

Quasi-Linear Size Zero Knowledge from Linear-Algebraic PCPs

Eli Ben-Sasson

eli@cs.technion.ac.il

Technion

Alessandro Chiesa

alexch@berkeley.edu

UC Berkeley

Ariel Gabizon

ariel.gabizon@gmail.com

Technion

Madars Virza

madars@mit.edu

MIT

January 9, 2016

Abstract

The seminal result that every language having an interactive proof also has a zero-knowledge interactive proof assumes the existence of one-way functions. Ostrovsky and Wigderson (ISTCS 1993) proved that this assumption is necessary: if one-way functions do not exist, then only languages in BPP have zero-knowledge interactive proofs.

Ben-Or et al. (STOC 1988) proved that, nevertheless, every language having a multi-prover interactive proof also has a zero-knowledge multi-prover interactive proof, unconditionally. Their work led to, among many other things, a line of work studying zero knowledge without intractability assumptions. In this line of work, Kilian, Petrank, and Tardos (STOC 1997) defined and constructed zero-knowledge probabilistically checkable proofs (PCPs).

While PCPs with quasilinear-size proof length, but without zero knowledge, are known, no such result is known for zero knowledge PCPs. In this work, we show how to construct “2-round” PCPs that are zero knowledge and of length $\tilde{O}(K)$ where K is the number of queries made by a malicious polynomial time verifier. Previous solutions required PCPs of length at least K^6 to maintain zero knowledge. In this model, which we call *duplex PCP* (DPCP), the verifier first receives an oracle string from the prover, then replies with a message, and then receives another oracle string from the prover; a malicious verifier can make up to K queries in total to both oracles.

Deviating from previous works, our constructions do not invoke the PCP Theorem as a blackbox but instead rely on certain algebraic properties of a specific family of PCPs. We show that if the PCP has a certain linear algebraic structure — which many central constructions can be shown to possess, including [BFLS91,ALMSS98,BS08] — we can add the zero knowledge property at virtually no cost (up to additive lower order terms) while introducing only minor modifications in the algorithms of the prover and verifier. We believe that our linear-algebraic characterization of PCPs may be of independent interest, as it gives a simplified way to view previous well-studied PCP constructions.

Keywords: zero knowledge, probabilistically-checkable proofs

Contents

1	Introduction	3
1.1	Our contributions	4
2	Preliminaries	6
2.1	Probabilistically checkable proofs	6
2.2	Probabilistically checkable proofs of proximity	7
2.3	Zero knowledge PCPs	7
2.4	Reed–Muller and Reed–Solomon codes	8
3	Duplex PCPs	10
4	Main theorem	11
4.1	Proof sketch	11
4.2	Roadmap of the rest of the paper	12
5	Linear algebraic CSPs and their canonical PCPs	14
5.1	Linear algebraic constraint satisfaction problems	14
5.2	A canonical PCP for linear algebraic CSPs	14
6	Zero-knowledge duplex PCPs from randomizable linear algebraic CSPs	17
6.1	Randomizable linear algebraic CSPs	17
6.2	Construction of zero-knowledge duplex PCPs	17
7	From NTIME to randomizable linear algebraic CSPs	22
7.1	Algebraic problems and group preservation	22
7.2	Algebraic problems naturally reduce to linear algebraic CSPs	23
7.3	From group-preserving algebraic problems to randomizable linear algebraic CSPs	24
7.4	An efficient reduction from NTIME to group-preserving algebraic problems	25
7.5	Combining the two reductions	27
8	Proof of Theorem 4.1	30
	Acknowledgments	32
	References	33

1 Introduction

We continue the study of proof systems that provide soundness and zero knowledge, simultaneously and unconditionally (i.e., no intractability assumptions are needed to achieve the two), as we now explain.

Interactive proofs. An *interactive proof* [BM88, GMR85] for a language \mathcal{L} is a pair of interactive algorithms (P, V) , where P is known as the prover and V as the verifier, that satisfies the following: (i) (completeness) for every instance x in \mathcal{L} , $P(x)$ can make $V(x)$ accept with probability 1; (ii) (soundness) for every instance x not in \mathcal{L} , every prover \tilde{P} can make $V(x)$ accept with at most a small probability ϵ . Shamir [Sha92] showed the expressive power of interactive proofs by proving that $\mathbf{IP} = \mathbf{PSPACE}$, i.e., all and only languages in \mathbf{PSPACE} have interactive proofs.

Zero knowledge. An interactive proof is *zero knowledge* [GMR85] if the verifier, even if malicious, cannot learn any information about an instance x in \mathcal{L} , by interacting with the prover, besides the fact x is in \mathcal{L} : for any efficient verifier \tilde{V} there exists an efficient simulator S such that $S(x)$ is “indistinguishable” from the view of \tilde{V} while interacting with $P(x)$. Depending on the choice of definition for indistinguishability, one gets different flavors of zero knowledge.

If indistinguishability is required to hold for efficient deciders only, then one gets *computational zero knowledge*; \mathbf{CZK} denotes the corresponding complexity class. A seminal result in cryptography says that if one-way functions exist then $\mathbf{CZK} = \mathbf{IP}$, i.e., every language having an interactive proof also has a computational zero-knowledge interactive proof [GMR85, IY87, BGG⁺88]. If indistinguishability is required to hold for all deciders, then one gets *statistical zero knowledge*; if instead the simulator’s output and the verifier’s view are the same distribution (and not merely close to each other), then one gets *perfect zero knowledge*. These stronger notions determine the corresponding complexity classes \mathbf{SZK} and \mathbf{PZK} , both of which are contained in $\mathbf{AM} \cap \mathbf{coAM}$; of course, $\mathbf{PZK} \subseteq \mathbf{SZK} \subseteq \mathbf{CZK}$.

Unfortunately, zero knowledge cannot be achieved unconditionally for non-trivial languages: Ostrovsky and Wigderson [OW93] proved that if one-way functions do not exist then \mathbf{CZK} equals an average-case variant of \mathbf{BPP} .

Other types of proof systems. Due to the limitations of interactive proofs with respect to zero knowledge that holds unconditionally, researchers have explored other types of proof systems, as an alternative to interactive proofs.

- **MIP.** Ben-Or et al. [BOGKW88] first studied statistical zero knowledge, and proved that it can be achieved in a new model, *multi-prover interactive proof* (MIPs), where the verifier interacts with multiple provers that are not allowed to communicate while interacting with the verifier (though they may share a random string before such an interaction begins). More precisely, Ben-Or et al. prove that every language having a multi-prover interactive proof also has a perfect zero-knowledge multi-prover interactive proof (again, without relying on intractability assumptions). The result of [BOGKW88] was subsequently improved in a number of papers [LS95, BFL91, DFK⁺92].
- **PCP.** Kilian et al. [KPT97] study statistical zero knowledge in the model of *probabilistically checkable proofs* (PCPs) [BFLS91, AS98, ALM⁺98], where the verifier has oracle access to a string. Essentially, the oracle string can be thought of as a stateless prover: the answer to a query depends only on the query itself, but not any other queries that were previously made. Building on results implicit in [DFK⁺92], Kilian et al. showed two main theorems. First, every language in \mathbf{NEXP} has a PCP that is statistical zero knowledge against verifiers that make at most any polynomial number of queries to the PCP. Second, every language in \mathbf{NP} has, for every constant $c > 0$, a PCP that is statistically zero knowledge against verifiers that make at most $k(n) := n^c$ queries to the PCP.

Subsequent works [IKOS09, IMS12, MX13, IMSX15] provided simplifications (giving alternative constructions or simplifying that of [KPT97]) and limitations (showing that for languages in \mathbf{NP} one cannot efficiently sample the oracle if one seeks statistical zero knowledge against verifiers that make at most a polynomial number of queries).

- **IPCP.** Goyal et al. [GIMS10] study statistical zero knowledge in the model of *interactive PCPs* (IPCPs) [KR08], where the verifier interacts with two provers of which one is restricted to be an oracle. Goyal et al. prove that every language in \mathbf{NP} has a constant-round interactive PCP that is statistical zero knowledge against verifiers that make at most any polynomial number of queries to the PCP, and where both provers’ strategies can be implemented efficiently as a function of the instance and the witness.

Researchers have also studied other proof systems in which various forms of zero knowledge hold unconditionally; these include *linear IPs* [BCI⁺13], *PCPs of proximity* [IW14], and others.

A limitation of prior work. PCPs with quasilinear-size proof length, but without zero knowledge, are known: for every language \mathcal{L} in $\text{NTIME}(T(n))$, there is a PCP with proof length $\tilde{O}(T(n))$ and query complexity $O(1)$ [Din07, BS08, BGH⁺06, Mie08]. On the other hand, no such result for statistical zero knowledge PCPs is known: even when applied to PCPs of length $\tilde{O}(T(n))$, [KPT97]’s result and followup improvements yields a proof length that is polynomial in $T(n) \cdot k(n)$, where $k(n)$, known as the *knowledge bound*, is a bound on the number of queries by any verifier (see Section 4.1 for further discussion). We thus ask the following question: are there statistical zero knowledge PCPs with proof length quasilinear in $T(n) + k(n)$?

1.1 Our contributions

We do not answer the above question in the PCP model, but we give a positive answer in a closely related model that can be thought of as a “2-round PCP”, which we call *duplex PCP* (DPCP). At a high level, a DPCP works as follows: the prover first sends an oracle string π_0 to the verifier, just as in a PCP; then, the verifier sends a message ρ to the prover; finally, the prover answers with a second oracle string π_1 ; the verifier may query both oracles, and then accept or reject. In other words, a DPCP is merely a 2-round interactive proof in which the prover sends oracle strings rather than messages. We prove the following theorem:

Theorem 1.1 (see Theorem 4.1 for formal statement). *For every language \mathcal{L} in $\text{NTIME}(T) \cap \text{NP}$ and polynomially-bounded knowledge bound k there exists a DPCP system satisfying the following:*

- *the proof length (in fact, also the prover running time) is quasilinear in $n + T(n) + k(n)$;*
- *the query complexity is polynomial in $\log(T(n) + k(n))$;*
- *the verifier running time is polynomial in $n + \log(T(n) + k(n))$;*
- *perfect zero knowledge holds against any verifier that makes at most $k(n)$ adaptive queries (in total to both oracles);*
- *the soundness error is $1/2$ (and can be reduced by repetition to $2^{-\lambda}$ while preserving perfect zero knowledge, provided that the number of queries does not exceed $k(n)$).*

Moreover, similarly to the PCPs of [KPT97], the DPCP system that we construct is in fact not only sound but is also a *proof of knowledge* [BG93]; however, in contrast to [KPT97], the DPCP verifier is *non-adaptive*, in the sense that the query locations depend only on the verifier’s random tape.

Perhaps the main difference between our construction and prior work is the techniques that we use. While previous works use the PCP Theorem as a black box, compiling a PCP into a zero knowledge PCP by using *locking schemes* [KPT97], we use certain *algebraic* properties of a specific family of PCPs to guarantee zero knowledge. In comparison to the generic approach, we are more specific, but the addition of zero knowledge essentially comes “for free” when compared to the corresponding constructions without zero knowledge. (In contrast, [KPT97] achieves a proof length of $\Omega(k(n)^6 \cdot l(n)^c)$, for some large enough c , when starting from a PCP with proof length $l(n)$.)

DPCP vs IPCP. Duplex PCPs are an alternative to interactive PCPs that combine PCPs and interaction. In a DPCP, the verifier gets an oracle string from the prover, replies with a message, and then gets another oracle string from the prover; in an IPCP, the verifier gets an oracle string from the prover, and then engages in an interactive proof with him.

Both [GIMS10] and our work are similar in that both address aspects that we do not know how to address in the PCP model, and resort to studying alternative models, i.e., IPCP and DPCP respectively. The two works however give different flavors of results: [GIMS10] obtain IPCPs that are zero knowledge against verifiers that ask at most any polynomial number of queries $k(n)$ but their oracle is of polynomial size in $k(n)$ (actually, of exponential size but with a polynomial-size circuit describing it); on the other hand, our work obtains DPCPs that are zero knowledge against verifiers that ask at most a fixed polynomial number of queries $k(n)$ and our oracles are of quasilinear size in $k(n)$.

Finally, we note that our construction can be also cast as an IPCP, because the knowledge bound $k(n)$ holds only for the first oracle, i.e., perfect zero knowledge is preserved even if the verifier reads the second oracle in full. This provides a result on a 2-round IPCP incomparable to [GIMS10]’s 4-round IPCP.

On the minimal computational gap between prover and verifier needed for zero knowledge. IPs and MIPs assume a computational gap between prover and verifier. The prover is allowed (and often assumed) to be computationally unbounded and the verifier is polynomially bounded. An intriguing corollary of our theorem is that the computational gap between prover and verifier can be drastically reduced, to a mere polylogarithmic one. Namely, suppose that we wish to create zero-knowledge proof systems in which the verifier runs in time $\text{tv}(n)$; in the model

above, as long as $\text{tp}(n) > \text{tv}(n) \cdot (\log \text{tv}(n))^c$ for an absolute constant c , then perfect zero knowledge with a small soundness error can be obtained under no intractability assumptions. (See Corollary 4.2 for a formal statement.)

2 Preliminaries

Functions and distributions. We use $f: D \rightarrow R$ to denote a function with domain D and range R ; given a subset \tilde{D} of D , we use $f|_{\tilde{D}}$ to denote the restriction of f to \tilde{D} . Given a distribution \mathcal{D} , we write $x \leftarrow \mathcal{D}$ to denote that x is sampled according to \mathcal{D} .

Distances. A distance measure is a function $\Delta: \Sigma^n \times \Sigma^n \rightarrow [0, 1]$ such that for all $x, y, z \in \Sigma^n$: (i) $\Delta(x, x) = 0$, (ii) $\Delta(x, y) = \Delta(y, x)$, and (iii) $\Delta(x, y) \leq \Delta(x, z) + \Delta(z, y)$. For example, the *relative Hamming distance* over alphabet Σ is a distance measure: $\Delta_{\Sigma}^{\text{Ham}}(x, y) := |\{i \mid x_i \neq y_i\}|/n$. We extend Δ to distances of strings to sets: given $x \in \Sigma^n$ and $S \subseteq \Sigma^n$, we define $\Delta(x, S) := \min_{y \in S} \Delta(x, y)$ (or 1 if S is empty). We say that a string x is ϵ -close to another string y if $\Delta(x, y) \leq \epsilon$, and ϵ -far from y if $\Delta(x, y) > \epsilon$; similar terminology applies for a string x and a set S .

Fields and polynomials. We denote by \mathbb{F} a finite field, by \mathbb{F}_q the field of size q , and by \mathcal{F} the set of all finite fields. We denote by $\mathbb{F}[X_1, \dots, X_m]$ the ring of polynomials in m variables over \mathbb{F} ; given a polynomial P in $\mathbb{F}[X_1, \dots, X_m]$, $\deg_{X_i}(P)$ is the degree of P in the variable X_i ; the total degree of P is the sum of all of these individual degrees. For any set $S \subseteq \mathbb{F}$ let $Z_S(X)$ be the unique monic polynomial of degree $|S|$ that vanishes on S : $Z_S(X) := \prod_{\alpha \in S} (X - \alpha)$.

Linear spaces. Given $n \in \mathbb{N}$, a subset S of \mathbb{F}^n is an \mathbb{F} -linear space if $\alpha x + \beta y \in S$ for all $\alpha, \beta \in \mathbb{F}$ and $x, y \in S$.

