{F}} (z - a)$, computed using $O(|\mathbb{F}|)$ constant and \times gates arranged in $O(\log |\mathbb{F}|)$ layers. Notice that

$$\mathcal{T}(z) = \begin{cases} 1 & \text{if } z = 0 \\ 0 & \text{if } z \neq 0 \end{cases}$$

Notation 32 (Mask Checker $\mathcal{C}_{0\text{-check}}$). Let $M \in \mathbb{N}$. The mask checker $\mathcal{C}_{0\text{-check}} : \mathbb{F}^M \times \mathbb{F}^M \rightarrow \mathbb{F}$ is defined as follows. $\mathcal{C}_{0\text{-check}}(y, z) = \prod_{i, j \in [M]} \mathcal{T}(y_i \times z_j)$, computed using a multiplication tree of size $O(M^2)$ and depth $O(\log M)$ (on top of the multiplication trees used to compute \mathcal{T}). Notice that

$$\mathcal{C}_{0\text{-check}}(y, z) = 1 \Leftrightarrow \mathcal{T}(y_i, z_j) = 1 \quad \forall i, j \in [M] \Leftrightarrow y = 0^M \vee z = 0^M.$$

The mask decoder \mathcal{C}_{dec} . As noted above, the mask decoder simply decodes the masking inputs used in $\widehat{\mathcal{C}}_{0\text{-check}}$, and checks that all decoded values are 0. Decoding is done using the decoder Dec^{in} of the internal encoding procedure of [FRR⁺10]. We note that this encoding procedure is linear, and so decoding is done simply by computing the inner product of the input encoding with some fixed vector (e.g., the all-1 vector). In particular, for each encoding length \widehat{n} , the corresponding decoding circuit Dec^{in} can be implemented using $O(\widehat{n})$ gates arranged in $O(\log \widehat{n})$ layers.

Notation 33 (Mask Decoder \mathcal{C}_{dec}). Let $M_0 \in \mathbb{N}$, and let \widehat{n} denote the encoding length of $(\text{Enc}^{\text{in}}, \text{Dec}^{\text{in}})$. The mask decoder $\mathcal{C}_{\text{dec}} : (\mathbb{F}^{\widehat{n}})^{M_0} \rightarrow \mathbb{F}$, on input $\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_{M_0})$ (where $\mathbf{r}_i \in \mathbb{F}^{\widehat{n}}$ for every $1 \leq i \leq M_0$) outputs $\prod_{i \in [M_0]} \mathcal{T}(\text{Dec}^{\text{in}}(\mathbf{r}_i))$. Notice that \mathcal{C}_{dec} outputs 1 if and only if all \mathbf{r}_i ’s are well-formed, otherwise it outputs 0. \mathcal{C}_{dec} is computed using $O(M_0)$ -many \times gates, arranged in a tree of depth $O(\log M_0)$ (on top of the sub-circuits $\mathcal{T} \circ \text{Dec}^{\text{in}}$).

The SAT-Respecting LRCC. Having defined the three components of the SAT-respecting LRCC, we are now ready to describe the compiler itself, which is given in Figure 2.

Remark 35 (Setting the parameters). Let $M^* = M^*(\sigma)$ denote the maximal number of masking inputs used in a gadget of Comp^* , and $S_0(M')$ denote the size of $\mathcal{C}_{0\text{-check}}$ on inputs of length M' . Then $M(\sigma) = \sigma \cdot M^*$ and $M_0(\sigma) = M^* \cdot S_0(M) = M^* \cdot S_0(\sigma \cdot M^*)$.

The following claim summarizes the properties of Construction 34. Roughly, it states that if the internal encoding scheme used in Construction 34 is leakage resilient against a leakage family \mathcal{LEAK}_E , then Construction 34 is a SAT-respecting LRCC against a slightly weaker leakage family \mathcal{LEAK} . Formally,

Claim 36 (SAT-respecting LRCC over \mathbb{F}). Let $\mathcal{LEAK}, \mathcal{LEAK}_E$ be families of functions, $S(n) : \mathbb{N} \rightarrow \mathbb{N}$ be a size function, and $\epsilon(n) : \mathbb{N} \rightarrow \mathbb{R}^+$. Let $\mathbf{E}^{\text{in}} = (\text{Enc}^{\text{in}}, \text{Dec}^{\text{in}})$ be a linear, onto, $(\mathcal{LEAK}_E, \epsilon(n))$ -leakage-resilient encoding scheme with parameters n, σ and $\widehat{n} = \widehat{n}(n, \sigma)$, such that

²⁰This check assumes that any masking input is a valid encoding – according to Enc^{in} – of some masking value (i.e., that the internal encoding scheme is onto). This is indeed the case for the internal encoding scheme used in [FRR⁺10].

Construction 34 (SAT-Respecting LRCC). Let σ be a security parameter, and $M = M(\sigma)$, $M_0 = M_0(\sigma) : \mathbb{N} \rightarrow \mathbb{N}$ (the value of these parameters will be set in Remark 35 below). Let $(\text{Comp}^*, (\text{Enc}^*, \text{Dec}^*))$ denote the LRCC of [FRR⁺10], and let $(\text{Enc}^{\text{in}} : \mathbb{F}^n \times \{1\}^* \rightarrow \mathbb{F}^{\widehat{n}_{\text{in}}(n, \sigma)}, \text{Dec}^{\text{in}})$ denote the internal encoding scheme which it employs. The SAT-respecting LRCC $(\text{Comp}, E = (\text{Enc}, \text{Dec}))$ is defined as follows.

The Encoding Scheme. $\text{Enc} : \mathbb{F}^n \times \{1\}^* \rightarrow \mathbb{F}^{2\widehat{n}_{\text{in}} + 2\widehat{M} + 2\widehat{M}_0}$ on input $x, 1^\sigma$ outputs an encoding $(\widehat{x}_1, \widehat{x}_2)$, where $\widehat{x}_i \leftarrow \text{Enc}^{\text{in}}((x, 0^{M+M_0}), 1^\sigma)$, and $\widehat{n}_{\text{in}}, \widehat{M}, \widehat{M}_0$ denote the lengths of encodings of messages of length m, M, M_0 , respectively. Dec on input $(\widehat{x}_1, \widehat{x}_2), 1^\sigma$ computes $\text{Dec}^{\text{in}}(\widehat{x}_1, 1^\sigma)$, and discards the last $M + M_0$ field elements. We use $\widehat{n} = \widehat{n}(n, \sigma)$ to denote the length of encodings output by Enc . For $(\widehat{x}_1, \widehat{x}_2) \leftarrow \text{Enc}(x, 1^\sigma)$, we interpret $\widehat{x}_i = (\widehat{x}_i^{\text{in}}, \mathbf{r}^i, \mathbf{r}^{i,0})$, where $\widehat{x}_i^{\text{in}}$ encodes x , and $\mathbf{r}^i, \mathbf{r}^{i,0}$ encode $0^M, 0^{M_0}$, respectively.^a

The Compiler Comp. On input a circuit $C : \mathbb{F}^n \rightarrow \mathbb{F}$, the compiler Comp outputs the circuit $\widehat{C} : \mathbb{F}^{\widehat{n}(n, |C|)} \rightarrow \mathbb{F}$ that on input $(\widehat{x}_1^{\text{in}}, \mathbf{r}^1, \mathbf{r}^{1,0}), (\widehat{x}_2^{\text{in}}, \mathbf{r}^2, \mathbf{r}^{2,0})$ operates as follows.

- Let $C_1 = C_2 = C$ denote two copies of C , let $\widehat{C}_i = \text{Comp}^*(C_i)$ for $i = 1, 2$, and $\widehat{C}_{0\text{-check}} = \text{Comp}^*(C_{0\text{-check}})$.
- \widehat{C} computes

$$f := \mathcal{T} \left(\widehat{C}_1(\widehat{x}_1^{\text{in}}, \mathbf{r}^1) - \widehat{C}_2(\widehat{x}_2^{\text{in}}, \mathbf{r}^2) \right) \times \widehat{C}_{0\text{-check}}((\mathbf{r}^1, \mathbf{r}^2), \mathbf{r}^{1,0}) \times C_{\text{dec}}(\mathbf{r}^{1,0}).$$

(Notice that $f = 1$ if and only if (1) $\widehat{C}_1(\widehat{x}_1^{\text{in}}, \mathbf{r}^1) = \widehat{C}_2(\widehat{x}_2^{\text{in}}, \mathbf{r}^2)$, and (2) the masking inputs used in at least one of them, as well as in $\widehat{C}_{0\text{-check}}$, are well-formed. Otherwise, $f = 0$.)

- \widehat{C} outputs

$$(1 - f((\widehat{x}_1^{\text{in}}, \mathbf{r}^1, \mathbf{r}^{1,0}), (\widehat{x}_2^{\text{in}}, \mathbf{r}^2, \mathbf{r}^{2,0}))) + f((\widehat{x}_1^{\text{in}}, \mathbf{r}^1, \mathbf{r}^{1,0}), (\widehat{x}_2^{\text{in}}, \mathbf{r}^2, \mathbf{r}^{2,0})) \cdot \widehat{C}_1(\widehat{x}_1^{\text{in}}, \mathbf{r}^1)$$

(Notice that the output is $\widehat{C}_1(\widehat{x}_1^{\text{in}}, \mathbf{r}^1, \mathbf{r}^{1,0})$ if $f = 1$, otherwise it is 1.)

^aWe note that $\mathbf{r}^{2,0}$ is not used by the leakage-resilient circuit, but is included in \widehat{x}_2 because the same internal encoding scheme Enc^{in} is used to generate both \widehat{x}_1 and \widehat{x}_2 .

Figure 2: SAT-Respecting LRCC [IWY16]

$\mathcal{LEAK}_E = \mathcal{LEAK} \circ \mathcal{SHALLOW}(7, O(\hat{n}^4(1, S(n)) \cdot S(n)))$. Then there exists a SAT-respecting, $(\mathcal{LEAK}, 4\epsilon(n) \cdot (\hat{n}(1, S(n)) + 1) \cdot S(n), S(n))$ -LRCC over \mathbb{F} . Moreover, for every $C : \mathbb{F}^n \rightarrow \mathbb{F}$, the compiled circuit \hat{C} has size $|\hat{C}| = O(|\mathbb{F}|^2 \cdot \hat{n}^4(1, S(n)) \cdot |C|^2)$.

Proof. The SAT-respecting property follows from Lemma 38. The leakage-resilience property follows from Lemma 39. As for the size $|\hat{C}|$ of the leakage-resilient version \hat{C} of C , in each of the copies \hat{C}_1, \hat{C}_2 each gate is replaced with a size- $O(\hat{n}^2(1, S(n)))$ gadget (this is the size of gadgets generated by Comp^* , see Fact 37 below), So $|\hat{C}_1|, |\hat{C}_2| \leq O(\hat{n}^2(1, S(n)) \cdot |C|)$. Since each gadget uses at most $O(\hat{n}(1, S(n)))$ masking inputs (see Fact 37 below), then $\mathcal{C}_{0\text{-check}}$ contains $O(\hat{n}^2(1, S(n)) \cdot |C|^2)$ “binarization” components \mathcal{T} , each of size at most $O(|\mathbb{F}|)$, arranged in a tree with $O(\hat{n}^2(1, S(n)) \cdot |C|^2)$ multiplications, so $|\mathcal{C}_{0\text{-check}}| \leq O(|\mathbb{F}| \cdot \hat{n}^2(1, S(n)) \cdot |C|^2)$, and consequently $|\hat{\mathcal{C}}_{0\text{-check}}| \leq O(|\mathbb{F}| \cdot \hat{n}^4(1, S(n)) \cdot |C|^2)$. Finally, \mathcal{C}_{dec} contains a decoding sub-circuit for each of the $O(\hat{n}(1, S(n)))$ masking inputs used in the $O(|\mathbb{F}| \cdot \hat{n}^2(1, S(n)) \cdot |C|^2)$ gadgets of $\hat{\mathcal{C}}_{0\text{-check}}$. Because E is linear, each of these decoding sub-circuits consists of $\hat{n}(1, S(n)) \times$ gates followed by $\hat{n}(1, S(n)) +$ gates. In addition, \mathcal{C}_{dec} contains a “binarization” component \mathcal{T} of size $O(|\mathbb{F}|)$ for each decoding sub-circuit, followed by $O(|\mathbb{F}| \cdot \hat{n}^3(1, S(n)) \cdot |C|^2) \times$ gates, so overall $|\mathcal{C}_{\text{dec}}| \leq O(|\mathbb{F}|^2 \cdot \hat{n}^3(1, S(n)) \cdot |C|^2)$. ■

The proof of Claim 36 used the following fact regarding the LRCC of [FRR⁺10].

Fact 37. Each gadget generated by the LRCC (Comp^*, E^*) of [FRR⁺10] has size at most $O(\hat{n}^2(1, S(n)))$, and uses at most $O(\hat{n}(1, S(n)))$ masking inputs.

Lemma 38. If E^{in} is linear and onto, then Construction 34 is SAT-respecting.

