

A Zero-Knowledge PCP Theorem

Tom Gur *

Jack O'Connor[†]

Nicholas Spooner[‡]

Abstract

We show that for every polynomial q^* there exist polynomial-size, constant-query, non-adaptive PCPs for NP which are perfect zero knowledge against (adaptive) adversaries making at most q^* queries to the proof. In addition, we construct exponential-size constant-query PCPs for NEXP with perfect zero knowledge against any polynomial-time adversary. This improves upon both a recent construction of perfect zero-knowledge PCPs for #P (STOC 2024) and the seminal work of Kilian, Petrank and Tardos (STOC 1997).

*University of Cambridge. Email: tom.gur@cl.cam.ac.uk. Supported by ERC Starting Grant 101163189 and UKRI Future Leaders Fellowship MR/X023583/1.

[†]University of Cambridge. Email: jack.oconnor@cl.cam.ac.uk.

[‡]Cornell University. Email: nspooner@cornell.edu.

Contents

1	Introduction	3
1.1	Techniques	4
1.2	Open problems	7
1.3	Related work	7
2	Preliminaries	9
2.1	Coding theory	9
2.2	PCPs	9
2.3	Zero-knowledge PCPs	11
2.4	Low-degree testing	12
3	Zero-knowledge proof composition	12
3.1	Locally computable proofs	12
3.2	Alphabet reduction of PZK-PCPs	14
3.3	Proof composition	15
4	Robust PCPPs for polynomial summation	17
5	Robust PZK-PCPPs for polynomial summation	20
6	Robust PZK-PCPs for NP and NEXP	25
7	PZK-PCP for NP and NEXP with constant query complexity	30
	References	31

1 Introduction

The PCP theorem [AS98; ALMSS92] states that for any language in NP, there exists a polynomial-size proof that can be checked probabilistically by reading only a constant number of bits from the proof; or, succinctly,

$$\text{NP} \subseteq \text{PCP}[\log n, 1],$$

where $\text{PCP}[r, q]$ is the class of all languages that admit a PCP verifier that uses $O(r)$ random bits and reads $O(q)$ bits of the proof (note that logarithmic randomness complexity implies polynomial proof length).

While PCPs originated from the study of zero-knowledge proofs [GMR89], there seems to be an intrinsic tension between the two notions: PCPs achieve locality by encoding NP witnesses in a manner that *spreads* global information throughout the proof, whereas zero-knowledge proofs aim to *hide* all information except for the validity of the statement. Moreover, PCPs are fundamentally non-interactive objects, whereas interaction is often crucial for zero knowledge.

Indeed, one must take care in even *defining* zero knowledge PCPs (ZK-PCPs), as it is impossible to achieve non-trivial zero knowledge against a malicious verifier that reads the entire proof. In their seminal work on ZK-PCPs, Kilian, Petrank and Tardos [KPT97] identify two regimes of interest:

- (a) *Polynomial-size* PCPs that are zero-knowledge against a verifier that makes at most a *fixed* polynomial number of queries q^* to the PCP. In this regime, they construct ZK-PCPs for NP with polylogarithmic query complexity. We refer to q^* as the “query bound”.
- (b) *Exponential-size* PCPs that are zero-knowledge against *any polynomial-time* verifier. In this regime, they construct ZK-PCPs for NEXP with polynomial query complexity.

However, these constructions fall short of a “zero-knowledge PCP theorem” in a fundamental way: the query complexity is polylogarithmic, as opposed to $O(1)$, which is a characteristic property of the PCP theorem.

A major obstacle to achieving $O(1)$ query complexity via the [KPT97] approach is that the technique used to obtain zero knowledge leads to an inherently *adaptive* honest verifier (i.e., which makes multiple rounds of queries to the proof). Known query-reduction methods apply only to non-adaptive PCPs. Aside from its theoretical interest, non-adaptivity is crucial for some applications [IWY16]. Finally, these constructions only achieve statistical zero knowledge (SZK-PCP) and not perfect zero knowledge (PZK-PCP).

Building on techniques developed in [BCFGRS17; CFGS18; CCGOS23], a recent work [GOS24] constructed exponential-size PZK-PCPs with polynomially many non-adaptive queries for $\#P$. Our first result shows that such PZK-PCPs exist for NEXP, with constant query complexity.

Theorem 1. *There exist PCPs for NEXP of exponential length with a non-adaptive verifier that reads $O(1)$ bits of the proof, which are perfect zero-knowledge against any efficient adversary.*

Our second result “scales down” the above to obtain polynomial-size non-adaptive PZK-PCPs for NP with constant query complexity.¹

Theorem 2 (“Zero-knowledge PCP theorem”). *For any $q^* \leq 2^{\text{poly}(n)}$, there exist PCPs for NP of length $\text{poly}(q^*, n)$ with a non-adaptive verifier that reads $O(1)$ bits of the proof, which are perfect zero-knowledge against any adversary reading at most q^* bits of the proof; i.e.,*

$$\text{NP} \subseteq \text{PZK-PCP}[\log n, 1].$$

¹For our definition of the class PZK-PCP, see Definition 2.12.

1.1 Techniques

PZK-PCPs for nondeterministic computation. This paper builds on the prior work of [GOS24], which proves a weaker version of Theorem 1 that only captures (decision) $\#P$, by constructing non-adaptive PZK-PCPs for the sumcheck problem.

Readers familiar with the PCP literature may wonder why this construction does not lead immediately to a ZK-PCP for NEXP; indeed, the first construction of a PCP for NEXP is essentially a reduction to sumcheck [BFLS91]. However, as we discuss below, even with a PZK-PCP for the sumcheck problem, the BFLS construction is not zero knowledge.

We start by briefly reviewing the BFLS construction, which is a PCP for the NEXP-complete problem Oracle-3SAT, defined as follows.

Definition 1 (Oracle 3-SAT). *Let $B: \{0, 1\}^{r+3s+3} \rightarrow \{0, 1\}$ be a 3-CNF. We say that B is implicitly satisfiable if there exists $A \in \{0, 1\}^s \rightarrow \{0, 1\}$ such that for all $z \in \{0, 1\}^r, b_1, b_2, b_3 \in \{0, 1\}^s$, $B(z, b_1, b_2, b_3, A(b_1), A(b_2), A(b_3)) = 1$. Let Oracle-3SAT be the language of implicitly satisfiable 3-CNFs.*

The BFLS PCP consists of two parts. First, the *multilinear extension* of the witness A ; i.e., a multilinear polynomial \hat{A} over some finite field \mathbb{F} such that $\hat{A}(x) = A(x)$ for all $x \in \{0, 1\}^s$. Second, a sumcheck PCP² for the following claim³:

$$\sum_{\substack{z \in \{0, 1\}^r \\ b_1, b_2, b_3 \in \{0, 1\}^s}} \hat{B}(z, b_1, b_2, b_3, \hat{A}(b_1), \hat{A}(b_2), \hat{A}(b_3)) = 2^{r+3s}, \quad (1)$$

where \hat{B} is an arithmetisation of the circuit B .

A natural first step is to use the zero-knowledge sumcheck PCP of [GOS24] in place of the standard sumcheck PCP, which hides hard-to-compute information about partial sums. However, this *does not* yet yield a ZK-PCP, because the first part of the construction contains an encoding of the witness A , which violates the zero-knowledge condition.

To hide A , we would like to use the *sumcheck commitment scheme*, introduced by [CFGS22] (see also [CFS17]) for their construction of an *interactive* PCP (IPCP) for NEXP. A sumcheck commitment to the polynomial $\hat{A}(\vec{X})$ is a random polynomial $C(\vec{X}, \vec{Y})$ of individual degree $d' \geq 2$ in each Y -variable such that

$$\sum_{c \in \{0, 1\}^k} C(\alpha, c) = \hat{A}(\alpha)$$

for all $\alpha \in \mathbb{F}^s$. It was shown in [CFGS22] that this commitment perfectly hides \hat{A} against all adversaries that make fewer than 2^k queries to C .

The construction of [BFLS91] uses the sumcheck protocol to reduce checking (1) to three random queries to \hat{A} . Building on this idea, the ZK-IPCP construction of [CFGS22] uses the *interactive* sumcheck protocol to open the sumcheck commitment to \hat{A} at those points. Finally, to ensure that those three evaluations do not leak information about the witness, \hat{A} is chosen to be a random multiquartic—rather than the unique multilinear—extension of A .

However, the strategy above strongly relies on the *interactivity* of the ZK-IPCP. Indeed, observe that if we were to “unroll” this interaction into a PCP, we would simply write down all of \hat{A} ! We must therefore

²Strictly speaking, a PCP of *proximity*; we will ignore this distinction for this overview.

³To mitigate some technical issues with soundness, the actual construction uses a slightly different summand polynomial.

establish (1) *without* opening the sumcheck commitment. That is, we would like to directly check:

$$\sum_{\substack{z \in \{0,1\}^r \\ b_1, b_2, b_3 \in \{0,1\}^s}} \hat{B} \left(z, b_1, b_2, b_3, \sum_{c \in \{0,1\}^k} C(b_1, c), \sum_{c \in \{0,1\}^k} C(b_2, c), \sum_{c \in \{0,1\}^k} C(b_3, c) \right) = 2^{r+3s},$$

To do this, we first “pull out” the three inner summations. There are various ways this can be achieved; we follow a linearisation approach. First observe that (assuming \hat{A} takes boolean values on $\{0, 1\}^s$), the LHS of (1) is equal to

$$\sum_{a_1, a_2, a_3 \in \{0,1\}} \sum_{\substack{z \in \{0,1\}^r \\ b_1, b_2, b_3 \in \{0,1\}^s}} \hat{B}(z, b_1, b_2, b_3, a_1, a_2, a_3) \cdot \prod_{i=1}^3 (\hat{A}(b_i) - (1 - a_i)),$$

since the product expression “zeroes out” any term of the sum for which some $a_i \neq \hat{A}(b_i)$. Next, observe that for any $b_1, b_2, b_3, a_1, a_2, a_3$, provided that \mathbb{F} has characteristic different from 2,

$$\prod_{i=1}^3 (\hat{A}(b_i) - (1 - a_i)) = \prod_{i=1}^3 \sum_{c \in \{0,1\}^k} \left(C(b_i, c) - \frac{1 - a_i}{2^k} \right) = \sum_{c_1, c_2, c_3 \in \{0,1\}^k} \prod_{i=1}^3 \left(C(b_i, c_i) - \frac{1 - a_i}{2^k} \right).$$

Taken together, we see that checking (1) is equivalent to checking

$$\sum_{c_1, c_2, c_3 \in \{0,1\}^k} \sum_{a_1, a_2, a_3 \in \{0,1\}} \sum_{\substack{z \in \{0,1\}^r \\ b_1, b_2, b_3 \in \{0,1\}^s}} \hat{B}(z, b_1, b_2, b_3, a_1, a_2, a_3) \prod_{i=1}^3 \left(C(b_i, c_i) - \frac{1 - a_i}{2^k} \right) = 2^{r+3s},$$

which can be proven in zero knowledge using the [GOS24] ZK-PCPP for sumcheck. In particular, the PCPP simulator requires only polynomially many evaluations of the summand to simulate polynomially many queries to the proof. By setting $k = \omega(\log n)$ we can simulate those evaluations by lazily simulating C as a uniformly random polynomial (via an algorithm of [BCFGRS17]), and then evaluating the summand directly. The query complexity of the verifier is $\text{poly}(n)$.

Our result for NP is obtained in a similar way, scaling the parameters appropriately. In particular, for a given adversary query bound q^* , we can set $k = O(\log q^*)$. The query complexity of the verifier is then $\text{poly}(\log n, \log q^*)$. We also point out that the honest prover in our construction for NP is efficient, given a valid witness as input.

Proof composition and zero knowledge. In order to obtain constant query ZK-PCPs for NP, we would like to apply the proof composition paradigm [BGHSV06] to our ZK-PCPs. This involves composing a robust outer PCP with an inner PCP of proximity to obtain a PCP which inherits the randomness complexity of the former and query complexity of the latter. To do this, we first need to strengthen our ZK-PCPs to satisfy *robust* soundness (i.e., the local view of the verifier must be far from an accepting view, with high probability). Then we show that proof composition *preserves* the zero knowledge of the outer PCP. To the best of our knowledge, this is the first composition theorem for ZK-PCPs.

Our first step is to obtain robust ZK-PCPs for NP with polylogarithmic query complexity. While the PZK-PCP of [GOS24] (upon which our PZK-PCP builds) is an algebraic construction, it does *not* have constant robust soundness as written. We present a modification of the [GOS24] PCP for sumcheck, following the “query bundling” approach of [BGHSV06], that has constant robust soundness.

The key challenge is to show that this modification preserves zero knowledge. In fact, we will show that a much more general class of “local” transformations preserve zero knowledge. This class includes not only query bundling, but also the subsequent steps of alphabet reduction and proof composition. To capture this class formally, we define a new notion, *locally computable proofs*.

Definition 2 (Locally computable algorithms (informal; see Definition 3.1)). *Let \mathcal{A} and \mathcal{A}_0 be randomized algorithms. We say that \mathcal{A} is ℓ -locally computable from \mathcal{A}_0 if there exists an efficient, deterministic oracle algorithm f making at most ℓ queries to its oracle such that, for every input x , the following two distributions are identically distributed:*

$$\mathcal{A}(x), \quad (f^{\pi_0} \mid \pi_0 \leftarrow \mathcal{A}_0(x)).$$

We show that if the PCP prover algorithm \mathcal{P} is ℓ -locally computable from a zero-knowledge PCP prover \mathcal{P}_0 , then provided ℓ is asymptotically smaller than the query bound on the ZK-PCP, \mathcal{P} inherits the zero knowledge guarantee of \mathcal{P}_0 . Intuitively, this is true because if a proof π is locally computable from a proof π_0 , and π_0 is zero knowledge, then we can apply f to the simulator for π_0 to obtain a simulator for π . This notion is surprisingly versatile, and allows us to prove zero knowledge in an array of distinct settings:

- **Robustification of [GOS24].** We show that our modified [GOS24] construction is locally computable from the original construction, and thus inherits zero knowledge.
- **Alphabet reduction.** Recall that alphabet reduction allows us to transform a robust PCP over a large alphabet into a boolean PCP, while maintaining robustness. It is performed by encoding each symbol of the PCP with a good error correcting code. More formally, if π is a distribution over robust PCP proofs over the alphabet $\{0, 1\}^a$ for some $a \in \mathbb{N}$, and $\text{ECC}: \{0, 1\}^a \rightarrow \{0, 1\}^b$ is a systematic binary error-correcting code of constant relative distance and rate, then we can obtain a robust boolean PCP by defining a new proof $\tau(\alpha) := \text{ECC}(\pi(\alpha))$ for every proof index α , and writing τ over the alphabet $\{0, 1\}$. Then τ is 1-locally computable from π by the following function:

$$f^\pi(\alpha, i) = \text{ECC}(\pi(\alpha))_i,$$

where $i \in [b]$. We note that prior work on alphabet reduction for ZKPCPs [HVW22] required a much more complex construction because their ZKPCP achieves only a quadratic gap between the honest verifier’s query complexity and the query bound.

- **Proof composition.** Recall that composition of an outer PCP system $(\mathcal{P}_{\text{out}}, \mathcal{V}_{\text{out}})$ for \mathcal{L} with an inner PCP of proximity $(\mathcal{P}_{\text{in}}, \mathcal{V}_{\text{in}})$ for circuit evaluation proceeds as follows. Let $\pi_{\text{out}} \leftarrow \mathcal{P}_{\text{out}}$. For every choice $r \in \{0, 1\}^{r_{\text{out}}}$ of \mathcal{V}_{out} ’s randomness, the composed prover computes \mathcal{V}_{out} ’s query set $Q(r)$ and an “inner proof” $\pi_r := \mathcal{P}_{\text{in}}(\mathcal{V}_{\text{out}}, \pi_{\text{out}}|_{Q(r)})$, which attests to the fact that if \mathcal{V}_{out} performed its verification of π_{out} using randomness r , then it would have accepted. Each inner proof is a function of at most ℓ_{out} many locations of π_{out} , where ℓ_{out} is the query complexity of \mathcal{V}_{out} . The composed proof is given by $(\pi_{\text{out}}, (\pi_r)_{r \in \{0, 1\}^{r_{\text{out}}}})$.