Languages and relations. We denote by \mathcal{R} a relation consisting of pairs $(\mathfrak{x}, \mathfrak{w})$, where \mathfrak{x} is the *instance* and \mathfrak{w} is the *witness*. We denote by $\text{Lan}(\mathcal{R})$ the language corresponding to \mathcal{R} , and by $\mathcal{R}|_{\mathfrak{x}}$ the set of witnesses in \mathcal{R} for \mathfrak{x} .

Complexity classes. We write complexity classes in bold capital letters: **NP**, **PSPACE**, **NEXP**, and so on. We take a “relation-centric” point of view: we view **NTIME** as a class of relations rather than as the class of the corresponding languages; we thus may write things like “let \mathcal{R} be in **NP**”. If \mathcal{R} is in **NTIME**(T), we fix an arbitrary machine $M_{\mathcal{R}}$ that decides \mathcal{R} in time $T(n)$, i.e., $M_{\mathcal{R}}(\mathfrak{x}, \mathfrak{w})$ always halts after $T(|\mathfrak{x}|)$ steps and $M_{\mathcal{R}}(\mathfrak{x}, \mathfrak{w}) = 1$ if and only if $(\mathfrak{x}, \mathfrak{w}) \in \mathcal{R}$; we then say that $M_{\mathcal{R}}$ decides \mathcal{R} (or $\text{Lan}(\mathcal{R})$). Throughout, we assume that $T(n) \geq n$.

Codes. An error correcting code C is a set of functions $w: H \rightarrow \Sigma$, where H, Σ are finite sets. The message length of C is $n := \log_{|\Sigma|} |C|$, its block length is $\ell := |H|$, its rate is $\rho := n/\ell$, its (minimum) distance is $d := \min\{\Delta(w, z) \mid w, z \in C, w \neq z\}$ when Δ is the (absolute) Hamming distance, and its (minimum) relative distance is $\delta := d/\ell$. Given a code family \mathcal{C} , we denote by $\text{Rel}(\mathcal{C})$ the relation that naturally corresponds to \mathcal{C} , i.e., $\{(C, w) \mid C \in \mathcal{C}, w \in C\}$. A code C is linear if Σ is a finite field and C is a Σ -linear space in Σ^{ℓ} ; we denote by $\dim(C)$ the dimension of C when viewed as a linear space. A code C is t -wise independent if, for every subset I of $[\ell]$ with cardinality t , the distribution of $w|_I$ (viewed as a string) for a random $w \in C$ equals the uniform distribution on Σ^t .

Random shifts. We later use the following folklore claim about distance preservation for random shifts in linear spaces; for completeness, we include its short proof.

Claim 2.1. *Let n be in \mathbb{N} , \mathbb{F} a finite field, S an \mathbb{F} -linear space in \mathbb{F}^n , and $x, y \in \mathbb{F}^n$. If x is ϵ -far from S , then $\alpha x + y$ is $\epsilon/2$ -far from S , with probability $1 - |\mathbb{F}|^{-1}$ over a random $\alpha \in \mathbb{F}$. (Distances are relative Hamming distances.)*

Proof. Suppose, by way of contradiction, that there exist $\alpha_1, \alpha_2 \in \mathbb{F}$ and $y_1, y_2 \in S$ with $\alpha_1 \neq \alpha_2$ such that, for every $i \in \{1, 2\}$, $\alpha_i x + y_i$ is $\epsilon/2$ close to S . Then, by the triangle inequality, $z := y_1 - y_2$ is ϵ -close to $(\alpha_1 x + y_1) - (\alpha_2 x + y_2) = (\alpha_1 - \alpha_2)x$. We conclude that x is ϵ -close to $\frac{1}{\alpha_1 - \alpha_2} z \in S$, a contradiction. \square

2.1 Probabilistically checkable proofs

A *PCP system* [BFLS91, AS98, ALM⁺98] for a relation \mathcal{R} is a tuple $\text{PCP} = (P, V)$ that works as follows.

- The *prover* P is a probabilistic algorithm that, given as input an instance-witness pair $(\mathfrak{x}, \mathfrak{w})$ with $n := |\mathfrak{x}|$, outputs a proof $\pi: D(n) \rightarrow \Sigma(n)$, where both $D(n)$ and $\Sigma(n)$ are finite sets.
- The *verifier* V is a probabilistic oracle algorithm that, given as input an instance \mathfrak{x} with $n := |\mathfrak{x}|$ and with oracle access to a proof $\pi: D(n) \rightarrow \Sigma(n)$, queries π at a few locations and then outputs a bit.

The system PCP has (perfect) completeness and soundness error $e(n)$ if the following two conditions hold. (Below, we explicitly denote the prover’s and verifier’s randomness as r_P and r_V .)

Completeness: For every instance-witness pair (\mathbf{x}, \mathbf{w}) in the relation \mathcal{R} , $\Pr_{r_P, r_V} [V^{P(\mathbf{x}, \mathbf{w}; r_P)}(\mathbf{x}; r_V) = 1] = 1$.

Soundness: For every instance \mathbf{x} not in the language $\text{Lan}(\mathcal{R})$ and proof $\pi: D(n) \rightarrow \Sigma(n)$, $\Pr_{r_V} [V^\pi(\mathbf{x}; r_V) = 1] \leq e(n)$.

A relation \mathcal{R} belongs to the complexity class **PCP** $[a, l, q, e, \text{tp}, \text{tv}]$ if there is a PCP system for \mathcal{R} in which:

- the answer alphabet (i.e., $\Sigma(n)$) is $a(n)$,
- the proof length over that alphabet (i.e., $|D(n)|$) is at most $l(n)$,
- the verifier queries the proof in at most $q(n)$ locations,
- the soundness error is $e(n)$,
- the prover runs in time $\text{tp}(n)$, and
- the verifier runs in time $\text{tv}(n)$.

Finally, we write $\text{na} = \text{Yes}$ in the square brackets if the queries to the oracles are non-adaptive (i.e., the queried locations only depend on the verifier's inputs); otherwise we write $\text{na} = \text{No}$.

2.2 Probabilistically checkable proofs of proximity

A *PCPP system* [BGH⁺04, DR04] for a relation \mathcal{R} is a tuple $\text{PCPP} = (P, V)$ that works as follows.

- The *prover* P is a probabilistic algorithm that, given as input an instance-witness pair (\mathbf{x}, \mathbf{w}) with $n := |\mathbf{x}|$, outputs a proof $\pi: D(n) \rightarrow \Sigma(n)$, where both $D(n)$ and $\Sigma(n)$ are finite sets.
- The *verifier* V is a probabilistic oracle algorithm that, given as input an instance \mathbf{x} with $n := |\mathbf{x}|$ and with oracle access to a witness \mathbf{w} and proof $\pi: D(n) \rightarrow \Sigma(n)$, queries \mathbf{w} and π at a few locations and then outputs a bit.

The system PCPP has (perfect) completeness, soundness error e , distance measure Δ , and proximity parameter d if the following two conditions hold. (Below, we explicitly denote the prover's and verifier's randomness as r_P and r_V .)

Completeness: For every instance-witness pair (\mathbf{x}, \mathbf{w}) in the relation \mathcal{R} , $\Pr_{r_P, r_V} [V^{(\mathbf{w}, P(\mathbf{x}, \mathbf{w}; r_P))}(\mathbf{x}; r_V) = 1] = 1$.

Soundness: For every instance-witness pair (\mathbf{x}, \mathbf{w}) , perhaps not in the language, such that $\Delta(\mathbf{w}, \mathcal{R}|_{\mathbf{x}}) \geq d(n)$ and proof $\pi: D(n) \rightarrow \Sigma(n)$, $\Pr_{r_V} [V^{(\mathbf{w}, \pi)}(\mathbf{x}; r_V) = 1] \leq e(n)$.

A relation \mathcal{R} belongs to the complexity class **PCPP** $[a, l, q, \Delta, d, e, \text{tp}, \text{tv}]$ if there is a PCPP system for \mathcal{R} in which:

- the answer alphabet (i.e., $\Sigma(n)$) is $a(n)$,
- the proof length over that alphabet (i.e., $|D(n)|$) is at most $l(n)$,
- the verifier queries the two oracles (codeword and proof) in at most $q(n)$ locations (in total),
- the distance measure is Δ ,
- the proximity parameter is $d(n)$,
- the soundness error is $e(n)$,
- the prover runs in time $\text{tp}(n)$, and
- the verifier runs in time $\text{tv}(n)$.

Finally, we write $\text{na} = \text{Yes}$ in the square brackets if the queries to the oracles are non-adaptive (i.e., the queried locations only depend on the verifier's inputs); otherwise we write $\text{na} = \text{No}$.

2.3 Zero knowledge PCPs

The notion of zero knowledge for PCPs was first considered in [DFK⁺92, KPT97]. A PCP system $\text{PCP} = (P, V)$ for a relation \mathcal{R} has *perfect zero knowledge with knowledge bound* k if there exists an expected-polynomial-time probabilistic algorithm S such that, for every k -query polynomial-time probabilistic oracle algorithm \tilde{V} , the following two distribution families are identical:

$$\{S(\tilde{V}, \mathbf{x})\}_{(\mathbf{x}, \mathbf{w}) \in \mathcal{R}} \quad \text{and} \quad \{\text{PCPView}(\tilde{V}, P, \mathbf{x}, \mathbf{w})\}_{(\mathbf{x}, \mathbf{w}) \in \mathcal{R}},$$

where $\text{PCPView}(\tilde{V}, \pi, \mathbf{x}, w)$ is the view of \tilde{V} in its execution when given input \mathbf{x} and oracle access to $\pi := P(\mathbf{x}, w)$. The definition of statistical and computational zero knowledge (with knowledge bound k) are similar: rather than identical, the two distribution families are required to be statistically and computationally close (as $|\mathbf{x}|$ grows), respectively.

A relation \mathcal{R} belongs to the complexity class $\text{PCP}_{\text{pzk}}[a, l, q, e, \text{tp}, \text{tv}, k]$ if there exists a PCP system for \mathcal{R} that (i) puts \mathcal{R} in $\text{PCP}[a, l, q, e, \text{tp}, \text{tv}]$, and (ii) has perfect zero knowledge with knowledge bound k ; as for PCP , we add the symbol na in the square brackets of PCP_{pzk} to denote if the queries to the proof are non-adaptive. The complexity classes PCP_{szk} and PCP_{czk} are similarly defined for statistical and computational zero knowledge.

The KPT result. Kilian, Petrank, and Tardos proved the following theorem:

Theorem 2.2 ([KPT97]). *For every polynomial time function $T: \mathbb{N} \rightarrow \mathbb{N}$, polynomial security function $\lambda: \mathbb{N} \rightarrow \mathbb{N}$, and polynomial knowledge bound function $k: \mathbb{N} \rightarrow \mathbb{N}$,*

$$\text{NTIME}(T) \subseteq \text{PCP}_{\text{szk}} \left[\begin{array}{lll} \text{answer alphabet} & a & = \mathbb{F}_{2^{\text{poly}(\lambda)}} \\ \text{proof length} & l & = \text{poly}(T, k) \\ \text{query complexity} & q & = \text{poly}(\lambda) \\ \text{soundness error} & e & = 2^{-\lambda} \\ \text{prover time} & \text{tp} & = \text{poly}(\lambda, T) \\ \text{verifier time} & \text{tv} & = \text{poly}(\lambda, T, k) \\ \text{knowledge bound} & k & \\ \text{non-adaptive queries} & \text{na} & = \text{No} \end{array} \right].$$

Remark 2.3. We make two remarks: (i) $\text{na} = \text{No}$ because [KPT97]’s construction relies on adaptively querying the proof; (ii) inspection of [KPT97]’s construction reveals that $l(n) \geq \text{poly}(T(n)) \cdot k(n)^6$.

2.4 Reed–Muller and Reed–Solomon codes

We define Reed–Muller and Reed–Solomon codes, as well as their “vanishing” variants [BS08]; all of these are linear codes. We then state a theorem about PCPPs for certain families of RS codes.

RM codes. Let \mathbb{F} be a finite field, H, V subsets of \mathbb{F} , m a positive integer, and ϱ a constant in $(0, 1]$; ϱ is called the *fractional degree*. The Reed–Muller code with parameters $\mathbb{F}, H, m, \varrho$ is $\text{RM}[\mathbb{F}, H, m, \varrho] := \{w: H^m \rightarrow \mathbb{F} \mid \max_{i \in [m]} \deg_{X_i}(w) < \varrho |H|\}$; its message length is $n = (\varrho |H|)^m$, block length is $\ell = |H|^m$, rate is $\rho = \varrho^m$, and relative distance is $\delta = 1 - \varrho$. The vanishing Reed–Muller code with parameters $\mathbb{F}, H, m, \varrho, V$ is $\text{VRM}[\mathbb{F}, H, m, \varrho, V] := \{w \in \text{RM}[\mathbb{F}, H, m, \varrho] \mid w(V^m) = \{0\}\}$; it is a subcode of $\text{RM}[\mathbb{F}, H, m, \varrho]$. We use the following folklore claim, whose correctness can be proved by induction on m :

Claim 2.4. *Let \mathbb{F} be a finite field, H, V subsets of \mathbb{F} with $H \cap V = \emptyset$, m a positive integer, and t a positive integer not exceeding $|H| - |V|$. Then $\text{VRM}[\mathbb{F}, H, m, \frac{|V|+t}{|H|}, V]$ is t -wise independent.*

RS codes. Let \mathbb{F} be a finite field, H, V subsets of \mathbb{F} , and ϱ a constant in $(0, 1]$. The Reed–Solomon code with parameters \mathbb{F}, H, ϱ is $\text{RS}[\mathbb{F}, H, \varrho] := \text{RM}[\mathbb{F}, H, 1, \varrho]$. The vanishing Reed–Solomon code with parameters $\mathbb{F}, H, \varrho, V$ is $\text{VRS}[\mathbb{F}, H, \varrho, V] := \{w \in \text{RS}[\mathbb{F}, H, \varrho] \mid w(V) = \{0\}\}$.

Two RS code families and their PCPPs. Given $\ell, \chi: \mathbb{N} \rightarrow \mathbb{N}$ and $\varrho: \mathbb{N} \rightarrow (0, 1)$, we denote by:

- $\mathcal{RS}_{\ell, \varrho, \chi}^*$ the set of Reed–Solomon codes $\text{RS}[\mathbb{F}, H, \varrho]$ for which \mathbb{F} has characteristic 2, H has cardinality at most ℓ , and the vanishing polynomial Z_H is computable by a χ -size \mathbb{F} -arithmetic circuit; and
- $\mathcal{VRS}_{\ell, \varrho, \chi}^*$ the set of vanishing Reed–Solomon codes $\text{VRS}[\mathbb{F}, H, \varrho, V]$ for which \mathbb{F} has characteristic 2, H has cardinality at most ℓ , and the vanishing polynomials Z_H and Z_V are computable by χ -size \mathbb{F} -arithmetic circuits.

The following theorem is from [BS08, BCGT13a] (the prover running time is shown in [BCGT13a] and the other parameters in [BS08]).