Proof. Assume that $\hat{C}(\hat{x}) = 0$ for some $\hat{x} \in \mathbb{F}^{\hat{n}}$, and denote $\hat{x} = ((\hat{x}_1^*, \mathbf{r}^1, \mathbf{r}^{1,0}), (\hat{x}_2^*, \mathbf{r}^2, \mathbf{r}^{2,0}))$. We show that there exists an $x \in \mathbb{F}^n$ such that $C(x) = 0$. Since $\hat{C}(\hat{x}) = 0$ then by the definition of \hat{C} , we have (1) $f = 1$, and (2) $\hat{C}_1(\hat{x}_1^*, \mathbf{r}^1) = 0$. By (1), $\hat{C}_1(\hat{x}_1^*, \mathbf{r}^1) = \hat{C}_2(\hat{x}_2^*, \mathbf{r}^2)$, and $\mathcal{C}_{\text{dec}}, \mathcal{C}_{0\text{-check}}$ output 1. Notation 33 then guarantees that $\mathbf{r}^{1,0}$ is well-formed which – by the correctness of Comp^* – guarantees that $\hat{\mathcal{C}}_{0\text{-check}}$ emulates $\mathcal{C}_{0\text{-check}}$. (Here, we also use the fact that \mathcal{C}_{dec} is independent of all other components of, and inputs to, \hat{C} .) Moreover, since the encoding scheme is onto then $\mathbf{r}^1, \mathbf{r}^2$ define inputs to $\mathcal{C}_{0\text{-check}}$ which cause it to output 1 (because $\hat{\mathcal{C}}_{0\text{-check}}$ outputs 1, and its masking inputs are well-formed). Notation 32 then guarantees that at least one of $\mathbf{r}^1, \mathbf{r}^2$ is well-formed. Assuming (without loss of generality) that \mathbf{r}^1 is well-formed, then the correctness of Comp^* guarantees that \hat{C}_1 emulates C , so $0 = \hat{C}_1(\hat{x}_1^*, \mathbf{r}^1) = C(x)$, where $x \in \mathbb{F}^n$ is obtained by computing $x = \text{Dec}^{\text{in}}(\hat{x}_1^*)$ (x is well-defined because E^{in} is onto). ■

Lemma 39. If E^{in} is $(\mathcal{LEAK}_E, \epsilon(n))$ -leakage resilient, then for every function class \mathcal{LEAK} such that $\mathcal{LEAK} \circ \mathcal{SHALLOW}(7, O(\hat{n}^5(S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$, every circuit $C : \mathbb{F}^n \rightarrow \mathbb{F}$ of size $|C| \leq S(n)$, and every $x, y \in \mathbb{F}^n$ such that $C(x) = C(y)$, it holds that $[\hat{C}, \hat{x}]$ and $[\hat{C}, \hat{y}]$ are $(\mathcal{LEAK}, 4\epsilon(n) \cdot (\hat{n}(1, S(n)) + 1) \cdot S(n))$ -leakage resilient, where $\hat{x} \leftarrow \text{Enc}(x, 1^{|C|})$ and $\hat{y} \leftarrow \text{Enc}(y, 1^{|C|})$.

4.2.1 Proof of Lemma 39

In this section we prove the leakage-resilience property of Construction 34. The analysis follows a (by now standard) proof paradigm for gadget-based LRCCs, but is more complex (compared to, e.g., the analysis in [FRR⁺10]) because of the mask-checker and mask-decoder components.

High-Level Description of the Leakage-Resilience Argument. Recall that the goal is to prove that leakage functions in some leakage class \mathcal{LEAK} cannot distinguish between the wire values of \widehat{C} when evaluated on (encodings of) two different inputs x, y . This is achieved by reduction to the leakage resilience of the underlying encoding scheme E against leakage classes in a somewhat larger leakage class \mathcal{LEAK}_E . The wire bundles of \widehat{C} are divided into two sets: *internal* bundles that are part of the internal computations in a gadget, and *external* bundles that connect two gadgets. The proof uses a hybrid argument in which we first replace all the internal bundles from carrying the real encodings to “simulated” encodings, and then replace all external bundles from encodings of the real values to encodings of random values. We then prove leakage resilience by showing that each pair of adjacent hybrids are computationally indistinguishable.

More specifically, consider two hybrid distributions $\mathcal{H}, \mathcal{H}'$ in which we replace some external bundle i . Proving that $\mathcal{H} \approx \mathcal{H}'$ is by reduction to the leakage resilience of the underlying encoding scheme, where we show that if \mathcal{H} and \mathcal{H}' are distinguishable by a function $\ell \in \mathcal{LEAK}$, then there exists a function $\ell_E \in \mathcal{LEAK}_E$ that can distinguish between encodings of two different values v, v' . The idea is that given an encoding e of either v or v' , ℓ_E would use e to generate *the entire hybrid distribution* (which would be either \mathcal{H} or \mathcal{H}' , depending on whether e encodes v or v' , respectively), then evaluate ℓ . It turns out that the only wires of $\mathcal{H}, \mathcal{H}'$ that depend on e are the internal wire bundles of the (at most two) gadgets that “touch” bundle i , in the sense that bundle i is either an input or output bundle of the gadget. Thus, while all other wire bundles can be hard-wired into ℓ_E , it would still need to generate the internal wire bundles of the gadget(s) that touch bundle i . In particular, the computational complexity of ℓ_E would be higher than that of ℓ , which causes a loss in leakage resilience. It is therefore imperative that the internal wire values of the gadgets could be generated from their inputs and outputs by a function in a low complexity class. These functions are called *local reconstructors*. The crucial point here is that the local reconstructor is given not only the gadget’s inputs, but also its *outputs* (otherwise, it would be impossible to generate the internal wires in a low complexity class, because some of the gadgets perform complex computations). Before delving deeper into the details of the analysis, we first set the needed terminology regarding gadgets and their local reconstructors.

Properties of Gadgets. The analysis will use two properties of gadgets. The first is *local reconstructibility*: for every pair of “legal” input and output encodings, the internal wires of the gadget (as determined by the input-output pair, and its masking inputs), can be simulated in a low complexity class.

Definition 40 (Local reconstructibility). *Let \mathcal{G} be a gadget and \mathcal{WF} denote its set of well-formed masking inputs. A pair (\mathbf{x}, \mathbf{y}) of encodings is plausible for \mathcal{G} if $\mathcal{G}(\mathbf{x}, \mathbf{m}) = \mathbf{y}$ for some $\mathbf{m} \in \mathcal{WF}$. For $\epsilon > 0$, and families $\mathcal{LEAK}, \mathcal{F}_{\mathcal{G}}$ of functions, \mathcal{G} is $(\mathcal{LEAK}, \epsilon)$ -reconstructible by $\mathcal{F}_{\mathcal{G}}$ if the following holds. There exists a distribution REC over functions rec such that:*

- $\text{Supp}(\text{REC}) \subseteq \mathcal{F}_{\mathcal{G}}$.
- *rec takes as input \mathcal{G} ’s inputs and output, and simulates the masking inputs, and internal wires, of \mathcal{G} .*

- The following distributions are $(\mathcal{LEAK}, \epsilon)$ -leakage resilient for any plausible pair (\mathbf{x}, \mathbf{y}) for \mathcal{G} : (1) $\text{rec}(\mathbf{x}, \mathbf{y})$ for $\text{rec} \leftarrow \text{REC}$; and (2) the actual distribution of the wires of \mathcal{G} (as determined by the distribution of the masking inputs), conditioned on \mathbf{x}, \mathbf{y} .

We will need the following result of [FRR⁺10] which shows that all gadgets in their LRCC (Comp^*, E^*) are locally reconstructible in a low complexity class.

Lemma 41 (Gadgets are locally reconstructible [FRR⁺10]). ²¹ Let $\sigma \in \mathbb{N}$ be a security parameter, let $n : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ be a length parameter, let $\epsilon(n) : \mathbb{N} \rightarrow \mathbb{R}^+$, and let $\mathcal{LEAK}, \mathcal{LEAK}_E$ be families of functions such that $\mathcal{LEAK}_E = \mathcal{LEAK} \circ \text{SHALLOW}(3, O(\widehat{n}(1, \sigma)))$. Let (Comp^*, E^*) denote the LRCC of [FRR⁺10], and let E^{in} denote the internal encoding scheme which it uses. If E^{in} is $(\mathcal{LEAK}_E, \epsilon(n))$ -leakage resilient, then all gadgets used in (Comp^*, E^*) are $(\mathcal{LEAK}, \widehat{n}(1, \sigma) \cdot \epsilon(n))$ -reconstructible by $\text{SHALLOW}(2, O(\widehat{n}^2(1, \sigma)))$.

The second property is that gadgets are *re-randomizing* in the sense that the encodings at the output of each gadget are uniform subject to encoding the “correct” value. Formally,

Definition 42 (Gadget Re-Randomization). A gadget \mathcal{G} with set \mathcal{WF} of well-formed masking inputs is re-randomizing if for every standard input $\mathbf{x} = \text{Enc}(x)$, when the masking input is sampled $\mathbf{m} \leftarrow \mathcal{WF}$ then $\mathcal{G}(\mathbf{x}, \mathbf{m})$ is random subject to encoding the correct output (as determined by x , and the gate which \mathcal{G} emulates).

The Hybrid Argument. Let $C : \mathbb{F}^n \rightarrow \mathbb{F}$ be an arithmetic circuit of size $|C| \leq S(n)$, and let $x, y \in \mathbb{F}^n$. We show that for every $\ell \in \mathcal{LEAK}$,

$$\text{SD} \left(\ell \left[\widehat{C}, \widehat{x} \right], \ell \left[\widehat{C}, \widehat{y} \right] \right) \leq \epsilon'(n)$$

where $\epsilon'(n) := 4\epsilon(n) \cdot (\widehat{n}(1, S(n)) + 1) \cdot S(n)$. We bound the statistical distance using a hybrid argument. We define:

$$\begin{aligned} \mathcal{H}^x &:= \left(\left[\widehat{C}_1, \widehat{x}_1, \mathbf{r}^1 \right], \left[\widehat{C}_2, \widehat{x}_2, \mathbf{r}^2 \right], \left[\widehat{C}_{0\text{-check}}, \mathbf{r}^1, \mathbf{r}^2, \mathbf{r}^{1,0} \right], \left[\mathcal{C}_{\text{dec}}, \mathbf{r}^{1,0} \right] \right) \\ \mathcal{H}^y &:= \left(\left[\widehat{C}_1, \widehat{y}_1, \mathbf{r}^1 \right], \left[\widehat{C}_2, \widehat{y}_2, \mathbf{r}^2 \right], \left[\widehat{C}_{0\text{-check}}, \mathbf{r}^1, \mathbf{r}^2, \mathbf{r}^{1,0} \right], \left[\mathcal{C}_{\text{dec}}, \mathbf{r}^{1,0} \right] \right) \\ \mathcal{H}^{y,x} &:= \left(\left[\widehat{C}_1, \widehat{y}_1, \mathbf{r}^1 \right], \left[\widehat{C}_2, \widehat{x}_2, \mathbf{r}^2 \right], \left[\widehat{C}_{0\text{-check}}, \mathbf{r}^1, \mathbf{r}^2, \mathbf{r}^{1,0} \right], \left[\mathcal{C}_{\text{dec}}, \mathbf{r}^{1,0} \right] \right) \end{aligned}$$

and notice that \mathcal{H}^x are the wire values of \widehat{C} on input an encoding of x , \mathcal{H}^y are the wire values of \widehat{C} on input an encoding of y , and $\mathcal{H}^{y,x}$ is a hybrid distribution, consisting of the wire values of \widehat{C} when the first copy \widehat{C}_1 has input (an encoding according to Enc^{in}) of x , whereas the second copy \widehat{C}_2 has input (an encoding according to Enc^{in}) of y . We note that $\mathcal{H}^{y,x}$ is only used in the proof – it is never obtained in an actual evaluation of \widehat{C} . The proof proceeds by showing that for every $\ell \in \mathcal{LEAK}$, $\text{SD}(\ell(\mathcal{H}^x), \ell(\mathcal{H}^{y,x}))$ and $\text{SD}(\ell(\mathcal{H}^{y,x}), \ell(\mathcal{H}^y))$ are upper bounded by $2\epsilon(n) \cdot (\widehat{n}(1, S(n)) + 1) \cdot S(n)$, through a sequence of hybrids.

²¹Faust et al. [FRR⁺10] actually prove a stronger result, where some of the gadgets have local reconstructors in lower complexity classes than the one stated here, and indistinguishability holds regardless of the leakage resilience of E . For clarity reasons, we chose to give a simplified and weaker version of their results which nonetheless suffices for our needs.

Bounding $SD(\ell(\mathcal{H}^x), \ell(\mathcal{H}^{y,x}))$ for a leakage function ℓ . The difference between these distributions is that in \mathcal{H}^x both copies $\widehat{C}_1, \widehat{C}_2$ of (the leakage-resilient version of) C are evaluated on input x , whereas in $\mathcal{H}^{y,x}$ the first is evaluated on y (and the second on x). We bound the statistical distance through a sequence of hybrids, defined as follows.

$\mathcal{H}_{\text{in}}^x$: this hybrid distribution replaces the internal wires of gadgets, and is obtained by: (1) evaluating \widehat{C} honestly on $\widehat{x} \leftarrow \text{Enc}(x, 1^{|C|})$; (2) picking local reconstructors for all gadgets of \widehat{C}_1 (such reconstructors exist by Lemma 41), and re-computing their internal wires using these reconstructors; and (3) re-evaluating $\widehat{C}_{0\text{-check}}$ on the new masking inputs generated for the gadgets of \widehat{C}_1 , by re-computing the internal wires of all gadgets \mathcal{G}' of $\widehat{C}_{0\text{-check}}$ whose inputs are masking inputs used in gadgets of \widehat{C}_1 – but without changing the masking inputs these \mathcal{G}' gadgets use, or their outputs.²² (Crucially, since re-evaluating $\widehat{C}_{0\text{-check}}$ does not influence its masking inputs, this does not influence the computation in \mathcal{C}_{dec} , so there is no need to re-evaluate it.)