Hence the composed proof is ℓ_{out} -locally computable from π_{out} , by the following function:

$$f^{\pi_{\text{out}}}(\mathcal{O}, i) = \begin{cases} \pi_{\text{out}}(i) & \text{if } \mathcal{O} = \pi_{\text{out}} \\ \mathcal{P}_{\text{in}}(\mathcal{V}_{\text{out}}, \pi_{\text{out}}|_{Q(r)})_i & \text{if } \mathcal{O} = \pi_r, \text{ for some } r \in \{0, 1\}^{r_{\text{out}}} \end{cases}.$$

Note that the composed PCP is locally computable from the outer PCP, so only the outer PCP needs to be zero knowledge to ensure the composed PCP is zero knowledge. Therefore we can employ the existing (non-ZK) PCP of proximity for circuit evaluation of [BGHSV06] as the inner PCPP.

1.2 Open problems

This work shows that for any polynomial q^* , any language in NP has a polynomial-sized proof that can be probabilistically checked by probing only $O(1)$ bits, but where any set of q^* bits carries no information about the witness. This can be viewed as a zero-knowledge PCP theorem that matches parameters of the original PCP theorem [AS98; ALMSS92]. Since then, stronger versions of the PCP theorem have been shown. It is tempting to ask whether zero-knowledge PCPs can match the strongest constructions of standard PCPs.

In particular, one of the most immediate open questions is whether it is possible to obtain ZK-PCPs with nearly-linear length. Optimising the proof length of PCPs received much attention for decades after the first proof of the PCP theorem, where the current state of the art achieves quasilinear proof length [BS08; Din07]. We ask whether the same can be obtained for zero-knowledge PCPs.

Open Problem 1 (Short ZK-PCPs). *Do there exist $O(1)$ -query ZK-PCPs for NP with proof length $\tilde{O}(n)$?*

We remark that our algebraic zero-knowledge techniques are based on Reed-Muller arithmetisation and the sumcheck protocol, whereas (non-ZK) constructions of quasilinear length PCPs are based on Reed-Solomon arithmetisation and combinatorial gap amplification, hence new ideas are necessary for such strengthening of our theorem.

Even more ambitiously, one could ask whether it is possible to transform any construction of a (non-ZK) PCP to a ZK-PCP while preserving its parameters.

Open Problem 2 (PCP to ZK-PCP transformation). *Is there a black-box transformation that imbues a PCP construction with zero knowledge?*

We note that prior to [GOS24], all works on ZK-PCPs (see below) followed this approach, but those transformations either introduce adaptivity or only achieve weak ZK guarantees, and none preserve query complexity. Similar transformations are known, e.g., for multi-prover interactive proofs [BGKW88] and their quantum analogues [GSY19; MS24] whereas in other models, such as in zero-knowledge streaming interactive proofs [CDGH23], we have a zero knowledge sumcheck protocol but no generic transformation is known.

We remark that our techniques make whitebox use of the structure of the [BFLS91] PCP, and rely strongly on the [GOS24] ZK-PCP for sumcheck; hence, obtaining a generic transformation would require new ideas.

The quantum PCP conjecture. Finally, we highlight a connection between zero-knowledge PCPs and one of the most important open problems in quantum complexity theory: the quantum PCP (QPCP) conjecture.

Most classical constructions of PCPs rely on an encoding of the NP witness via a locally-testable and (relaxed) locally-decodable code [BGHSV06; GRR20]. Even though there is growing evidence that good quantum LTCs may exist [AE15; LLZ22; ABN23; DLV24], quantum codes *cannot* be locally decodable due to the no-cloning theorem. This is one of the main barriers towards applying algebraic and coding-theoretic techniques to QPCPs.

Quantum codes are fundamentally tied to zero knowledge. It is a well-known fact that the erasure of a subset of qubits of a codeword is correctable if and only if the reduced density matrix on that subset is independent of the encoded state. In other words, roughly speaking, a quantum code has good distance if and only if it satisfies a quantum analogue of the PZK property for PCPs. Thus PZK-PCPs are perhaps the closest classical analogue of QPCPs, and studying them may help to shed light on the QPCP conjecture.

1.3 Related work

The first zero-knowledge PCPs appeared in the work of Kilian, Petrank and Tardos [KPT97]. Later works [IMS12; ISVW13; IW14; IMSX15] simplified this construction, and extended it to PCPs of proximity and

the closely related notion of zero-knowledge locally testable codes (LTCs). These constructions rely on an adaptive honest verifier, and hence it is unclear how to use proof composition to improve their parameters. [IMS12; IMSX15] showed that PCPs which are zero knowledge against any efficient adversary and where the proof oracle is described by a polynomial-sized circuit exist only for languages in SZK.

Another line of work, motivated by cryptographic applications, focuses on obtaining SZK-PCPs for NP with a *non-adaptive* honest verifier from leakage resilience. These results come with caveats, achieving either a weaker notion of zero knowledge known as *witness indistinguishability* [IWY16], or simulation against adversaries making only quadratically many more queries than the honest verifier [HVW22]. See [Wei22] for a survey of this line of work.

In related models that allow for interaction, zero-knowledge proofs are easier to construct. We know that $\text{PZK-MIP} = \text{MIP} (= \text{NEXP})$ [BGKW88], where MIP is the class of languages with a multi-prover interactive proofs. The quantum analogue of this result, $\text{PZK-MIP}^* = \text{MIP}^* (= \text{RE})$, is also known to hold [CFG22; GSY19; MS24]. The constructions in this work draw inspiration from a similar result for interactive PCPs (IPCPs) [KR08], an interactive generalisation of PCPs (and special case of IOPs): $\text{PZK-IPCP} = \text{IPCP} = \text{NEXP}$ [BCFGRS17; CFS17].

2 Preliminaries

Throughout \mathbb{F} is a finite field. CIRCUIT VALUE denotes the P-complete language $\{(C, x) \in \{0, 1\}^* : C \text{ is an encoding of a boolean circuit such that } C(x) = 1\}$. For an oracle π , we use $\text{Dom}(\pi)$ to denote the domain of π . For $t \in \mathbb{N}$, and oracles π_1, \dots, π_t , we denote their concatenation by (π_1, \dots, π_t) , and use the notation (π_i, q) to denote the index $q \in \text{Dom}(\pi_i)$. For a pair language \mathcal{L} , for every string $x \in \{0, 1\}^n$, we denote the set $\mathcal{L}[x] := \{y : (x, y) \in \mathcal{L}\}$. For a relation \mathcal{R} , we denote the language corresponding to \mathcal{R} by $\mathcal{L}(\mathcal{R}) := \{x : (x, w) \in \mathcal{R}\}$.

Algorithms. We write $\mathcal{A}^\pi(x)$ to denote the output of \mathcal{A} when given input x (explicitly) and oracle access to π . We use the abbreviation PPT to denote probabilistic polynomial-time. If an oracle algorithm makes at most $q^*(n)$ oracle queries we say that it is q^* -query-bounded. We generally omit the internal randomness of an algorithm from probability statements; that is, we write $\Pr[\mathcal{A}(x) = 0]$ to mean $\Pr_{r \leftarrow \{0, 1\}^n}[\mathcal{A}(x; r) = 0]$.

Polynomials. For $d \in \mathbb{N}$, we write $\mathbb{F}^{\leq d}[X_1, \dots, X_m]$ for the ring of polynomials in m variables over \mathbb{F} of total degree at most d . For $\vec{d} := (d_1, \dots, d_m) \in \mathbb{N}^m$, we will write $\mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ for the ring of polynomials in m variables over \mathbb{F} of individual degree d_i in the i -th variable. For a product set $S = S_1 \times \dots \times S_m \subseteq \mathbb{F}^m$, $\vec{a} \in S$, and $\vec{d} = (|S_1| - 1, \dots, |S_m| - 1)$, we denote by $L_{S, \vec{a}}$ the unique element of $\mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ such that for all $\vec{b} \in S$,

$$L_{S, \vec{a}}(\vec{b}) = \begin{cases} 1 & \text{if } \vec{b} = \vec{a}, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Vector reversal and truncation. For a vector $\vec{\alpha} = (\alpha_1, \dots, \alpha_n) \in \mathbb{F}^n$, we denote by $\vec{\alpha}_{\text{rev}}$ the reversed vector $(\alpha_n, \alpha_{n-1}, \dots, \alpha_1)$. For any $i \in [n]$, we denote the truncation of $\vec{\alpha}$ to its first i entries by $\vec{\alpha}_i := (\alpha_1, \dots, \alpha_i)$.

2.1 Coding theory

Let Σ be an alphabet and let $\ell \in \mathbb{N}$. A code C is a subset $C \subseteq \Sigma^\ell$. Given two strings $x, y \in \Sigma^n$, we denote the relative Hamming distance between x and y by $\Delta(x, y) := |\{i \in [n] : x_i \neq y_i\}|/n$. We say that a vector x is ε -far from a set $S \subseteq \Sigma^n$ if $\min_{y \in S} \Delta(x, y) \geq \varepsilon$.

Reed–Muller codes. The Reed–Muller (RM) code is the code consisting of evaluations of multivariate low-degree polynomials over a finite field. Given a finite field \mathbb{F} , and positive integers m and d we denote by $\text{RM}[\mathbb{F}, m, d]$ the linear code consisting of evaluations of m -variate polynomials over \mathbb{F} of total degree at most d .

2.2 PCPs

Definition 2.1 (PCP). A probabilistically checkable proof (PCP) for a relation \mathcal{R} consists of a prover \mathcal{P} and a PPT verifier \mathcal{V} such that the following holds.

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

$$\Pr_{\pi \leftarrow \mathcal{P}(x, w)}[\mathcal{V}^\pi(x) = 1] = 1.$$

2. **Soundness.** For every $x \notin \mathcal{L}(\mathcal{R})$ and every oracle π^* ,

$$\Pr[\mathcal{V}^{\pi^*}(x) = 1] \leq \frac{1}{2}.$$

We say that a PCP is **efficient** if $\mathcal{P}(x, w)$ can be computed efficiently for any $(x, w) \in \mathcal{R}$.

We formally define PCPs for relations (rather than languages) to allow us to discuss efficient PCP provers in the NP setting.

Definition 2.2 (PCPP). For $\delta: \mathbb{N} \rightarrow [0, 1]$, a **probabilistically checkable proof of proximity (PCPP)** for a pair language \mathcal{L} with proximity parameter δ consists of a prover \mathcal{P} and a verifier \mathcal{V} such that the following holds for every pair of strings (x, y) .

1. **Completeness.** If $(x, y) \in \mathcal{L}$,

$$\Pr_{\pi \leftarrow \mathcal{P}(x, y)} [\mathcal{V}^{(y, \pi)}(x) = 1] = 1.$$

2. **Soundness.** If y is δ -far from the set $\mathcal{L}[x] := \{y : (x, y) \in \mathcal{L}\}$, then for every oracle π^* ,

$$\Pr[\mathcal{V}^{(y, \pi^*)}(x) = 1] \leq \frac{1}{2}.$$

Definition 2.3 (Non-adaptive PCP verifiers). A **non-adaptive PCP verifier** is an algorithm of the form $\mathcal{V}^\pi(x; \mu) = D(x, \pi|_{Q(x; \mu)}; \mu)$, where $D: \{0, 1\}^n \times \Sigma^q \times \{0, 1\}^{r(n)} \rightarrow \{0, 1\}$ is the **decision algorithm**, and $Q: \{0, 1\}^n \times \{0, 1\}^{r(n)} \rightarrow \text{Dom}(\pi)^q$ is the **query algorithm**.

Definition 2.4 (Accepting view of a PCP verifier). For non-adaptive q -query PCP verifiers of the form $\mathcal{V}^\pi(x; \mu) = D(x, \pi|_{Q(x; \mu)}; \mu)$, where $D: \{0, 1\}^n \times \Sigma^q \times \{0, 1\}^{r(n)} \rightarrow \{0, 1\}$, and $Q: \{0, 1\}^n \times \{0, 1\}^{r(n)} \rightarrow \text{Dom}(\pi)^q$, we define, for each choice of randomness $\mu \in \{0, 1\}^{r(n)}$ the **set of accepting views of \mathcal{V}** to be:

$$\text{Acc}(\mathcal{V}(x; \mu)) := \{a \in \Sigma^q : D(x, a; \mu) = 1\}.$$

Definition 2.5 (Accepting view of a PCP of proximity verifier). For non-adaptive q -query PCP of proximity verifiers of the form $\mathcal{V}^{(y, \pi)}(x; \mu) = D(x, (y, \pi)|_{Q(x; \mu)}; \mu)$, where $D: \{0, 1\}^n \times \Sigma^q \times \{0, 1\}^{r(n)} \rightarrow \{0, 1\}$, and $Q: \{0, 1\}^n \times \{0, 1\}^{r(n)} \rightarrow \text{Dom}((y, \pi))^q$, we define, for each choice of randomness $\mu \in \{0, 1\}^{r(n)}$, the **set of accepting views of \mathcal{V}** to be:

$$\text{Acc}(\mathcal{V}(x; \mu)) := \{a \in \Sigma^q : D(x, a; \mu) = 1\}.$$

In both of the above definitions when it is clear from context we will omit μ and write $\text{Acc}(\mathcal{V}(x))$.

Definition 2.6 (Robust soundness). A non-adaptive PCP for a relation \mathcal{R} has **robustness** ρ if for every $x \notin \mathcal{L}(\mathcal{R})$, it holds that, for every oracle π^* ,

$$\Pr_{\mu} [\Delta(\pi^*|_{Q(x)}, \text{Acc}(\mathcal{V}(x; \mu))) \leq \rho] \leq \frac{1}{2}.$$

A non-adaptive PCP of proximity for a language \mathcal{L} has **robustness** ρ if for every (x, y) such that y is δ -far from $\mathcal{L}[x]$, it holds that, for every oracle π^* ,

$$\Pr_{\mu} [\Delta((y, \pi^*)|_{Q(x)}, \text{Acc}(\mathcal{V}(x; \mu))) \leq \rho] \leq \frac{1}{2}.$$

Definition 2.7. For $\rho: \mathbb{N} \rightarrow [0, 1]$, a PCP has **expected robustness** ρ if for every $x \notin \mathcal{L}(\mathcal{R})$, we have for every oracle π^* ,

$$\mathbb{E}_{\mu} [\Delta(\pi^*|_{Q(x)}, \text{Acc}(\mathcal{V}(x; \mu)))] \geq \rho(|x|).$$

A PCP of proximity has **expected robustness** ρ if for every (x, y) such that y is δ -far from $\mathcal{L}[x]$, it holds that, for every oracle π^* ,

$$\mathbb{E}_{\mu} [\Delta((y, \pi^*)|_{Q(x)}, \text{Acc}(\mathcal{V}(x; \mu)))] \geq \rho(|x|).$$

The following proposition relating robustness and expected robustness appears as Proposition 2.10 in [BGHSV06].