Theorem 2.5. For every $\lambda, s, \chi: \mathbb{N} \rightarrow \mathbb{N}$, and $\varrho: \mathbb{N} \rightarrow (0, 1)$,

$$\text{Rel}(\mathcal{RS}_{\ell, \varrho, \chi}^*), \text{Rel}(\mathcal{VRS}_{\ell, \varrho, \chi}^*) \in \mathbf{PCPP} \left[\begin{array}{ll} \text{answer alphabet} & \mathbf{a} = \mathbb{F}_{2^{s+\log \ell}} \\ \text{proof length} & \mathbf{l} = \tilde{O}(\ell) \\ \text{query complexity} & \mathbf{q} = \lambda \cdot \text{polylog}(\ell) \\ \text{distance measure} & \Delta = \Delta_a^{\text{Ham}} \\ \text{proximity parameter} & \mathbf{d} = \varrho/2 \\ \text{soundness error} & \mathbf{e} = 2^{-\lambda} \\ \text{prover time} & \mathbf{tp} = \text{poly}(s + \chi) \cdot \tilde{O}(\ell) \\ \text{verifier time} & \mathbf{tv} = \lambda \cdot \text{poly}(s + \chi + \log \ell) \\ \text{non-adaptive queries} & \mathbf{na} = \text{Yes} \end{array} \right].$$

3 Duplex PCPs

We define duplex PCPs, and then define notions of zero knowledge for this model. Our main theorem is the construction of a duplex PCP with certain parameters; see Section 4. The difference between a PCP and a duplex PCP is that all provers (both honest and malicious) produce two proof oracles rather than one: the prover produces a proof π_0 ; then the verifier sends a message ρ to the prover; then the prover produces another proof π_1 ; finally the verifier queries both π_0 and π_1 and either accepts or rejects. (Thus, a PCP is a special case of a duplex PCP, but not vice versa.) More precisely, a *duplex PCP system* for a relation \mathcal{R} is a tuple $\text{DPCP} = (P, V)$ that works as follows.

- The *prover* P is a pair (P_0, P_1) of probabilistic algorithms, with shared randomness, where: (a) given as input an instance-witness pair (\mathbf{x}, \mathbf{w}) with $n := |\mathbf{x}|$, P_0 outputs a proof $\pi_0: D_0(n) \rightarrow \Sigma(n)$; (b) given as input (\mathbf{x}, \mathbf{w}) and the verifier's message ρ (see below), P_1 outputs a proof $\pi_1: D_1(n) \rightarrow \Sigma(n)$. Here $D_0(n), D_1(n), \Sigma(n)$ are finite sets.
- The *verifier* V is a pair (V_0, V_1) of probabilistic algorithms, with shared randomness, where: (a) given as input an instance \mathbf{x} with $n := |\mathbf{x}|$, V_0 outputs a message ρ ; (b) given as input \mathbf{x} and with oracle access to proofs $\pi_0: D_0(n) \rightarrow \Sigma(n)$ and $\pi_1: D_1(n) \rightarrow \Sigma(n)$, V_1 queries π_0 and π_1 at a few locations and then outputs a bit.

The system DPCP has (perfect) completeness and soundness error $\epsilon(n)$ if the following two conditions hold. (Below, we explicitly denote the prover's and verifier's randomness as r_P and r_V .)

Completeness: For every instance-witness pair (\mathbf{x}, \mathbf{w}) in the relation \mathcal{R} ,

$$\Pr_{r_P, r_V} \left[V_1^{\pi_0, \pi_1}(\mathbf{x}; r_V) = 1 \mid \begin{array}{l} \pi_0 \leftarrow P_0(\mathbf{x}, \mathbf{w}; r_P) \\ \rho \leftarrow V_0(\mathbf{x}; r_V) \\ \pi_1 \leftarrow P_1(\mathbf{x}, \mathbf{w}, \rho; r_P) \end{array} \right] = 1 .$$

Soundness: For every instance \mathbf{x} not in the language $\text{Lan}(\mathcal{R})$ and pair of algorithms $\tilde{P} = (\tilde{P}_0, \tilde{P}_1)$,

$$\Pr_{r_V} \left[V_1^{\pi_0, \pi_1}(\mathbf{x}; r_V) = 1 \mid \begin{array}{l} \pi_0 \leftarrow \tilde{P}_0 \\ \rho \leftarrow V_0(\mathbf{x}; r_V) \\ \pi_1 \leftarrow \tilde{P}_1(\rho) \end{array} \right] \leq \epsilon(n) .$$

A relation \mathcal{R} belongs to the complexity class $\text{DPCP}[a, l, q, \epsilon, \text{tp}, \text{tv}]$ if there is a DPCP system for \mathcal{R} in which:

- the answer alphabet (i.e., $\Sigma(n)$) is $a(n)$,
- the proof length over that alphabet (i.e., $(|D_0(n)| + |D_1(n)|)$) is at most $l(n)$,
- the verifier queries the two proofs in at most $q(n)$ locations (in total),
- the soundness error is $\epsilon(n)$,
- the prover runs in time $\text{tp}(n)$, and
- the verifier runs in time $\text{tv}(n)$.

Finally, we write $\text{na} = \text{Yes}$ in the square brackets if the queries to the oracles are non-adaptive (i.e., the queried locations only depend on the verifier's inputs); otherwise we write $\text{na} = \text{No}$.

Zero knowledge. A DPCP system $\text{DPCP} = (P, V)$ for a relation \mathcal{R} has *perfect zero knowledge with knowledge bound* k if there exists an expected-polynomial-time probabilistic algorithm S such that for every pair of polynomial-time probabilistic oracle algorithms $\tilde{V} := (\tilde{V}_0, \tilde{V}_1)$ the following two distribution families are identical:

$$\{S(\tilde{V}, \mathbf{x})\}_{(\mathbf{x}, \mathbf{w}) \in \mathcal{R}} \quad \text{and} \quad \{\text{DPCPView}(k, \tilde{V}, P, \mathbf{x}, \mathbf{w})\}_{(\mathbf{x}, \mathbf{w}) \in \mathcal{R}} ,$$

where $\text{DPCPView}(k, \tilde{V}, P, \mathbf{x}, \mathbf{w})$ is the view of \tilde{V}_1 in its execution when given input \mathbf{x} and when allowed to make a total of $k(n)$ adaptive queries to π_0, π_1 , where $\pi_0 := P_0(\mathbf{x}, \mathbf{w})$ and $\pi_1 := P_1(\mathbf{x}, \mathbf{w}, \tilde{V}_0^{\pi_0}(\mathbf{x}))$. (As above, P_0, P_1 share the same randomness r_P ; ditto for \tilde{V}_0, \tilde{V}_1 .) The definition of statistical and computational zero knowledge (with knowledge bound k) are similar: rather than identical, the two distribution families are required to be statistically and computationally close (as $|\mathbf{x}|$ grows), respectively.

4 Main theorem

The main result of this paper is the following.

Theorem 4.1. *For every polynomially-bounded time function $T: \mathbb{N} \rightarrow \mathbb{N}$ and polynomially-bounded knowledge bound function $k: \mathbb{N} \rightarrow \mathbb{N}$ with $k \leq T$,*

$$\text{NTIME}(T) \subseteq \text{DPCP}_{\text{pzk}} \left[\begin{array}{ll} \text{answer alphabet} & \mathbf{a} = \mathbb{F}_{2^{O(\log(T+k))}} \\ \text{proof length} & \mathbf{l} = \tilde{O}(T+k) \\ \text{query complexity} & \mathbf{q} = \text{polylog}(T+k) \\ \text{soundness error} & \mathbf{e} = 1/2 \\ \text{prover time} & \mathbf{tp} = \text{poly}(n) \cdot \tilde{O}(T+k) \\ \text{verifier time} & \mathbf{tv} = \text{poly}(n + \log(T+k)) \\ \text{knowledge bound} & \mathbf{k} \\ \text{non-adaptive queries} & \mathbf{na} = \text{Yes} \end{array} \right].$$

The soundness error can be reduced from $e = 1/2$ to $e = 2^{-\lambda}$ via parallel repetition, with a corresponding λ -fold increase in the query complexity q and verifier running time tv ; perfect zero knowledge is preserved so long as the query complexity does not exceed the knowledge bound k .

A corollary. The above theorem says that, fixing T , the prover running time is quasilinear in the knowledge bound k , while the verifier running time is polylogarithmic in k . This yields an intriguing corollary: in the duplex PCP model, we only need a polylogarithmic computational overhead of the prover over the verifier in order to guarantee perfect zero knowledge. We state this formally next.

Corollary 4.2. *For every polynomial time function $T: \mathbb{N} \rightarrow \mathbb{N}$ and relation $\mathcal{R} \in \text{NTIME}(T)$, there is a constant c such that, for every function $tv: \mathbb{N} \rightarrow \mathbb{N}$ with $tv(n) \geq n \cdot (\log T(n))^c$, there is a DPCP system with:*

- *completeness 1 and soundness $2^{-tv(n)/\text{polylog}(T(n))}$;*
- *perfect zero knowledge;*
- *the verifier running time is $tv(n)$ and prover running time is $tp(n) := \max\{T(n) \cdot (\log T(n))^c, tv(n) \cdot (\log tv(n))^c\}$. The verifier has no limitations other than a bound on its running time (its query complexity can be as large as $tv(n)$).*

4.1 Proof sketch

Let \mathcal{R} be a relation in NP, and let (x, w) be an instance-witness pair in \mathcal{R} . The prover and verifier both know x , while the prover also knows w . The prover wishes to convince the verifier that he knows a witness w for x , in such a way that the verifier does not learn anything about w (beyond what can be inferred from the prover's claim).

The KPT approach. We introduce our ideas by contrasting them with those of [KPT97]. Suppose that the prover wishes to convince the verifier by sending him a PCP proof $\pi = \pi(w)$ such that any k values in π do not reveal anything about w . Loosely speaking, [KPT97] (building on [DFK⁺92]) provide a probabilistic transformation that maps the PCP proof π to a new proof π' , in which each bit of π is “hidden” amongst many bits of π' . The main tool employed in the transformation is a *locking scheme*, and its use imposes certain limitations: (i) the new proof π' is $\text{poly}(k)$ larger than the original one (k^6 by inspection of [DFK⁺92, KPT97]); (ii) zero knowledge holds only statistically, but not perfectly, because a malicious verifier can be “lucky” and obtain information on the bit of π being locked with fewer queries to π' than expected.

Our approach (ideally). We take a different approach: apply a “local” PCP to a “random” witness, as we now explain. Suppose that $\pi = \pi(w)$ is (t, k) -local, i.e., any k positions of the PCP proof π jointly depend on at most t positions of the witness w . Note that, even if π is (t, k) -local, a single bit of π can still leak information about w . So suppose further that the relation \mathcal{R} is t -randomizable: given $(x, w) \in \mathcal{R}$, one can efficiently sample a witness w' from a t -wise independent subset of the set of witnesses for x . In such a case, the prover can produce a zero-knowledge PCP as follows: (1) sample a witness w' from the t -wise independent subset; then (2) send to the verifier the PCP proof $\pi = \pi(w')$. Indeed, the locality of π ensures that seeing any k indices of π reveals nothing about w , because these k indices are a function of t random bits. In sum, if we had a (t, k) -local PCP for a t -randomizable relation \mathcal{R} , then we could obtain a PCP for \mathcal{R} that is zero knowledge against verifiers that ask at most k queries.

Our approach (in reality). Unfortunately, we do not know how to obtain local PCPs for randomizable relations. However, we are able to obtain “partially local” *duplex* PCPs for certain randomizable relations, and also show that **NTIME** can be efficiently reduced to these randomizable relations, as we now explain.

Our starting point are *algebraic PCPs*: certain PCPs that prove satisfiability of *algebraic problems* (APs) [PS94]. Numerous known PCP constructions can be viewed as algebraic PCPs. Informally, in this work we make two basic observations: (i) algebraic PCPs exist for certain randomizable relations; and (ii) an algebraic PCP proof can be split in two parts, one part is local, while the other part is not local but enjoys convenient linear algebraic properties that, nevertheless, enable us to hide information about the witness, in the duplex PCP model. (Recall that, in the duplex PCP model, the prover produces a proof π_0 ; then the verifier sends a message ρ to the prover; then the prover produces another proof π_1 ; finally the verifier queries both π_0 and π_1 and either accepts or rejects.)

In more detail, from a technical viewpoint, we proceed as follows. First, we introduce a family of constraint satisfaction problems (CSPs) called *linear algebraic CSPs*, and show that **NTIME** is efficiently reducible to *randomizable* linear algebraic CSPs. The reduction consists of two parts: we go through an intermediary that we call *group preserving algebraic problems* (GAPs), a special case of APs that we believe to be of independent interest for the study of algebraic PCPs. Second, we construct a duplex PCP system for randomizable linear algebraic CSPs that is zero knowledge against verifiers that ask at most a certain number of queries.

A technical piece: zero-knowledge duplex PCPP for low-degreeness. Later sections address all of the above steps (see Section 4.2 for a roadmap of these), and for now we only sketch one of these steps. Above we mention that an algebraic PCP proof has two parts: a local part, and a non-local part. This latter part of the proof arises from a central component of many PCP proofs: a *PCP of proximity* (PCPP) [BGH⁺05, DR04] that facilitates low-degree testing. Informally, given a function $f: H \rightarrow \mathbb{F}$ and an integer d , a *PCPP for degree d* is a proof $\pi(f)$ that f is ϵ -close to an evaluation of a polynomial degree at most d . We explain how to transform a PCPP for low-degreeness into a duplex PCPP for low-degreeness that is zero knowledge against verifiers that make at most t queries.

The set C of functions $f: H \rightarrow \mathbb{F}$ that are evaluations of a polynomial of degree at most d is a subspace of $\mathbb{F}^{|H|}$. The basic idea is that, in order for the prover to convince the verifier that a function f is close to C , it suffices for the prover to convince the verifier that a *random offset* of f is close to C : one can verify that, for any $u: H \rightarrow \mathbb{F}$, if f is ϵ -far from C , then $\alpha f + u$ is $\epsilon/2$ -far from C , with probability $1 - |\mathbb{F}|^{-1}$ over a random $\alpha \in \mathbb{F}$. Hence, we can let the duplex PCP work as follows: (i) the prover samples a witness w' from the t -wise independent subset, chooses a random $u \in C$, and sends $\pi_0 := (w', u)$ to the verifier; (ii) the verifier sends to the prover a random $\alpha \in \mathbb{F}$; (iii) the prover sends $\pi_1 = (v, \pi(v))$ to the verifier, where $v := \alpha w' + u$ and $\pi(v)$ is a PCPP for low-degreeness of v ; (iv) the verifier runs the PCPP verifier on (v, π) to check that v is close to C , and then checks that $v_i = \alpha w'_i + u_i$ for a few random indices i in $\{1, \dots, |H|\}$.

Let us discuss the various properties of the duplex PCPP.

- **COMPLETENESS:** If $w \in C$, then $\alpha w' + u \in C$; therefore, the prover convinces the verifier.
- **ZERO-KNOWLEDGE:** If the verifier asks at most t queries, then he learns nothing about w because: $\pi_0 = (w', u)$ contains w' sampled from a t -wise independent subset and u random in C ; $\pi_1 = (v, \pi(v))$ is running the PCPP on a vector v that is random in C .
- **SOUNDNESS:** If v does equal $\alpha \cdot w + u$, then the verifier rejects with high probability because v is far from C (and the PCPP verifier rejects π with high probability). If instead v does not equal $\alpha \cdot w + u$, then the fact that v is close to C does not prove anything about whether w is also close. So, in this case, we need to reason about the success probability of the verifier’s linearity tests: if these pass with enough probability, then with high probability v is close to $\alpha w + u$, which again suffices for our purpose. Overall, soundness holds.

Next, we discuss how the technical sections are organized, and how they come together to yield our main theorem.

4.2 Roadmap of the rest of the paper

The rest of the paper is dedicated to turn the above intuition into a more formal proof. To do so, we introduce various intermediate steps, as follows.

- In Section 5, we introduce *linear algebraic CSPs* (a family of constraint satisfaction problems), and then describe how to obtain a *canonical PCP* for any linear algebraic CSP.