$\mathcal{H}_{\text{ext}}^x$: this distribution replaces the external wires (i.e., wires connecting gadgets), and is obtained as follows:

- *Generating the wires of \widehat{C}_2* : encode $\widehat{x} = ((\widehat{x}_1, \mathbf{r}^1, \mathbf{r}^{1,0}), (\widehat{x}_2, \mathbf{r}^2, \mathbf{r}^{2,0})) \leftarrow \text{Enc}(x, 1^{|C|})$, and honestly evaluate \widehat{C}_2 on \widehat{x}_2 with masking inputs \mathbf{r}^2 .
- *Generating the wires of \widehat{C}_1* : pick a random encoding $\text{out} \leftarrow \text{Enc}^{\text{in}}(1, 1^{|C|})$ for the output of \widehat{C}_1 , and honestly compute the wires of the output decoder of \widehat{C}_1 . (As discussed in Section 4.1, \widehat{C}_1 contains an output decoder, which is needed because the computations in \widehat{C}_1 are executed over encodings.) Then, pick a random input $z \in_R \mathbb{F}^m$ for \widehat{C}_1 and encode $\widehat{z}_1 \leftarrow \text{Enc}^{\text{in}}(z, 1^{|C|})$. Next, pick random encodings (according to Enc^{in}) for the outputs of all gadgets (except the gadgets whose outputs are the inputs of the output decoder, since the outputs of these gadgets have already been fixed). This effectively determines the standard inputs, and outputs, of all gadgets of \widehat{C}_1 . Next, pick local reconstructors for all gadgets of \widehat{C}_1 , and use them to compute the internal wires of the gadgets. The reconstructors determine the (possibly ill-formed) masking inputs of the gadgets, which we denote by $\mathbf{r}^{1'}$. $\mathbf{r}^{1'}$ together with \mathbf{r}^2 form the standard inputs of $\widehat{C}_{0\text{-check}}$.
- *Generating the wires of $\widehat{C}_{0\text{-check}}$* : Evaluate $\widehat{C}_{0\text{-check}}$ on $\mathbf{r}^{1'}, \mathbf{r}^2$, with masking inputs $\mathbf{r}^{1,0}$.
- *Generating the wires of \mathcal{C}_{dec}* : Evaluate \mathcal{C}_{dec} on $\mathbf{r}^{1,0}$.
- Use the outputs of $\widehat{C}_1, \widehat{C}_2, \widehat{C}_{0\text{-check}}, \mathcal{C}_{\text{dec}}$ to generate the flag f , and the output of \widehat{C} .
- $\mathcal{H}_{\text{ext}}^x$ consists of the concatenation of all these wire values.

$\mathcal{H}_{\text{ext}}^{y,x}$: this hybrid is generated similarly to $\mathcal{H}_{\text{ext}}^x$, except that instead of evaluating \widehat{C} on an encoding of x , we use the internal encoding scheme to generate encodings of $(\widehat{y}_1, \mathbf{r}^1, \mathbf{r}^{1,0})$ and $(\widehat{x}_2, \mathbf{r}^2, \mathbf{r}^{2,0})$ (where $\widehat{y}_1, \widehat{x}_2$ encode y, x , respectively), and use them as inputs to $\widehat{C}_1, \widehat{C}_2$, respectively.

$\mathcal{H}_{\text{in}}^{y,x}$: this hybrid is generated similarly to $\mathcal{H}_{\text{in}}^x$, except that instead of evaluating \widehat{C} on an encoding of x , we use $(\widehat{y}_1, \mathbf{r}^1, \mathbf{r}^{1,0})$ and $(\widehat{x}_2, \mathbf{r}^2, \mathbf{r}^{2,0})$ as inputs to $\widehat{C}_1, \widehat{C}_2$, respectively.

We now bound the statistical distance between the outputs of leakage functions on each pair of adjacent hybrids. The following notation will be useful.

²²Lemma 44 below shows that this is possible.

Notation 43. We say that a gadget \mathcal{G}' of $\widehat{\mathcal{C}}_{0\text{-check}}$ is connected to a gadget \mathcal{G}_1 $\widehat{\mathcal{C}}_1$ (alternatively, a gadget \mathcal{G}_2 of $\widehat{\mathcal{C}}_2$) if an input of \mathcal{G}' is a masking input of \mathcal{G}_1 (alternatively, \mathcal{G}_2).

Bounding $\text{SD}(\ell(\mathcal{H}^x), \ell(\mathcal{H}_{\text{in}}^x))$: we show that $\text{SD}(\ell(\mathcal{H}^x), \ell(\mathcal{H}_{\text{in}}^x)) \leq \epsilon(n) \cdot \widehat{n}(1, S(n)) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$ such that $\mathcal{LEAK}' \circ \text{SHALLOW}(5, O(\widehat{n}^5(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$. We use a hybrid argument, replacing the internal wires of the $M \leq S(n)$ gadgets of $\widehat{\mathcal{C}}_1$ one at a time. We define hybrids $\mathcal{H}_0, \dots, \mathcal{H}_M$ where \mathcal{H}_i is obtained by evaluating $\widehat{\mathcal{C}}$ on (an encoding of) x , then recomputing the internal wires of the first i gadgets of $\widehat{\mathcal{C}}_1$ using their local reconstructors. Additionally, we recompute the internal wires of gadgets \mathcal{G}' of $\widehat{\mathcal{C}}_{0\text{-check}}$ connected to one of these i gadgets, but *without changing the masking inputs of \mathcal{G}'* (this is possible by Lemma 44 below). Then $\mathcal{H}_0 = \mathcal{H}^x$, $\mathcal{H}_M = \mathcal{H}_{\text{in}}^x$. To show that $\text{SD}(\ell(\mathcal{H}^x), \ell(\mathcal{H}_{\text{in}}^x)) \leq \epsilon(n) \cdot \widehat{n}(1, S(n)) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$, we show that for every $m \in [M]$, and any $\ell \in \mathcal{LEAK}'$, it holds that $\text{SD}(\ell(\mathcal{H}_m), \ell(\mathcal{H}_{m-1})) \leq \epsilon(n) \cdot \widehat{n}(1, S(n))$. Denote the m 'th gadget by \mathcal{G} . We use Lemma 41 to show that this follows from the leakage resilience of the internal encoding scheme E_{in} . For this, we will need to generate – in a low complexity class – the entire hybrid distributions given only the (either real or reconstructed) internal wires of \mathcal{G} . The problem is that since the internal wires of \mathcal{G} contain also its masking inputs, changing them affects also computations in $\widehat{\mathcal{C}}_{0\text{-check}}$. The following lemma from [IWY16] states that the influence of modifying the masking inputs of \mathcal{G} can be blocked, specifically: (1) that it only affects the internal wires of gadgets \mathcal{G}' of $\widehat{\mathcal{C}}_{0\text{-check}}$ connected to \mathcal{G} , and more importantly (2) these internal wire values can be reconstructed *without* changing the masking inputs used in \mathcal{G}' .

Lemma 44 (Restating of Lemma 3.13 of [IWY15]). *Let \mathcal{G}' be a gadget of $\widehat{\mathcal{C}}_{0\text{-check}}$ connected to gadgets of $\widehat{\mathcal{C}}_1, \widehat{\mathcal{C}}_2$. Then for any fixed well-formed inputs r^1, r^2 from gadgets of $\widehat{\mathcal{C}}_1, \widehat{\mathcal{C}}_2$ (respectively), and any fixed well-formed masking inputs m for \mathcal{G}' , the following holds. For any $r^{1'}$, the internal wires of \mathcal{G}' on input $r^{1'}, r^2, m$ can be computed in $\text{SHALLOW}(2, O(\widehat{n}^2(1, S(n))))$ given $r^{1'}$ and the output out of \mathcal{G}' when evaluated on input r^1, r^2 and masking inputs m .*

Using an averaging argument, we can fix all wires of $\widehat{\mathcal{C}}_2$, and all wires of $\widehat{\mathcal{C}}_1$ except the internal wires of \mathcal{G} (its input and output wires *can* be fixed). Lemma 44 (which can be used because $\widehat{\mathcal{C}}_2$ and $\widehat{\mathcal{C}}_{0\text{-check}}$ use well-formed masking inputs) shows that we can further fix all wires of $\widehat{\mathcal{C}}_{0\text{-check}}$ and \mathcal{C}_{dec} , except for the internal wires of all gadgets \mathcal{G}' of $\widehat{\mathcal{C}}_{0\text{-check}}$ connected to \mathcal{G} (crucially, the masking inputs of these gadgets *can* be fixed).

Consequently, given the wire values of \mathcal{G} (either the real values \mathcal{W}_R , or the reconstructed wires \mathcal{W}_S), we can generate the entire hybrid distribution (either \mathcal{H}_{m-1} or \mathcal{H}_m , respectively) in $\text{SHALLOW}(2, O(\widehat{n}^5(1, S(n)) \cdot S(n)))$ by recomputing the internal wires of each of the (at most) $O(\widehat{n}^3(1, S(n)) \cdot S(n))$ gadgets of $\widehat{\mathcal{C}}_{0\text{-check}}$ connected to \mathcal{G} . (Here, we use Fact 37 – \mathcal{G} uses at most $\widehat{n}(1, S(n))$ masking inputs; $\widehat{\mathcal{C}}_2$ has at most $O(\widehat{n}(1, S(n)) \cdot S(n))$ gadgets, each using at most $\widehat{n}(1, S(n))$ masking inputs; and each masking input of \mathcal{G} is connected to every masking input used in $\widehat{\mathcal{C}}_2$.) Using Lemma 44, the internal wires of each \mathcal{G}' can be computed in $\text{SHALLOW}(2, O(\widehat{n}^2(1, S(n))))$, and all these computations can be done in parallel. Therefore, the hybrid distributions can be generated in $\text{SHALLOW}(2, O(\widehat{n}^5(1, S(n)) \cdot S(n)))$. By the assumption of Claim 36, E_{in} is $(\mathcal{LEAK}_E, \epsilon(n))$ -leakage resilient, which by Lemma 41 implies that $\mathcal{W}_R, \mathcal{W}_S$ are $(\mathcal{LEAK}'', \epsilon(n) \cdot \widehat{n}(1, S(n)))$ -leakage resilient for any class \mathcal{LEAK}'' of leakage functions such that $\mathcal{LEAK}'' \circ \text{SHALLOW}(3, O(\widehat{n}(1, S(n)))) \subseteq \mathcal{LEAK}_E$. We now claim that $\mathcal{H}_{m-1}, \mathcal{H}_m$ are $(\mathcal{LEAK}', \epsilon(n) \cdot \widehat{n}(1, S(n)))$ -leakage resilient for every family \mathcal{LEAK}' of leakage functions such that $\mathcal{LEAK}' \circ \text{SHALLOW}(2, O(\widehat{n}^5(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}''$, i.e., for any \mathcal{LEAK}' such that $\mathcal{LEAK}' \circ \text{SHALLOW}(5, O(\widehat{n}^5(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$. Indeed, this follows from the follow-

ing lemma of [FRR⁺10] (by choosing \mathcal{F} to be a singleton containing the function described above which generates $\mathcal{H}_{m-1}, \mathcal{H}_m$ from $\mathcal{W}_R, \mathcal{W}_S$).

Lemma 45 ([FRR⁺10]). *Let $n \in \mathbb{N}$, let $\mathcal{W}_R, \mathcal{W}_S$ be distributions over \mathbb{F}^n , let $\mathcal{LEAK}, \mathcal{F}$ be families of functions, and let $\epsilon > 0$. Let \mathcal{D} be a distribution over functions in \mathcal{F} of input length n . For $f \leftarrow \mathcal{D}$, let $\mathcal{W}'_R := f(\mathcal{W}_R), \mathcal{W}'_S := f(\mathcal{W}_S)$. If $\mathcal{W}_R, \mathcal{W}_S$ are $(\mathcal{LEAK}, \epsilon)$ -leakage resilient, then $\mathcal{W}'_R, \mathcal{W}'_S$ are $(\mathcal{LEAK}', \epsilon)$ -leakage resilient for any family \mathcal{LEAK}' of leakage functions such that $\mathcal{LEAK}' \circ \mathcal{F} \subseteq \mathcal{LEAK}$.*

Using Lemma 45, for any $\ell \in \mathcal{LEAK}'$ we have

$$\text{SD}(\ell(\mathcal{H}_m), \ell(\mathcal{H}_{m-1})) = \text{SD}(\ell(\mathcal{W}_S), \ell(\mathcal{W}_R)) \leq \epsilon(n) \cdot \widehat{n}(1, S(n)).$$

Bounding $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^x), \ell(\mathcal{H}_{\text{in}}^x))$: We show that $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^x), \ell(\mathcal{H}_{\text{in}}^x)) \leq \epsilon(n) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$ such that $\mathcal{LEAK}' \circ \text{SHALLOW}(4, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$. We again use a hybrid argument, this time replacing the $M \leq S(n)$ inputs bundles of \widehat{C}_1 (i.e., bundles corresponding to input wires of C), and bundles at the output of gadgets of \widehat{C}_1 (except for bundles that are used as input of the output decoder of \widehat{C}_1). We define hybrids $\mathcal{H}_0, \dots, \mathcal{H}_M$, where \mathcal{H}_i is generated as follows. We evaluate \widehat{C} on a random encoding of x . Then, we replace the first i bundles with random encodings of random values, except if one of these bundles corresponds to the output wire of C , in which case we replace it with a random encoding of 1. Finally, we recompute the internal wires of the gadgets of \widehat{C}_1 using their gadget reconstructors (Lemma 41), and recompute (using Lemma 44) the internal wires (except the masking inputs and outputs) of all gadgets \mathcal{G}' of $\widehat{C}_{0\text{-check}}$ connected to one of these gadgets. In particular, since the inputs of $\widehat{C}_{0\text{-check}}$ contain also masking inputs used in gadgets of \widehat{C}_1 , this re-computation of wires of $\widehat{C}_{0\text{-check}}$ uses the masking inputs generated by the reconstructors (for any gadget of \widehat{C}_1 whose internal wires were re-constructed). Then $\mathcal{H}_0 = H_{\text{in}}^x$ and $\mathcal{H}_M = \mathcal{H}_{\text{ext}}^x$. Therefore, to prove that $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^x), \ell(\mathcal{H}_{\text{in}}^x)) \leq \epsilon(n) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$, it suffices to prove that $\text{SD}(\ell(\mathcal{H}_m), \ell(\mathcal{H}_{m-1})) \leq \epsilon(n)$ for all $m \in [M]$. Let \mathcal{G}_o (\mathcal{G}_i) denote the gadget whose output (input) is the m' th bundle. (If the m' th bundle is an input bundle, then we consider only the gadget \mathcal{G}_i .) Using an averaging argument, we can fix all wires in $\mathcal{H}_m, \mathcal{H}_{m-1}$ except for: the m' th bundle; the masking inputs, outputs, and internal wires of \mathcal{G}_o ; the masking inputs, and internal wires, of \mathcal{G}_i , as well as its input wire corresponding to the m' th bundle; and the internal wires of all gadgets \mathcal{G}' of $\widehat{C}_{0\text{-check}}$ connected to \mathcal{G}_o or \mathcal{G}_i (except for the masking inputs and the output of \mathcal{G}' , which *can* be fixed, see Lemma 44).