Proposition 2.8 ([BGHSV06]). *If a PCPP has expected robustness ρ , then for every $\varepsilon \leq \rho$, it has robust-soundness error $1 - \varepsilon$ with robustness parameter $\rho - \varepsilon$.*

2.3 Zero-knowledge PCPs

Definition 2.9 (View). For a PCP $(\mathcal{P}, \mathcal{V})$ and a (possibly malicious) verifier \mathcal{V}^* , $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x, w)$ denotes the view of \mathcal{V}^* with input x and oracle access to $\pi \leftarrow \mathcal{P}(x, w)$. That is, $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x, w)$ comprises \mathcal{V}^* 's random coins and all answers to \mathcal{V}^* 's queries to π .

Definition 2.10 (Perfect zero-knowledge PCP). We say that a PCP system $(\mathcal{P}, \mathcal{V})$ for a relation \mathcal{R} is **perfect zero knowledge** (a PZK-PCP) with query bound q^* if for every (possibly malicious) adaptive q^* -query-bounded verifier \mathcal{V}^* there exists a PPT algorithm $\text{Sim}_{\mathcal{V}^*}$ (the simulator), such that for every $(x, w) \in \mathcal{R}$, $\text{Sim}_{\mathcal{V}^*}(x)$ is distributed identically to $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x, w)$.

All of our constructions satisfy the stronger notion of *black-box* zero knowledge, defined next.

Definition 2.11 (Black-box perfect zero-knowledge PCP). We say that a PCP system $(\mathcal{P}, \mathcal{V})$ for a relation \mathcal{R} is **black-box perfect zero knowledge** with query bound q^* if there exists a PPT algorithm Sim (the simulator), such that for every (possibly malicious) adaptive q^* -query-bounded verifier \mathcal{V}^* , and for every $(x, w) \in \mathcal{R}$, $\text{Sim}^{\mathcal{V}^*}(x)$ is distributed identically to $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x, w)$.

We define the complexity class PZK-PCP and its robust version.

Definition 2.12 (PZK-PCP, rPZK-PCP). For functions $r, q, q^*: \mathbb{N} \rightarrow \mathbb{N}$, we define the complexity class $\text{PZK-PCP}_{q^*, \Sigma}[r(n), q(n)]$ to be the set of all languages that admit perfect zero-knowledge PCPs written over the alphabet Σ with randomness complexity $O(r(n))$, query complexity $O(q(n))$ and query bound $q^*(n)$. Further, we define the complexity class $\text{rPZK-PCP}_{q^*, \Sigma}[r(n), q(n)]$ to be the set of all languages that admit perfect zero-knowledge PCPs (with the same parameters) with robustness $\Omega(1)$.

We define the complexity class

$$\text{PZK-PCP}[r, q] := \bigcap_{q^*} \text{PZK-PCP}_{q^*, \{0,1\}}[r + \log q^*, q]$$

where the intersection ranges over all functions $q^*(n)$ that are bounded by $2^{p(n)}$ for some polynomial p .

That is, $\mathcal{L} \in \text{PZK-PCP}[r, q]$ if for any query bound q^* (that is at most exponential in n), \mathcal{L} has a perfect zero-knowledge PCP of length $\text{poly}(2^r, q^*)$ over the boolean alphabet with query complexity $O(q)$.

We also define zero-knowledge PCPs of proximity.

Definition 2.13 (View of PCP of proximity). For a PCP of proximity $(\mathcal{P}, \mathcal{V})$ and a (possibly malicious) verifier \mathcal{V}^* , $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x, y)$ denotes the view of \mathcal{V}^* with explicit input x and oracle access to both implicit input y and $\pi \leftarrow \mathcal{P}(x)$. That is, $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x, w)$ comprises \mathcal{V}^* 's random coins and all answers to \mathcal{V}^* 's queries to (y, π) .

Definition 2.14 (Black-box perfect zero-knowledge PCP of proximity). We say that a PCP of proximity system $(\mathcal{P}, \mathcal{V})$ for a pair language \mathcal{L} is **black-box perfect zero knowledge** with query bound q^* if there exists an algorithm Sim , such that for every (possibly malicious) adaptive q^* -query-bounded verifier \mathcal{V}^* , and for every $(x, y) \in \mathcal{L}$, $\text{Sim}^{(\mathcal{V}^*, y)}(x)$ is distributed identically to $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x, y)$, and moreover, Sim runs in time $\text{poly}(|x|, q^*)$.

2.4 Low-degree testing

Definition 2.15 (Vector-valued Reed–Muller code). The vector-valued Reed–Muller code consists of tuples of evaluations of multivariate low-degree polynomials over a finite field. Given a finite field \mathbb{F} , and positive integers m , d and k we denote by $\text{RM}^k[\mathbb{F}, m, d]$ the linear code over the alphabet \mathbb{F}^k consisting of k -tuples of evaluations of m -variate polynomials over \mathbb{F} of total degree at most d . More formally,

$$\text{RM}^k[\mathbb{F}, m, d] := \{(p_1(\vec{x}), \dots, p_k(\vec{x})) \in \mathbb{F}^k : \vec{x} \in \mathbb{F}^m, \forall i \in [k], p_i(\vec{X}) \in \mathbb{F}^{\leq d}[X_1, \dots, X_m]\}.$$

The following theorem appears as Proposition 5.7 in [Par21].

Theorem 2.16 (Robust vector-valued low-degree test [Par21]). For $\delta > 0$, provided $|\mathbb{F}| > 25k$ there exists a test that given oracle access to a function $F: \mathbb{F}^m \rightarrow \mathbb{F}^k$, makes $|\mathbb{F}|$ queries to F , runs in time $\text{poly}(|\mathbb{F}|, m, d, k)$, and:

- if $F \in \text{RM}^k[\mathbb{F}, m, d]$ then the test accepts with probability 1;
- if F is ε -far from $\text{RM}^k[\mathbb{F}, m, d]$ then the expected distance of the tester's view from any accepting view is at least $\Omega(\varepsilon)$.

3 Zero-knowledge proof composition

In this section, we will prove that PCP proof composition (as in [BGHSV06]) preserves zero knowledge. Towards this end, we introduce the notion of *locally computable proofs*, which are proofs in which each symbol can be deterministically and efficiently computed from a local view of another proof. For example, in proof composition, the each inner proof is locally computable from the outer (zero-knowledge) proof.

In Section 3.1, we formally define locally computable proofs and show that if a PCP is locally computable from a PZK-PCP, then the locally computable PCP also exhibits perfect zero knowledge. In Section 3.2, we show how to convert a PZK-PCP over an arbitrary alphabet into a boolean PZK-PCP. In Section 3.3, we show that proof composition preserves perfect zero knowledge. Specifically, if the outer PCP of the composition is perfect zero knowledge, then the composed PCP will also be perfect zero knowledge.

3.1 Locally computable proofs

We introduce the notion of ℓ -locally computable proofs. This is a new notion that we consider to be of independent interest and will be useful towards showing that query bundling, alphabet reduction and proof composition preserve zero knowledge.

Definition 3.1. Let \mathcal{A} and \mathcal{A}_0 be randomized algorithms, and let $\ell: \mathbb{N} \rightarrow \mathbb{N}$. We say that \mathcal{A} is ℓ -locally computable from \mathcal{A}_0 on $C \subseteq \{0, 1\}^*$ if there exists an efficient, deterministic oracle algorithm f making at most ℓ queries to its oracle such that, for every $x \in C$, the following two distributions are identical:

$$\mathcal{A}(x); \quad (f^{\pi_0} \mid \pi_0 \leftarrow \mathcal{A}_0(x)).$$

Lemma 3.2. Let $(\mathcal{P}_0, \mathcal{V}_0)$ be a PZK-PCP for a relation \mathcal{R} with query bound q^* , and let $(\mathcal{P}, \mathcal{V})$ be a PCP for \mathcal{L} such that \mathcal{P} is ℓ -locally computable from \mathcal{P}_0 on \mathcal{L} . Then $(\mathcal{P}, \mathcal{V})$ is perfect zero knowledge with query bound q^*/ℓ .

While we state and prove Lemma 3.2 for the case of PCPs, the analysis straightforwardly extends to the case of PCPs of proximity.

The intuition behind Lemma 3.2 is that if $\pi_0 \leftarrow \mathcal{P}_0$ is zero knowledge, then since any local view of π_0 is efficiently simulatable and $\pi \leftarrow \mathcal{P}$ is computable from a local view of π_0 , we can combine these two functionalities to obtain a simulator for \mathcal{P} . More formally, let \mathcal{V}^* be a (possibly adaptive) malicious verifier of $(\mathcal{P}, \mathcal{V})$. Our goal is to construct a simulator for the view of \mathcal{V}^* and our strategy for doing so consists of two parts. In Construction 3.3, we construct a hybrid simulator $\mathcal{A}_{\mathcal{V}^*}$ for \mathcal{V}^* , which is given oracle access to the (zero-knowledge) proof $\pi_0 \leftarrow \mathcal{P}_0$. This hybrid simulator uses the local computability of $\pi \leftarrow \mathcal{P}$ and its oracle access to π_0 to simulate \mathcal{V}^* 's view of π . Construction 3.4 is our simulator for \mathcal{V}^* . As $\mathcal{A}_{\mathcal{V}^*}$ is itself a malicious verifier of the proof π_0 , we can invoke the zero-knowledge property of π_0 to simulate the view of $\mathcal{A}_{\mathcal{V}^*}$. Our simulator for \mathcal{V}^* works by running $\mathcal{A}_{\mathcal{V}^*}$ and answering its queries to π_0 using the simulator for π_0 .

Construction 3.3. A hybrid simulator for any malicious verifier \mathcal{V}^* of the PCP $(\mathcal{P}, \mathcal{V})$, a PCP which is ℓ -locally computable from $(\mathcal{P}_0, \mathcal{V}_0)$ by a function f . Receives oracle access to $\pi_0 \leftarrow \mathcal{P}_0$.

$\mathcal{A}_{\mathcal{V}^*}^{\pi_0}(x, r)$:

1. Initialise an empty list T .
2. Run \mathcal{V}^* on random coins r . Every time \mathcal{V}^* makes a query α to π , compute $\beta := f^{\pi_0}(\alpha)$, add (α, β) to T and feed β to \mathcal{V}^* as the response from π .
3. Output (r, T) .

Construction 3.4. A perfect zero-knowledge simulator for a PCP system $(\mathcal{P}, \mathcal{V})$ which is ℓ -locally computable from a PZK-PCP $(\mathcal{P}_0, \mathcal{V}_0)$. Let $\overline{\text{Sim}}_{\mathcal{V}^*}$ denote the simulator for $(\mathcal{P}_0, \mathcal{V}_0)$ for a malicious verifier \mathcal{V}^* .

$\text{Sim}_{\mathcal{V}^*}(x)$:

1. Run $\overline{\text{Sim}}_{\mathcal{A}_{\mathcal{V}^*}}(x)$ (that is, simulate the view of $\mathcal{A}_{\mathcal{V}^*}$ (Construction 3.3)) to obtain a query-answer set T_0 for π_0 and random coins r_0 .
2. Run $\mathcal{A}_{\mathcal{V}^*}(x, r)$ (Construction 3.3) using T_0 to answer its queries to π_0 to obtain (r, T) .
3. Output (r, T) .

Proof of Lemma 3.2. We show that Construction 3.4 is a perfect zero-knowledge simulator for $(\mathcal{P}, \mathcal{V})$. We first analyse Construction 3.3, and show that its output is identically distributed to $\text{View}_{\mathcal{V}^*, \mathcal{P}}$. Then we analyse Construction 3.4, and show that provided $\mathcal{A}_{\mathcal{V}^*}$ does not break the query bound on $\pi_0 \leftarrow \mathcal{P}_0$, we can simulate its view of π_0 . Combining these two facts yields the result.

Let \mathcal{V}^* be an arbitrary (possibly adaptive) malicious verifier of $(\mathcal{P}, \mathcal{V})$, and let π and π_0 be random variables denoting the output of \mathcal{P} and \mathcal{P}_0 respectively. Consider the hybrid simulator $\mathcal{A}_{\mathcal{V}^*}^{\pi_0}$ defined in Construction 3.3. As \mathcal{P} is ℓ -locally computable by f , we have that π is identically distributed to $(f^{\pi_0}(\alpha))_{\alpha \in \text{Dom}(\pi)}$. Thus the output of Construction 3.3 is distributed identically to $\text{View}_{\mathcal{V}^*, \mathcal{P}}$.

Now we analyse Construction 3.4. As $(\mathcal{P}_0, \mathcal{V}_0)$ is a PZK-PCP, and $\mathcal{A}_{\mathcal{V}^*}$ is a verifier of $\pi_0 \leftarrow \mathcal{P}_0$, there exists a simulator $\overline{\text{Sim}}_{\mathcal{A}_{\mathcal{V}^*}}$ for $\mathcal{A}_{\mathcal{V}^*}^{\pi_0}$, whose output is identically distributed to $\text{View}_{\mathcal{A}_{\mathcal{V}^*}, \mathcal{P}_0}(x, w)$ provided that $\mathcal{A}_{\mathcal{V}^*}$ makes at most q^* queries to π_0 . Note that for every query made by \mathcal{V}^* , $\mathcal{A}_{\mathcal{V}^*}$ makes at most ℓ queries to π_0 . Thus, the output of $\overline{\text{Sim}}_{\mathcal{A}_{\mathcal{V}^*}}$ is identically distributed to $\text{View}_{\mathcal{A}_{\mathcal{V}^*}, \mathcal{P}_0}(x, w)$ provided that \mathcal{V}^* makes at most q^*/ℓ queries. Therefore, as $\mathcal{A}_{\mathcal{V}^*}(x)$ is receiving answers to its queries to π_0 identically distributed to the real proof, by the first part, its output is identically distributed to $\text{View}_{\mathcal{V}^*, \mathcal{P}}$. \square

3.2 Alphabet reduction of PZK-PCPs

In this section, we use the notion of ℓ -locally computable proofs to show that standard alphabet reduction techniques preserve zero knowledge. In particular, we will show this for the alphabet reduction construction in [BGHSV06, Lemma 2.13]. For the sake of clarity, we include a self-contained presentation of their construction below (Construction 3.6).

Lemma 3.5 (Zero-knowledge alphabet reduction). *If a language \mathcal{L} has a PZK-PCP over the alphabet $\{0, 1\}^a$, query complexity q , randomness complexity r , query bound q^* , decision complexity d and robust-soundness error s with robustness parameter ρ , then \mathcal{L} has a boolean PZK-PCP with query complexity $O(a \cdot q)$, randomness complexity r , query bound q^* , decision complexity $d + O(a \cdot q)$ and robust-soundness error s with robustness parameter $\Omega(\rho)$.*

Following [BGHSV06], the construction relies on the existence of good error-correcting codes, which have constant relative distance and rate. To obtain the desired decision complexity, a code which is also computable by linear-sized circuits is used, such as [Spi96].

Construction 3.6. A boolean PZK-PCP for the language \mathcal{L} , given a PZK-PCP $(\mathcal{P}_a, \mathcal{V}_a)$ over the alphabet $\{0, 1\}^a$ for \mathcal{L} , where the verifier \mathcal{V}_a has decision algorithm D_a and query algorithm Q_a . Let $\text{ECC}: \{0, 1\}^a \rightarrow \{0, 1\}^b$ where $b = O(a)$ be a systematic binary error-correcting code of constant relative minimum distance, computable by an explicit circuit of size $O(a)$.

Proof:

1. Run $\mathcal{P}_a(x)$ to obtain a proof π_a .
2. Define the proof oracle τ , by $\tau(\gamma) := \text{ECC}(\pi_a(\gamma))$, for all $\gamma \in \text{Dom}(\pi_a)$.
3. Output (π_a, τ) .

Verifier:

1. Compute the query set $Q := Q_a(x) \subseteq \text{Dom}(\pi_a)$.
2. Run $D_a(x, \pi_a|_Q)$ and reject if it rejects.
3. For each $\gamma \in Q$, compute $\text{ECC}(\pi_a(\gamma))$, query $\tau(\gamma)$ and reject if $\text{ECC}(\pi_a(\gamma)) \neq \tau(\gamma)$.
4. If all of the above checks pass, then accept.