- In Section 6, we introduce *randomizable* linear algebraic CSPs, a subfamily of linear algebraic CSPs; then we show that, for every randomizable linear algebraic CSP, we can convert the CSP's canonical PCP into a corresponding zero-knowledge duplex PCP, incurring only little overheads.
- In Section 7, we show an efficient reduction from **NTIME** to randomizable linear algebraic CSPs; along the way, we introduce a family of algebraic problems, having special symmetry properties, that we believe to be of independent interest (e.g., for studying other questions about PCPs).

Combining (i) the efficient reduction from **NTIME** to randomizable linear algebraic CSPs together with (ii) the zero-knowledge duplex PCP for such problems yields Theorem 4.1. In Section 8 we provide details about how these components are combined.

5 Linear algebraic CSPs and their canonical PCPs

We introduce *linear algebraic CSPs*, a family of constraint satisfaction problems; then we describe how to obtain a *canonical PCP* for any linear algebraic CSP.

5.1 Linear algebraic constraint satisfaction problems

A constraint satisfaction problem asks whether, for a given “local” function g , there exists an input α such that $g(\alpha)$ is an “accepting” output. For example, in the case of 3-SAT with n variables and m clauses, the function g maps $\{0, 1\}^n$ to $\{0, 1\}^m$, and $g(\alpha)$ indicates which clauses are satisfied by $\alpha \in \{0, 1\}^n$; hence α yields an accepting output if (and only if) $g(\alpha) = 1^m$. Below we introduce a family of constraint satisfaction problems whose domain and range are linear-algebraic objects, namely, linear error correcting codes.

We begin by providing the notion of locality that we use for g ; we also provide two other notions, one for the efficiency of computing a single coordinate of g ’s output, and another for measuring g ’s “pseudorandomness”.

Definition 5.1. *Let $g: \Sigma^n \rightarrow \Sigma^m$ be a function. We say that g is:*

- q -local if for every $j \in [m]$ there exists $I_j \subseteq [n]$ with $|I_j| \leq q$ such that $g(\alpha)[j]$ (the j -th coordinate of $g(\alpha)$) depends only on $\alpha|_{I_j}$ (the restriction of α to I_j);
- c -efficient if there is a time c algorithm that, given j and $\alpha|_{I_j}$, computes the set I_j and value $g(\alpha)[j]$;
- (γ, ϵ) -sampling if $\Pr[I_j \cap I \neq \emptyset \mid j \leftarrow [m]] \leq \gamma$ for every $I \subseteq [n]$ with $|I|/n \leq \epsilon$.

Next we introduce \mathcal{R}_{LA} , the relation of **linear algebraic CSPs**:

Definition 5.2 (\mathcal{R}_{LA}). *Given functions $f: \mathbb{N} \rightarrow \mathcal{F}$, $\ell, q, c: \mathbb{N} \rightarrow \mathbb{N}$, and $\rho, \delta, \gamma, \epsilon: \mathbb{N} \rightarrow (0, 1]$, the relation*

$$\mathcal{R}_{\text{LA}}[f, \ell, \rho, \delta, q, c, \gamma, \epsilon]$$

consists of instance-witness pairs $(\mathfrak{x}, \mathfrak{w})$ satisfying the following.

- *The instance \mathfrak{x} is a tuple $(1^n, C_\circ, C_\bullet, g)$ where:*
 - C_\circ, C_\bullet are linear error correcting codes with block lengths $\ell_\circ(n), \ell_\bullet(n)$ at most $\ell(n)$, each with rate at most $\rho(n)$ and relative distance at least $\delta(n)$ over the same field $f(n)$;
 - $g: f(n)^{\ell_\circ(n)} \rightarrow f(n)^{\ell_\bullet(n)}$ is a $q(n)$ -local, $c(n)$ -efficient, $(\gamma(n), \epsilon(n))$ -sampling function;
 - $C_\bullet \cup g(C_\circ)$ has relative distance at least $\delta(n)$ (though may not be a linear space).
- *The witness \mathfrak{w} is a tuple $(\alpha_\circ, \alpha_\bullet)$ where $\alpha_\circ \in f(n)^{\ell_\circ(n)}$ and $\alpha_\bullet \in f(n)^{\ell_\bullet(n)}$.*
- *The instance \mathfrak{x} and witness \mathfrak{w} jointly satisfy the following: $\alpha_\circ \in C_\circ$, $\alpha_\bullet \in C_\bullet$, and $g(\alpha_\circ) = \alpha_\bullet$.*

Finally we need notation for referring to codes appearing in instances of \mathcal{R}_{LA} :

Definition 5.3. *Given $\mathcal{R} \subseteq \mathcal{R}_{\text{LA}}$, we denote by*

- $\mathcal{C}_{\mathcal{R}, \circ}$ the set of codes C for which there is an instance $\mathfrak{x} = (1^n, C_\circ, C_\bullet, g)$ in the relation \mathcal{R} with $C = C_\circ$;
- $\mathcal{C}_{\mathcal{R}, \bullet}$ the set of codes C for which there is an instance $\mathfrak{x} = (1^n, C_\circ, C_\bullet, g)$ in the relation \mathcal{R} with $C = C_\bullet$.

5.2 A canonical PCP for linear algebraic CSPs

We show how to construct a “canonical” PCP system for instances in \mathcal{R}_{LA} (the relation of linear algebraic CSPs). At a high level, a canonical PCP proof for a \mathcal{R}_{LA} -instance \mathfrak{x} consists of a witness $\mathfrak{w} = (\alpha_\circ, \alpha_\bullet)$ concatenated with two PCPP proofs π_\circ, π_\bullet , showing that $\alpha_\circ, \alpha_\bullet$ are close to C_\circ, C_\bullet respectively. The canonical PCP verifier first checks the two PCPP proofs and then checks that $g(\alpha_\circ)[j] = \alpha_\bullet[j]$ for a uniformly random $j \in [\ell_\bullet]$.

Definition 5.4. *Given (i) a relation $\mathcal{R} \subseteq \mathcal{R}_{\text{LA}}$, (ii) a PCPP system $\text{PCPP}_\circ = (P_\circ, V_\circ)$ for $\text{Rel}(\mathcal{C}_{\mathcal{R}, \circ})$, and (iii) a PCPP system $\text{PCPP}_\bullet = (P_\bullet, V_\bullet)$ for $\text{Rel}(\mathcal{C}_{\mathcal{R}, \bullet})$, the canonical PCP system for the triple $(\mathcal{R}, \text{PCPP}_\circ, \text{PCPP}_\bullet)$ is the PCP system $\text{PCP} = (P, V)$ constructed as follows.*

- **Prover.** Given $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}_{\text{LA}}$, the PCP prover P outputs $\pi := (\mathbf{w}, \pi_\circ, \pi_\bullet)$ where $\pi_\circ := P_\circ(C_\circ, \alpha_\circ)$ and $\pi_\bullet := P_\bullet(C_\bullet, \alpha_\bullet)$. In other words, the PCP prover outputs a PCP proof that is the concatenation of the witness $\mathbf{w} = (\alpha_\circ, \alpha_\bullet)$ and a pair of PCPP proofs, the first proving that $\alpha_\circ \in C_\circ$ and the second proving that $\alpha_\bullet \in C_\bullet$.
- **Verifier.** Given \mathbf{x} and oracle access to a PCP proof $\pi = (\mathbf{w}, \pi_\circ, \pi_\bullet)$, the PCP verifier V works as follows:
 - (proximity) check that $V_\circ^{(\alpha_\circ, \pi_\circ)}(C_\circ)$ and $V_\bullet^{(\alpha_\bullet, \pi_\bullet)}(C_\bullet)$ both accept;
 - (consistency) check that $g(\alpha_\circ)[j] = \alpha_\bullet[j]$ for a uniformly random $j \in [\ell_\bullet]$.

The next lemma says that the above construction is a PCP system when \mathcal{R}_{LA} 's parameters are sufficiently ‘‘good’’.

Lemma 5.5 ($\mathcal{R}_{\text{LA}} \rightarrow \text{PCP}$). *Suppose that \mathcal{R} is a relation that satisfies the following conditions:*

- (i) $\mathcal{R} \subseteq \mathcal{R}_{\text{LA}}[f_1, \ell_1, \rho_1, \delta_1, q_1, c_1, \gamma_1, \epsilon_1]$ with $\epsilon_1 < \min\{\frac{\delta_1}{2}, \delta_1 - \gamma_1\}$;
- (ii) $\text{Rel}(\mathcal{C}_{\mathcal{R}, \circ}), \text{Rel}(\mathcal{C}_{\mathcal{R}, \bullet}) \in \text{PCPP}[\mathbf{a}_2, l_2, \mathbf{q}_2, \Delta_{\mathbf{a}_2}^{\text{Ham}}, d_2, \mathbf{e}_2, \text{tp}_2, \text{tv}_2, \text{na}_2]$ with $\mathbf{a}_2 = f_1$ and $d_2 \leq \epsilon_1$.

Then there is a canonical PCP system for a triple $(\mathcal{R}, \text{PCPP}_\circ, \text{PCPP}_\bullet)$ that yields

$$\mathcal{R} \in \text{PCP} \left[\begin{array}{ll} \text{answer alphabet} & \mathbf{a} = f_1 (= \mathbf{a}_2) \\ \text{proof length} & l = 2l_2(\ell_1) + 2\ell_1 \\ \text{query complexity} & \mathbf{q} = 2\mathbf{q}_2(\ell_1) + q_1 + 1 \\ \text{soundness error} & \mathbf{e} = \max\{1 - \delta_1 + \gamma_1 + \epsilon_1, \mathbf{e}_2\} \\ \text{prover time} & \text{tp} = 2\text{tp}_2(\ell_1) \\ \text{verifier time} & \text{tv} = 2\text{tv}_2(\ell_1) + c_1 + \log \ell_1 \\ \text{non-adaptive queries} & \text{na} = \text{na}_2 \end{array} \right].$$

Proof of Lemma 5.5. First, we show that the canonical PCP system satisfies completeness and soundness; afterwards, we discuss the efficiency parameters achieved by it.

Completeness. Consider an instance-witness pair (\mathbf{x}, \mathbf{w}) in the relation \mathcal{R} . Parse the instance \mathbf{x} as $(1^n, C_\circ, C_\bullet, g)$ and the witness \mathbf{w} as $(\alpha_\circ, \alpha_\bullet)$. Since $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$, we have that $\alpha_\circ \in C_\circ$, $\alpha_\bullet \in C_\bullet$, and $g(\alpha_\circ) = \alpha_\bullet$. Therefore, the PCP proof $(\mathbf{w}, \pi_\circ, \pi_\bullet)$ generated by the PCP prover is accepted by the PCP verifier with probability 1: the PCPP verifiers $V_\circ^{(\alpha_\circ, \pi_\circ)}(C_\circ)$ and $V_\bullet^{(\alpha_\bullet, \pi_\bullet)}(C_\bullet)$ always accept and $g(\alpha_\circ)[j] = \alpha_\bullet[j]$ for every $j \in [\ell_\bullet]$.

Soundness. Consider an instance \mathbf{x} not in the language $\text{Lan}(\mathcal{R})$ and a PCP proof $\tilde{\pi} = (\tilde{\mathbf{w}}, \tilde{\pi}_\circ, \tilde{\pi}_\bullet)$. Parse the instance \mathbf{x} as $(1^n, C_\circ, C_\bullet, g)$ and the witness $\tilde{\mathbf{w}}$, inside $\tilde{\pi}$, as $(\tilde{\alpha}_\circ, \tilde{\alpha}_\bullet)$. We prove that V accepts $\tilde{\pi}$ with probability at most $\max\{1 - \delta_1 + \gamma + \epsilon_1, \mathbf{e}_2\}$, by considering the following three cases.

- *Case 1:* $\tilde{\alpha}_\circ$ is ϵ_1 -far in relative Hamming distance from C_\circ . The canonical PCP verifier's proximity test fails, because $\Delta_{\mathbf{a}}^{\text{Ham}}(\tilde{\alpha}_\circ, C_\circ) \geq \epsilon_1 \geq d_2$, and so the PCPP verifier $V_\circ^{(\tilde{\alpha}_\circ, \tilde{\pi}_\circ)}(C_\circ)$ accepts with probability at most \mathbf{e}_2 .
- *Case 2:* $\tilde{\alpha}_\bullet$ is ϵ_1 -far in relative Hamming distance from C_\bullet . This case is analogous to the previous one.
- *Case 3:* there exist $\alpha_\circ \in C_\circ$ and $\alpha_\bullet \in C_\bullet$ with $\Delta_{\mathbf{a}}^{\text{Ham}}(\alpha_\circ, \tilde{\alpha}_\circ) \leq \epsilon_1$ and $\Delta_{\mathbf{a}}^{\text{Ham}}(\alpha_\bullet, \tilde{\alpha}_\bullet) \leq \epsilon_1$.

First, since ϵ_1 is less than $\delta_1/2$ (the unique decoding radius of C_\circ and C_\bullet), the codewords α_\circ and α_\bullet are unique.

Next, we claim that $\alpha'_\bullet := g(\alpha_\circ)$ and $g(\tilde{\alpha}_\circ)$ are γ_1 -close. Indeed, since g is (γ_1, ϵ_1) -sampling, α_\circ and $\tilde{\alpha}_\circ$ differ in at most $\epsilon_1 \cdot \ell_\circ(n)$ positions, and so at most $\gamma_1 \cdot \ell_\bullet(n)$ positions of $g(\tilde{\alpha}_\circ)$ depend on an index where α_\circ and $\tilde{\alpha}_\circ$ differ.

Next, we claim that $\Delta_{\mathbf{a}}^{\text{Ham}}(\alpha_\bullet, \alpha'_\bullet) \geq \delta_1$. Indeed, we have that $\alpha_\bullet \neq \alpha'_\bullet$ because otherwise $(\alpha_\circ, \alpha_\bullet)$ would be a satisfying assignment for \mathbf{x} (contradicting the assumption that $\mathbf{x} \notin \text{Lan}(\mathcal{R})$); moreover, we also have that $C_\bullet \cup g(C_\circ)$ has relative distance at least δ_1 .

We now use the triangle inequality, along with the above observations, to obtain that

$$\begin{aligned} \delta_1 &\leq \Delta_{\mathbf{a}}^{\text{Ham}}(\alpha_\bullet, \alpha'_\bullet) \\ &\leq \Delta_{\mathbf{a}}^{\text{Ham}}(\alpha_\bullet, \tilde{\alpha}_\bullet) + \Delta_{\mathbf{a}}^{\text{Ham}}(\tilde{\alpha}_\bullet, g(\tilde{\alpha}_\circ)) + \Delta_{\mathbf{a}}^{\text{Ham}}(g(\tilde{\alpha}_\circ), \alpha'_\bullet) \\ &\leq \epsilon_1 + \Delta_{\mathbf{a}}^{\text{Ham}}(\tilde{\alpha}_\bullet, g(\tilde{\alpha}_\circ)) + \gamma_1. \end{aligned}$$

Thus, $\Delta_{\mathbf{a}}^{\text{Ham}}(\tilde{\alpha}_\bullet, g(\tilde{\alpha}_\circ)) \geq \delta_1 - (\gamma_1 + \epsilon_1)$, and so the canonical PCP verifier's consistency check passes with probability at most $1 - \delta_1 + \gamma_1 + \epsilon_1$.

We conclude that V accepts $\tilde{\pi}$ with probability at most $\max\{1 - \delta_1 + \gamma_1 + \epsilon_1, e_2\}$.