Let b denote the value encoded by the m' th bundle in H_{in}^x . Let $\mathcal{W}_R \leftarrow \text{Enc}^{\text{in}}(b, 1^{S(n)})$, and $\mathcal{W}_S \leftarrow \text{Enc}^{\text{in}}(r, 1^{S(n)})$ for a random r (except when the m' th bundle corresponds to the output wire of C , in which case we set $\mathcal{W}_S \leftarrow \text{Enc}^{\text{in}}(1, 1^{S(n)})$). Then $\mathcal{W}_R, \mathcal{W}_S$ are distributed identically to the m' th bundle in $\mathcal{H}_{m-1}, \mathcal{H}_m$, respectively. We define a distribution \mathcal{F} over $\text{SHALLOW}(4, O(\widehat{n}^4(1, S(n)) \cdot S(n)))$ as follows. Sampling a function $f \leftarrow \mathcal{F}$ is done by sampling $\text{rec}_o, \text{rec}_i$ from the distribution over reconstructors for $\mathcal{G}_o, \mathcal{G}_i$, respectively (see Definition 40). The function f has all the hard-wired values of \mathcal{H}_{m-1} hard-wired into it. On input $\mathbf{e} \in \mathbb{F}^{\widehat{n}(1, S(n))}$, f performs the following: (1) evaluates rec_o on the (hard-wired) inputs of \mathcal{G}_o , and the output \mathbf{e} (this reconstructs the masking inputs, and internal wires, of \mathcal{G}_o); (2) evaluates rec_i on \mathbf{e} as one of the inputs, and the other (hard-wired) input and output of \mathcal{G}_i ; and finally (3) for every gadget \mathcal{G}' of $\widehat{C}_{0\text{-check}}$ connected to \mathcal{G}_o or \mathcal{G}_i , generates its internal wires using Lemma 44 (without changing the output or masking inputs of \mathcal{G}'). Then $f \in \text{SHALLOW}(4, O(\widehat{n}^4(1, S(n)) \cdot S(n)))$. Indeed, by Lemma 41, $\text{rec}_o, \text{rec}_i \in \text{SHALLOW}(2, O(\widehat{n}^2(1, S(n))))$ (and given \mathbf{e} , they can be evaluated in parallel). Moreover, given the internal wires of $\mathcal{G}_o, \mathcal{G}_i$, the wires that need to be computed in a gadget \mathcal{G}' connected to them are computable in $\text{SHALLOW}(2, O(\widehat{n}^2(1, S(n))))$ (by Lemma 44, and

since the masking inputs of \widehat{C}_2 and $\widehat{C}_{0\text{-check}}$ are well-formed). We conclude by noting that there are at most $O(\widehat{n}^2(1, S(n))) \cdot S(n)$ such gadgets, and they can be evaluated in parallel.

Denote $\mathcal{W}'_R := f(\mathcal{W}_R), \mathcal{W}'_S := f(\mathcal{W}_S)$ for $f \leftarrow \mathcal{F}$, then $\mathcal{W}'_R \equiv \mathcal{H}_{m-1}$ and $\mathcal{W}'_S \equiv \mathcal{H}_m$ because the gadgets are re-randomizing (which, in particular, guarantees that these equivalences hold *despite* the fact some of the values in the computation were fixed in advance). By the assumption of Lemma 39, $\mathcal{W}_R, \mathcal{W}_S$ – which are encodings according to E^{in} – are $(\mathcal{LEAK}_E, \epsilon(n))$ -leakage resilient. Therefore, by Lemma 45, $\mathcal{W}'_R, \mathcal{W}'_S$ (and consequently also $\mathcal{H}_{m-1}, \mathcal{H}_m$) are $(\mathcal{LEAK}', \epsilon(n))$ -leakage resilient for any \mathcal{LEAK}' such that $\mathcal{LEAK}' \circ \mathcal{SHALLOW}(4, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$.

Bounding $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^x), \ell(\mathcal{H}_{\text{ext}}^{y,x}))$: We show that $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^x), \ell(\mathcal{H}_{\text{ext}}^{y,x})) = 0$ for every ℓ . Indeed, $\mathcal{H}_{\text{ext}}^x \equiv \mathcal{H}_{\text{ext}}^{y,x}$ because the hybrids are independent of the input for \widehat{C}_1 (since the input is re-sampled as a fresh $z \in_R \mathbb{F}^n$ in both hybrids).

Bounding $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^{y,x}), \ell(\mathcal{H}_{\text{in}}^{y,x}))$: We show that $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^{y,x}), \ell(\mathcal{H}_{\text{in}}^{y,x})) \leq \epsilon(n) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$ where $\mathcal{LEAK}' \circ \mathcal{SHALLOW}(4, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$. The proof is similar to the proof that $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^x), \ell(\mathcal{H}_{\text{in}}^x)) \leq \epsilon(n) \cdot S(n)$, because the argument was independent of the actual inputs used in $\widehat{C}_1, \widehat{C}_2$ (as long as both hybrids use the same input in each copy).

Bounding $\text{SD}(\ell(\mathcal{H}_{\text{in}}^{y,x}), \ell(\mathcal{H}^{y,x}))$: We show that $\text{SD}(\ell(\mathcal{H}_{\text{in}}^{y,x}), \ell(\mathcal{H}^{y,x})) \leq \epsilon(n) \cdot \widehat{n}(1, S(n)) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$, where $\mathcal{LEAK}' \circ \mathcal{SHALLOW}(5, O(\widehat{n}^5(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$. The proof is similar to the proof that $\text{SD}(\ell(\mathcal{H}^x), \ell(\mathcal{H}_{\text{in}}^x)) \leq \epsilon(n) \cdot \widehat{n}(1, S(n)) \cdot S(n)$, because the argument was independent of the actual inputs used in $\widehat{C}_1, \widehat{C}_2$.

Bounding $\text{SD}(\ell(\mathcal{H}^y), \ell(\mathcal{H}^{y,x}))$ for a leakage function ℓ . The argument here follows the same blueprint as the one used to bound $\text{SD}(\ell(\mathcal{H}^x), \ell(\mathcal{H}^{y,x}))$, but is more involved because we now need to switch the input of the *second copy* \widehat{C}_2 . In particular, the wire values in this hybrid argument will no longer correspond to the values in an *actual* evaluation of \widehat{C}_2 . While the computations in \widehat{C}_1 will still be executed honestly, we will no longer be able to claim that reconstructing the internal wires of gadgets of \widehat{C}_2 (or the external wires connecting gadgets of \widehat{C}_2) does not affect the computations in $\widehat{C}_{0\text{-check}}$ and \mathcal{C}_{dec} . This is because the manner in which $\widehat{C}_{0\text{-check}}$ uses the masking inputs from $\widehat{C}_1, \widehat{C}_2$ is *not symmetric*, and in particular, resampling the masking inputs used in \widehat{C}_2 (as is done by the local reconstructors of gadgets of \widehat{C}_2) *will* affect the computations in $\widehat{C}_{0\text{-check}}$, and will necessitate reevaluating it. This, in turn, will affect the masking inputs used in $\widehat{C}_{0\text{-check}}$, which will affect the computations in \mathcal{C}_{dec} . However, we cannot simply re-evaluate $\widehat{C}_{0\text{-check}}$ and \mathcal{C}_{dec} in the hybrid argument. Indeed, this would require evaluating circuits of large depth, and the leakage resilience guarantee will therefore deteriorate significantly. Instead, Ishai et al. [IWY16] use an alternative method of locally reconstructing the needed wire values of $\widehat{C}_{0\text{-check}}, \mathcal{C}_{\text{dec}}$. Specifically, they show a low-depth reconstructor for the gadgets \mathcal{G}' of $\widehat{C}_{0\text{-check}}$ connected to gadgets of \widehat{C}_2 , which generates the “correct” distribution if the input of \mathcal{G}' from \widehat{C}_1 , and \mathcal{G}' 's output, are well formed. In particular, this implies that the internal wires of \mathcal{G}' can be reconstructed without modifying its output. They also show a low-depth reconstructor for each decoding circuit of \mathcal{C}_{dec} , that generates the “correct” distribution if its inputs are well formed (again, without changing the output of these decoding circuits). Specifically, we will use the following results from [IWY16].

Lemma 46 (Local reconstructors for $\widehat{C}_{0\text{-check}}, \mathcal{C}_{\text{dec}}$, restatement of Lemmas 3.17 and 3.18 of [IWY15]). *There exists a distribution REC over $\mathcal{SHALLOW}(2, O(\widehat{n}^2(1, S(n))))$ such that the following holds for any gadget \mathcal{G}' of $\widehat{C}_{0\text{-check}}$ connected to a gadget of \widehat{C}_1 or \widehat{C}_2 . For every plausible pair $((\mathbf{r}^{1,0}, \mathbf{r}^{2,0}), \mathbf{c})$ for \mathcal{G}' such that $\mathbf{r}^{1,0}$ is well formed, when $\text{rec} \leftarrow \text{REC}$ then $\text{rec}(\mathbf{r}^{1,0}, \mathbf{r}^{2,0}, \mathbf{c})$ is distributed identically to the wire values of \mathcal{G}' in a real execution, conditioned on the input-output pair $((\mathbf{r}^{1,0}, \mathbf{r}^{2,0}), \mathbf{c})$. In particular, \mathbf{c} and the masking inputs computed by rec are well-formed.*

Moreover, for a fixed $\text{rec} \in \text{Supp}(\text{REC})$, and any fixed well-formed $\mathbf{r}^{1,0}, \mathbf{c}$, there exists a function $\text{rec}_{\text{Dec}}^{\mathbf{r}^{1,0}, \mathbf{c}, \text{rec}} \in \text{SHALLOW}(2, O(\widehat{n}(1, S(n))))$ such that the following holds for any input $\mathbf{r}^{2,0}$ of \mathcal{G}' . If $((\mathbf{r}^{1,0}, \mathbf{r}^{2,0}), \mathbf{c})$ is a plausible pair for \mathcal{G}' , then $\text{rec}_{\text{Dec}}^{\mathbf{r}^{1,0}, \mathbf{c}, \text{rec}}(\mathbf{r}^{2,0})$ generates the wire values of the decoding sub-circuits of \mathcal{C}_{dec} whose inputs are the masking inputs of \mathcal{G}' which rec generated. In particular, because the masking inputs generated by rec are well-formed, the outputs of these decoding sub-circuits are 0.

The proof now follows using a hybrid argument. The hybrids are similar to the ones used to bound $\text{SD}(\ell(\mathcal{H}^x), \ell(\mathcal{H}^{y,x}))$, except that we will need to regenerate the internal wires of $\widehat{\mathcal{C}}_{0\text{-check}}, \mathcal{C}_{\text{dec}}$ using the reconstructors of Lemma 46. We now define the hybrids.

$\mathcal{H}_{\text{in}}^y$: this hybrid distribution replaces the internal wires of gadgets (similar to $\mathcal{H}_{\text{in}}^x$), and is obtained by: (1) evaluating \widehat{C} honestly on $\widehat{y} \leftarrow \text{Enc}(y, 1^{|C|})$; (2) picking local reconstructors for all gadgets of \widehat{C}_2 , and re-computing their internal wires using these reconstructors; (3) re-computing the internal wires of $\widehat{C}_{0\text{-check}}$ using $\text{rec} \leftarrow \text{REC}$, where REC is the distribution of Lemma 46; and (4) re-computing the internal wires of \mathcal{C}_{dec} using the functions $\text{rec}_{\text{Dec}}^{\mathbf{r}^{1,0}, \mathbf{c}, \text{rec}}$ defined in Lemma 46 (here, $\mathbf{r}^{1,0}, \mathbf{c}$ are determined by the re-computed wires values of $\widehat{C}_{0\text{-check}}$).