Proof of Lemma 3.5. As Construction 3.6 is entirely unchanged from [BGHSV06], a proof of completeness, soundness and all of the parameters (apart from the query bound) can be found in Lemma 2.13 of [BGHSV06]. Here we argue zero knowledge.

We show zero knowledge by demonstrating that \mathcal{P} , as defined in Construction 3.6, is 1-locally computable from the PZK-PCP \mathcal{P}_a , which has query bound q^* . Hence, by Lemma 3.2, the PCP \mathcal{P} inherits its zero-knowledge property from \mathcal{P}_a . Towards this end, observe that \mathcal{P} is 1-locally computable by the function $f: \text{Dom}(\pi) \rightarrow \{0, 1\}^b$, defined as

$$f^{\pi_a}(\mathcal{O}, \alpha, i) = \begin{cases} \pi_a(\alpha) & \text{if } \mathcal{O} = \pi_a \\ \text{ECC}(\pi_a(\alpha))_i & \text{if } \mathcal{O} = \tau \end{cases}.$$

Thus, by Lemma 3.2, Construction 3.6 is perfect zero knowledge, with query bound q^* . \square

3.3 Proof composition

We extend the proof composition theorem in [BGHSV06, Theorem 2.7] to account for zero knowledge. Recall that proof composition is performed with a robust outer PCP and an inner PCP of proximity. We show that if proof composition is carried out with a robust outer PCP that is zero knowledge, then the composed PCP will be zero knowledge. Crucially, the inner PCP of proximity is *not* required to be zero knowledge.

Theorem 3.7 (Zero-knowledge proof composition). *Suppose that for functions $r_{out}, r_{in}, d_{out}, d_{in}, q_{in}, q^*: \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon_{out}, \varepsilon_{in}, \rho_{out}, \delta_{in}: \mathbb{N} \rightarrow [0, 1]$ the following hold:*

- *Language \mathcal{L} has a robust boolean PZK-PCP $(\mathcal{P}_{out}, \mathcal{V}_{out})$ with randomness complexity r_{out} , query complexity q_{out} , decision complexity d_{out} , robust-soundness error $1 - \varepsilon_{out}$, robustness parameter ρ_{out} and zero-knowledge query bound q^* .*
- *CIRCUIT VALUE has a boolean PCPP $(\mathcal{P}_{in}, \mathcal{V}_{in})$ where \mathcal{P}_{in} is deterministic and efficient, with randomness complexity r_{in} , query complexity q_{in} , decision complexity d_{in} , proximity parameter δ_{in} , and soundness error $1 - \varepsilon_{in}$.*
- *$\delta_{in}(d_{out}(n)) \leq \rho_{out}(n)$ for every n .*

Then \mathcal{L} has a PZK-PCP, denoted $(\mathcal{P}_{comp}, \mathcal{V}_{comp})$ (cf. Construction 3.8) with:

- *randomness complexity $r_{out}(n) + r_{in}(d_{out}(n))$;*
- *query complexity $q_{in}(d_{out}(n))$;*
- *decision complexity $d_{in}(d_{out}(n))$;*
- *soundness error $1 - \varepsilon_{out}(n) \cdot \varepsilon_{in}(d_{out}(n))$; and*
- *query bound q^*/q_{out} ;*

Construction 3.8 is entirely unchanged from [BGHSV06] (apart from the stipulation that the outer PCP be zero knowledge), and is reproduced here for completeness.

Construction 3.8 (Composed PCP). A composed PZK-PCP $(\mathcal{P}_{\text{comp}}, \mathcal{V}_{\text{comp}})$ for the language \mathcal{L} , given a robust boolean outer PZK-PCP $(\mathcal{P}_{\text{out}}, \mathcal{V}_{\text{out}})$ for \mathcal{L} (with randomness complexity r_{out} and query complexity q_{out}) and an inner boolean PCPP $(\mathcal{P}_{\text{in}}, \mathcal{V}_{\text{in}})$ for CIRCUIT VALUE (with query complexity q_{in}). Let D_{out} and Q_{out} denote the decision and query algorithms for \mathcal{V}_{out} respectively. Both parties receive common input $x \in \{0, 1\}^n$.

Proof:

1. Run $\mathcal{P}_{\text{out}}(x)$ to obtain a proof π_{out} .
2. For each choice of randomness $r \in \{0, 1\}^{r_{\text{out}}}$, compute the query set $I_{\text{out},r} := Q_{\text{out}}(x; r)$.
3. Compile D_{out} into a circuit $C_{\text{out}}: \{0, 1\}^n \times \{0, 1\}^{\ell_{\text{out}}} \times \{0, 1\}^{r_{\text{out}}} \rightarrow \{0, 1\}$, and for each $r \in \{0, 1\}^{r_{\text{out}}}$, define the restricted circuit $C_{\text{out},r} := C_{\text{out}}(x, \cdot, r)$.
4. For each $r \in \{0, 1\}^{r_{\text{out}}}$, run $\mathcal{P}_{\text{in}}(C_{\text{out},r}, \pi_{\text{out}}|_{I_{\text{out},r}})$ to obtain a proof π_r .
5. Output $\pi := (\pi_{\text{out}}, (\pi_r)_{r \in \{0, 1\}^{r_{\text{out}}}})$.

Verifier:

1. Compile D_{out} into a circuit $C_{\text{out}}: \{0, 1\}^n \times \{0, 1\}^{\ell_{\text{out}}} \times \{0, 1\}^{r_{\text{out}}} \rightarrow \{0, 1\}$.
2. Sample $r \leftarrow \{0, 1\}^{r_{\text{out}}}$, and define the restricted circuit $C_{\text{out},r} := C_{\text{out}}(x, \cdot, r)$.
3. Run $\mathcal{V}_{\text{in}}^{\pi_{\text{out}}, \pi_r}(C_{\text{out},r})$, treating π_{out} as the input oracle and π_r as the proof oracle. Accept if and only if \mathcal{V}_{in} accepts.

In order to show that proof composition preserves zero-knowledge, we will show that $\mathcal{P}_{\text{comp}}$ (defined in Construction 3.8) is locally computable from \mathcal{P}_{out} .

Claim 3.9. Let $(\mathcal{P}_{\text{comp}}, \mathcal{V}_{\text{comp}})$ be the composed PCP system defined in Construction 3.8, and let $(\mathcal{P}_{\text{out}}, \mathcal{V}_{\text{out}})$ be the perfect zero-knowledge outer PCP of $(\mathcal{P}_{\text{comp}}, \mathcal{V}_{\text{comp}})$. $\mathcal{P}_{\text{comp}}$ is q_{out} -locally computable from \mathcal{P}_{out} .

Proof. Clearly queries to the outer proof are 1-locally computable. To locally compute queries to π_r for any $r \in \{0, 1\}^{r_{\text{out}}}$, run Steps 2 to 4 of the prover in Construction 3.8 to obtain a proof π_r , then output $\pi_r(\alpha)$. Correctness is clear. In the worst case, we make $\max_{r \in \{0, 1\}^{r_{\text{out}}}} |I_{\text{out},r}| = q_{\text{out}}$ many queries to π_{out} . Finally, the efficiency and determinism follow from the efficiency and determinism of \mathcal{P}_{in} and \mathcal{V}_{out} (which is deterministic for any fixed choice of randomness). \square

Proof of Theorem 3.7. The randomness complexity, query complexity and decision complexity follow straightforwardly from the construction. A proof that Construction 3.8 satisfies completeness and the prescribed soundness error can be found in [BGHSV05]. The construction is perfect zero-knowledge with query bound q^*/q_{out} by Lemma 3.2 and Claim 3.9. \square

We finish this section by instantiating the inner PCPP for CIRCUIT VALUE to prove a corollary of Theorem 3.7, in which the parameters of the composed ZKPCP only depend on the parameters of the outer ZKPCP. We shall need Theorem 3.3 in [BGHSV05], which we restate below, setting the parameter $t(n) = 2/\varepsilon$, while noting that the honest prover strategy is efficient.

Theorem 3.10 ([BGHSV05]). For any $\varepsilon > 0$, CIRCUIT VALUE has a PCPP $(\mathcal{P}, \mathcal{V})$ where \mathcal{P} is deterministic and efficient, with randomness complexity $\log n + O(\log^\varepsilon(n))$, query complexity $O(1/\varepsilon)$, proximity parameter $\Theta(\varepsilon)$ and soundness error $1/2$.

Corollary 3.11. $\text{rPZK-PCP}_{q^*, \{0, 1\}}[r, q] \subseteq \text{PZK-PCP}_{q^*/q, \{0, 1\}}[r + \log n, 1]$.

Proof. Let $\mathcal{L} \in \text{rPZK-PCP}_{q^*, \{0,1\}}[r, q]$. That is, there is a PZK-PCP $(\mathcal{P}, \mathcal{V})$ for \mathcal{L} with randomness complexity $r(n)$, query complexity $q(n)$, query bound $q^*(n)$ and constant robustness ρ . Theorem 3.10 guarantees the existence of a PCPP for CIRCUIT VALUE with proximity parameter $k\varepsilon$, for some universal constant k and any $\varepsilon > 0$. Set $\varepsilon := \rho/k$. Then by Theorem 3.7 (note that our choice of ε ensures that the third bullet of the hypothesis is met), \mathcal{L} has a PZK-PCP with randomness complexity $r(n) + \log(n) + O(\log^\varepsilon(n))$ and query complexity $O(1/\varepsilon)$, as required. \square

4 Robust PCPPs for polynomial summation

In this section, we construct a PCP of proximity for the Sum pair language, defined below, which achieves constant *robust* soundness.

Definition 4.1. *We define the language Sum to be the set of all pairs $((\mathbb{F}, 1^m, 1^d, H, \gamma), F)$ where \mathbb{F} is a finite field, $m, d \in \mathbb{N}$, $H \subseteq \mathbb{F}$, $\gamma \in \mathbb{F}$, and $F \in \text{RM}[\mathbb{F}, m, d]$ with $\sum_{\vec{b} \in H^m} F(\vec{b}) = \gamma$.*

Note that d and m are specified in unary, so that the verifier and simulator for our PCPP can run in time $\text{poly}(\log |\mathbb{F}|, m, d, |H|)$. Below, for readability, we will write explicit inputs of Sum as $(\mathbb{F}, m, d, H, \gamma)$; this should be understood as $(\mathbb{F}, 1^m, 1^d, H, \gamma)$.

Lemma 4.2. *Let \mathbb{F} be a finite field, and let $m, d \in \mathbb{N}$, $H \subseteq \mathbb{F}$ and $\delta > 0$ be such that $\delta > \frac{md}{|\mathbb{F}|}$ and $d \geq |H| + 1$. Then Construction 4.3 is a robust PCP of proximity for Sum, over the alphabet \mathbb{F}^{m-1} , with proximity parameter δ and robustness parameter $\rho = \Omega(\delta)$. The verifier makes $O(|\mathbb{F}|)$ -many queries to F and $|\mathbb{F}|$ -many queries to π . The proof length is $|\mathbb{F}|^{m-1}$.*

While the robust PCP of [BGHSV06] uses (a variant of) sumcheck, to the best of our knowledge Lemma 4.2 has not previously been shown explicitly. Our proof here is substantially simpler than that of [BGHSV06], in part because the structure of our PCP permits a simpler “bundling” of symbols. We will build on this construction when constructing robust PZK-PCPPs for Sum, in Section 5.

Construction 4.3. A robust PCP of proximity RSC for Sum with proximity parameter $\delta > 0$ over the alphabet \mathbb{F}^{m-1} . Both parties receive the explicit input $(\mathbb{F}, m, d, H, \gamma, \delta)$, and oracle access to the evaluation table of the implicit input $F: \mathbb{F}^m \rightarrow \mathbb{F}$.

Proof:

1. For each $i \in [m-1]$, define the i -th sumcheck polynomial to be the i -variate polynomial $g_i(\vec{X}) \in \mathbb{F}^{\leq d}[X_1, \dots, X_i]$ given by

$$g_i(x_1, \dots, x_i) := \sum_{\vec{b} \in H^{m-i}} F(x_1, \dots, x_i, \vec{b}).$$

2. Define $\pi: \mathbb{F}^{m-1} \rightarrow \mathbb{F}^{m-1}$ by

$$\pi(c_1, \dots, c_{m-2}, \alpha) := (g_1(\alpha), g_2(c_1, \alpha), \dots, g_{m-1}(c_1, \dots, c_{m-2}, \alpha)),$$

for every $(c_1, \dots, c_{m-2}, \alpha) \in \mathbb{F}^{m-1}$.

3. Output π .

Verifier:

1. Sample $\vec{c} := (c_1, \dots, c_{m-1}) \leftarrow \mathbb{F}^{m-1}$ uniformly at random.
2. For all $\alpha \in \mathbb{F}$, query $\pi(c_1, \dots, c_{m-2}, \alpha) := (g_1(\alpha), g_2(c_1, \alpha), \dots, g_{m-1}(c_1, \dots, c_{m-2}, \alpha))$.
3. For all $\alpha \in \mathbb{F}$, query $F(c_1, \dots, c_{m-1}, \alpha)$.
4. For each $i \in [m-1]$, reject if $g_i(c_1, \dots, c_{i-1}, X)$ is not a polynomial of degree at most d .
5. Check that

$$\sum_{b \in H} g_1(b) = \gamma.$$

6. For each $i \in [m-2]$, check that

$$\sum_{b \in H} g_{i+1}(c_1, \dots, c_i, b) = g_i(c_1, \dots, c_i).$$

7. Check that

$$\sum_{b \in H} F(\vec{c}, b) = g_{m-1}(\vec{c}).$$

8. Perform a robust low total degree test (Theorem 2.16) on F , with proximity parameter $\delta_{\text{RM}} = \min(\delta, 1/5)$, and reject if the test fails.
9. If none of the above checks fail, then accept.

We now prove that Construction 4.3 is robust, but first outline our strategy for doing so. Conceptually speaking, Construction 4.3 consists of two tests: the sumcheck tests in Steps 4 to 7 (which includes a check that the proof itself has low-degree structure), and the low-degree test on the input oracle in Step 8. The low-degree test is indeed robust (by Theorem 2.16), so the key technical component in proving the robustness of Construction 4.3 is Lemma 4.4, which asserts that the sumcheck tests (Steps 4 to 7) reject robustly, under the assumption the input oracle is close to a low-degree polynomial. Since the verifier makes roughly the same number of queries when performing both of these tests, this implies the overall construction is robust.

Lemma 4.4 holds due to the assumed low-degree structure of F and our specific choice of bundling. If the view of the sumcheck verifier is rejecting, then, since both the proof π and the input oracle F are close to low degree, at least one of F, g_1, \dots, g_{m-1} needs to be modified to equal a different polynomial. By the

Schwartz-Zippel lemma, this requires a substantial fraction of modifications. Our bundling strategy ensures that these modifications form a substantial fraction of the view of the verifier. We now proceed with a formal proof.