Other parameters. The remaining parameters are straightforward to establish. The canonical PCP does not change the alphabet, so $a = f_1$ (which also equals a_2). The proof length, and the running times of the prover and verifier are the sum of the same measures of the sub-components: the PCP proof has $l = 2l_2(\ell_1) + 2\ell_1$ symbols, is produced in time $t_p = 2t_{p_2}(\ell_1)$, and is verified in time $t_v = 2t_{v_2}(\ell_1) + c_1 + O(1)$. The canonical PCP verifier makes $q_1 + 1$ queries on top of those of the PCPP verifiers, so its query complexity is $q = 2q_2(\ell_1) + q_1 + 1$. The $q_1 + 1$ additional queries are non-adaptive; so if the PCPP verifiers are non-adaptive, so is the canonical PCP verifier (i.e., $na = na_2$). \square

6 Zero-knowledge duplex PCPs from randomizable linear algebraic CSPs

We introduce *randomizable linear algebraic CSPs*, a subfamily of linear algebraic CSPs. Then we show that, for every randomizable linear algebraic CSP, we can convert the CSP's canonical PCP into a corresponding zero-knowledge duplex PCP, incurring only little overheads.

6.1 Randomizable linear algebraic CSPs

The definition below specifies the notion of randomizability for linear algebraic CSPs.

Definition 6.1 (\mathcal{R}_{RLA}). *The relation $\mathcal{R}_{\text{RLA}}[f, \ell, \rho, \delta, q, c, \gamma, \epsilon, t, r]$ is the sub-relation of $\mathcal{R}_{\text{LA}}[f, \ell, \rho, \delta, q, c, \gamma, \epsilon]$ obtained by restricting it to instances that are t -randomizable in time r . An instance $\mathfrak{x} = (1^n, C_\circ, C_\bullet, g)$ is $t(n)$ -randomizable in time $r(n)$ if: (i) there exists a $t(n)$ -wise independent subcode $C' \subseteq C_\circ$ such that if $(w_\circ, g(w_\circ))$ satisfies \mathfrak{x} , then, for every w'_\circ in $C' + w_\circ := \{w' + w_\circ \mid w' \in C'\}$, the witness $(w'_\circ, g(w'_\circ))$ satisfies \mathfrak{x} ; and (ii) one can sample, in time $r(n)$, three uniformly random elements in C', C_\circ and C_\bullet respectively.*

Remark 6.2. A linear code $C \subseteq \mathbb{F}^\ell$ is t -wise independent if and only if its dual distance is at least t , i.e., its dual code $C^\perp := \{w \in \mathbb{F}^\ell \mid \forall z \in C, \sum_{i=1}^\ell w_i z_i = 0\}$ has distance at least t .

Remark 6.3. A random codeword in a linear code $C \subseteq \mathbb{F}^\ell$ can be sampled with $r := \ell \cdot \dim(C)$ field operations in \mathbb{F} , by using a generating matrix for C . In certain cases, a random codeword can be sampled more efficiently; e.g., in the case of Reed–Muller codes, $\ell \cdot \text{polylog}(\dim(C))$ field operations are sufficient (via a suitable use of FFTs).

6.2 Construction of zero-knowledge duplex PCPs

We construct a zero-knowledge duplex PCP system for randomizable linear algebraic CSPs. The duplex PCP system does little more than invoking, as a subroutine, the canonical PCP system for the linear algebraic CSP; hence, the efficiency of the duplex PCP and of the canonical PCP system are closely related. The construction demonstrates that “adding zero knowledge to an algebraic PCP” is cheap, provided that one moves from the PCP model to the (more general) duplex PCP model. More precisely, we prove the following theorem.

Theorem 6.4 ($\mathcal{R}_{\text{RLA}} \rightarrow \text{DPCP}_{\text{pzk}}$). *Suppose that \mathcal{R} is a relation that satisfies the following conditions:*

- (i) $\mathcal{R} \subseteq \mathcal{R}_{\text{RLA}}[f_1, \ell_1, \rho_1, \delta_1, q_1, c_1, \gamma_1, \epsilon_1, t_1, r_1]$ with $\epsilon_1 < \min\{\frac{\delta_1}{2}, \delta_1 - \gamma_1\}$ and ℓ_1, c_1, r_1 polynomially bounded;
- (ii) $\text{Rel}(\mathcal{C}_{\mathcal{R}, \circ}), \text{Rel}(\mathcal{C}_{\mathcal{R}, \bullet}) \in \text{PCPP}[\mathfrak{a}_2, \mathfrak{l}_2, \mathfrak{q}_2, \Delta_{\mathfrak{a}_2}^{\text{Ham}}, \mathfrak{d}_2, \mathfrak{e}_2, \mathfrak{tp}_2, \mathfrak{tv}_2, \mathfrak{na}_2]$ with $\mathfrak{a}_2 = f_1$ and $\mathfrak{d}_2 \leq \epsilon_1/4$.

Then there is a duplex PCP system for \mathcal{R} that yields

$$\mathcal{R} \in \text{DPCP}_{\text{pzk}} \left[\begin{array}{ll} \text{answer alphabet} & \mathfrak{a} = f_1 (= \mathfrak{a}_2) \\ \text{proof length} & \mathfrak{l} = 2\mathfrak{l}_2(\ell_1) + 6\ell_1 \\ \text{query complexity} & \mathfrak{q} = 2\mathfrak{q}_2(\ell_1) + q_1 + 7 \\ \text{soundness error} & \mathfrak{e} = \max\{1 - \delta_1 + \gamma_1 + \epsilon_1, (1 - |f_1|^{-1}) \cdot \max\{\mathfrak{e}_2, \epsilon_1/4\} + |f_1|^{-1}\} \\ \text{prover time} & \mathfrak{tp} = 2\mathfrak{tp}_2(\ell_1) + (c_1 + 5)\ell_1 + r_1 \\ \text{verifier time} & \mathfrak{tv} = 2\mathfrak{tv}_2(\ell_1) + c_1 + \log \ell_1 \\ \text{knowledge bound} & \mathfrak{k} = t_1/q_1 \\ \text{non-adaptive queries} & \mathfrak{na} = \mathfrak{na}_2 \end{array} \right].$$

Proof. We prove the claim by constructing a suitable duplex PCP system $\text{DPCP} = (P, V)$ for the relation \mathcal{R} . Recall that: the prover P is a pair of algorithms (P_0, P_1) , and the verifier V is also a pair of algorithms (V_0, V_1) ; moreover, an instance \mathfrak{x} of \mathcal{R} is of the form $(1^n, C_\circ, C_\bullet, g)$, while a witness \mathfrak{w} of \mathcal{R} is of the form $(\alpha_\circ, \alpha_\bullet)$; finally, randomizability implies that there is a $t(n)$ -wise independent subcode $C' \subseteq C_\circ$ such that if $(w_\circ, g(w_\circ))$ satisfies \mathfrak{x} then so does the witness $(w'_\circ, g(w'_\circ))$, for every w'_\circ in $C' + w_\circ$.

We now describe the construction of the duplex PCP system $\text{DPCP} = (P, V)$:

- $P_0(\mathfrak{x}, \mathfrak{w}) \rightarrow \pi_0$

Sample uniformly random $v_\circ \in C_\circ, v_\bullet \in C_\bullet, u' \in C'$; compute $w_\circ := u' + \alpha_\circ, w_\bullet := g(w_\circ)$ and output $\pi_0 := (w_\circ \| v_\circ \| w_\bullet \| v_\bullet)$.

- $V_0(\mathbf{x}) \rightarrow \rho$
Sample uniformly random $\rho_o, \rho_\bullet \in f_1$, and output $\rho := (\rho_o, \rho_\bullet)$.
- $P_1(\mathbf{x}, \mathbf{w}, \rho) \rightarrow \pi_1$
Compute $z_o := \rho_o w_o + v_o$ and $z_\bullet := \rho_\bullet w_\bullet + v_\bullet$; compute $\pi_o := P_o(C_o, z_o)$ and $\pi_\bullet := P_\bullet(C_\bullet, z_\bullet)$; and output $\pi_1 := (z_o \| z_\bullet \| \pi_o \| \pi_\bullet)$. (Essentially, this step corresponds to running the canonical PCP prover with respect to a uniformly random pair (z_o, z_\bullet) in (C_o, C_\bullet) .)
- $V_1^{\pi_o, \pi_1}(\mathbf{x}) \rightarrow b$
Conduct the following tests (and reject if any of them fails):
 - (proximity) check that $V_o^{(z_o, \pi_o)}(C_o)$ and $V_\bullet^{(z_\bullet, \pi_\bullet)}(C_\bullet)$ both accept;
 - (consistency) check that $g(w_o)[j] = w_\bullet[j]$ for a random $j \in [\ell_\bullet]$;
 - (linearity) check that $z_o[i] = \rho_o w_o[i] + v_o[i]$ and $z_\bullet[k] = \rho_\bullet w_\bullet[k] + v_\bullet[k]$ for random $i \in [\ell_o(n)]$ and $k \in [\ell_\bullet(n)]$.
(Essentially the first two steps correspond to running the canonical PCP verifier on modified inputs, while the third step consists of two linearity tests.)

Having described the duplex PCP system, we now show that it satisfies completeness, soundness and zero-knowledge; afterwards, we discuss the efficiency parameters achieved by it.

Completeness. Consider an instance-witness pair (\mathbf{x}, \mathbf{w}) in the relation \mathcal{R} . Since $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$, we have that $\alpha_o \in C_o$, $\alpha_\bullet \in C_\bullet$, and $g(\alpha_o) = \alpha_\bullet$. Since $w_o \in C' + \alpha_o$ and \mathcal{R} is randomizable, we have that $(w_o, w_\bullet) := (w_o, g(w_o))$ satisfies \mathbf{x} ; thus V_1 's consistency check passes with probability 1. Since the codes C_o and C_\bullet are linear and $w_o, v_o \in C_o$, $w_\bullet, v_\bullet \in C_\bullet$, we have that $z_o := \rho_o w_o + v_o \in C_o$ and $z_\bullet := \rho_\bullet w_\bullet + v_\bullet \in C_\bullet$; thus the PCPP verifiers $V_o^{(z_o, \pi_o)}(C_o)$ and $V_\bullet^{(z_\bullet, \pi_\bullet)}(C_\bullet)$ accept with probability 1. Finally, by construction of z_o and z_\bullet , V_1 's linearity tests also accept with probability 1. We conclude that the duplex PCP system described above has perfect completeness.

Soundness. Consider an instance \mathbf{x} not in the language $\text{Lan}(\mathcal{R})$. Fix an arbitrary proof string $\tilde{\pi}_0 = (\tilde{w}_o \| \tilde{v}_o \| \tilde{w}_\bullet \| \tilde{v}_\bullet)$, and let the proof string $\tilde{\pi}_1 = (\tilde{z}_o \| \tilde{z}_\bullet \| \tilde{\pi}_o \| \tilde{\pi}_\bullet)$ depend arbitrarily on the verifier message $\rho = (\rho_o, \rho_\bullet)$. We prove this by distinguishing between the three cases below.

- *Case 1: \tilde{w}_o is ϵ_1 -far in relative Hamming distance from C_o .*

Claim 2.1 implies that $z'_o := \rho_o \tilde{w}_o + \tilde{v}_o$ is $\epsilon_1/2$ -far from C_o , with probability $1 - |f_1|^{-1}$ over a random choice of ρ_o . Let $\theta := \Delta_a^{\text{Ham}}(z'_o, \tilde{z}_o)$ and $\eta := \Delta_a^{\text{Ham}}(\tilde{z}_o, C_o)$. By the triangle inequality, $\theta + \eta \geq \Delta_a^{\text{Ham}}(z'_o, C_o) \geq \epsilon_1/2$; hence, at least one of the inequalities $\theta \geq \epsilon_1/4$ and $\eta \geq \epsilon_1/4$ holds. In the former case, V_1 's first linearity test accepts with probability at most $1 - \epsilon_1/4$; in the latter case, the PCPP verifier $V_o^{(\tilde{z}_o, \tilde{\pi}_o)}(C_o)$ for V_1 's first proximity test accepts with probability at most e_2 , as $\Delta_a^{\text{Ham}}(\tilde{z}_o, C_o) \geq \epsilon_1/4 \geq d_2$.

- *Case 2: \tilde{w}_\bullet is ϵ_1 -far in relative Hamming distance from C_\bullet .*

This case is analogous to the previous one.

- *Case 3: there exist $w_o \in C_o$ and $w_\bullet \in C_\bullet$ with $\Delta_a^{\text{Ham}}(w_o, \tilde{w}_o) \leq \epsilon_1$ and $\Delta_a^{\text{Ham}}(w_\bullet, \tilde{w}_\bullet) \leq \epsilon_1$.*

In this case we follow the very end of the soundness analysis in Lemma 5.5's proof, replacing $\tilde{\alpha}_o, \tilde{\alpha}_\bullet$ there with $\tilde{w}_o, \tilde{w}_\bullet$, and conclude that the verifier accepts with probability at most $1 - \delta_1 + \gamma_1 + \epsilon_1$.

Summing up, in the first case the verifier's acceptance probability is at most $(1 - |f_1|^{-1}) \cdot \max\{e_2, \epsilon_1/4\} + |f_1|^{-1}$; similarly for the second case. In the third case the rejection probability is $1 - \delta_1 + \gamma_1 + \epsilon_1$, that of the canonical PCP consistency verifier. This completes the soundness analysis.

Zero knowledge. We construct a simulator S that yields perfect zero knowledge with knowledge bound k . Consider an instance-witness pair (\mathbf{x}, \mathbf{w}) in the relation \mathcal{R} , and a malicious verifier $\tilde{V} = (\tilde{V}_0, \tilde{V}_1)$ making at most k adaptive queries. The output of the simulator S when given as input \tilde{V} and \mathbf{x} , denoted $S(\tilde{V}, \mathbf{x})$, has to be identically distributed to $\text{DPCPView}(k, \tilde{V}, P, \mathbf{x}, \mathbf{w})$, which is the view of \tilde{V}_1 in its execution when given input \mathbf{x} and when allowed to make a total of $k(n)$ adaptive queries to π_0, π_1 , where $\pi_0 := P_0(\mathbf{x}, \mathbf{w})$ and $\pi_1 := P_1(\mathbf{x}, \mathbf{w}, V_0^{\pi_0}(\mathbf{x}))$. In fact, we will prove a

stronger statement: the output of the simulator continues to exactly match the view of the verifier, interacting with the honest prover, even if the verifier is allowed unbounded access to π_1 , provided that \tilde{V} makes at most k queries to π_0 .

We now discuss how S works. At a high level, S treats \tilde{V} as a black box, running it once without rewinding; along the way, S samples suitable answers for each query (as discussed below); when \tilde{V} halts, S outputs all the answers and \tilde{V} 's message (which together form the view of the verifier). The simulator S runs in strict polynomial time, without ever aborting. We now describe how S answers each query.

The simulator S maintains a proof string $\pi^S = (\pi_0^S, \pi_1^S)$ where $\pi_0^S = (w_\circ^S, v_\circ^S, w_\bullet^S, v_\bullet^S)$ and $\pi_1^S = (z_\circ^S, z_\bullet^S, \pi_\circ^S, \pi_\bullet^S)$. Initially, π^S is unspecified at all locations. The simulator maintains two sets, $A_\circ \subseteq [\ell_\circ]$ and $A_\bullet \subseteq [\ell_\bullet]$, that are initially empty. At any stage of the simulation, A_\circ and A_\bullet contain the set of indices of w_\circ^S, v_\circ^S and w_\bullet^S, v_\bullet^S that are currently specified, respectively.