$\mathcal{H}_{\text{ext}}^y$: this distribution replaces the external wires (similar to $\mathcal{H}_{\text{ext}}^x$), and is obtained as follows:

- *Generating the wires of \widehat{C}_1* : encode $\widehat{y} = ((\widehat{y}_1, \mathbf{r}^1, \mathbf{r}^{1,0}), (\widehat{y}_2, \mathbf{r}^2, \mathbf{r}^{2,0})) \leftarrow \text{Enc}(y, 1^{|C|})$, and honestly evaluate \widehat{C}_1 on \widehat{y}_1 with masking inputs \mathbf{r}^1 .
- *Generating the wires of \widehat{C}_2* : pick a random input $z \in_R \mathbb{F}^n$ for \widehat{C}_2 , and generate random encodings $\text{out} \leftarrow \text{Enc}^{\text{in}}(1, 1^{|C|})$, $\widehat{z}_1 \leftarrow \text{Enc}^{\text{in}}(z, 1^{|C|})$ for the output and input of \widehat{C}_2 . Next, pick random encodings (according to Enc^{in}) for the outputs of all gadgets (except the gadgets whose outputs are the inputs of the output decoder, since the outputs of these gadgets have already been fixed). Then, pick local reconstructors for all gadgets of \widehat{C}_2 , and use them to compute the internal wires of the gadgets. The reconstructors determine the (possibly ill-formed) masking inputs $\mathbf{r}^{2'}$ of the gadgets, which (together with \mathbf{r}^1) form the standard inputs of $\widehat{C}_{0\text{-check}}$.
- *Generating the wires of $\widehat{C}_{0\text{-check}}$* : for every gadget \mathcal{G}' of $\widehat{C}_{0\text{-check}}$ connected to a gadget of \widehat{C}_2 , pick a reconstructor for it according to the distribution REC (Lemma 46) and use it to compute the internal wires of \mathcal{G}' . These reconstructors determine the inputs to the decoding sub-circuits of \mathcal{C}_{dec} .
- *Generating the wires of \mathcal{C}_{dec}* : Use the functions of Lemma 46 to compute the internal wires of the decoding sub-circuits of the \mathcal{C}_{dec} .
- Use the outputs of $\widehat{C}_1, \widehat{C}_2, \widehat{C}_{0\text{-check}}, \mathcal{C}_{\text{dec}}$ to generate the flag f , and the output of \widehat{C} .
- $\mathcal{H}_{\text{ext}}^y$ consists of the concatenation of all these wire values.

$\mathcal{H}_{\text{ext},2}^{y,x}$: this hybrid is generated similarly to $\mathcal{H}_{\text{ext}}^y$, except that instead of evaluating \widehat{C} on an encoding of y , we use the internal encoding scheme to generate encodings of $(\widehat{y}_1, \mathbf{r}^1, \mathbf{r}^{1,0})$ and $(\widehat{x}_2, \mathbf{r}^2, \mathbf{r}^{2,0})$ (where $\widehat{y}_1, \widehat{x}_2$ encode y, x , respectively), and use them as inputs to $\widehat{C}_1, \widehat{C}_2$, respectively.

$\mathcal{H}_{\text{in},2}^{y,x}$: this hybrid is generated similarly to $\mathcal{H}_{\text{in}}^y$, except that instead of evaluating \widehat{C} on an encoding of y , we use $(\widehat{y}_1, \mathbf{r}^1, \mathbf{r}^{1,0})$ and $(\widehat{x}_2, \mathbf{r}^2, \mathbf{r}^{2,0})$ as inputs to $\widehat{C}_1, \widehat{C}_2$, respectively.

The indistinguishability of the hybrids now follows similarly to the proof bounding $\text{SD}(\mathcal{H}^x, \mathcal{H}^{y,x})$, and we only sketch the difference.

Bounding $\text{SD}(\ell(\mathcal{H}^y), \ell(\mathcal{H}_{\text{in}}^y))$: We show that $\text{SD}(\ell(\mathcal{H}^y), \ell(\mathcal{H}_{\text{in}}^y)) \leq \epsilon(n) \cdot \widehat{n}(1, S(n)) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$ such that $\mathcal{LEAK}' \circ \mathcal{SHALLOW}(7, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$. We define the hybrids $\mathcal{H}_0, \dots, \mathcal{H}_M$, where \mathcal{H}_i is obtained by: (1) evaluating \widehat{C} on (an encoding of) y , then recomputing the internal wires of the first i gadgets of \widehat{C}_2 using their local reconstructors; (2) *re-computing the internal wires of $\widehat{C}_{0\text{-check}}$* that are influenced by this re-computation of the first i gadgets of \widehat{C}_2 (i.e., gadgets of $\widehat{C}_{0\text{-check}}$ that are connected to one of these i gadgets); and (3) (using the functions of Lemma 46) *recomputing the internal wires of \mathcal{C}_{dec}* that were influenced by re-computing $\widehat{C}_{0\text{-check}}$. Then $\mathcal{H}_0 = \mathcal{H}^y$, $\mathcal{H}_M = \mathcal{H}_{\text{in}}^y$, and we show that $\text{SD}(\ell(\mathcal{H}_m), \ell(\mathcal{H}_{m-1})) \leq \epsilon(n) \cdot \widehat{n}(1, S(n))$ for every $m \in [M]$, and any $\ell \in \mathcal{LEAK}'$. Denote the m 'th gadget by \mathcal{G} . We can fix all wires of \widehat{C}_1 , all wires of \widehat{C}_2 except the internal wires \mathcal{W} of \mathcal{G} (its input and output wires *can* be fixed), and all the internal wires of $\widehat{C}_{0\text{-check}}, \mathcal{C}_{\text{dec}}$ that are influenced by \mathcal{W} . (These consist of the internal wires – but not the output – of any gadget \mathcal{G}' of $\widehat{C}_{0\text{-check}}$ connected to \mathcal{G} , and the internal wires – but not the outputs – in the decoding sub-circuits of \mathcal{C}_{dec} that decode masking inputs of \mathcal{G}' .)

We now describe a distribution \mathcal{F} over $\mathcal{SHALLOW}(4, O(\widehat{n}^4(1, S(n)) \cdot S(n)))$, where given the wire values of \mathcal{G} (either the real wire values \mathcal{W}_R or the reconstructed wires \mathcal{W}_S), $f \in \mathcal{F}$ generates the entire hybrid distribution (either \mathcal{H}_{m-1} or \mathcal{H}_m , respectively), as follows. For every gadget \mathcal{G}' of $\widehat{C}_{0\text{-check}}$ connected to \mathcal{G} , \mathcal{F} chooses a reconstructor $\text{rec}_{\mathcal{G}'} \leftarrow \text{REC}$ (see Lemma 46). The function f has all the hard-wired values of \mathcal{H}_m hard-wired into it. On input \mathcal{W} it evaluates the $\text{rec}_{\mathcal{G}'}$'s on the masking inputs of \mathcal{G} (as reported in \mathcal{W}) and the hard-wired values, to generate the internal wires (including the masking inputs) of \mathcal{G}' , and then uses the functions defined in Lemma 46 to reconstruct the internal wires of the decoding sub-circuits of \mathcal{C}_{dec} that decode the masking inputs of \mathcal{G}' . Then $f \in \mathcal{SHALLOW}(4, O(\widehat{n}^4(1, S(n)) \cdot S(n)))$ because by Lemma 46 the reconstructors for the \mathcal{G}' 's (the decoding sub-circuits, respectively) are computable in $\mathcal{SHALLOW}(2, O(\widehat{n}^2(1, S(n))))$ ($\mathcal{SHALLOW}(2, O(\widehat{n}(1, S(n))))$, respectively); the reconstructors of all the (at most) $O(\widehat{n}^2(1, S(n)) \cdot S(n))$ gadgets \mathcal{G}' can be computed in parallel; and the reconstructors for all of the (at most) $O(\widehat{n}^3(1, S(n)) \cdot S(n))$ decoding sub-circuits (each \mathcal{G}' is connected to $O(\widehat{n}(1, S(n)))$ decoding sub-circuits) can be computed in parallel.

Since $\mathcal{W}_R, \mathcal{W}_S$ are $(\mathcal{LEAK}', \epsilon(n) \cdot \widehat{n}(1, S(n)))$ -leakage resilient for any class \mathcal{LEAK}' of leakage functions such that $\mathcal{LEAK}' \circ \mathcal{SHALLOW}(3, O(\widehat{n}(1, \sigma))) \subseteq \mathcal{LEAK}_E$ (this follows from Lemma 41 because E^{in} is $(\mathcal{LEAK}_E, \epsilon(n))$ -leakage resilient), Lemma 45 guarantees that $\mathcal{H}_{m-1}, \mathcal{H}_m$ are $(\mathcal{LEAK}'', \epsilon(n) \cdot \widehat{n}(1, S(n)))$ -leakage resilient for every family \mathcal{LEAK}'' of leakage functions such that $\mathcal{LEAK}'' \circ \mathcal{SHALLOW}(4, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}'$, i.e., for any \mathcal{LEAK}'' such that $\mathcal{LEAK}'' \circ \mathcal{SHALLOW}(7, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$.

Bounding $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^y), \ell(\mathcal{H}_{\text{in}}^y))$: We show that $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^y), \ell(\mathcal{H}_{\text{in}}^y)) \leq \epsilon(n) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$ such that $\mathcal{LEAK}' \circ \mathcal{SHALLOW}(6, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$. The proof is by a hybrid argument in which we replace the input bundles of \widehat{C}_2 , and the bundles at the output of gadget of \widehat{C}_2 , one at a time. More specifically, we define hybrids $\mathcal{H}_0, \dots, \mathcal{H}_M$, where \mathcal{H}_i is generated from $\mathcal{H}_{\text{in}}^y$, by: (1) replacing the first i bundles with random encodings of random values (except for the bundle corresponding to the output of \widehat{C}_2 , which is set to a random encoding of 1); (2) recomputing the internal wires of the first i gadgets of \widehat{C}_2 using the gadget reconstructors; (3) recomputing the internal wires of $\widehat{C}_{0\text{-check}}$ that are influenced by this re-computation of the first i gadgets of \widehat{C}_2 ; and (4) (using the functions of Lemma 46) recomputing the internal wires of \mathcal{C}_{dec} that were influenced by re-computing $\widehat{C}_{0\text{-check}}$. Then $\mathcal{H}_0 = \mathcal{H}_{\text{in}}^y$ and $\mathcal{H}_M = \mathcal{H}_{\text{ext}}^y$, and we show that $\text{SD}(\ell(\mathcal{H}_m), \ell(\mathcal{H}_{m-1})) \leq \epsilon(n)$ for all $m \in [M]$ and $\ell \in \mathcal{LEAK}'$. We denote by \mathcal{G}_o (\mathcal{G}_i) the gadget

whose output (input) is the m' th bundle, and fix all wires in $\mathcal{H}_m, \mathcal{H}_{m-1}$ except for: the m' th bundle; the masking inputs, outputs, and internal wires of \mathcal{G}_o ; the masking inputs, and internal wires, of \mathcal{G}_i , as well as its input wire corresponding to the m' th bundle; the internal wires of all gadgets \mathcal{G}' of $\widehat{\mathcal{C}}_{0\text{-check}}$ connected to \mathcal{G}_o or \mathcal{G}_i ; and the internal wires of the decoding sub-circuits of \mathcal{C}_{dec} whose inputs are masking inputs of one of these \mathcal{G}' 's.

Let $\mathcal{W}_R \leftarrow \text{Enc}^{\text{in}}(b, 1^{S(n)})$ (where b is the value encoded by the m' th bundle in H_{in}^y), and $\mathcal{W}_S \leftarrow \text{Enc}^{\text{in}}(r, 1^{S(n)})$ for a random r (except if m corresponds to the output bundle, in which case $\mathcal{W}_S \leftarrow \text{Enc}^{\text{in}}(1, 1^{S(n)})$). We define a distribution \mathcal{F} over $\mathcal{SHALLOW}(6, O(\widehat{n}^4(1, S(n)) \cdot S(n)))$ as follows. Sampling a function $f \leftarrow \mathcal{F}$ is done by sampling $\text{rec}_o, \text{rec}_i$ from the distribution over reconstructors for $\mathcal{G}_o, \mathcal{G}_i$, respectively (see Definition 40), and sampling reconstructors $\text{rec}_{\mathcal{G}'}$ for the gadgets of $\widehat{\mathcal{C}}_{0\text{-check}}$ connected to \mathcal{G}_o or \mathcal{G}_i . The function f has all the hard-wired values of \mathcal{H}_{m-1} hard-wired into it. On input $e \in \mathbb{F}^{\widehat{n}(1, S(n))}$, f : (1) evaluates rec_o on the (hard-wired) inputs of \mathcal{G}_o , and the output e (this reconstructs the masking inputs, and internal wires, of \mathcal{G}_o); (2) evaluates rec_i on e as one of the inputs, and the other (hard-wired) input and output of \mathcal{G}_i ; (3) for every gadget \mathcal{G}' of $\widehat{\mathcal{C}}_{0\text{-check}}$ connected to \mathcal{G}_o or \mathcal{G}_i , uses $\text{rec}_{\mathcal{G}'}$ to generate its internal wires; and finally (4) uses the functions of Lemma 46 to generate the internal wires of the decoding sub-circuits that decode the masking inputs used in the \mathcal{G}' 's. Then $f \in \mathcal{SHALLOW}(6, O(\widehat{n}^4(1, S(n)) \cdot S(n)))$ because the reconstructors $\text{rec}_o, \text{rec}_i \in \mathcal{SHALLOW}(2, O(\widehat{n}^2(1, S(n))))$ (by Lemma 41) and can be evaluated in parallel, the $O(\widehat{n}^2(1, S(n)) \cdot S(n))$ reconstructors $\text{rec}_{\mathcal{G}'}$ (by Lemma 46), and can be evaluated in parallel, and the $O(\widehat{n}^3(1, S(n)) \cdot S(n))$ reconstructors of the decoding sub-circuits are each computable in $\mathcal{SHALLOW}(2, O(\widehat{n}(1, S(n))))$ (by Lemma 46) and can be evaluated in parallel.