Lemma 4.4. *Let $\tilde{F}: \mathbb{F}^m \rightarrow \mathbb{F}$ be δ_Σ -far from $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$, but δ_{RM} -close to $\text{RM}[\mathbb{F}, m, d]$ for some $\frac{md}{|\mathbb{F}|} < \delta_{RM} \leq \delta_\Sigma$ and $d \geq |H| + 1$. Then for all proofs π^* ,*

$$\mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} \left[\Delta \left((\pi^*(c_1, \dots, c_{m-2}, \alpha)_{\alpha \in \mathbb{F}}, \tilde{F}(c_1, \dots, c_{m-1}, \alpha)_{\alpha \in \mathbb{F}}), \text{Acc}(\mathcal{V}) \right) \right] \geq \frac{1}{2} \min(\delta_{RM}, 1 - 4\delta_{RM}).$$

Proof. Let $\hat{F} \in \mathbb{F}^{\leq d}[X_1, \dots, X_m]$ be the some polynomial of total degree at most d which is δ_{RM} -close to \tilde{F} , as guaranteed by the hypothesis. Note that by the condition that \tilde{F} is δ_Σ -far from $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$ and $\delta_{RM} \leq \delta_\Sigma$, it must be that $\sum_{\vec{b} \in H^m} \hat{F}(\vec{b}) \neq \gamma$. Before proceeding, we remind the reader of some notation: for any $\vec{c} = (c_1, \dots, c_{m-1}) \in \mathbb{F}^{m-1}$, and any $i \in [m-1]$, we denote $\vec{c}_i := (c_1, \dots, c_i)$. Also, for a polynomial $f \in \mathbb{F}^{\leq d}[X_1, \dots, X_i]$, we denote the restriction of f to the axis-parallel line defined by \vec{c}_{i-1} by $f|_{\vec{c}_{i-1}}(X) := f(\vec{c}_{i-1}, X)$.

First, fix some $\vec{c} = (c_1, \dots, c_{m-1}) \in \mathbb{F}^{m-1}$, and suppose that for some $i \in [m-1]$, $g_i(c_1, \dots, c_{i-1}, X)$ is δ_{RM} -far from every univariate polynomial of degree at most d . Then since the verifier inspects each $g_i(c_1, \dots, c_{i-1}, X)$ at every point in \mathbb{F} and rejects if $g_i(c_1, \dots, c_{i-1}, X)$ is not a polynomial of degree at most d (Step 4), at least a δ_{RM} -fraction of symbols of $\pi^*(c_1, \dots, c_{m-2}, \alpha)_{\alpha \in \mathbb{F}}$ would need to be modified in order for the verifier to accept. Thus $(\pi^*(c_1, \dots, c_{m-2}, \alpha)_{\alpha \in \mathbb{F}}, \tilde{F}(c_1, \dots, c_{m-1}, \alpha)_{\alpha \in \mathbb{F}})$ has distance at least $\delta_{RM}/2$ to any accepting view.

Hence we can proceed under the assumption that for all $i \in [m-1]$, $g_i(c_1, \dots, c_{i-1}, X)$ is δ_{RM} -close to some univariate polynomial $\hat{g}_i(X) \in \mathbb{F}^{\leq d}[X]$. Let E be the event that either

$$\sum_{b \in H} \hat{g}_1(b) \neq \gamma,$$

or

$$\sum_{b \in H} \hat{g}_{i+1}(b) \neq \hat{g}_i(c_i)$$

for some $i \in [m-2]$, or

$$\sum_{b \in H} \hat{F}(\vec{c}, b) \neq \hat{g}_{m-1}(c_{m-1}).$$

Note that since $\sum_{\vec{b} \in H^m} \hat{F}(\vec{b}) \neq \gamma$, by the soundness of sumcheck, $\Pr_{\vec{c}}[\neg E] \leq md/|\mathbb{F}|$.

Conditioned on E , in order for the verifier to accept, one of $g_1(X), g_2|_{\vec{c}_1}(X), \dots, g_{m-1}|_{\vec{c}_{m-2}}(X), \tilde{F}|_{\vec{c}}(X)$ must be modified to equal a degree d polynomial other than $\hat{g}_1(X), \hat{g}_2|_{\vec{c}_1}(X), \dots, \hat{g}_{m-1}|_{\vec{c}_{m-2}}(X), \hat{F}|_{\vec{c}}(X)$ respectively. This would require modifying $g_i|_{\vec{c}_{i-1}}(X)$, for some $i \in [m-1]$, in at least a $1 - d/|\mathbb{F}| - \Delta(g_i|_{\vec{c}_{i-1}}, \hat{g}_i|_{\vec{c}_{i-1}}) \geq 1 - d/|\mathbb{F}| - \delta_{RM}$ fraction of points, or modifying $\tilde{F}|_{\vec{c}}(X)$ in at least a $1 - d/|\mathbb{F}| - \Delta(\tilde{F}|_{\vec{c}}, \hat{F}|_{\vec{c}})$ fraction of points. Thus, the verifier's view of the proof would require modification in at least a $\frac{1}{2}(1 - \frac{d}{|\mathbb{F}|} - \max(\delta_{RM}, \Delta(\tilde{F}|_{\vec{c}}, \hat{F}|_{\vec{c}})))$ fraction of points.

Thus the expected distance of $(\pi^*(c_1, \dots, c_{m-2}, \alpha)_{\alpha \in \mathbb{F}}, \tilde{F}(c_1, \dots, c_{m-1}, \alpha)_{\alpha \in \mathbb{F}})$ to an accepting view is

at least

$$\begin{aligned}
\frac{1}{2} \cdot \mathbb{E}_{\tilde{c} \leftarrow \mathbb{F}^{m-1}} \left[1 - \frac{d}{|\mathbb{F}|} - \max(\delta_{\text{RM}}, \Delta(\tilde{F}|_{\tilde{c}}, \hat{F}|_{\tilde{c}})) \right] - \Pr_{\tilde{c}}[\neg E] &= \frac{1}{2} \left(1 - \frac{d}{|\mathbb{F}|} - \delta_{\text{RM}} \right) - \Pr_{\tilde{c}}[\neg E] \\
&\geq \frac{1}{2} \left(1 - \frac{d}{|\mathbb{F}|} - \delta_{\text{RM}} \right) - \frac{md}{|\mathbb{F}|} \\
&\geq \frac{1}{2} (1 - 4\delta_{\text{RM}}),
\end{aligned}$$

where we have used that $\mathbb{E}[\Delta(\tilde{F}|_{\tilde{c}}, \hat{F}|_{\tilde{c}})] \leq \delta_{\text{RM}}$ and that $\frac{md}{|\mathbb{F}|} \leq \delta_{\text{RM}}$ by hypothesis. \square

Now we employ Lemma 4.4 to prove that Construction 4.3 is robust, following the approach we outlined above.

Proof of Lemma 4.2. Completeness is straightforward; we prove robust soundness and efficiency.

Robust soundness. Suppose that \tilde{F} is δ -far from $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$. First, if F is δ_{RM} -far from $\text{RM}[\mathbb{F}, m, d]$, then by Theorem 2.16, the expected distance of the view of the low degree test (Step 8) from any accepting view is $\Omega(\delta)$ (since $\delta_{\text{RM}} \leq \delta$). As the low-degree test makes $O(|\mathbb{F}|)$ queries to F , and (apart from the low-degree test) the verifier in Construction 4.3 makes $|\mathbb{F}|$ queries to each of F and π , the queries made by the low-degree test form a constant fraction of the queries made overall by the verifier. Hence the expected distance of the view of the verifier from any accepting view is also $\Omega(\delta)$.

Second, if F is δ -far from $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$ but δ_{RM} -close to $\text{RM}[\mathbb{F}, m, d]$, then by Lemma 4.4, the expected distance of the view of the verifier (when ignoring the low-degree test) to any accepting view is at least $\frac{1}{2} \min(\delta_{\text{RM}}, (1 - 4\delta_{\text{RM}})) \geq \min(\delta/2, 1/10) = \Omega(\delta)$ (since $\delta_{\text{RM}} := \min(\delta, 1/5)$). Once again, the queries made by the verifier apart from the low-degree test form a constant fraction of the overall number of queries made by the verifier, so the expected distance of the view of the verifier from any accepting view is $\Omega(\delta)$. In other words, Construction 4.3 has expected robustness $\Omega(\delta)$.

By Proposition 2.8, for an appropriate choice of constant ε , Construction 4.3 has robustness $\Omega(\delta)$.

Efficiency. The proof length is clear from construction. The query complexity follows from the fact that the low-degree test (Step 8) makes $O(|\mathbb{F}|)$ queries to F , and the rest of Construction 4.3 makes $|\mathbb{F}|$ queries to each of F and π . \square

5 Robust PZK-PCPPs for polynomial summation

In this section, we construct PZK-PCPPs of proximity for Sum (Definition 4.1) with constant robust soundness. The construction makes black-box use of the robust sumcheck PCPP we obtain in Lemma 4.2 and builds on the PZK-PCP for $\#P$ constructed in [GOS24].

Lemma 5.1. *Let $\delta > 0$, \mathbb{F} be a finite field, $H \subseteq \mathbb{F}$, $\gamma \in \mathbb{F}$, and $m, d \in \mathbb{N}$ be such that $\frac{md}{|\mathbb{F}|} < \delta$ and $d \geq |H| + 1$. Construction 5.2 is a robust (black-box) perfect zero-knowledge PCP of proximity for $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$ over the alphabet \mathbb{F}^{m+1} with proximity parameter δ and robustness parameter $\rho = \Omega(\delta)$. The verifier complexity is $\text{poly}(|\mathbb{F}|, m, d)$, the verifier makes $O(|\mathbb{F}|)$ queries to both F and π , and the proof length is $O(|\mathbb{F}|^m)$.*

We note that (as in [GOS24]) the zero-knowledge property for this PCPP holds against all polynomial-time malicious verifiers regardless of their query complexity (in fact, under an appropriate definition, ZK holds against *all* verifiers regardless of their complexity).

Construction 5.2 closely follows [GOS24], with only a few modifications to achieve constant robustness. We provide a brief high-level description of the [GOS24] construction, and then present our construction in detail. We refer the reader to [GOS24] for further exposition.

The [GOS24] ZKPCP for Sum is a (standard) sumcheck PCPP for a “masked” version of the polynomial F ; specifically, a sumcheck PCPP for the claim

$$\sum_{\vec{a} \in H^m} F(\vec{a}) + \underbrace{Q(\vec{a}) - Q(\vec{a}_{\text{rev}})}_{\text{antisymmetric}} + \underbrace{\sum_{i=1}^m Z_H(a_i) T_i(\vec{a})}_{\text{nullstellensatz}} = \gamma \quad (2)$$

where

- Q is a uniformly random polynomial in $\mathbb{F}^{\leq d}[X_1, \dots, X_m]$;
- $Z_H := \prod_{a \in H} (X - a)$ is the nonzero univariate polynomial of minimal degree that is zero on H ; and
- T_i , for each $i \in [m]$, is a uniformly random polynomial in $\mathbb{F}^{\leq \vec{d}_i}[X_1, \dots, X_m]$, where \vec{d}_i taking value $d - |H|$ in the i -th coordinate and d in all other coordinates.

The polynomials Q, T_1, \dots, T_m are provided in the proof. It is not difficult to see that, for any functions Q, T_1, \dots, T_m , (2) is true if and only if $\sum_{\vec{a} \in H^m} F(\vec{a}) = \gamma$.

The mask polynomials make the PCPP zero knowledge. The purpose of the “antisymmetric” term $Q(\vec{X}) - Q(\vec{X}_{\text{rev}})$ is to hide the values of intermediate sums in the sumcheck proof, which are #P-hard to compute, while summing to zero over H^m to maintain soundness. The “nullstellensatz” term computes, by the combinatorial nullstellensatz [Alo99], a uniformly random polynomial N such that $N(\vec{h}) = 0$ for all $\vec{h} \in H$. By adding this term, we retain the antisymmetric structure of the mask on H^m (necessary for soundness) while maximising the randomness outside of H^m (necessary for simulation).

We are now ready to present our construction.

Construction 5.2. A robust PZK-PCP of proximity for $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$ with proximity parameter $\delta > 0$ over the alphabet \mathbb{F}^{m+1} . Both parties receive common input $(\mathbb{F}, m, d, H, \gamma, \delta)$, and oracle access to the evaluation table of $F: \mathbb{F}^m \rightarrow \mathbb{F}$.

Proof:

1. Sample a polynomial $Q \leftarrow \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ uniformly at random, where $\vec{d} = (d, \dots, d)$.
2. For each $i \in [m]$, sample $T_i \leftarrow \mathbb{F}^{\leq \vec{d}_i}[X_1, \dots, X_m]$ uniformly at random, where $\vec{d}_i = (d, \dots, d - |H|, \dots, d)$ is the vector which takes the value d in every coordinate except for the i -th, which takes the value $d - |H|$.
3. Define $\pi_P(\vec{x}) = (Q(\vec{x}), T_1(\vec{x}), \dots, T_m(\vec{x}))$ for all $\vec{x} \in \mathbb{F}^m$, written over the alphabet \mathbb{F}^{m+1} .
4. Define $R(\vec{X}) := Q(\vec{X}) - Q_{\text{rev}}(\vec{X}) + \sum_{i=1}^m Z_H(X_i)T_i(\vec{X})$, where $Z_H := \prod_{a \in H} (X - a)$, and $Q_{\text{rev}}(\vec{X}) := Q(\vec{X}_{\text{rev}})$.
5. Compute $\pi_\Sigma := \text{RSC}[(\mathbb{F}, m, d, H, \gamma, \delta), F + R]$, where $\text{RSC}[x, f]$ denotes the robust sumcheck PCP with explicit input x and implicit input f (Construction 4.3). Write this proof over the alphabet \mathbb{F}^{m+1} (rather than \mathbb{F}^{m-1}) by padding each symbol with two zeros at the end.
6. Output $\pi := (\pi_\Sigma, \pi_P)$, written over the alphabet \mathbb{F}^{m+1} .

Verifier:

1. Emulate the robust sumcheck verifier \mathcal{V}_Σ (Construction 4.3) on input $F + R$ and proof π_Σ . To query the input $F + R$ at some point $\vec{\alpha} \in \mathbb{F}^m$, query $F(\vec{\alpha}), \pi_P(\vec{\alpha}) := (Q(\vec{\alpha}), T_1(\vec{\alpha}), \dots, T_m(\vec{\alpha}))$ and $\pi_P(\vec{\alpha}_{\text{rev}}) := (Q(\vec{\alpha}_{\text{rev}}), T_1(\vec{\alpha}_{\text{rev}}), \dots, T_m(\vec{\alpha}_{\text{rev}}))$, then compute $(F + R)(\vec{\alpha}) := F(\vec{\alpha}) + Q(\vec{\alpha}) - Q(\vec{\alpha}_{\text{rev}}) + \sum_{i=1}^m Z_H(\alpha_i)T_i(\vec{\alpha})$. Reject if \mathcal{V}_Σ rejects.
2. Perform a robust low-degree test (Theorem 2.16) on F , with proximity parameter $\delta_{\text{RM}} := \min(\delta, 1/5)$, and reject if the test fails.
3. Perform a robust vector-valued low-degree test on π_P (Theorem 2.16), with proximity parameter $\varepsilon_P := \frac{1}{8}\delta_{\text{RM}}$ and degree parameter $d_P = md$, and reject if the test fails.
4. Accept if and only if none of the above tests rejected.

To prove the robustness of our construction, we show a technical lemma (see Lemma 5.3) which asserts that the PCPP is robust under the assumption that the input oracle is close to some Reed-Muller codeword that does not sum to γ over H^m (and indeed is far from all Reed-Muller codewords that do). This will imply that the whole construction is robust, since the low-degree test is robust.

We also need to argue that the zero-knowledge property holds. We accomplish this using the notion of locally computable proofs developed in Section 3.1. Essentially, we show that the PCP in Construction 5.2 is locally computable from the perfect zero-knowledge PCPP for sumcheck defined in [GOS24], and so it inherits zero knowledge from the GOS construction by Lemma 3.2.