We now discuss how S adaptively specifies locations in π^S . We distinguish between two parts of the simulation: before the point when \tilde{V} sends his message ρ , and only queries to π_0 are possible; and afterwards, when queries to both π_0 and π_1 are possible. When \tilde{V} requests an index $j \in [\ell_\circ]$ of w_\circ or v_\circ that is unspecified, the simulator will update both w_\circ^S and v_\circ^S at j . The value $v_\circ^S[j]$ is set to a random field element, unless its value is already determined by previously specified indices together with the linear constraint that v_\circ^S must be a projection of an element of C_\circ . This implies that *at any stage of the simulation, the specified part of v_\circ^S is a projection of a uniform element of C_\circ* . During the first part of the simulation, $w_\circ^S[j]$ is simply updated to a random field element. During the second part, it is updated according to the linear relation $\tilde{\rho}_\circ \cdot w_\circ^S + v_\circ^S = z_\circ^S$. The strings w_\bullet^S and v_\bullet^S are updated similarly. An important thing to note is that to update $w_\bullet^S[j]$ the simulator first queries all indices of w_\circ^S on which $w_\bullet^S[j]$ depends. In more detail, the simulator works as follows.

- *Simulating answers to $\pi_0 = (w_\circ \| v_\circ \| w_\bullet \| v_\bullet)$, before \tilde{V} outputs $\tilde{\rho} = (\tilde{\rho}_\circ, \tilde{\rho}_\bullet)$.*
 1. Initialize $A_\circ := \emptyset, A_\bullet := \emptyset$.
 2. For a query $j \in [\ell_\circ]$ to either $w_\circ[j]$ or $v_\circ[j]$ with $j \notin A_\circ$:
 - Check if $v_\circ^S[j]$ is determined by a linear constraint of C_\circ together with the values $\{v_\circ^S[j'] \mid j' \in A_\circ\}$ and, if so, update $v_\circ^S[j]$ to that value. Otherwise, update $v_\circ^S[j]$ to a random field element $\gamma \in f_1$.
 - Set $w_\circ^S[j]$ to a random field element $\beta \in f_1$.
 - Add j to A_\circ .
 3. For a query $j \in [\ell_\bullet]$ to either $w_\bullet[j]$ or $v_\bullet[j]$ with $j \notin A_\bullet$:
 - Check if $v_\bullet^S[j]$ is determined by a linear constraint of C_\bullet together with the values $\{v_\bullet^S[j'] \mid j' \in A_\bullet\}$ and, if so, update $v_\bullet^S[j]$ to that value. Otherwise, update $v_\bullet^S[j]$ to a random field element $\gamma \in f_1$.
 - Do the following: (i) compute the set $I_j \subseteq [\ell_\bullet]$ of locations on which $g(w_\bullet^S)[j]$ depends (see Definition 5.2); (ii) deduce $w_\circ^S|_{I_j}$ by querying each $i \in I_j$ according to Step 2; and (iii) set $w_\bullet^S[j] := g(w_\circ^S|_{I_j})$.
 - Add j to A_\bullet .
- *Simulating answers to $\pi_0 = (w_\circ \| v_\circ \| w_\bullet \| v_\bullet)$ and $\pi_1 = (z_\circ \| z_\bullet \| \pi_\circ \| \pi_\bullet)$, after \tilde{V} outputs $\tilde{\rho} = (\tilde{\rho}_\circ, \tilde{\rho}_\bullet)$.*
 4. After receiving $\tilde{\rho} = (\tilde{\rho}_\circ, \tilde{\rho}_\bullet)$, immediately do the following:
 - (a) sample a random $z_\circ^S \in C_\circ$ under the constraint “ $z_\circ^S[i] = \tilde{\rho}_\circ w_\circ^S[i] + v_\circ^S[i]$ for all $i \in A_\circ$ ”;
 - (b) sample a random $z_\bullet^S \in C_\bullet$ under the analogous constraint;
 - (c) compute $\pi_\circ^S := P_\circ(C_\circ, z_\circ^S)$;
 - (d) compute $\pi_\bullet^S := P_\bullet(C_\bullet, z_\bullet^S)$.
 5. All queries to $z_\circ, z_\bullet, \pi_\circ, \pi_\bullet$ are answered according to the values specified in Step 4.
 6. For a query $j \in [\ell_\circ]$ to either $w_\circ[j]$ or $v_\circ[j]$ with $j \notin A_\circ$:
 - Check if $v_\circ^S[j]$ is determined by a linear constraint of C_\circ together with the values $\{v_\circ^S[j'] \mid j' \in A_\circ\}$ and, if so, update $v_\circ^S[j]$ to that value. Otherwise, update $v_\circ^S[j]$ to a random field element $\gamma \in f_1$.
 - If $\tilde{\rho}_\circ = 0$, set $w_\circ^S[j]$ uniformly at random, otherwise set $w_\circ^S[j] := (z_\circ^S[j] - v_\circ^S[j])/\tilde{\rho}_\circ$.
 - Add j to A_\circ .
 7. For a query $j \in [\ell_\bullet]$ to either $w_\bullet[j]$ or $v_\bullet[j]$ with $j \notin A_\bullet$:
 - Do the following: (i) compute the set $I_j \subseteq [\ell_\bullet]$ of locations on which $g(w_\bullet^S)[j]$ depends (see Definition 5.2); (ii) deduce $w_\circ^S|_{I_j}$ by querying each $i \in I_j$ according to Step 6; and (iii) set $w_\bullet^S[j] := g(w_\circ^S|_{I_j})$.
 - Set $v_\bullet^S[j] := z_\bullet^S[j] - \tilde{\rho}_\bullet \cdot w_\bullet^S[j]$.

- Add j to A_\bullet .

We claim that the above simulation achieves perfect zero-knowledge, that is, $S(\tilde{V}, \mathbf{x})$ is identically distributed to $\text{DPCPView}(k, \tilde{V}, P, \mathbf{x}, \mathbf{w})$. We show that the distribution of answers provided by the simulation to \tilde{V} is the same as the distribution of answers obtained by \tilde{V} from the oracles provided by the honest prover.

For $i \in [k]$ we fix the following notation:

- Let $A_{\bullet(i)} \subseteq [\ell_\bullet]$ be the set of indices of either w_\bullet or v_\bullet that have been queried by \tilde{V} by step i of the interaction with the honest prover. Similarly, let $A_{\bullet^S(i)} \subseteq [\ell_\bullet]$ be the set of indices of either w_\bullet or v_\bullet that have been queried by \tilde{V} by step i of the interaction with the simulator.
- Let $A_{\circ(i)} \subseteq [\ell_\circ]$ be the set of indices of either w_\circ or v_\circ that have been queried by \tilde{V} by step i of the interaction with the honest prover; together with elements of $[\ell_\circ]$ influencing an index $j \in A_{\bullet(i)}$ by application of g , i.e., $j' \in [\ell_\circ]$ such that $g(a)[j']$ depends on $a_{j'}$ for some $j' \in A_{\bullet(i)}$. Analogously, let $A_{\circ^S(i)} \subseteq [\ell_\circ]$ be the set of indices of either w_\circ or v_\circ that have been queried by \tilde{V} by step i of the interaction with the simulator; together with elements of $[\ell_\circ]$ influencing an index $j \in A_{\bullet^S(i)}$ by application of g , i.e., $j' \in [\ell_\circ]$ such that $g(a)[j']$ depends on $a_{j'}$ for some $j' \in A_{\bullet^S(i)}$.
- Let $\pi_{0(i)} := (w_{\circ(i)}, v_{\circ(i)}, w_{\bullet(i)}, v_{\bullet(i)})$ be the projection of π_0 where w_\circ and v_\circ are restricted to $A_{\circ(i)}$ and w_\bullet and v_\bullet are restricted to $A_{\bullet(i)}$.
- Let $\pi_{\circ^S(i)} := (w_{\circ^S(i)}, v_{\circ^S(i)}, w_{\bullet^S(i)}, v_{\bullet^S(i)})$ be the projection of π_\circ^S where w_{\circ^S} and v_{\circ^S} are restricted to $A_{\circ^S(i)}$ and w_{\bullet^S} and v_{\bullet^S} are restricted to $A_{\bullet^S(i)}$. That is, $\pi_{\circ^S(i)}$ is the part of π_\circ^S updated by the simulator up to step i .

Note that the random variables $A_{\bullet(i)}$ and $A_{\circ(i)}$ depend on the randomness r_P of the honest prover, and the randomness $r_{\tilde{V}}$ of \tilde{V} . The random variables $A_{\bullet^S(i)}$ and $A_{\circ^S(i)}$ depend on the randomness r_S of the simulator and $r_{\tilde{V}}$. The core of the proof is the following lemma:

Lemma 6.5. *For every $i \in [k]$, we have $(r_{\tilde{V}}, \pi_{\circ^S(i)}, \pi_1^S) \equiv (r_{\tilde{V}}, \pi_{0(i)}, \pi_1)$.*

The lemma implies that $S(\tilde{V}, \mathbf{x}) \equiv \text{DPCPView}(k, \tilde{V}, P, \mathbf{x}, \mathbf{w})$ because the view of \tilde{V} is a deterministic function of $r_{\tilde{V}}$ and the parts of the proof seen by \tilde{V} are identically distributed (by the lemma). We now prove the lemma.

Proof. We prove a stronger claim by induction on i . We will have that for every $i \in [k]$ the following equivalences hold: (i) $A_{\circ^S(i)} \equiv A_{\circ(i)}$; (ii) $A_{\bullet^S(i)} \equiv A_{\bullet(i)}$; and (iii) $(r_{\tilde{V}}, \pi_{\circ^S(i)}, \pi_1^S) \equiv (r_{\tilde{V}}, \pi_{0(i)}, \pi_1)$.

The claim holds for $i = 0$. For any $i > 0$ we assume the claim for $i - 1$ and prove it for i as follows. Note that $A_{\circ(i)}$ is a deterministic function f of the verifier randomness and the values of locations queried at previous steps. That is,

$$A_{\circ^S(i)} \equiv f(r_{\tilde{V}}, \pi_{\circ^S(i-1)}, \pi_1^S) \equiv f(r_{\tilde{V}}, \pi_{0(i-1)}, \pi_1) \equiv A_{\circ(i)} .$$

Similarly, $A_{\bullet(i)}$ is a deterministic function f' of the verifier randomness and the values of locations queried at previous steps. That is,

$$A_{\bullet^S(i)} \equiv f'(r_{\tilde{V}}, \pi_{\circ^S(i-1)}, \pi_1^S) \equiv f'(r_{\tilde{V}}, \pi_{0(i-1)}, \pi_1) \equiv A_{\bullet(i)} .$$

We are left to prove that $(r_{\tilde{V}}, \pi_{\circ^S(i)}, \pi_1^S) \equiv (r_{\tilde{V}}, \pi_{0(i)}, \pi_1)$. We do this by going over the sequence of corresponding variables and show that they are identically distributed conditioned on fixing the previous corresponding variables to the same values. Clearly, the first components are identically distributed: $r_{\tilde{V}} \equiv r_{\tilde{V}}$. We now show that the same holds for the rest.

- $(w_{\circ^S(i)} | r_{\tilde{V}}) \equiv (w_{\circ(i)} | r_{\tilde{V}})$. The variable $w_{\circ^S(i)} | r_{\tilde{V}}$ is simply a uniform sequence of elements of size $|A_{\circ(i)}|$. Because w_\circ is t -wise independent and $|A_{\circ(i)}| \leq t$, we have that $w_{\circ(i)} | r_{\tilde{V}}$ is also uniformly random.
- $(v_{\circ^S(i)} | (r_{\tilde{V}}, w_{\circ^S(i)})) \equiv (v_{\circ(i)} | (r_{\tilde{V}}, w_{\circ(i)}))$. This holds because both random variables are a projection of uniformly random element of C_\circ to $A_{\circ(i)}$.
- $(w_{\bullet^S(i)} | (r_{\tilde{V}}, w_{\circ^S(i)}, v_{\circ^S(i)})) \equiv (w_{\bullet(i)} | (r_{\tilde{V}}, w_{\circ(i)}, v_{\circ(i)}))$. This holds because $w_{\bullet^S(i)}$ is a deterministic function g' of $w_{\circ^S(i)}$: $w_{\bullet(i)} = g'(w_{\circ(i)})$ and we are conditioning on $w_{\circ^S(i)}$ and $w_{\circ(i)}$ being fixed to identical values.

- $(v_{\bullet(i)}^S | (r_{\tilde{V}}, w_{\circ(i)}^S, v_{\circ(i)}^S, w_{\bullet(i)}^S)) \equiv (v_{\bullet(i)} | (r_{\tilde{V}}, w_{\circ(i)}, v_{\circ(i)}, w_{\bullet(i)}))$. This holds because both random variables are a projection of uniformly random element of C_{\bullet} to $A_{\bullet(i)}$.
- $(z_{\circ}^S | (r_{\tilde{V}}, w_{\circ(i)}^S, v_{\circ(i)}^S, w_{\bullet(i)}^S, v_{\bullet(i)}^S)) \equiv (z_{\circ} | (r_{\tilde{V}}, w_{\circ(i)}, v_{\circ(i)}, w_{\bullet(i)}, v_{\bullet(i)}))$. This holds because both are a uniform element a of C_{\circ} conditioned on $\rho_0 \cdot w_{\circ}[j] + v_{\circ}[j] = a[j]$ for every $j \in A_{\circ(i)}$.
- $(z_{\bullet}^S | (r_{\tilde{V}}, w_{\circ(i)}^S, v_{\circ(i)}^S, w_{\bullet(i)}^S, v_{\bullet(i)}^S, z_{\circ}^S)) \equiv (z_{\bullet} | (r_{\tilde{V}}, w_{\circ(i)}, v_{\circ(i)}, w_{\bullet(i)}, v_{\bullet(i)}, z_{\circ}))$. This holds because both are a uniform element a of C_{\bullet} conditioned on $\rho_1 \cdot w_{\bullet}[j] + v_{\bullet}[j] = a[j]$ for every $j \in A_{\bullet(i)}$.

Recall that $\pi_{0(i)} = (w_{\circ(i)} \| v_{\circ(i)} \| w_{\bullet(i)} \| v_{\bullet(i)})$ and $\pi_1 = (z_{\circ} \| z_{\bullet} \| \pi_{\circ} \| \pi_{\bullet})$ and similarly for $\pi_{0(i)}^S$ and π_1^S ; the above then implies that: (i) $r_{\tilde{V}} \equiv r_{\tilde{V}}$; (ii) $(\pi_{0(i)}^S | r_{\tilde{V}}) \equiv (\pi_{0(i)} | r_{\tilde{V}})$; and (iii) $(\pi_1^S | (r_{\tilde{V}}, \pi_{0(i)}^S)) \equiv (\pi_1 | (r_{\tilde{V}}, \pi_{0(i)}))$, that is, $(r_{\tilde{V}}, \pi_{0(i)}^S, \pi_1^S) \equiv (r_{\tilde{V}}, \pi_{0(i)}, \pi_1)$. \square

We conclude the discussion of the simulator by examining the time complexity of the simulation. Most steps of the simulation require (a) sampling a random field element and, possibly, (b) solving a linear system with a polynomial number of equations. The only expensive part of the simulation is Step 4, because it requires sampling random codewords in C_{\circ} and C_{\bullet} , as well as computing PCPP proofs for these two codewords. Provided that ℓ_1, c_1, r_1 are polynomially bounded, the entire simulation also runs in polynomial time in the instance size n . (The definition of zero knowledge in Section 3 prescribes, as typically done, a simulator that runs in expected probabilistic polynomial time; our simulator runs in strict probabilistic polynomial time.)