Consequently, if $\mathcal{W}'_R := f(\mathcal{W}_R), \mathcal{W}'_S := f(\mathcal{W}_S)$ for $f \leftarrow \mathcal{F}$, then $\mathcal{W}'_R \equiv \mathcal{H}_{m-1}, \mathcal{W}'_S \equiv \mathcal{H}_m$ and by Lemma 45, $\mathcal{W}'_R, \mathcal{W}'_S$ are $(\mathcal{LEAK}', \epsilon(n))$ -leakage resilient for any \mathcal{LEAK}' such that $\mathcal{LEAK}' \circ \mathcal{SHALLOW}(6, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$.

Bounding SD $(\ell(\mathcal{H}_{\text{ext}}^y), \ell(\mathcal{H}_{\text{ext},2}^{y,x}))$: it holds that $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^y), \ell(\mathcal{H}_{\text{ext},2}^{y,x})) = 0$ for every ℓ because the hybrids are independent of the input for $\widehat{\mathcal{C}}_2$ (since the input is re-sampled as a fresh $z \in_R \mathbb{F}^n$ in both).

Bounding SD $(\ell(\mathcal{H}_{\text{ext},2}^{y,x}), \ell(\mathcal{H}_{\text{in},2}^{y,x}))$: We show that $\text{SD}(\ell(\mathcal{H}_{\text{ext},2}^{y,x}), \ell(\mathcal{H}_{\text{in},2}^{y,x})) \leq \epsilon(n) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$ such that $\mathcal{LEAK}' \circ \mathcal{SHALLOW}(6, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$. The proof is similar to the proof that $\text{SD}(\ell(\mathcal{H}_{\text{ext}}^y), \ell(\mathcal{H}_{\text{in}}^y)) \leq \epsilon(n) \cdot S(n)$, because the argument was independent of the actual inputs used in $\widehat{\mathcal{C}}_1, \widehat{\mathcal{C}}_2$ (as long as both hybrids use the same input in each copy).

Bounding SD $(\ell(\mathcal{H}_{\text{in},2}^{y,x}), \ell(\mathcal{H}^{y,x}))$: We show that $\text{SD}(\ell(\mathcal{H}_{\text{in},2}^{y,x}), \ell(\mathcal{H}^{y,x})) \leq \epsilon(n) \cdot \widehat{n}(1, S(n)) \cdot S(n)$ for all $\ell \in \mathcal{LEAK}'$ such that $\mathcal{LEAK}' \circ \mathcal{SHALLOW}(7, O(\widehat{n}^4(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$. The proof is similar to the proof that $\text{SD}(\ell(\mathcal{H}^y), \ell(\mathcal{H}_{\text{in}}^y)) \leq \epsilon(n) \cdot \widehat{n}(1, S(n)) \cdot S(n)$, because the argument was independent of the actual inputs used in $\widehat{\mathcal{C}}_1, \widehat{\mathcal{C}}_2$.

Bounding SD $(\ell(\mathcal{H}^x), \ell(\mathcal{H}^y))$ for $\ell \in \mathcal{LEAK}'$. From this analysis, we can now conclude using the union bound that for every $\ell \in \mathcal{LEAK}'$ such that

$$\mathcal{LEAK}' \circ \mathcal{SHALLOW}(7, O(\widehat{n}^5(1, S(n)) \cdot S(n))) \subseteq \mathcal{LEAK}_E$$

it holds that

$$\text{SD}(\ell(\mathcal{H}^x), \ell(\mathcal{H}^y)) \leq 4\epsilon(n) \cdot S(n) \cdot (\widehat{n}(1, S(n)) + 1).$$

4.3 A SAT-Respecting LRCC Against “Useful” Leakage

Ishai et al. [IWY16] use Construction 34 to devise a SAT-respecting LRCC which they later employ in their WI-PCP construction (described in Section 4.4). This is done in two steps. First, since the LRCC will be used to compile verification circuits of NP relations, we need the compiler to be *boolean*. Second, we need to instantiate the internal encoding scheme E such that it would resist leakage computable by functions that: (1) apply the PCP-prover algorithm, and then (2) restrict the output to a small subset of bits. Indeed, this is exactly the “leakage” on the witness which a query-bounded verifier (even a malicious one) obtains by querying the proof. We now provide more details on each of these steps.

Step (1): A SAT-Respecting LRCC for Boolean Circuits. The high-level idea is to transform the boolean circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$ into a functionally-equivalent arithmetic circuit C_3 over \mathbb{F}_3 (i.e., the field with 3 elements), use Construction 34 over the field \mathbb{F}_3 to compile C_3 into its leakage-resilient version \widehat{C}_3 , and then output the boolean circuit \widehat{C} that emulates \widehat{C}_3 using boolean operations. This is an over-simplified description of the compiler, where the actual construction needs to address several subtleties. We now describe each of these steps, and the subtleties that arise, in more detail.

From Boolean to Arithmetic Operations. The circuit C_3 is obtained from C by representing each boolean operation using an appropriate polynomial over \mathbb{F}_3 in the natural way. While C_3 is guaranteed to be functionally equivalent to C on binary strings, two issues arise concerning the SAT-respecting property. First, satisfiability over \mathbb{F}_2 means that $C(x) = 1$ for some x , whereas satisfiability over \mathbb{F}_3 means $C_3(y) = 0$ for some y . In particular, we want the leakage-resilient circuit to output 1 only if there exists an x such that $C(x)$ outputs 1, whereas the SAT-respecting property guarantees only that if $\widehat{C}_3(\widehat{y}) = 0$ for some \widehat{y} , then $C_3(y) = 0$ for some y . Therefore, we need to “translate” a 1-output of C into a 0-output of C_3 , and a 0-output of C into a non-0-output of C_3 . Thus, we will have that $\widehat{C}_3(\widehat{y}) = 0$ for some \widehat{y} only if $C_3(y) = 0$ for some y . This brings us to the second issue: while we would like to use the function-equivalence of C and C_3 to claim that if $C_3(y) = 0$ for some $y \in \mathbb{F}_3^n$ then $C(x) = 1$ for some $x \in \mathbb{F}_2^n$, this is not necessarily the case, because C_3 's inputs are from \mathbb{F}_3^n and might *not* correspond to an input in \mathbb{F}_2^n . To overcome this issue, we add to C_3 an “input checker” component which checks that each of its n inputs is a bit. We denote this “enhanced” version of C_3 – which flips the output and checks validity of the inputs – by C'_3 .

From Arithmetic Back To Boolean. Once we generate the leakage-resilient version \widehat{C}'_3 of C'_3 , we need to represent it using a *boolean* circuit. (Indeed, the original circuit C was boolean, and its leakage-resilient version should also be a boolean circuit.) We do so by replacing each field element with a binary string representing it, and implementing each gate over \mathbb{F}_3 by a boolean sub-circuit. We also flip the output of \widehat{C}'_3 , so the resultant circuit would again be functionally-equivalent to C . There are two important points that we need to handle. First, while we can represent field elements using any (injective) encoding scheme, to preserve the SAT-respecting property it must also be onto. Otherwise, the boolean circuit could potentially be satisfied using invalid encodings, namely ones that do not encode any field element, and thus the computation in the boolean circuit would not correspond to a computation in \widehat{C}'_3 . In particular, the boolean sub-circuits implementing gates over \mathbb{F}_3 must be defined for all possible encodings – even ones that would not be used in an honest evaluation of the circuit.

Second, these boolean sub-circuits should have small depth and size. The reason is the reduction from the leakage resilience of the final circuit \widehat{C} to the leakage resilience of \widehat{C}'_3 . More

specifically, the reduction proceeds by assuming that a leakage function ℓ in some class \mathcal{LEAK} can distinguish between the wire values $[\widehat{C}, \widehat{x}]$ of \widehat{C} on an encoding \widehat{x} of some input x , and its wire values $[\widehat{C}, \widehat{x}']$ on the encoding \widehat{x}' of some other input x' such that $C(x) = C(x')$. It then uses this to break the leakage resilience of \widehat{C}'_3 for some leakage function ℓ_3 in the leakage class \mathcal{LEAK}_3 against which the arithmetic LRCC is secure. This is done as follows: given the wire values of \widehat{C}'_3 on some input (these values are elements of \mathbb{F}_3), the reduction first replaces the field elements with the corresponding encodings. The resultant values constitute only *part* of the wire values of \widehat{C} . Specifically, these are the values of the wires *between the sub-circuits emulating the gates over \mathbb{F}_3* , whereas \widehat{C} contains also the internal wires of these sub-circuits, namely wires which do not appear in \widehat{C}'_3 . Thus, the leakage function ℓ_3 must first generate these missing wires, and only then can it evaluate ℓ . In particular, $\mathcal{LEAK} \subset \mathcal{LEAK}_3$, where the difference between the two classes depends on the complexity of the boolean sub-circuits implementing gates over \mathbb{F}_3 . Fortunately, these sub-circuits are both shallow and small.

In summary, Ishai et al. show [IWY15, Proposition 3.31] the following boolean LRCC.

Claim 47 (Boolean SAT-respecting LRCC). *Let $\mathcal{LEAK}, \mathcal{LEAK}_E$ be families of functions, $S(n) : \mathbb{N} \rightarrow \mathbb{N}$ be a size function, and $\epsilon(n) : \mathbb{N} \rightarrow \mathbb{R}^+$. Let $E^{\text{in}} = (\text{Enc}^{\text{in}}, \text{Dec}^{\text{in}})$ be a linear, onto, $(\mathcal{LEAK}_E, \epsilon(n))$ -leakage-resilient encoding scheme with parameters n, σ and $\widehat{n} = \widehat{n}(n, \sigma)$, such that $\mathcal{LEAK}_E = \mathcal{LEAK} \circ \text{BOOL}(33, O(\widehat{n}^6(1, S(n)) \cdot S^2(n)))$. Then there exist constants $c, c' > 0$ for which there exists a SAT-respecting, $(\mathcal{LEAK}, c' \cdot \epsilon(n) \cdot (\widehat{n}(1, c \cdot S(n)) + 1) \cdot c \cdot S(n), S(n))$ -LRCC over $\{0, 1\}$. Moreover, for every $C : \{0, 1\}^n \rightarrow \{0, 1\}$, the compiled circuit \widehat{C} has size $|\widehat{C}| = O(\widehat{n}^6(1, c \cdot S(n)) \cdot |C|^2)$.*

Step (2): Leakage Resilience Against “Useful” Leakage. The second component of the construction is an internal encoding scheme E_{in} – resisting leakage from a “useful” class of leakage functions – with which we instantiate Claim 47. More specifically, the leakage class consists of \mathcal{AC}^0 circuits (namely, constant-depth, polynomial-sized boolean circuits over unbounded fan-in and fan-out \wedge, \vee, \neg gates), augmented with a sublinear number of \oplus gates of unbounded fan-in and fan-out. Formally,

Notation 48 ($\mathcal{L}_{n,d,s,\oplus}^m$ leakage family). *Let $n, d, s \in \mathbb{N}$ be length, depth and size parameters (respectively), and let $t \in \mathbb{N}$ be a parity gate bound. The family $\mathcal{L}_{n,d,s,\oplus}$ consists of all functions computable by a boolean circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$ of size at most s and depth d , with unbounded fan-in and fan-out $\wedge, \vee, \neg, \oplus$ gates, out of which at most t are \oplus gates. We denote $\mathcal{L}_{d,s,\oplus} = \cup_{n \in \mathbb{N}} \mathcal{L}_{n,d,s,\oplus}$.*

For a length parameter $m \in \mathbb{N}$, and a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, let $f_i(x_1, \dots, x_n), i \in [m]$ denote the i 'th output bit of f . We denote: $\mathcal{L}_{n,d,s,\oplus}^m = \{f : \{0, 1\}^n \rightarrow \{0, 1\}^m : \forall 1 \leq i \leq m, f_i \in \mathcal{L}_{n,d,s,\oplus}\}$, and $\mathcal{L}_{d,s,\oplus}^m := \cup_{n \in \mathbb{N}} (\mathcal{L}_{n,d,s,\oplus}^m)$.