Throughout the remainder of this section we will use the following notation. We denote the non-ZK robust sumcheck PCP verifier (Construction 4.3) by \mathcal{V}_Σ . We denote the verifier that performs Step 1 of Construction 5.2 by \mathcal{V}_1 . That is, \mathcal{V}_1 is given $F: \mathbb{F}^m \rightarrow \mathbb{F}$ as input, and runs \mathcal{V}_Σ on input oracle $F + R: \mathbb{F}^m \rightarrow \mathbb{F}$ (computing queries to $F + R$ via F and the proof π_P) and proof oracle π_Σ . \mathcal{V}_1 then outputs the output of \mathcal{V}_Σ . For any $\vec{c} \in \mathbb{F}^{m-1}$, we denote the view of \mathcal{V}_Σ , on randomness $\vec{c} \in \mathbb{F}^{m-1}$ and input $G: \mathbb{F}^m \rightarrow \mathbb{F}$, by $\nu_\Sigma(\vec{c}) := (\pi_\Sigma(\vec{c}_{m-2}, \alpha)_{\alpha \in \mathbb{F}}, G(\vec{c}, \alpha)_{\alpha \in \mathbb{F}})$ and we denote the view of \mathcal{V}_1 , on randomness $\vec{c} \in \mathbb{F}^{m-1}$ and input $F: \mathbb{F}^m \rightarrow \mathbb{F}$, by $\nu(\vec{c}) := (\pi_\Sigma(\vec{c}_{m-2}, \alpha)_{\alpha \in \mathbb{F}}, \pi_P(\vec{c}, \alpha)_{\alpha \in \mathbb{F}}, \pi_P(\alpha, \vec{c}_{\text{rev}})_{\alpha \in \mathbb{F}}, F(\vec{c}, \alpha)_{\alpha \in \mathbb{F}})$.

Lemma 5.3. Let $\tilde{F}: \mathbb{F}^m \rightarrow \mathbb{F}$ be δ_Σ -far from $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$, but δ_{RM} -close to $\text{RM}[\mathbb{F}, m, d]$ for $\delta_\Sigma >$

$\delta_{RM} \geq \frac{md}{|\mathbb{F}|}$, $\delta_{RM} < 1/5$ and $d \geq |H| + 1$. Then for all proofs (π_Σ^*, π_P^*) ,

$$\mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} \left[\Delta \left((\pi_\Sigma^*(\vec{c}_{m-2}, \alpha)_{\alpha \in \mathbb{F}}, \pi_P^*(\vec{c}, \alpha)_{\alpha \in \mathbb{F}}, \pi_P^*(\alpha, \vec{c}_{\text{rev}})_{\alpha \in \mathbb{F}}, \tilde{F}(\vec{c}, \alpha)_{\alpha \in \mathbb{F}}), \text{Acc}(\mathcal{V}) \right) \right] = \Omega(\delta_{RM}).$$

Here we view $(\pi_\Sigma^*(\vec{c}_{m-2}, \alpha)_{\alpha \in \mathbb{F}}, \pi_P^*(\vec{c}, \alpha)_{\alpha \in \mathbb{F}}, \pi_P^*(\alpha, \vec{c}_{\text{rev}})_{\alpha \in \mathbb{F}}, \tilde{F}(\vec{c}, \alpha)_{\alpha \in \mathbb{F}}) \in (\mathbb{F}^{m+1})^{4\mathbb{F}}$ as a string over the alphabet \mathbb{F}^{m+1} , of length $4|\mathbb{F}|$, via some natural embedding of \mathbb{F} into \mathbb{F}^{m+1} . The metric Δ refers to Hamming distance over \mathbb{F}^m .

Proof. First, suppose that π_P^* is ε_P -far (where $\varepsilon_P := \frac{1}{8}\delta_{RM}$, see Step 3), with respect to the alphabet \mathbb{F}^{m+1} , from $\text{RM}^{m+1}[\mathbb{F}, m, m \cdot d]$. In this case, then by the robustness of the vector-valued low-degree test (Theorem 2.16), the view of the verifier in Step 3 has expected distance $\Omega(\varepsilon_P) = \Omega(\delta_{RM})$ to any accepting view. Since this test forms a constant fraction of the overall view of the verifier, the expected distance of the view of the verifier from any accepting view would be $\Omega(\delta_{RM})$.

Hence we can assume that π_P^* is ε_P -close, with respect to the alphabet \mathbb{F}^{m+1} , to some vector of m -variate polynomials each of degree at most md , which we denote by $(\hat{Q}(\vec{X}), \hat{T}_1(\vec{X}), \dots, \hat{T}_m(\vec{X}))$. We denote the ‘‘corrected low-degree’’ proof by

$$\pi_{\hat{P}}(\vec{x}) := (\hat{Q}(\vec{x}), \hat{T}_1(\vec{x}), \dots, \hat{T}_m(\vec{x}))$$

and note that by definition of $\pi_{\hat{P}}$ we have $\Delta(\pi_{\hat{P}}, \pi_P) \leq \varepsilon_P$, where this distance is with respect to \mathbb{F}^{m+1} . Lastly, we denote \mathcal{V}_1 's view of the ‘‘corrected’’ proof $\pi_{\hat{P}}$ by

$$\hat{\nu}(\vec{c}) := (\pi_\Sigma(\vec{c}_{m-2}, \alpha)_{\alpha \in \mathbb{F}}, \hat{F}(\vec{c}, \alpha)_{\alpha \in \mathbb{F}}, \pi_{\hat{P}}(\vec{c}, \alpha)_{\alpha \in \mathbb{F}}, \pi_{\hat{P}}(\alpha, \vec{c}_{\text{rev}})_{\alpha \in \mathbb{F}}).$$

We show next that the verifier's view of the corrected proof $\hat{\nu}(\vec{c})$ is far (in expectation) from any accepting view. We accomplish this step by appealing to the robustness of the sumcheck PCPP (Construction 4.3), and by arguing that \mathcal{V}_1 inherits this robustness. Then as the view of the actual proof $\nu(\vec{c})$ is close (in expectation) to the view of the corrected proof $\hat{\nu}(\vec{c})$, $\nu(\vec{c})$ must be far (in expectation) from any accepting view.

By Lemma 4.4, if \mathcal{V}_Σ has oracle access to the evaluation table of a function $G: \mathbb{F}^m \rightarrow \mathbb{F}$, which is δ_Σ -far from $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$, but δ_{RM} -close to $\text{RM}[\mathbb{F}, m, d]$, then for all proofs π_Σ , we have

$$\mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} [\Delta(\nu_\Sigma(\vec{c}), \text{Acc}(\mathcal{V}_\Sigma))] \geq \frac{1}{2} \min(\delta_{RM}, 1 - 4\delta_{RM}) = \frac{1}{2}\delta_{RM}, \quad (3)$$

where the final equality holds since $\delta_{RM} \leq 1/5$. Now consider the verifier \mathcal{V}_1 , which uses its access to F and π_P to run \mathcal{V}_Σ on input oracle $F + R$ and proof π_Σ . Each location of (F, π_P) is used to compute at most two locations of $F + R$. This means that if \mathcal{V}_Σ 's view of $F + R$ is an absolute distance of k away from accepting, then \mathcal{V}_1 's view of (F, π_P) must be at least an absolute distance of $k/2$ away from accepting in the worst case.

Combining this observation with the fact that the full view $\hat{\nu}(\vec{c})$ of \mathcal{V}_1 is twice the length of the full view $\nu_\Sigma(\vec{c})$ of \mathcal{V}_Σ , we see that the expected relative distance degrades by at most a factor of $1/4$. Thus, for all proofs π_Σ ,

$$\mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} [\Delta(\hat{\nu}(\vec{c}), \text{Acc}(\mathcal{V}_1))] \geq \frac{1}{8}\delta_{RM}.$$

Next, we upper bound the expected distance of $\nu(\vec{c})$ and $\hat{\nu}(\vec{c})$. Since $\Delta(\pi_{\hat{P}}, \pi_P) \leq \varepsilon_P$, we have that

$$\mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} [\Delta(\pi_{\hat{P}}(\vec{c}, \alpha)_{\alpha \in \mathbb{F}}, \pi_P(\vec{c}, \alpha)_{\alpha \in \mathbb{F}})] \leq \varepsilon_P = \frac{1}{8}\delta_{RM}.$$

Since $\pi_{\hat{P}}(\vec{c}, \alpha)_{\alpha \in \mathbb{F}}$ and $\pi_P(c, \alpha)_{\alpha \in \mathbb{F}}$ are substrings of $\hat{\nu}(\vec{c})$ and $\nu(\vec{c})$ (respectively) of half the length, and the other components of the two views are identical, it follows that

$$\mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} [\Delta(\hat{\nu}(\vec{c}), \nu(\vec{c}))] \leq \frac{1}{16} \delta_{\text{RM}}.$$

Finally, we put everything together to obtain the result:

$$\begin{aligned} \mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} [\Delta(\nu(\vec{c}), \text{Acc}(\mathcal{V}_1))] &\geq \mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} [|\Delta(\hat{\nu}(\vec{c}), \text{Acc}(\mathcal{V}_1)) - \Delta(\nu(\vec{c}), \hat{\nu}(\vec{c}))|] \\ &\geq \left| \mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} [\Delta(\hat{\nu}(\vec{c}), \text{Acc}(\mathcal{V}_1))] - \mathbb{E}_{\vec{c} \leftarrow \mathbb{F}^{m-1}} [\Delta(\nu(\vec{c}), \hat{\nu}(\vec{c}))] \right| \\ &\geq \frac{1}{8} \delta_{\text{RM}} - \frac{1}{16} \delta_{\text{RM}} \\ &= \frac{1}{16} \delta_{\text{RM}}, \end{aligned}$$

where the first inequality is a consequence of the reverse triangle inequality, and the second is due to linearity of expectation and the fact that the $\mathbb{E}|X| \geq |\mathbb{E}[X]|$ for any random variable X . \square

We now prove Lemma 5.1. We prove robustness by splitting the argument into two cases: one in which the input is far from even being low-degree (in which case the low-degree test guarantees robustness), and the other in which the input is close to low-degree, but far from summing to the correct value (in which case Lemma 5.3 guarantees robustness). As alluded to earlier, we prove zero knowledge using the fact that this robust PCP is locally computable from the PZK-PCP constructed in [GOS24].

Proof of Lemma 5.1. Completeness is straightforward. Here we argue robust soundness and perfect zero-knowledge.

Robust soundness. Suppose that \tilde{F} is δ -far from $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$. First, if \tilde{F} is δ_{RM} -far from $\text{RM}[\mathbb{F}, m, d]$, then by the robustness of the low-degree test (Theorem 2.16), the expected distance of the verifier's view of \tilde{F} to any accepting view is $\Omega(\delta_{\text{RM}}) = \Omega(\delta)$. As the low-degree test makes $O(|\mathbb{F}|)$ many queries to \tilde{F} , and the verifier makes $O(|\mathbb{F}|)$ many queries to both π_{Σ} and π_P , this forms a constant fraction of the view of the verifier. Therefore the expected distance of the view of the verifier from any accepting view is also $\Omega(\delta)$.

Second, if \tilde{F} is δ -far from $\text{Sum}[\mathbb{F}, m, d, H, \gamma]$, but δ_{RM} -close to $\text{RM}[\mathbb{F}, m, d]$, then by Lemma 5.3, the view of the verifier (when ignoring the low-degree test) to any accepting view is $\Omega(\delta_{\text{RM}}) = \Omega(\delta)$. Once again, this forms a constant fraction of the number of queries made by the verifier overall, so the expected distance of the view of the verifier from any accepting view is $\Omega(\delta)$. In other words, Construction 5.2 has expected robustness $\Omega(\delta)$.

Thus by Proposition 2.8, for an appropriate choice of constant, Construction 5.2 has robustness $\Omega(\delta)$.

Zero knowledge. We show that Construction 5.2 is perfect zero knowledge by showing that it is $(m+1)$ -locally computable from the perfect zero-knowledge PCP described in Construction 8.9 of [GOS24]. The proof generated by Construction 5.2 is $(m+1)$ -locally computable by the following function:

$$f^{(\pi'_{\Sigma}, \pi'_Q, \pi'_{T_1}, \dots, \pi'_{T_m})}(\mathcal{O}, \vec{\alpha}) = \begin{cases} (\pi'_Q(\vec{\alpha}), \pi'_{T_1}(\vec{\alpha}), \dots, \pi'_{T_m}(\vec{\alpha})) & \text{if } \mathcal{O} = \pi_P \\ (\pi'_{\Sigma}(\alpha_{m-1}), \pi'_{\Sigma}(\alpha_1, \alpha_{m-1}), \dots, \pi'_{\Sigma}(\vec{\alpha}_{i-1}, \alpha_{m-1}), \dots, \pi'_{\Sigma}(\vec{\alpha})) & \text{if } \mathcal{O} = \pi_{\Sigma} \end{cases},$$

where $(\pi'_{\Sigma}, \pi'_Q, \pi'_{T_1}, \dots, \pi'_{T_m})$ is a proof sampled by the prover of Construction 8.9 of [GOS24]. Thus, by Lemma 3.2 (in this work), and Lemma 8.11 from [GOS24], Construction 5.2 is perfect zero knowledge.

Efficiency. The proof length and verifier complexity are clear from construction. For every query that \mathcal{V}_Σ makes to its input, the verifier in Construction 5.2 makes three queries to F and π_P . For every query \mathcal{V}_Σ makes to its proof, the verifier in Construction 5.2 makes one query to π_Σ . By Lemma 5.1, \mathcal{V}_Σ makes $O(|\mathbb{F}|)$ queries to its input and $|\mathbb{F}|$ queries to π , and the desired query complexity follows. \square

6 Robust PZK-PCPs for NP and NEXP

In this section we show how to build robust PZK-PCPs for NP and NEXP using a robust PZK-PCPP for sumcheck as a building block.

Definition 6.1 (Oracle 3-SAT). *Let $B: \{0, 1\}^{r+3s+3} \rightarrow \{0, 1\}$ be a 3-CNF. We say that B is implicitly satisfiable if there exists $A \in \{0, 1\}^s \rightarrow \{0, 1\}$ such that for all $z \in \{0, 1\}^r, b_1, b_2, b_3 \in \{0, 1\}^s$, $B(z, b_1, b_2, b_3, A(b_1), A(b_2), A(b_3)) = 1$. Let Oracle-3SAT be the language of implicitly satisfiable 3-CNFs.*

Theorem 6.2 (Cook-Levin). *For any $T(n)$ -time nondeterministic Turing machine M , $T = \Omega(n)$, there is a polynomial-time reduction R_M such that for any input $x \in \{0, 1\}^n$,*

$$R_M(x, T) \in \text{Oracle-3SAT} \Leftrightarrow \exists w, M(x, w) = 1.$$

Moreover, $R_M(x, T)$ is a formula in $O(\log T(n))$ variables of size $\text{poly}(n, \log T(n))$.

Theorem 6.3. *The following inclusions hold:*

- (i) *for any $q^* \leq 2^{\text{poly}(n)}$, $\text{NP} \subseteq \text{rPZK-PCP}_{q^*, \Sigma(n)}[\log n + \log q^*, \text{poly}(\log n + \log q^*)]$, where $|\Sigma(n)| = \text{poly}(n, q^*)$, and*
- (ii) *for any $q^* \leq 2^{\text{poly}(n)}$, $\text{NEXP} \subseteq \text{rPZK-PCP}_{q^*, \Sigma(n)}[\text{poly}(n) + \log q^*, \text{poly}(n)]$, where $|\Sigma(n)| = \text{poly}(2^n, q^*)$.*

Proof. Fix a language $\mathcal{L} \in \text{NTIME}(T)$, where for (i), $T = \text{poly}(n)$, and for (ii), $T = 2^{\text{poly}(n)}$. Let \mathbb{F} be a field of characteristic 2 of size to be fixed, H a subfield of \mathbb{F} of size $\Theta(\log T)$. We will fix an arithmetisation \hat{B} of $1 - B$ over \mathbb{F} ; i.e., for all $c \in \{0, 1\}^{r+3s+3}$, $\hat{B}(c) = 1 - B(c)$. The PCP constructions which prove (i) and (ii) differ only in the choice of \hat{B} . For (i), note that since $r + 3s + 3 = O(\log n)$, the evaluation table of B is computable in polynomial time. Hence we choose \hat{B} to be the unique multilinear extension of B ; this can be evaluated in polynomial time, and the total degree d_B of \hat{B} is $O(\log n)$. For (ii), we will take \hat{B} to be the function computed by the arithmetic formula corresponding to the boolean formula B ; this has total degree $d_B = O(|B|) = \text{poly}(n)$. In both cases we choose \mathbb{F} so that $d_B \ll |\mathbb{F}| = \text{poly}(\log T)$.