Other parameters. The remaining parameters are straightforward to establish. The duplex PCP construction does not change the alphabet, so $a = f_1$ (which also equals a_2). The proof length, and the running times of the prover and verifier are the sum of the same measures of the duplex PCP's components: the PCP proofs have $l = 2l_2(\ell_1) + 6\ell_1$ symbols in total, are produced in time $tp = 2tp_2(\ell_1) + (c_1 + 5)\ell_1 + r_1$, and are verified in time $tv = 2tv_2(\ell_1) + c_1 + \log \ell_1$. The duplex PCP verifier makes $q_1 + 7$ queries on top of those made by the two PCPP verifiers, so its query complexity is $q = 2q_2(\ell_1) + q_1 + 7$. The $q_1 + 7$ additional queries are non-adaptive; so if the PCPP verifiers are non-adaptive, so is the duplex PCP verifier (i.e., $na = na_2$). \square

Remark 6.6 (Necessity of duplex PCP). Claim 2.1 suggests a way to modify our duplex PCP construction so that the prover sends the oracle proof strings in one message (as in a standard PCP) rather than over two rounds of interaction. Namely, assuming w_{\circ} is ϵ -far from C_{\circ} , there is at most one verifier message ρ_{\circ} that makes z_{\circ} be $\epsilon/2$ -close to C_{\circ} . Hence the verifier can fix in advance r distinct $\rho_{\circ}^{(1)}, \dots, \rho_{\circ}^{(r)}$ and have the prover send, in addition to π_0 , the oracle strings $\pi_{\circ}^{(1)}, \dots, \pi_{\circ}^{(r)}$, where each $\pi_{\circ}^{(j)}$ is a PCPP proof of proximity for $z_{\circ}^{(j)} := \rho_{\circ}^{(j)} w_{\circ} + v_{\circ}$. The (honest) verifier now samples $j \in [r]$ and continues as before with $\pi_{\circ} := \pi_{\circ}^{(j)}$. This modification yields perfect completeness; also it yields small soundness error, because at most one of $z_{\circ}^{(1)}, \dots, z_{\circ}^{(r)}$ is $\epsilon/2$ -close to C_{\circ} .

Unfortunately, such a modification does not achieve zero knowledge against malicious verifiers, as we now explain. If the PCP proofs of proximity $\pi_{\circ}^{(j)}$ are computed via a linear transformations applied to $z_{\circ}^{(j)}$ (as is the case for the PCPPs mentioned in Section 2.4), then a malicious verifier can use $\pi_{\circ}^{(1)}$ and $\pi_{\circ}^{(2)}$ to recover individual positions of $\pi'_{\circ} := P_{\circ}(C_{\circ}, w_{\circ})$, as follows: $\pi'_{\circ}[i] := (\rho_{\circ}^{(1)} \cdot \pi_{\circ}^{(2)}[i] - \rho_{\circ}^{(2)} \cdot \pi_{\circ}^{(1)}[i]) / (\rho_{\circ}^{(1)} - \rho_{\circ}^{(2)})$. Certain entries of π'_{\circ} can (and do, for the aforementioned PCPPs) depend on more than t locations of $w_{\circ} := u' + \alpha_{\circ}$. Thus a security proof that makes black-box use of the PCPP cannot claim that the contribution from u' is uniformly random: a linear combination of more than t random variables from a t -wise independent ensemble is not necessarily random. Hence we cannot prove that the original assignment α_{\circ} is hidden by u' . In fact, the PCPPs considered in this paper reveal positive amount of knowledge in this case.

However, the above modification is zero knowledge against *honest* verifiers: an honest execution of the modified verifier above involves querying only one of the proof strings $\pi_{\circ}^{(i)}$ so that there is a corresponding duplex PCP interaction that is zero knowledge (as established in the proof above).

7 From NTIME to randomizable linear algebraic CSPs

We show an efficient reduction from NTIME to randomizable linear algebraic CSPs; we proceed as follows.

- $\mathcal{R}_{\text{AP}} \& \mathcal{R}_{\text{GAP}}$. In Section 7.1, we define *algebraic problems*, implicit in several influential works on PCPs and IP [BFL91, LFKN92, BFLS91, ALM⁺98] and explicitly defined in [PS94, Sze99, HS00]. Afterward, we define *group-preserving algebraic problems*, a new “symmetric” variant of algebraic problems that not only are powerful enough to efficiently capture NTIME but are also naturally “randomizable”, as discussed below.
- $\mathcal{R}_{\text{AP}} \rightarrow \mathcal{R}_{\text{LA}}$. In Section 7.2 (see Lemma 7.3), we show that algebraic problems are a sublanguage of linear algebraic CSPs. This observation shows that the techniques of this paper could potentially be applied to many PCP systems (e.g., those in [BFL91, LFKN92, BFLS91, ALM⁺98, Sze99, HS00, BGH⁺05, BGH⁺06, BS08, BCGT13b, BV14] to name a few) and also provides a “warm up” for the next item.
- $\mathcal{R}_{\text{GAP}} \rightarrow \mathcal{R}_{\text{RLA}}$. In Section 7.3 (see Lemma 7.5), we show an efficient reduction from group-preserving algebraic problems to randomizable linear algebraic CSPs. In other words, the property of group preservation allows the corresponding linear algebraic CSPs to be randomizable.
- $\text{NTIME} \rightarrow \mathcal{R}_{\text{GAP}}$. In Section 7.4 (see Lemma 7.6), we show an efficient reduction from NTIME to group-preserving algebraic problems.
- $\text{NTIME} \rightarrow \mathcal{R}_{\text{RLA}}$. In Section 7.5 (see Theorem 7.9), we explain how to combine the above to obtain the efficient reduction from NTIME to randomizable linear algebraic CSPs.

7.1 Algebraic problems and group preservation

The definition below of **algebraic problems** is essentially due to [PS94] (though the term “algebraic problem” is from [HS00]); variants of it appear in later works such as [Spi95, Sze99, HS00, BS08, BGH⁺06, BCGT13a, BV14].

Definition 7.1 (\mathcal{R}_{AP}). *Given functions $F: \mathbb{N} \rightarrow \mathcal{F}$, and $h, m, \eta, d, \sigma: \mathbb{N} \rightarrow \mathbb{N}$, the relation*

$$\mathcal{R}_{\text{AP}}[F, h, m, \eta, d, \sigma]$$

consists of instance-witness pairs $(\mathfrak{x}, \mathfrak{w})$ satisfying the following.

- *The instance \mathfrak{x} is a tuple $(1^n, H, Q, \vec{N})$ where:*
 - *H is a subset of $F(n)$ with cardinality $h(n)$;*
 - *Q is a polynomial in $F(n)[X_1, \dots, X_{m(n)}, Y_1, \dots, Y_{\eta(n)}]$ such that (i) it has degree less than $h(n)$ in each variable X_i , (ii) it has total degree at most $d(n)$ when viewed as a polynomial in the variables $Y_1, \dots, Y_{\eta(n)}$ with coefficients in $F(n)[X_1, \dots, X_{m(n)}]$, (iii) it can be evaluated by an arithmetic circuit of size $\sigma(n)$;*
 - *$\vec{N} = (N_1, \dots, N_{\eta(n)})$ and each $N_i: F(n)^{m(n)} \rightarrow F(n)^{m(n)}$ is an invertible affine function.*
- *The witness \mathfrak{w} is a polynomial A in $F(n)[X_1, \dots, X_{m(n)}]$ with degree less than $h(n)$ in each variable X_i .*
- *The instance \mathfrak{x} and witness \mathfrak{w} jointly satisfy the following:*

$$\text{for every } \alpha \in H^{m(n)}, (Q \circ A \circ \vec{N})(\alpha) = 0 \quad (1)$$

where

$$(Q \circ A \circ \vec{N})(X) := Q(X_1, \dots, X_{m(n)}, A(N_1(X_1, \dots, X_{m(n)})), \dots, A(N_{\eta(n)}(X_1, \dots, X_{m(n)}))) \quad (2)$$

Next, we define **group-preserving algebraic problems**, a family of algebraic problems in which the set H is a subgroup of $F(n)$ and the neighbor functions act on the product group $H^{m(n)}$. The additional symmetry enables a reduction to randomizable linear algebraic CSPs, which give rise to zero knowledge duplex PCPs. We believe that group-preserving algebraic problems may find applications in the study of PCPs beyond their use in this paper.

Definition 7.2 (\mathcal{R}_{GAP}). The relation $\mathcal{R}_{\text{GAP}}[F, h, m, \eta, d, \sigma]$ is the sub-relation of $\mathcal{R}_{\text{AP}}[F, h, m, \eta, d, \sigma]$ obtained via restriction to instances that are group preserving. An instance $\mathfrak{x} = (1^n, H, Q, \vec{N})$ is group preserving if: (i) H is an additive or a multiplicative subgroup of $F(n)$; (ii) each $N_i: F(n)^{m(n)} \rightarrow F(n)^{m(n)}$ in \vec{N} can be identified with an element χ_i in $H^{m(n)}$ such that $N_i(x) = \chi_i \odot x$, where \odot denotes the group operation of the product group $H^{m(n)}$.

We also write $\mathcal{R}_{\text{GAP}}[F, h, m, \eta, d, \sigma, +]$ to denote the further restriction to instances that are additively group preserving (i.e., H is an additive subgroup); similarly, we write $\mathcal{R}_{\text{GAP}}[F, h, m, \eta, d, \sigma, \times]$ to denote the restriction to instances that are multiplicatively group preserving.

7.2 Algebraic problems naturally reduce to linear algebraic CSPs

The next lemma says that algebraic problems naturally reduce to linear algebraic CSPs, i.e., there is an efficient reduction from the relation $\mathcal{R}_{\text{AP}}[F, h, m, \eta, d, \sigma]$ to the relation $\mathcal{R}_{\text{LA}}[f, \ell, \rho, \delta, q, c, \gamma, \epsilon]$, for suitable parameter choices.

Lemma 7.3 ($\mathcal{R}_{\text{AP}} \rightarrow \mathcal{R}_{\text{LA}}$). For every $F: \mathbb{N} \rightarrow \mathcal{F}$, $h, m, \eta, d, \sigma: \mathbb{N} \rightarrow \mathbb{N}$, $\epsilon: \mathbb{N} \rightarrow (0, 1)$, and $\mathcal{R} \subseteq \mathcal{R}_{\text{AP}}[F, h, m, \eta, d, \sigma]$ there exist a relation \mathcal{R}' and algorithms inst , wit_1 , wit_2 satisfying the following conditions:

- **EFFICIENT REDUCTION.** For every instance \mathfrak{x} , letting $\mathfrak{x}' := \text{inst}(\mathfrak{x})$:
 - for every witness \mathfrak{w} , if $(\mathfrak{x}, \mathfrak{w}) \in \mathcal{R}$ then $(\mathfrak{x}', \text{wit}_1(\mathfrak{x}, \mathfrak{w})) \in \mathcal{R}'$;
 - for every witness \mathfrak{w}' , if $(\mathfrak{x}', \mathfrak{w}') \in \mathcal{R}'$ then $(\mathfrak{x}, \text{wit}_2(\mathfrak{x}, \mathfrak{w}')) \in \mathcal{R}$.

Moreover, inst runs in time $\text{poly}(|\mathfrak{x}|)$, wit_1 in time $\text{poly}(|\mathfrak{x}|) \cdot \tilde{O}(|\mathfrak{w}| \cdot \eta \cdot \sigma)$, and wit_2 in time $\text{poly}(|\mathfrak{x}|) \cdot \tilde{O}(|\mathfrak{w}'|)$.

- **LINEAR ALGEBRAIC CSP.** The relation \mathcal{R}' is a subset of

$$\mathcal{R}_{\text{LA}} \left[\begin{array}{l} \text{field} \\ \text{block length} \\ \text{rate} \\ \text{relative distance} \\ \text{map locality} \\ \text{map efficiency} \\ \text{map sampling I} \\ \text{map sampling II} \end{array} \begin{array}{l} f = F \\ \ell = |F|^m \\ \rho = \left(\frac{hd}{|F|}\right)^m \\ \delta = 1 - \frac{hd}{|F|} \\ q = \eta \\ c = \sigma + \eta \\ \gamma = \eta\epsilon \\ \epsilon \end{array} \right].$$

- **RM CODES.** If $\mathfrak{x} = (1^n, H, Q, \vec{N})$ then $\text{inst}(\mathfrak{x}) = (1^n, C_\circ, C_\bullet, g)$ with

- $C_\circ = \text{RM} \left[F(n), F(n), m(n), \frac{h(n)}{|F(n)|} \right]$ and
- $C_\bullet = \text{VRM} \left[F(n), F(n), m(n), \frac{h(n)d(n)}{|F(n)|}, H \right]$.

Proof of Lemma 7.3. Let $\mathfrak{x} = (1^n, H, Q, \vec{N})$ be an instance of $\mathcal{R}_{\text{AP}}[F, h, m, \eta, d, \sigma]$, and construct $\mathfrak{x}' := \text{inst}(\mathfrak{x}) = (1^n, C_\circ, C_\bullet, g)$ with C_\circ, C_\bullet as above and g being the function that maps $F(n)[X_1, \dots, X_{m(n)}]$ to $F(n)^{F(n)^{m(n)}}$ as follows: given A in $F(n)[X_1, \dots, X_{m(n)}]$ and $\omega \in F(n)^{m(n)}$, the ω -th coordinate of $g(A)$ equals to $(Q \circ A \circ \vec{N})(\omega)$. We first argue that \mathfrak{x}' is an instance of $\mathcal{R}_{\text{LA}}[f, \ell, \rho, \delta, q, c, \gamma, \epsilon]$.

First, C_\circ and C_\bullet are linear error correcting codes with block length at most $\ell := |F|^m$, rate at most $\rho := \max\left\{\left(\frac{h}{|F|}\right)^m, \left(\frac{hd}{|F|}\right)^m\right\}$, and relative distance at least $\delta := \min\left\{1 - \frac{h}{|F|}, 1 - \frac{hd}{|F|}\right\}$ over the same field F . (See Section 2.4.)

By construction, the function g is q -local with $q := \eta$ and c -efficient with $c := \sigma + \eta$; moreover, g is (γ, ϵ) -sampling with $\gamma := \eta\epsilon$, as we now explain. (See Definition 5.1 for definitions of these properties.) For every $\omega \in F^m$, I_ω denotes the set of indices in F^m that $g(\cdot)[\omega]$ depends on; for the g above, I_ω equals $\{N_1(\omega), \dots, N_\eta(\omega)\}$. For every $\omega' \in F^m$ and $\omega \in F^m$, if $\omega' \in I_\omega$ then $\omega \in \{N_1^{-1}(\omega'), \dots, N_\eta^{-1}(\omega')\}$. Hence, the number of ω 's with $\omega' \in I_\omega$ is at most η , because each N_i is invertible. We deduce that $\Pr[I_\omega \cap I \neq \emptyset \mid \omega \leftarrow F^m] \leq (\eta \cdot |I|) / |F|^m \leq \eta\epsilon$.

Finally, $C_\bullet \cup g(C_\circ)$ has relative distance at least δ because it is a subset of $\text{RM}[F, F, m, \frac{hd}{|F|}]$. This claim is immediate for C_\bullet ; for $g(C_\circ)$, it follows from the fact that $Q \circ A \circ \vec{N}$ has, in each variable, a degree that is at most a multiplicative factor of d larger than the degree of A .