The encodings scheme we use encodes elements $\gamma \in \mathbb{F}_3$ as binary strings whose sum mod 3 is γ . Formally:

Notation 49. *For $\gamma \in \{0, 1, 2\}$ and $n \in \mathbb{N}$, \mathcal{U}_γ^n denotes the uniform distribution over $\{v \in \{0, 1\}^{3n} : \#_1(v) \equiv \gamma \pmod{3}\}$, where $\#_1(v)$ denotes the number of 1's in v .*

Definition 50. *The encodings scheme $E_3 = (\text{Enc}_3, \text{Dec}_3)$ is defined as follows. For every $\gamma \in \mathbb{F}_3$, $\text{Enc}_3(\gamma, 1^n)$ samples from \mathcal{U}_γ^n , and $\text{Dec}_3(v)$ returns $(\#_1(v) \pmod{3})$. We note that E_3 is linear and onto.*

Remark 51. *Enc_3 can be computed efficiently by repeating the following procedure n^2 times. Pick $v \in \{0, 1\}^{3n}$ uniformly at random, compute $t := \#_1(v)$, and if $t = \gamma$ then return v . If all iterations fail, return a fixed $v_\gamma \in \{0, 1\}^{3n}$ such that $\#_1(v_\gamma) = \gamma$. Then the output of Enc_3 is thus statistically close to \mathcal{U}_γ^n .*

Ishai et al. [IWY16] show that the encoding scheme of Definition 50 resists leakage from \mathcal{AC}^0 circuits augmented with few \oplus gates:

Corollary 52 (Corollary 3.44 in [IWY15]). *For every constant depth parameter $d \in \mathbb{N}$ there exist constants $c, \epsilon \in (0, 1)$, such that for every constant $l \in \mathbb{N}$ there exists a minimal length parameter $n_0 \in \mathbb{N}$ such that for every $n \geq n_0$ the encoding scheme $\text{Enc}_3(\cdot, 1^n)$ of Definition 50 is $(\mathcal{L}_{3n, d, n^l, \oplus n^\epsilon}, 2^{-n^c})$ -leakage resilient.*

Instantiating Claim 47 with the encoding scheme of Definition 50 as the internal encoding scheme, and using Corollary 52, [IWY16] show the existence of a boolean LRCC resisting leakage from \mathcal{AC}^0 circuits with few \oplus gates:

Theorem 53 (Boolean SAT-respecting LRCC for \mathcal{AC}^0 circuits with \oplus gates, Theorem 3.37 in [IWY15]). *Let $n \in \mathbb{N}$ be an input length parameter. For every positive constants d, c , polynomials $m = m(n)$, $t = t(n)$, and polynomial size bound $s = s(n)$, there exists a polynomial $l(n)$, such that the following holds. There exists a SAT-respecting $(\mathcal{L}_{l, d, l^c, \oplus t}^m, 2^{-n^c}, s(n))$ -LRCC over $\{0, 1\}$, which on input a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$ of size $|C| \leq s(n)$ outputs a circuit \widehat{C} of size $|\widehat{C}| \leq l(n)$.*

4.4 The Witness-Indistinguishable PCP

In this section, we describe the WI-PCPs for NP of [IWY16], which rely on the boolean SAT-respecting LRCC of Theorem 53. They use also a PCP system $(\mathcal{P}', \mathcal{V}')$ for the language 3SAT of all satisfiable 3CNF formulas, in which the prover algorithm can be implemented in a low complexity class.

The high-level idea of the construction for an NP-relation \mathcal{R} with verification circuit C is that given input x , instead of verifying that $C_x(\cdot) := C(x, \cdot)$ is satisfiable (which holds if and only if $x \in \mathcal{L}$ for the corresponding NP-language \mathcal{L}), the verifier will check that the *leakage-resilient version* \widehat{C}_x is satisfiable. For this, the prover and verifier will represent \widehat{C}_x as a 3CNF formula φ in the natural way,²³ then use the PCP system $(\mathcal{P}', \mathcal{V}')$ to verify that $\varphi \in 3\text{SAT}$. The construction is described in Figure 3.

The following theorem (which is a combination of [IWY15, Proposition 4.4] and [IWY15, Corollary 4.10]) asserts the connection between the properties of the LRCC and the resultant PCP system $(\mathcal{P}, \mathcal{V})$.

Theorem 55. *Let $n \in \mathbb{N}$ be a length parameter, $q^* = q^*(n)$, $S = S(n)$ be query and size functions, $\epsilon, \epsilon' \in [0, 1]$, and \mathcal{LEAK} be a family of leakage functions. Assume that Construction 54 is instantiated with:*

- *A boolean SAT-respecting $(\mathcal{LEAK}, \epsilon, S)$ -LRCC such that there exists a polynomial $g(\cdot)$ for which $|\text{Comp}(C)| \leq g(|C|)$ for every circuit C , and*
- *A PCP system $(\mathcal{P}', \mathcal{V}')$ for 3SAT with proofs of length $\text{len}(n)$, such that for every $(\varphi, \mathcal{W}) \in 3\text{SAT}$, every subset \mathcal{Q} of q^* bits of an honestly-generated proof $\pi = \pi(\varphi, \mathcal{W})$ is computable from \mathcal{W} by a function $f_{\varphi, \mathcal{Q}} \in \mathcal{LEAK}$.*

Then for every NP-relation \mathcal{R} with verification circuit C of size $|C| \leq S$, the PCP system $(\mathcal{P}, \mathcal{V})$ is a (q^, ϵ^*) -WI-PCP for \mathcal{R} , where $\epsilon^* = O\left(\epsilon \cdot q^* \cdot \text{len}^{2q^*}(t)\right) + e^{-\Omega(q^* \cdot \text{len}^{q^*}(t))}$ and $t = O(g(|C|))$. Moreover, if \mathcal{V}' is non-adaptive, the so is \mathcal{V} .*

²³That is, φ will have a variable for each wire of \widehat{C}_x . It will contain, for each gate g of \widehat{C}_x , a sub-formula verifying that the output wire of g is consistent with the input wires and the operation of g , and it will also check that the output wire of \widehat{C}_x is 1. A satisfying assignment for φ is the wire values of \widehat{C}_x when evaluated on (an encoding of) a witness w for x .

Construction 54 (A Witness-Indistinguishable PCP). Let $\mathcal{R} = \mathcal{R}(x, w)$ be an NP-relation with verification circuit C .^a

Building blocks:

- A PCP system $(\mathcal{P}', \mathcal{V}')$ for 3SAT.
- A boolean LRCC $(\text{Comp}, E = (\text{Enc}, \text{Dec}))$.

Prover algorithm. \mathcal{P} , on input $(x, w) \in \mathcal{R}$, operates as follows:

- Computes $\widehat{C}_x(\cdot) = \text{Comp}(C_x)$, where C_x denotes the circuit C with x hard-wired into it.
- Samples a random encoding $\widehat{w} \leftarrow \text{Enc}(w, 1^{|C_x|})$, and evaluates \widehat{C}_x on \widehat{w} to generate the wire values \mathcal{W} of \widehat{C}_x .
- Constructs the 3CNF φ_x representing \widehat{C}_x .
- Computes a PCP $\pi = \mathcal{P}'(\varphi_x, \mathcal{W})$ for the claim “ $\varphi_x \in 3\text{SAT}$ ”, and outputs π .

Verifier algorithm. \mathcal{V} is given input x and oracle access to π . It computes $\widehat{C}_x(\cdot) = \text{Comp}(C_x)$, and constructs the 3CNF formula φ_x . Then, \mathcal{V} emulates \mathcal{V}' with input φ_x and oracle access to π , and outputs whatever \mathcal{V}' outputs.

^aWe note that \mathcal{R} is actually associated with a family $\{C_n\}$ of circuits, where C_n is applied to inputs $x \in \{0, 1\}^n$. Somewhat abusing notation, we refer to all these circuits simply as “ C ”.

Figure 3: Witness-Indistinguishable PCPs from SAT-Respecting LRCCs [IWY16]

Furthermore, the system is (q^*, ϵ) -WI against non-adaptive (possibly malicious) verifiers. Moreover, proofs generated by \mathcal{P} have length $\text{len}(t)$, and if \mathcal{V} has query complexity $q(n)$ and tosses $r(n)$ coins, then \mathcal{V} has query complexity $q(t)$, tosses $r(t)$ coins.

Proof. We first analyze the parameters of the system. The wire assignment \mathcal{W} to \widehat{C}_x has size $|\mathcal{W}| = |\widehat{C}_x| \leq g(|C_x|) \leq g(|C|)$, and $|\varphi_x| = O(|\widehat{C}_x|) \leq O(g(|C|))$. Therefore, the internal PCP system $(\mathcal{P}', \mathcal{V}')$ is emulated using inputs of size $t = g(|C|)$.

Completeness follows directly from the completeness of the building blocks.

Soundness follows from a combination of the soundness of $(\mathcal{P}', \mathcal{V}')$ and the SAT-respecting property of (Comp, E) . Indeed, if $x \notin \mathcal{L}$, then C_x is not satisfiable, and so (by the SAT-respecting property) \widehat{C}_x is not satisfiable, i.e., $\varphi_x \notin \text{3SAT}$. Therefore, by the soundness of $(\mathcal{P}', \mathcal{V}')$ we have that $\Pr[\mathcal{V}^{\pi^*}(x) = 1] = \Pr[\mathcal{V}'^{\pi^*}(\varphi_x) = 1]$ is negligible.

Witness-indistinguishability. Let $x \in \mathcal{L}$, φ_x be the 3CNF formula representing \widehat{C}_x , and w_1, w_2 be two witnesses for x . We first show witness indistinguishability against *non-adaptive* verifiers. Let \mathcal{V}^* be a non-adaptive q^* -query-bounded verifier, and let $\pi_i \leftarrow \mathcal{P}(x, w_i)$ for $i = 1, 2$. Since \mathcal{V}^* 's entire view can be generated from the oracle answers to its queries (and this cannot increase the statistical distance), it suffices to show that $\text{SD}(\pi_1|_{\mathcal{Q}}, \pi_2|_{\mathcal{Q}}) \leq \epsilon$ for every set \mathcal{Q} of queries of \mathcal{V}^* such that size $|\mathcal{Q}| \leq q^*$, where $\pi_i|_{\mathcal{Q}}$ denotes the restriction of π_i to the entries in \mathcal{Q} . Since $|C| = |C_x| \leq S$, the leakage resilience of the LRCC guarantees that $\text{SD}(\ell[\widehat{C}_x, \widehat{w}_1], \ell[\widehat{C}_x, \widehat{w}_2]) \leq \epsilon$ for every $\ell \in \mathcal{LEAK}$. We conclude the proof by noting that $f_{\varphi_x, \mathcal{Q}} \in \mathcal{LEAK}$, and $\pi_i|_{\mathcal{Q}} = f_{\varphi_x, \mathcal{Q}}[\widehat{C}_x, \widehat{w}_i]$. We have shown that $(\mathcal{P}, \mathcal{V})$ is (q^*, ϵ) -witness indistinguishable against *non-adaptive* verifiers \mathcal{V}^* . Using Theorem 56 below, this implies that $(\mathcal{P}, \mathcal{V})$ is (q^*, ϵ^*) -WI (even against *adaptive* verifiers), for $\epsilon^* = O(\epsilon \cdot q^* \cdot \text{len}^{2q^*}(t)) + e^{-\Omega(q^* \cdot \text{len}^{q^*}(t))}$. ■

The proof of Theorem 55 used the following theorem, which is implicit in [CDD⁺01] (see also [IWY15, Theorem 4.11]).

Theorem 56 (Implicit in [CDD⁺01]). *Let $(\mathcal{P}, \mathcal{V})$ be a PCP system that is (q^*, ϵ) -WI against non-adaptive verifiers, with proofs of length len . Then $(\mathcal{P}, \mathcal{V})$ is (q^*, ϵ^*) -WI against adaptive verifiers, where $\epsilon^* = O(\epsilon \cdot q^* \cdot \text{len}^{2q^*}) + e^{-\Omega(q^* \cdot \text{len}^{q^*})}$.*

Theorem 23 now follows as a corollary of Theorem 55, using the SAT-respecting LRCC of Theorem 53, and a PCP system of [AS92], whose prover algorithm can be implemented in a low complexity class (the analysis of the prover complexity is due to [IWY15, Appendix B]):

Theorem 57 (PCPs for NP, [AS92]). *3SAT has a PCP system $(\mathcal{P}, \mathcal{V})$ with soundness error $1/2$ with an honest verifier that queries $O(\log^2 n)$ proof bits. The proofs have length $\text{poly}(n)$, where every proof bit can be generated by an \mathcal{AC}^0 circuit with a single \oplus gate of unbounded fan-in.*

We are now ready to prove Theorem 23.

Proof (of Theorem 23): We instantiate Construction 54 with the PCP system $(\mathcal{P}', \mathcal{V}')$ of Theorem 57 and the LRCC of Theorem 53. By Theorem 57, there exist constants $d, c \in \mathbb{N}$ such that every bit in a proof generated by \mathcal{P}' is computable from the NP witness in $\mathcal{L}_{d, n^c, \oplus 1}$, where n is the witness length, and the proofs have length $n^{c'}$, for some constant c' .

Let \mathcal{R} be an NP-relation with verification circuit C , then $|C| = n^{c'}$ for some constant c' . We instantiate Theorem 53 with parameters $d^* = d, s^* = |C|, n^* = n, t^* = 1, m^* = q^*,$ ²⁴ and $c^* \geq c$

²⁴We use the superscript $*$ to denote the parameters of Theorem 53, and note that s^*, t^*, m^* are in $\text{poly}(n)$.

which is a sufficiently large constant whose value is set below, to obtain an LRCC (Comp, E). We compute $\widehat{C} = \text{Comp}(C)$, where $|\widehat{C}| \leq l(n)$ which, because $|C| \leq s^*$, is $(\mathcal{L}_{d,l^c(n),\oplus 1}^{q^*}, 2^{-n^{c^*}})$ -LR (where $l(n)$ is the polynomial whose existence is guaranteed by Theorem 53).