Let $m_1 := r / \log |H|$, $m_2 := s / \log |H|$. Let $\gamma_1: H^{m_1} \rightarrow \{0, 1\}^r$, $\gamma_2: H^{m_2} \rightarrow \{0, 1\}^s$ be the lexicographic orderings of the elements of H^{m_1}, H^{m_2} respectively. By [GKR15, Claim 4.2], for $i \in [2]$ the unique minimal-degree extension $\hat{\gamma}_i$ of γ_i has individual degree $|H| - 1$ and can be evaluated in time $\text{poly}(|H|, m_i, \log |\mathbb{F}|)$.

For a polynomial $\hat{A}: \mathbb{F}^{m_2} \rightarrow \mathbb{F}$, we define $g_{\hat{A}}: \mathbb{F}^{m_1+3m_2+3} \rightarrow \mathbb{F}$:

$$g_{\hat{A}}(z, b_1, b_2, b_3, a_1, a_2, a_3) := \hat{B}(\gamma_1(z), \gamma_2(b_1), \gamma_2(b_2), \gamma_2(b_3), a_1, a_2, a_3) \cdot \prod_{i=1}^3 (\hat{A}(b_i) + a_i - 1).$$

Next, for a polynomial $\hat{C}: \mathbb{F}^{m_2+k} \rightarrow \mathbb{F}$, we define⁴ $h_{\hat{C}}: \mathbb{F}^{m_1+3m_2+3+3k} \rightarrow \mathbb{F}$:

$$h_{\hat{C}}(z, b_1, b_2, b_3, a_1, a_2, a_3, c_1, c_2, c_3) := \hat{B}(\gamma_1(z), \gamma_2(b_1), \gamma_2(b_2), \gamma_2(b_3), a_1, a_2, a_3) \\ \cdot \prod_{i=1}^3 \left(\hat{C}(b_i, c_i) + L_{H^k, 0^k}(c_i) \cdot (a_i - 1) \right).$$

Observe that $h_{\hat{C}}$ is a polynomial of total degree $d = O((d_B + k) \cdot |H| + \deg(C))$ in $O(k + \log T / \log \log T)$ variables, and that if $\hat{A}(X) := \sum_{c \in H^k} \hat{C}(X, c)$,

$$\sum_{c_1, c_2, c_3 \in H^k} h_{\hat{C}}(z, b_1, b_2, b_3, a_1, a_2, a_3, c_1, c_2, c_3) \equiv g_{\hat{A}}(z, b_1, b_2, b_3, a_1, a_2, a_3). \quad (4)$$

We are now ready to specify the construction. First, we choose the parameter k . In the construction, $\deg(C)$ will be $O((\log T + k)|H|)$. Denote by Sim' the simulator guaranteed by Lemma 5.1. Let \tilde{q} be a polynomial such that, for all $((\mathbb{F}, m, d, H, \gamma), F) \in \text{Sum}$, and for all PCPP verifiers \mathcal{V}^* that make q^* queries to the input F and the proof, Sim' makes $p(q^*, m, d, |H|, \log |\mathbb{F}|)$ queries to F . Define $\tilde{q} := \tilde{q}(q^*, m_1 + 3m_2 + 3 + k, d, |H|, \log |\mathbb{F}|) = \text{poly}(q^*, \log T, k)$. Choose $k = O(\log q^* + \log \log T)$ large enough such that $k > \log \tilde{q} / \log |H|$.

⁴Here, since H is a subfield of \mathbb{F} , we cannot divide $a_i - 1$ by $|H|^k$ as suggested in Section 1.1. Instead we use a Lagrange polynomial so that the term $a_i - 1$ appears only once in the sum over H^k .

Construction 6.4. A PZK-PCP for Oracle-3SAT.

Proof:

1. Let $A: \{0, 1\}^s \rightarrow \{0, 1\}$ be a satisfying assignment for B . Choose a uniformly random polynomial $\hat{C}: \mathbb{F}^{m_2+k}$ of individual degree $2(|H| - 1)$ such that for all $b \in H^{m_2}$, $\sum_{c \in H^k} \hat{C}(b, c) = A(\gamma_2(b))$. Compute the full evaluation table π_C of \hat{C} .
2. For each $\vec{\tau} \in \mathbb{F}^{m_1+3m_2+3}$, let $\pi_{\vec{\tau}}$ be a PZK-PCPP (Lemma 5.1) for the claim

$$\sum_{\substack{z \in H^{m_1} \\ b_1, b_2, b_3 \in H^{m_2}}} \sum_{a_1, a_2, a_3 \in \{0, 1\}} \sum_{c_1, c_2, c_3 \in H^k} L_{H^{m_1+3m_2} \times \{0, 1\}^3, (z, \vec{b}, \vec{a})}(\vec{\tau}) \cdot h_{\hat{C}}(\vec{\tau}, c_1, c_2, c_3) = 0. \quad (5)$$

3. Output $(\pi_C, (\pi_{\vec{\tau}})_{\vec{\tau} \in \mathbb{F}^{m_1+3m_2+3}})$.

Verifier:

1. Perform a low total degree test (Theorem 2.16) on Π_C with $\varepsilon = 1/100$ and $\delta = 1/2$. If the test rejects, then reject.
2. Choose $\vec{\tau} = (\zeta, \nu_1, \nu_2, \nu_3, \xi) \in \mathbb{F}^{m_1} \times (\mathbb{F}^{m_2})^3 \times \mathbb{F}^3$ uniformly at random.
3. Run the verifier for the sumcheck PZK-PCPP on $\pi_{\vec{\tau}}$, yielding a claim

$$L_{H^{m_1+3m_2} \times \{0, 1\}^3, \beta}(\vec{\tau}) \cdot h_{\hat{C}}(\vec{\tau}, \eta_1, \eta_2, \eta_3) = \gamma \quad (6)$$

for some $\beta \in \mathbb{F}^{m_1+3m_2+3}$, $\eta_1, \eta_2, \eta_3 \in \mathbb{F}^{3k}$, $\gamma \in \mathbb{F}$. Query \hat{C} at (ν_i, η_i) for $i \in \{1, 2, 3\}$ in order to compute $h_{\hat{C}}(\vec{\tau}, \eta_1, \eta_2, \eta_3)$. Accept if (6) is true, else reject.

We “balance” the verifier’s queries by repeating Steps 1 and 3 sufficiently many times so that each test accounts for at least $1/3$ of the verifier’s view.

Proof length, alphabet size and query complexity. The proof length is dominated by the sumcheck proofs $\pi_{\vec{\tau}}$; each such proof is of size $|\mathbb{F}|^{O(m_1+m_2+k+3)} = \text{poly}(T(n), q^*)$, and there are $|\mathbb{F}|^{O(m_1+m_2+3)} = \text{poly}(T(n))$ of them. The proof alphabet is $\mathbb{F}^{m_1+m_2+k+4}$, so the alphabet is of size $|\mathbb{F}|^{O(m_1+m_2+k+3)} = \text{poly}(T(n), q^*)$. The query complexity is $\text{poly}(\log T(n), \log q^*)$.

Completeness and robust soundness. Completeness and soundness rely on the following key observation.

Claim 6.5. *The following are equivalent:*

- (i) $B \in \text{Oracle-3SAT}$;
- (ii) *there exists a polynomial $\hat{A}: \mathbb{F}^m \rightarrow \mathbb{F}$ such that for all $z \in H^{m_1}$, $b_1, b_2, b_3 \in H^{m_2}$, $a_1, a_2, a_3 \in \{0, 1\}$, $g_{\hat{A}}(z, b_1, b_2, b_3, a_1, a_2, a_3) = 0$;*
- (iii) *there exists a polynomial $\hat{C}: \mathbb{F}^{m_2+k} \rightarrow \mathbb{F}$ such that for all $z \in H^{m_1}$, $b_1, b_2, b_3 \in H^{m_2}$, $a_1, a_2, a_3 \in \{0, 1\}$, $\sum_{c_1, c_2, c_3 \in H^k} h_{\hat{C}}(z, b_1, b_2, b_3, a_1, a_2, a_3, c_1, c_2, c_3) = 0$.*

Proof. The equivalence of (iii) and (ii) is a direct consequence of Eq. 4. It remains to prove that (i) and (ii) are equivalent.

Note first that for any $z \in \{0, 1\}^r, b_1, b_2, b_3 \in \{0, 1\}^s, a_1, a_2, a_3 \in \{0, 1\}, g_{\hat{A}}(z, b_1, b_2, b_3, a_1, a_2, a_3) = 0$ if and only if either $B(z, b_1, b_2, b_3, a_1, a_2, a_3) = 1$ (since \hat{B} extends $1 - B$) or, for some $i \in [3], \hat{A}(b_i) = 1 - a_i$.

(i) \Rightarrow (ii). Suppose that $B \in \text{Oracle-3SAT}$, and let \hat{A} be the multilinear extension of a satisfying assignment A , and fix $z, b_1, b_2, b_3, a_1, a_2, a_3$. If $B(z, b_1, b_2, b_3, a_1, a_2, a_3) = 1$ then we are done, so suppose not. Then since A is a satisfying assignment, there exists $i \in [3]$ such that $\hat{A}(b_i) = A(b_i) \neq a_i$, so $\hat{A}(b_i) = 1 - a_i$ since $A(b_i) \in \{0, 1\}$.

(ii) \Rightarrow (i). Fix some polynomial \hat{A} such that $g_{\hat{A}}$ is zero on $\{0, 1\}^{r+3s+3}$. Let A be the assignment given by, for each $b \in \{0, 1\}^s$,

$$A(b) = \begin{cases} \hat{A}(b) & \text{if } \hat{A}(b) \in \{0, 1\}, \text{ or} \\ 0 & \text{otherwise.} \end{cases}$$

Fix $z \in \{0, 1\}^r, b_1, b_2, b_3 \in \{0, 1\}^s$; we show that $B(z, b_1, b_2, b_3, A(b_1), A(b_2), A(b_3)) = 1$. Indeed, let $a_1, a_2, a_3 \in \{0, 1\}$ be such that $B(z, b_1, b_2, b_3, a_1, a_2, a_3) = 0$; if no such assignment exists then we are already done. Since $B(z, b_1, b_2, b_3, a_1, a_2, a_3) = 0$, it must be that there exists $i \in [3]$ such that $\hat{A}(b_i) = 1 - a_i$; and since $1 - a_i \in \{0, 1\}, A(b_i) = \hat{A}(b_i)$. Hence $A(b_i) \neq a_i$. Thus if $A(b_i) = a_i$ for all i , then $B(z, b_1, b_2, b_3, a_1, a_2, a_3) = 1$, from which the implication follows. \square

For completeness, it suffices to observe that the claim (5) is true for \hat{C} as in Claim 6.5. We proceed to show robust soundness. Suppose that $B \notin \text{Oracle-3SAT}$.

Let $\varepsilon := 1/100$. First, suppose that π_C is ε -far from $\text{RM}[\mathbb{F}, m_2 + k, (m_2 + k)d]$. Then the verifier's view in the low-degree test is on average $\Omega(\varepsilon)$ -far from accepting by Theorem 2.16.

Otherwise, let $\hat{C} \in \text{RM}[\mathbb{F}, m_2 + k]$ be the unique closest codeword to π_C ; then $d(\hat{C}, \pi_C) \leq \varepsilon$. It follows that $d(h_{\hat{C}}, h_{\pi_C}) \leq 3\varepsilon$, since a uniformly random evaluation of $h_{\hat{C}}$ depends on three uniformly random evaluations of \hat{C} .

Since $B \notin \text{Oracle-3SAT}$, by Claim 6.5 there exists $z \in H^{m_1}, b_1, b_2, b_3 \in H^{m_2}, a_1, a_2, a_3 \in \{0, 1\}$ such that $\sum_{c_1, c_2, c_3 \in H^k} h_{\hat{C}}(z, b_1, b_2, b_3, a_1, a_2, a_3, c_1, c_2, c_3) \neq 0$. It follows that

$$\sum_{\substack{z \in H^{m_1} \\ b_1, b_2, b_3 \in H^{m_2}}} \sum_{a_1, a_2, a_3 \in \{0, 1\}} L_{H^{m_1+3m_2} \times \{0, 1\}^3, (z, \vec{b}, \vec{a})}(\vec{X}) \cdot \sum_{c_1, c_2, c_3 \in H^k} h_{\hat{C}}(\vec{X}, c_1, c_2, c_3)$$

is a nonzero polynomial in $m_1 + m_2 + 3$ variables \vec{X} of individual degree $O(d_B \cdot |H|)$. Thus, with probability $1 - O((m_1 + m_2 + 3)d_B \cdot |H|/|\mathbb{F}|)$ over the choice of $\vec{\tau}$, the claim in Eq. 6 is false. By the expected robustness guarantee of the PZK-PCPP (Lemma 5.1), the verifier's view of the sumcheck PCP in this case is $\Omega(\varepsilon)$ -far on average from an accepting view. By a union bound, the overall expected robustness is $\Omega(\varepsilon)$.

Zero knowledge. Our simulator makes use of an algorithm due to Ben-Sasson et al. [BCFGRS17] for efficiently lazily sampling random multivariate polynomials.

Lemma 6.6 ([BCFGRS17], Corollary 4.10). *There exists a probabilistic algorithm PolySim such that, for every finite field \mathbb{F} , $m, d \in \mathbb{N}$, set $S = \{(\alpha_1, \beta_1), \dots, (\alpha_\ell, \beta_\ell)\} \subseteq \mathbb{F}^m \times \mathbb{F}$, and $(\alpha, \beta) \in \mathbb{F}^m \times \mathbb{F}$,*

$$\Pr \left[\text{PolySim}(\mathbb{F}, m, d, S, \alpha) = \beta \right] = \Pr_{Q \leftarrow \mathbb{F}^{\leq d}[X_1, \dots, X_m]} \left[Q(\alpha) = \beta \mid \begin{array}{c} Q(\alpha_1) = \beta_1 \\ \vdots \\ Q(\alpha_\ell) = \beta_\ell \end{array} \right].$$

Moreover PolySim runs in time $\text{poly}(\log |\mathbb{F}|, m, d, \ell)$.

We describe the construction of the simulator below, and then prove that its output is distributed identically to $\text{View}_{\mathcal{V}^*, \mathcal{P}}$.

Construction 6.7. A simulator Sim for Construction 6.4.

On query $\vec{\alpha}$ to oracle \mathcal{O} :

1. If $\mathcal{O} = \pi_C$, sample $\beta \leftarrow \text{PolySim}(\mathbb{F}, m+k, d, S, \vec{q})$, where S is the set of prior queries to π_C and their answers, and respond with β .
2. If $\mathcal{O} = \pi_{\vec{r}}$ for some $\vec{r} \in \mathbb{F}^{m_1+3m_2+3}$:
 - (a) If this is the first query to $\pi_{\vec{r}}$, start a new instance $\text{Sim}'_{\vec{r}}$ of the sumcheck PZK-PCPP simulator (Lemma 5.1).
 - (b) Answer the query using $\text{Sim}'_{\vec{r}}$. This sub-simulator may make queries to the summand polynomial (5). Each such query can be efficiently answered by making 3 queries to π_C , which are answered as in Step 1.