Let φ denote the 3CNF representing \widehat{C} . Then by Theorem 57 (and using the fact that $|\mathcal{W}| = |\widehat{C}| \leq l(n)$), every bit of a proof generated by \mathcal{P}' for φ can be generated from a wire assignment \mathcal{W} in $\mathcal{L}_{d,l^c(n),\oplus 1}^{q^*}$. Therefore, Theorem 55 guarantees that the system $(\mathcal{P}, \mathcal{V})$ of Construction 54 is a non-adaptive WI-PCP system for \mathcal{R} , with $\left(q^*, O\left(2^{-n^{c^*}} \cdot q^* \cdot \text{len}^{2q^*}(l(n)) \right) + e^{-\Omega(q^* \cdot \text{len}^{q^*}(l(n)))} \right)$ -WI (where $\text{len}(l(n))$ denotes the proof length), and soundness error $1/2$ with an honest verifier that queries $O(\log^2(l(n))) = \text{polylog}(n) \leq \text{polylog}(q^*)$ proof bits. We set c^* to be sufficiently large, such that the statistical LR error satisfies²⁵

$$O\left(2^{-n^{c^*}} \cdot q^* \cdot \text{len}^{2q^*}(l(n)) \right) + e^{-\Omega(q^* \cdot \text{len}^{q^*}(l(n)))} = \text{negl}(q^*) \leq \text{negl}(\kappa).$$

We conclude the proof by noting that soundness can be amplified to $\text{negl}(\kappa)$ with only a $\text{poly}(\kappa)$ blowup in the query complexity of the honest verifier. ■

5 Discussion

The works of [IWY16, HVW21a] show a connection between ZK-PCPs and the seemingly-unrelated field of leakage-resilient cryptography, and use it to circumvent an inherent limitation of previous constructions – that the *honest* verifier is adaptive. Specifically, using tools from the leakage-resilience literature, [IWY16, HVW21a] put forth two new paradigms of constructing ZK-PCPs, yielding PCPs with ZK against malicious verifiers, in which the honest verifier is *non-adaptive*. In the context of cryptographic applications of ZK-PCPs, non-adaptive verification translates into fewer communication rounds. The paradigm of [HVW21a] also extends to ZK-PCPs of *proximity*.

Despite this recent progress, several interesting questions remain open. The obvious open problem is to obtain ZK-PCPs and ZK-PCPPs with an exponential query gap as in [KPT97, IW14] but which can be verified non-adaptively. One possible approach is to design a ZK-PCP variant over a large alphabet with negligible soundness error and an honest verifier that makes fewer queries than [IKOS07] (hopefully, polylogarithmic). Another interesting research direction is to extend the techniques of [IWY16, HVW21a] to other related proof systems, such as interactive oracle proofs. Finding further applications of ZK-PCPs and ZK-PCPPs is also an interesting question. Finally, though in this survey we have focused on other parameters of ZK-PCPs, reducing the proof length is a fascinating open problem worthy of study. Whereas the locking-scheme-based ZK-PCPs of [KPT97, IW14] inherently incur a polynomial blowup in proof length, another advantage of the leakage-resilience-based approach is that it opens up the possibility of reducing the proof length of ZK-PCPs, potentially even matching the proof length of non-ZK PCPs.

Acknowledgements

I am extremely grateful to Yuval Ishai for introducing me to the world of ZK-PCPs. I am also grateful to all my collaborators on these works - Yuval Ishai, Carmit Hazay, Guang Yang, and Muthu

²⁵Such a constant exists since we assume that $q^* = \text{poly}(n)$.

Venkatasubramaniam – for many enjoyable discussions, and the fruitful collaborations whose results are surveyed here.

The author is supported by the BIU Center for Research in Applied Cryptography and Cyber Security in conjunction with the Israel National Cyber Bureau in the Prime Minister’s Office.

References

- [ALM⁺92] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and hardness of approximation problems. In *33rd Annual Symposium on Foundations of Computer Science, 1992*, pages 14–23. IEEE Computer Society, 1992.
- [ARW17] Amir Abboud, Aviad Rubinfeld, and R. Ryan Williams. Distributed PCP theorems for hardness of approximation in P. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 25–36. IEEE Computer Society, 2017.
- [AS92] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs; A new characterization of NP. In *33rd Annual Symposium on Foundations of Computer Science, 1992*, pages 2–13. IEEE Computer Society, 1992.
- [BCF⁺16] Eli Ben-Sasson, Alessandro Chiesa, Michael A. Forbes, Ariel Gabizon, Michael Riabzev, and Nicholas Spooner. On probabilistic checking in perfect zero knowledge. *Electron. Colloquium Comput. Complex.*, page 156, 2016.
- [BCG⁺13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: verifying program executions succinctly and in zero knowledge. In *Advances in Cryptology - CRYPTO, 2013. Proceedings, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 90–108. Springer, 2013.
- [BCGV16] Eli Ben-Sasson, Alessandro Chiesa, Ariel Gabizon, and Madars Virza. Quasi-linear size zero knowledge from linear-algebraic PCPs. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, 2016, Proceedings, Part II*, volume 9563 of *Lecture Notes in Computer Science*, pages 33–64. Springer, 2016.
- [BCL22] Jonathan Bootle, Alessandro Chiesa, and Siqi Liu. Zero-knowledge IOPs with linear-time prover and polylogarithmic-time verifier. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Proceedings, Part II*, volume 13276 of *Lecture Notes in Computer Science*, pages 275–304. Springer, 2022.
- [BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, 2016, Proceedings, Part II*, volume 9986 of *Lecture Notes in Computer Science*, pages 31–60, 2016.
- [BDG⁺18] Marshall Ball, Dana Dachman-Soled, Siyao Guo, Tal Malkin, and Li-Yang Tan. Non-malleable codes for small-depth circuits. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018*, pages 826–837. IEEE Computer Society, 2018.
- [BDKM16] Marshall Ball, Dana Dachman-Soled, Mukul Kulkarni, and Tal Malkin. Non-malleable codes for bounded depth, bounded fan-in circuits. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications*

of *Cryptographic Techniques, 2016, Proceedings, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 881–908. Springer, 2016.

- [BGH⁺04] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil P. Vadhan. Robust PCPs of proximity, shorter PCPs and applications to coding. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, 2004*, pages 1–10. ACM, 2004.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *Advances in Cryptology - CRYPTO 2001, Proceedings*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2001.
- [BIS19] Andrej Bogdanov, Yuval Ishai, and Akshayaram Srinivasan. Unconditionally secure computation against low-complexity leakage. In *Advances in Cryptology - CRYPTO, 2019, Proceedings, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 387–416. Springer, 2019.
- [BS08] Eli Ben-Sasson and Madhu Sudan. Short PCPs with polylog query complexity. *SIAM J. Comput.*, 38(2):551–607, 2008.
- [CCG⁺07] Hao Chen, Ronald Cramer, Shafi Goldwasser, Robbert de Haan, and Vinod Vaikuntanathan. Secure computation from random error correcting codes. In Moni Naor, editor, *Advances in Cryptology - EUROCRYPT 2007, 26th Annual International Conference on the Theory and Applications of Cryptographic Techniques, 2007, Proceedings*, volume 4515 of *Lecture Notes in Computer Science*, pages 291–310. Springer, 2007.
- [CDD⁺01] Ran Canetti, Ivan Damgård, Stefan Dziembowski, Yuval Ishai, and Tal Malkin. On adaptive vs. non-adaptive security of multiparty protocols. In *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, 2001, Proceeding*, volume 2045 of *Lecture Notes in Computer Science*, pages 262–279. Springer, 2001.
- [CDMW08] Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. Black-box construction of a non-malleable encryption scheme from any semantically secure one. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008*, volume 4948 of *Lecture Notes in Computer Science*, pages 427–444. Springer, 2008.
- [CDMW18] Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. A black-box construction of non-malleable encryption from semantically secure encryption. *J. Cryptol.*, 31(1):172–201, 2018.
- [DF12] Stefan Dziembowski and Sebastian Faust. Leakage-resilient circuits without computational assumptions. In *Theory of Cryptography - 9th Theory of Cryptography Conference, TCC 2012. Proceedings*, volume 7194 of *Lecture Notes in Computer Science*, pages 230–247. Springer, 2012.
- [DFK⁺92] Cynthia Dwork, Uriel Feige, Joe Kilian, Moni Naor, and Shmuel Safra. Low communication 2-prover zero-knowledge proofs for NP. In *Advances in Cryptology - CRYPTO, 1992, Proceedings*, volume 740 of *Lecture Notes in Computer Science*, pages 215–227. Springer, 1992.

- [DGR99] Scott E. Decatur, Oded Goldreich, and Dana Ron. Computational sample complexity. *SIAM J. Comput.*, 29(3):854–879, 1999.
- [Din06] Irit Dinur. The PCP theorem by gap amplification. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, 2006*, pages 241–250. ACM, 2006.
- [DR04] Irit Dinur and Omer Reingold. Assignment testers: Towards a combinatorial proof of the PCP-theorem. In *45th Symposium on Foundations of Computer Science FOCS, 2004, Proceedings*, pages 155–164. IEEE Computer Society, 2004.
- [FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st Annual Symposium on Foundations of Computer Science, 1990, Volume I*, pages 308–317. IEEE Computer Society, 1990.
- [FRR⁺10] Sebastian Faust, Tal Rabin, Leonid Reyzin, Eran Tromer, and Vinod Vaikuntanathan. Protecting circuits from leakage: the computationally-bounded and noisy cases. In *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*, pages 135–156. Springer, 2010.
- [GM10] Berndt M. Gammel and Stefan Mangard. On the duality of probing and fault attacks. *J. Electron. Test.*, 26(4):483–493, 2010.
- [GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, 1985*, pages 291–304. ACM, 1985.
- [HVW21a] Carmit Hazay, Muthuramakrishnan Venkatasubramanian, and Mor Weiss. ZK-PCPs from leakage-resilient secret sharing. In *2nd Conference on Information-Theoretic Cryptography, ITC 2021*, volume 199 of *LIPICs*, pages 6:1–6:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [HVW21b] Carmit Hazay, Muthuramakrishnan Venkatasubramanian, and Mor Weiss. ZK-PCPs from leakage-resilient secret sharing (full version). *IACR Cryptol. ePrint Arch.*, 2021. <https://eprint.iacr.org/2021/606>.
- [IKOS07] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC), 2007*, pages 21–30. ACM, 2007.
- [IMS12] Yuval Ishai, Mohammad Mahmoody, and Amit Sahai. On efficient zero-knowledge PCPs. In *Theory of Cryptography - 9th Theory of Cryptography Conference, TCC 2012. Proceedings*, volume 7194 of *Lecture Notes in Computer Science*, pages 151–168. Springer, 2012.
- [ISVW13] Yuval Ishai, Amit Sahai, Michael Viderman, and Mor Weiss. Zero knowledge LTCs and their applications. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 16th International Workshop, APPROX 2013, and 17th International Workshop, RANDOM 2013. Proceedings*, volume 8096 of *Lecture Notes in Computer Science*, pages 607–622. Springer, 2013.

- [ISW03] Yuval Ishai, Amit Sahai, and David A. Wagner. Private circuits: Securing hardware against probing attacks. In *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 463–481. Springer, 2003.
- [IW14] Yuval Ishai and Mor Weiss. Probabilistically checkable proofs of proximity with zero-knowledge. In *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC. Proceedings*, volume 8349 of *Lecture Notes in Computer Science*, pages 121–145. Springer, 2014.
- [IWY15] Yuval Ishai, Mor Weiss, and Guang Yang. Making the best of a leaky situation: Zero-knowledge PCPs from leakage-resilient circuits (full version). *IACR Cryptol. ePrint Arch.*, 2015. <http://eprint.iacr.org/2015/1055>.
- [IWY16] Yuval Ishai, Mor Weiss, and Guang Yang. Making the best of a leaky situation: Zero-knowledge PCPs from leakage-resilient circuits. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Proceedings, Part II*, volume 9563 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2016.
- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In S. Rao Kosaraju, Mike Fellows, Avi Wigderson, and John A. Ellis, editors, *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 723–732. ACM, 1992.
- [KPT97] Joe Kilian, Erez Petrank, and Gábor Tardos. Probabilistically checkable proofs with zero knowledge. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, 1997*, pages 496–505. ACM, 1997.
- [LS11] Shachar Lovett and Srikanth Srinivasan. Correlation bounds for poly-size AC^0 circuits with $n^{1-o(1)}$ symmetric gates. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011. Proceedings*, volume 6845 of *Lecture Notes in Computer Science*, pages 640–651. Springer, 2011.
- [Mic94] Silvio Micali. CS proofs (extended abstract). In *35th Annual Symposium on Foundations of Computer Science, 1994*, pages 436–453. IEEE Computer Society, 1994.
- [Mie09] Thilo Mie. Short PCPPs verifiable in polylogarithmic time with $O(1)$ queries. *Ann. Math. Artif. Intell.*, 56(3-4):313–338, 2009.
- [MR04] Silvio Micali and Leonid Reyzin. Physically observable cryptography (extended abstract). In *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Proceedings*, volume 2951 of *Lecture Notes in Computer Science*, pages 278–296. Springer, 2004.
- [MV13] Eric Miles and Emanuele Viola. Shielding circuits with groups. In *Symposium on Theory of Computing Conference, STOC, 2013*, pages 251–260. ACM, 2013.
- [NY90] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, 1990*, pages 427–437. ACM, 1990.

- [RRR16] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016*, pages 49–62. ACM, 2016.
- [Sha90] Adi Shamir. IP=PSPACE. In *31st Annual Symposium on Foundations of Computer Science, 1990, Volume I*, pages 11–15. IEEE Computer Society, 1990.
- [Tha22] Justin Thaler. Proofs, arguments, and zero-knowledge. *Manuscript*, 2022. <https://people.cs.georgetown.edu/jthaler/ProofsArgsAndZK.pdf>.
- [Wei16] Mor Weiss. Secure computation and probabilistic checking. *PhD Thesis*, 2016.