We prove perfect zero knowledge via a hybrid argument. Consider a hybrid oracle simulator $\widetilde{\text{Sim}}$ which behaves as Sim except that it answers queries to π_C by querying an external oracle.

We consider the following sequence of hybrid algorithms:

- H_0 : $\text{Sim}^{\mathcal{V}^*}(x)$.
 - H_1 : $\widetilde{\text{Sim}}^{\mathcal{V}^*, Z}(x)$, where $Z: \mathbb{F}^{m_2+k}$ is a uniformly random polynomial of individual degree $|H| - 1$.
 - H_2 : $\widetilde{\text{Sim}}^{\mathcal{V}^*, Z}(x)$, where $Z: \mathbb{F}^{m_2+k}$ is sampled as \hat{C} in Step 1 of the prover algorithm.
 - H_3 : $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x)$.
- $H_0 \equiv H_1$. This follows directly from the correctness of PolySim (Lemma 6.6).
 - $H_1 \equiv H_2$. This claim relies crucially on the following lemma, originally shown in [CFS17; CFGS22].

Lemma 6.8 ([CFS17, Corollary 5.3]). *Let \mathbb{F} be a finite field, $H \subseteq \mathbb{F}$, $d, d' \in \mathbb{N}$ with $d' \geq 2(|H| - 1)$. Let $Q \subseteq \mathbb{F}^{m+k}$ with $|Q| < |H|^k$. Let Z be chosen uniformly at random in $\mathbb{F}[X_1, \dots, X_m, Y_1, \dots, Y_k]$ such that $\deg_{X_i}(Z) \leq d$ for $1 \leq i \leq m$ and $\deg_{Y_i}(Z) \leq d'$ for all $1 \leq i \leq k$.*

The ensembles $(\sum_{\vec{y} \in H^k} Z(\vec{\alpha}, \vec{y}))_{\vec{\alpha} \in \mathbb{F}^m}$ and $(Z(\vec{q}))_{\vec{q} \in Q}$ are independent.

Since \mathcal{V}^* makes at most q^* queries to the proof, $\widetilde{\text{Sim}}$ makes at most $p(q^*)$ queries to its oracle. Then since $p(q^*) < |H|^k$ by assumption, by Lemma 6.8, the answers to those queries are statistically independent from $(\sum_{\vec{y} \in H^k} Z(\vec{\alpha}, \vec{y}))_{\vec{\alpha} \in \mathbb{F}^m}$. Thus for all v ,

$$\begin{aligned} \Pr[H_2 \rightarrow v] &= \Pr_Z \left[\widetilde{\text{Sim}}^{\mathcal{V}^*, Z}(x) \rightarrow v \mid \forall b \in H^{m_2}, \sum_{c \in H^k} \hat{C}(b, c) = A(\gamma_2(b)) \right] \\ &= \Pr_Z \left[\widetilde{\text{Sim}}^{\mathcal{V}^*, Z}(x) \rightarrow v \right]. \end{aligned}$$

- $H_2 \equiv H_3$. This follows from the perfect zero knowledge guarantee of the sumcheck PZK-PCPP. \square

7 PZK-PCP for NP and NEXP with constant query complexity

In this section, we combine our zero-knowledge proof composition theorem and our randomness-efficient, robust PZK-PCPs for NP and NEXP to obtain constant-query PZK-PCPs for NP and NEXP.

Theorem 7.1. *The following inclusions hold:*

- (i) $\text{NP} \subseteq \text{PZK-PCP}[\log n, 1]$, and
- (ii) $\text{NEXP} \subseteq \text{PZK-PCP}[\text{poly}(n), 1]$.

Moreover, when q^* is polynomial, NP has efficient PZK-PCPs with constant query complexity and query bound q^* .

Proof. Part (i). We prove this via a chain of inclusions as follows.

We start by performing alphabet reduction on the robust PZK-PCP for NP we constructed in Theorem 6.3. Let $q^* \leq 2^{\text{poly}(n)}$ be an arbitrary query bound; for simplicity, we will assume without loss of generality that $q^* \geq n$. By part (i) of Theorem 6.3, for any $\tilde{q} \leq 2^{\text{poly}(n)}$,

$$\text{NP} \subseteq \text{rPZK-PCP}_{\tilde{q}, \Sigma(n)}[\log n + \log \tilde{q}, \text{poly}(\log n + \log \tilde{q})],$$

where $|\Sigma(n)| = \text{poly}(n, \tilde{q})$. This can be viewed as a PCP over the alphabet $\{0, 1\}^a$, where $a := \log |\Sigma(n)| = O(\log n + \log \tilde{q})$. Therefore, we can apply alphabet reduction (Lemma 3.5) to obtain

$$\text{rPZK-PCP}_{\tilde{q}, \Sigma(n)}[\log n + \log \tilde{q}, \text{poly}(\log n + \log \tilde{q})] \subseteq \text{rPZK-PCP}_{\tilde{q}, \{0,1\}}[\log n + \log \tilde{q}, \text{poly}(\log n + \log \tilde{q})].$$

Next, we perform proof composition to obtain constant query complexity. Let $Q(n) = \text{poly}(\log n + \log \tilde{q})$ be the query complexity of the boolean rPZK-PCP and set $\tilde{q}(n)$ equal to a polynomial in q^* and n large enough to ensure that $\tilde{q}(n)/Q(n) \geq q^*(n)$ for all $n \in \mathbb{N}$; $\tilde{q} = O((q^*)^{1+\varepsilon})$ suffices for any $\varepsilon > 0$. Then by Corollary 3.11, $\text{rPZK-PCP}_{\tilde{q}, \{0,1\}}[r, q] \subseteq \text{PZK-PCP}_{\tilde{q}/Q, \{0,1\}}[r + \log n, 1]$. Thus, we have that

$$\text{rPZK-PCP}_{\tilde{q}, \{0,1\}}[\log n + \log \tilde{q}, \text{poly}(\log n + \log \tilde{q})] \subseteq \text{PZK-PCP}_{q^*, \{0,1\}}[\log n + \log q^*, 1].$$

By chaining the above inclusions together we have that $\text{NP} \subseteq \text{PZK-PCP}_{q^*, \{0,1\}}[\log n + \log q^*, 1]$. Since q^* was arbitrary, $\text{NP} \subseteq \text{PZK-PCP}[\log n, 1]$.

Part (ii). Let $q^* \leq 2^{\text{poly}(n)}$ be an arbitrary query bound. By part (ii) of Theorem 6.3, for any $\tilde{q} \leq 2^{\text{poly}(n)}$,

$$\text{NEXP} \subseteq \text{rPZK-PCP}_{\tilde{q}, \Sigma(n)}[\text{poly}(n) + \log q^*, \text{poly}(n)],$$

where $|\Sigma(n)| = \text{poly}(2^n, \tilde{q})$. As before, this can be viewed as a PCP over the alphabet $\{0, 1\}^a$, where $a := \log |\Sigma(n)| = O(n + \log \tilde{q})$. Therefore, we can apply alphabet reduction (Lemma 3.5) to obtain

$$\text{rPZK-PCP}_{\tilde{q}, \Sigma(n)}[\text{poly}(n) + \log q^*, \text{poly}(n)] \subseteq \text{rPZK-PCP}_{\tilde{q}, \{0,1\}}[\text{poly}(n), \text{poly}(n, \log \tilde{q})].$$

As before, we perform proof composition to obtain constant query complexity. Let $Q(n) = \text{poly}(n)$ be the query complexity of the boolean rPZK-PCP and set $\tilde{q} := q^*(n) \cdot Q(n)$ for all $n \in \mathbb{N}$, so $\tilde{q} = \text{poly}(q^*, n)$. Then by Corollary 3.11, $\text{rPZK-PCP}_{\tilde{q}, \{0,1\}}[r, q] \subseteq \text{PZK-PCP}_{\tilde{q}/Q, \{0,1\}}[r + \log n, 1]$, so

$$\text{rPZK-PCP}_{\tilde{q}, \{0,1\}}[\text{poly}(n), \text{poly}(n, \log \tilde{q})] \subseteq \text{PZK-PCP}_{q^*, \{0,1\}}[\text{poly}(n) + \log q^*, 1].$$

Since q^* was arbitrary, $\text{NP} \subseteq \text{PZK-PCP}[\text{poly}, 1]$. □

References

- [ABN23] A. Anshu, N. P. Breuckmann, and C. Nirkhe. “NLTS Hamiltonians from good quantum codes”. In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 2023, pp. 1090–1096.
- [AE15] D. Aharonov and L. Eldar. “Quantum locally testable codes”. In: *SIAM Journal on Computing* 44.5 (2015), pp. 1230–1262.
- [ALMSS92] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. “Proof verification and hardness of approximation problems”. In: *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*. 1992, pp. 14–23.
- [Al99] N. Alon. “Combinatorial Nullstellensatz”. In: *Combinatorics, Probability and Computing* 8 (1999), pp. 7–29.
- [AS98] S. Arora and S. Safra. “Probabilistic checking of proofs: a new characterization of NP”. In: *Journal of the ACM* 45.1 (1998). Preliminary version in FOCS ’92., pp. 70–122.
- [BCFGRS17] E. Ben-Sasson, A. Chiesa, M. A. Forbes, A. Gabizon, M. Riabzev, and N. Spooner. “Zero Knowledge Protocols from Succinct Constraint Detection”. In: *Proceedings of the 15th Theory of Cryptography Conference*. TCC ’17. 2017, pp. 172–206.
- [BFLS91] L. Babai, L. Fortnow, L. A. Levin, and M. Szegedy. “Checking computations in polylogarithmic time”. In: *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*. STOC ’91. 1991, pp. 21–32.
- [BGHSV05] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. Vadhan. “Short PCPs Verifiable in Polylogarithmic Time”. In: *Proceedings of the 20th Annual IEEE Conference on Computational Complexity*. CCC ’05. 2005, pp. 120–134.
- [BGHSV06] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. P. Vadhan. “Robust PCPs of Proximity, Shorter PCPs, and Applications to Coding”. In: *SIAM Journal on Computing* 36.4 (2006), pp. 889–974.
- [BGKW88] M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson. “Multi-prover interactive proofs: how to remove intractability assumptions”. In: *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*. STOC ’88. 1988, pp. 113–131.
- [BS08] E. Ben-Sasson and M. Sudan. “Short PCPs with Polylog Query Complexity”. In: *SIAM Journal on Computing* 38.2 (2008). Preliminary version appeared in STOC ’05., pp. 551–607.
- [CCGOS23] M. Chen, A. Chiesa, T. Gur, J. O’Connor, and N. Spooner. “Proof-carrying data from arithmetized random oracles”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’23. 2023, pp. 379–404.
- [CDGH23] G. Cormode, M. Dall’Agnol, T. Gur, and C. Hickey. “Streaming zero-knowledge proofs”. In: *arXiv preprint arXiv:2301.02161* (2023).
- [CFG18] A. Chiesa, M. A. Forbes, T. Gur, and N. Spooner. “Spatial Isolation Implies Zero Knowledge Even in a Quantum World”. In: *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science*. FOCS ’18. 2018, pp. 755–765.
- [CFG22] A. Chiesa, M. A. Forbes, T. Gur, and N. Spooner. “Spatial Isolation Implies Zero Knowledge Even in a Quantum World”. In: *Journal of the ACM* 69.2 (2022), pp. 1–44.

- [CFS17] A. Chiesa, M. A. Forbes, and N. Spooner. *A Zero Knowledge Sumcheck and its Applications*. Cryptology ePrint Archive, Report 2017/305. 2017.
- [Din07] I. Dinur. “The PCP theorem by gap amplification”. In: *Journal of the ACM* 54.3 (2007), p. 12.
- [DLV24] I. Dinur, T.-C. Lin, and T. Vidick. “Expansion of higher-dimensional cubical complexes with application to quantum locally testable codes”. In: *arXiv preprint arXiv:2402.07476* (2024).
- [GKR15] S. Goldwasser, Y. T. Kalai, and G. N. Rothblum. “Delegating Computation: Interactive Proofs for Muggles”. In: *Journal of the ACM* 62.4 (2015), 27:1–27:64.
- [GMR89] S. Goldwasser, S. Micali, and C. Rackoff. “The knowledge complexity of interactive proof systems”. In: *SIAM Journal on Computing* 18.1 (1989). Preliminary version appeared in STOC ’85., pp. 186–208.
- [GOS24] T. Gur, J. O’Connor, and N. Spooner. “Perfect Zero-Knowledge PCPs for #P”. In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024, pp. 1724–1730.
- [GRR20] T. Gur, G. Ramnarayan, and R. Rothblum. “Relaxed locally correctable codes”. In: *Theory of Computing* 16.1 (2020), pp. 1–68.
- [GSY19] A. B. Grilo, W. Slofstra, and H. Yuen. “Perfect Zero Knowledge for Quantum Multiprover Interactive Proofs”. In: *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019*. IEEE Computer Society, 2019, pp. 611–635.
- [HVW22] C. Hazay, M. Venkatasubramanian, and M. Weiss. “ZK-PCPs from Leakage-Resilient Secret Sharing”. In: *J. Cryptol.* 35.4 (2022), p. 23.
- [IMS12] Y. Ishai, M. Mahmoody, and A. Sahai. “On Efficient Zero-Knowledge PCPs”. In: *Proceedings of the 9th Theory of Cryptography Conference on Theory of Cryptography*. TCC ’12. 2012, pp. 151–168.
- [IMSX15] Y. Ishai, M. Mahmoody, A. Sahai, and D. Xiao. *On Zero-Knowledge PCPs: Limitations, Simplifications, and Applications*. Available at <http://www.cs.virginia.edu/~mohammad/files/papers/ZKPCPs-Full.pdf>. 2015.
- [ISVW13] Y. Ishai, A. Sahai, M. Viderman, and M. Weiss. “Zero Knowledge LTCs and Their Applications”. In: *Proceedings of the 16th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, and of the 17th International Workshop on Randomization and Computation*. APPROX-RANDOM ’13. 2013, pp. 607–622.
- [IW14] Y. Ishai and M. Weiss. “Probabilistically Checkable Proofs of Proximity with Zero-Knowledge”. In: *Proceedings of the 11th Theory of Cryptography Conference*. TCC ’14. 2014, pp. 121–145.
- [IWY16] Y. Ishai, M. Weiss, and G. Yang. “Making the Best of a Leaky Situation: Zero-Knowledge PCPs from Leakage-Resilient Circuits”. In: *Proceedings of the 13th Theory of Cryptography Conference*. TCC ’16-A. 2016, pp. 3–32.
- [KPT97] J. Kilian, E. Petrank, and G. Tardos. “Probabilistically checkable proofs with zero knowledge”. In: *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*. STOC ’97. 1997, pp. 496–505.

- [KR08] Y. Kalai and R. Raz. “Interactive PCP”. In: *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*. ICALP ’08. 2008, pp. 536–547.
- [LLZ22] A. Leverrier, V. Londe, and G. Zémor. “Towards local testability for quantum coding”. In: *Quantum* 6 (2022), p. 661.
- [MS24] K. Mastel and W. Slofstra. “Two Prover Perfect Zero Knowledge for MIP”. In: *STOC*. ACM, 2024, pp. 991–1002.
- [Par21] O. Paradise. “Smooth and strong PCPs”. In: *computational complexity* 30.1 (2021), p. 1.
- [Spi96] D. A. Spielman. “Linear-time encodable and decodable error-correcting codes”. In: *IEEE Transactions on Information Theory* 42.6 (1996). Preliminary version appeared in STOC ’95., pp. 1723–1731.
- [Wei22] M. Weiss. “Shielding Probabilistically Checkable Proofs: Zero-Knowledge PCPs from Leakage Resilience”. In: *Entropy* 24.7 (2022), p. 970.