We conclude the proof by explaining how one obtains the two witness maps $\text{wit}_1, \text{wit}_2$. For wit_1 , suppose that $\mathfrak{w} = A \in F[X_1, \dots, X_m]$ is a witness for \mathfrak{x} ; then one can verify that $\mathfrak{w}' := (\alpha_\circ, \alpha_\bullet)$, where $\alpha_\circ := A$ and $\alpha_\bullet := Q \circ A \circ \vec{N}$, is a witness for \mathfrak{x}' ; α_\bullet can be efficiently obtained by first computing the evaluation of A on F^m (via an FFT), then computing the evaluation of $Q \circ A \circ \vec{N}$ on F^m (via point-to-point computation), and finally interpolating (via an inverse FFT). Conversely, for wit_2 , suppose that $\mathfrak{w}' = (\alpha_\circ, \alpha_\bullet)$ is a witness for \mathfrak{x}' ; then one can verify that $\mathfrak{w} := \alpha_\circ$ is a witness for \mathfrak{x} . \square

Remark 7.4. By combining the reduction from algebraic problems to linear algebraic CSPs (Lemma 7.3) with the canonical PCP for linear algebraic CSPs (Lemma 5.5), one can recover several known PCP results; these include PCPs with polylogarithmic verification of [BFLS91, BGH⁺06], constant-query PCPs of [ALM⁺98], quasilinear-size PCPs of [BS08] and, with some extra work, constant-rate PCPs based on algebraic geometry codes of [BKK⁺13]. Since the focus of this paper is zero knowledge, we do not pursue this direction further here.

7.3 From group-preserving algebraic problems to randomizable linear algebraic CSPs

The following lemma gives an efficient reduction from group-preserving algebraic problems to randomizable linear algebraic CSPs.

Lemma 7.5 ($\mathcal{R}_{\text{GAP}} \rightarrow \mathcal{R}_{\text{RLA}}$). *For every $F: \mathbb{N} \rightarrow \mathcal{F}$, $h, m, \eta, d, \sigma, t: \mathbb{N} \rightarrow \mathbb{N}$, $\delta, \epsilon: \mathbb{N} \rightarrow (0, 1)$ with $|F| \geq \hat{h}$, where \hat{h} denotes the smallest integral multiple of h that is greater than $\frac{(h+t)d}{1-\delta}$, and for any $\mathcal{R} \subseteq \mathcal{R}_{\text{GAP}}[F, h, m, \eta, d, \sigma]$ there exist a relation \mathcal{R}' and algorithms $\text{inst}, \text{wit}_1, \text{wit}_2$ satisfying the following conditions:*

- **EFFICIENT REDUCTION.** *For every instance \mathfrak{x} , letting $\mathfrak{x}' := \text{inst}(\mathfrak{x})$:*
 - *for every witness \mathfrak{w} , if $(\mathfrak{x}, \mathfrak{w}) \in \mathcal{R}$ then $(\mathfrak{x}', \text{wit}_1(\mathfrak{x}, \mathfrak{w})) \in \mathcal{R}'$;*
 - *for every witness \mathfrak{w}' , if $(\mathfrak{x}', \mathfrak{w}') \in \mathcal{R}'$ then $(\mathfrak{x}, \text{wit}_2(\mathfrak{x}, \mathfrak{w}')) \in \mathcal{R}$.*

Moreover, inst runs in time $\text{poly}(|\mathfrak{x}|)$, wit_1 in time $\text{poly}(|\mathfrak{x}|) \cdot \tilde{O}(|\mathfrak{w}| \cdot \eta \cdot \sigma)$, and wit_2 in time $\text{poly}(|\mathfrak{x}|) \cdot \tilde{O}(|\mathfrak{w}'|)$.

- **RANDOMIZABLE LINEAR ALGEBRAIC CSP.** *The relation \mathcal{R}' is a subset of*

$$\mathcal{R}_{\text{RLA}} \left[\begin{array}{ll} \text{field} & f = F \\ \text{block length} & \ell = \hat{h}^m \\ \text{rate} & \rho = (1 - \delta)^m \\ \text{relative distance} & \delta \\ \text{map locality} & q = \eta \\ \text{map efficiency} & c = \sigma + \eta \\ \text{map sampling I} & \gamma = \eta\epsilon \\ \text{map sampling II} & \epsilon \\ \text{randomizability} & t \\ \text{randomize time} & r = \tilde{O}(\hat{h}^m) \end{array} \right].$$

- **RM CODES WITH GROUP STRUCTURE.** *If $\mathfrak{x} = (1^n, H, Q, \vec{N})$ then $\text{inst}(\mathfrak{x}) = (1^n, C_\circ, C_\bullet, g)$ with*
 - $C_\circ := \text{RM} \left[F(n), \hat{H}, m(n), \frac{h(n)+t(n)}{\hat{h}(n)} \right]$ *and*
 - $C_\bullet := \text{VRM} \left[F(n), \hat{H}, m(n), \frac{(h(n)+t(n))d(n)}{\hat{h}(n)}, H \right]$,

where \hat{H} is a union of cosets of the subgroup H with cardinality $\hat{h}(n)$.

Proof of Lemma 7.5. Let $\mathfrak{x} = (1^n, H, Q, \vec{N})$ be an instance of $\mathcal{R}_{\text{GAP}}[F, h, m, \eta, d, \sigma]$. We construct an instance $\mathfrak{x}' := \text{inst}(\mathfrak{x}) = (1^n, C_\circ, C_\bullet, g)$ of $\mathcal{R}_{\text{RLA}}[f, \ell, \rho, \delta, q, c, \gamma, \epsilon, t, r]$ as follows.

Let \hat{H} be a subset of F that is a union of cosets of H with $|\hat{H}| = \hat{h}$ and $\hat{H} \cap H = \emptyset$. (This can be done as follows: let S be a subset of the quotient group F^\odot/H with cardinality $|S| = \hat{h}/h$ that does not include 1_\odot , where F^\odot denotes the additive or multiplicative group of F , depending on whether H is additive or multiplicative, and 1_\odot is the identity in H ; then set $\hat{H} := \{x \odot y \mid x \in S, y \in H\}$.) Analogously to the proof of Lemma 7.3, we define:

- $C_\circ := \text{RM} \left[F(n), \hat{H}, m(n), \frac{h(n)+t(n)}{h(n)} \right];$
- $C_\bullet := \text{VRM} \left[F(n), \hat{H}, m(n), \frac{(h(n)+t(n))d(n)}{h(n)}, H \right];$
- g to be the function that maps $F(n)[X_1, \dots, X_{m(n)}]$ to $F(n)^{\hat{H}^{m(n)}}$ as follows: given A in $F(n)[X_1, \dots, X_{m(n)}]$ and $\omega \in \hat{H}^{m(n)}$, the ω -th coordinate of $g(A)$ equals to $(Q \circ A \circ \vec{N})(\omega)$. Note that g is well-defined, i.e., $g(A)$ is a function from $\hat{H}^{m(n)}$ to $F(n)$; this follows from the group preservation property of \times (see Definition 7.2): for every $\omega \in \hat{H}^m$ and $i \in [\eta]$, it holds that $N_i(\omega) \subseteq \hat{H}^m$ because \hat{H} is a union of cosets of H and N_i multiplies every coordinate of ω by an element of H .

We first argue that \mathfrak{x}' constructed above is an instance of $\mathcal{R}_{\text{RLA}}[f, \ell, \rho, \delta, q, c, \gamma, \epsilon, t, r]$.

First, analogously to the proof of Lemma 7.3, we note that C_\circ and C_\bullet are linear error correcting codes with block length at most $\ell := \hat{h}^m$, rate at most $\max\left\{\left(\frac{h+t}{h}\right)^m, \left(\frac{(h+t)d}{h}\right)^m\right\} = \rho$, and relative distance at least $\min\left\{1 - \frac{h+t}{h}, 1 - \frac{(h+t)d}{h}\right\} = \delta$ over the same field F ; also, we deduce that g is q -local with $q := \eta$, c -efficient with $c := \sigma + \eta$, and (γ, ϵ) -sampling with $\gamma := \eta\epsilon$.

Next, recalling notions from Definition 6.1, we note that \mathfrak{x}' is t -randomizable in time $r := \tilde{O}(\hat{h}^m)$ because: (i) $C' := \text{VRM}[F(n), \hat{H}, m, \frac{h+t}{h}, H]$ is a subcode of C_\circ and it is t -wise independent due to Claim 2.4 (C' satisfies the hypotheses because $H \cap \hat{H} = \emptyset$ and $\hat{h} - h \geq \frac{(h+t)d}{1-\delta} - h \geq t$); and (ii) one can sample random elements from C' , C_\circ and C_\bullet in time $\tilde{O}(\hat{h}^m)$ by using FFT algorithms for multipoint evaluation (sampling from C_\circ involves selecting random coefficients for the multi-variate polynomial and then evaluating it on \hat{H}^m , while sampling from C' and C_\bullet relies on Alon's combinatorial Nullstellensatz [Alo99] to reduce this case to the previous case, as also done in [BS08]).

We conclude the proof by observing that necessary witness maps $\text{wit}_1, \text{wit}_2$ exist. Analogously to Lemma 7.3, if $\mathfrak{w} = A \in F(n)[X_1, \dots, X_{m(n)}]$ is a witness for \mathfrak{x} then $\text{wit}_1(\mathfrak{x}, \mathfrak{w})$ outputs $\mathfrak{w}' := (A, Q \circ A \circ \vec{N})$, which is a witness for \mathfrak{x}' ; conversely, if $\mathfrak{w}' = (\alpha_\circ, \alpha_\bullet)$ is a witness for \mathfrak{x}' then $\text{wit}_2(\mathfrak{x}, \mathfrak{w}')$ outputs $\mathfrak{w} := \alpha_\circ$, which is a witness for \mathfrak{x} . \square

7.4 An efficient reduction from NTIME to group-preserving algebraic problems

The following lemma gives an efficient reduction from NTIME to group-preserving algebraic problems in which instances are over fields of characteristic 2 and preserve additive groups.

Lemma 7.6 (NTIME $\rightarrow \mathcal{R}_{\text{GAP}}$). *For every $h, m, T, s: \mathbb{N} \rightarrow \mathbb{N}$ with $h(n)^{m(n)} = \Omega(T(n) \log T(n))$, $s(n) \geq 2$, and $\mathcal{R} \in \text{NTIME}(T)$ there exist a relation \mathcal{R}' and algorithms $\text{inst}, \text{wit}_1, \text{wit}_2$ satisfying the following conditions:*

- **EFFICIENT REDUCTION.** *For every instance \mathfrak{x} , letting $\mathfrak{x}' := \text{inst}(\mathfrak{x})$:*
 - *for every witness \mathfrak{w} , if $(\mathfrak{x}, \mathfrak{w}) \in \mathcal{R}$ then $(\mathfrak{x}', \text{wit}_1(\mathfrak{x}, \mathfrak{w})) \in \mathcal{R}'$;*
 - *for every witness \mathfrak{w}' , if $(\mathfrak{x}', \mathfrak{w}') \in \mathcal{R}'$ then $(\mathfrak{x}, \text{wit}_2(\mathfrak{x}, \mathfrak{w}')) \in \mathcal{R}$.*

Moreover, inst runs in time $\text{poly}(n + \log h(n) + m(n))$ and $\text{wit}_1, \text{wit}_2$ run in time $\tilde{O}(T(n))$.

- **GROUP PRESERVING ALGEBRAIC PROBLEM.** *The relation \mathcal{R}' is a subset of*

$$\mathcal{R}_{\text{GAP}} \left[\begin{array}{ll} \text{field} & F = \mathbb{F}_{2^{\log T + O(\log \log T) + s}} \\ \text{size of subset} & h \\ \text{copies of subset} & m \\ \text{number of neighbors} & \eta = \text{polylog}(T) \\ \text{circuit total degree} & d = O(1) \\ \text{circuit size} & \sigma = \text{poly}(n + \log T + m) \\ + & \end{array} \right].$$

Proof. The proof follows the reduction of [BV14, Section 2], which we summarize via the following statement.

Theorem 7.7 ([BV14]). For every $T: \mathbb{N} \rightarrow \mathbb{N}$ and $s: \mathbb{N} \rightarrow \mathbb{N}$ with $T(n) \geq n$ and $s(n) \geq 2$, and every relation $\mathcal{R} \in \mathbf{NTIME}(T)$ there exist a relation \mathcal{R}' and algorithms $\text{inst}, \text{wit}_1, \text{wit}_2$ satisfying the following conditions:

(1) EFFICIENT REDUCTION. For every instance \mathbf{x} , letting $\mathbf{x}' := \text{inst}(\mathbf{x})$:

- for every witness \mathbf{w} , if $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$ then $(\mathbf{x}', \text{wit}_1(\mathbf{x}, \mathbf{w})) \in \mathcal{R}'$;
- for every witness \mathbf{w}' , if $(\mathbf{x}', \mathbf{w}') \in \mathcal{R}'$ then $(\mathbf{x}, \text{wit}_2(\mathbf{x}, \mathbf{w}')) \in \mathcal{R}$.

Moreover, inst runs in time $\text{poly}(n + s + \log T)$, wit_1 in time $T \cdot \text{poly}(n + \log T)$, and wit_2 in time $\tilde{O}(T)$.

(2) ALGEBRAIC PROBLEM. The relation \mathcal{R}' is a subset of

$$\mathcal{R}_{\text{AP}} \left[\begin{array}{ll} \text{field} & F = \mathbb{F}_{2^{\log T + O(\log \log T) + s}} \\ \text{size of subset} & h = \tilde{O}(T) \\ \text{copies of subset} & m = 1 \\ \text{number of neighbors} & \eta = \text{poly}(n + \log T) \\ \text{circuit total degree} & d = O(1) \\ \text{circuit size} & \sigma = \text{poly}(n + \log T) \end{array} \right].$$

(3) STRUCTURE OF WITNESS EXTRACTOR. There exists a sequence of elements $\alpha_1, \dots, \alpha_T \in F$ such that, for every $\mathbf{w}' = A \in F[X]$ that satisfies \mathbf{x}' , we have $A(\alpha_i) \in \{0, 1\}$ and the witness extractor wit_2 outputs $\text{wit}_2(\mathbf{x}, \mathbf{w}') := (A(\alpha_1), \dots, A(\alpha_T))$. The vector $(\alpha_1, \dots, \alpha_T)$ can be computed from \mathbf{x} in time $\tilde{O}(T)$.

(4) STRUCTURE OF NEIGHBOR FUNCTIONS. The neighborhood functions $\vec{N} = (N_1, \dots, N_\eta)$ are such that:

- each $N_j(X)$ is of the form $N_j(X) := X + \gamma_j$ where $\gamma_j \in H$; and
- $\dim(\text{span}(\gamma_1, \dots, \gamma_\eta)) \leq \log T + O(\log \log T + \log n)$.

Remark 7.8. While not explicitly stated or stated differently, all four claims of Theorem 7.7 can be deduced by inspecting [BV14, Theorem 4.2] and its proof and [BV14, Theorem 2.2] and its proof.

The statement of [BV14, Theorem 4.2] does not involve group-preserving algebraic problems, but inspection of Theorem 4.2 and its proof reveals that one can achieve $\mathcal{R}' \in \mathcal{R}_{\text{GAP}}$ with $H \subseteq F$ being a linear space over \mathbb{F}_2 and other parameters remaining the same; that is, one can establish the first two properties.

The third property follows by inspection of the proof of Theorem 2.2 there. The fourth property holds because $\{\gamma_1, \dots, \gamma_\eta\} = \{S_i(\beta) \mid i \in \{1, 2, 3\}, \beta \in \{0, 1\}^{\log T}\}$ where the boolean circuits S_i are as stated in [BV14, Theorem 2.2]. The image of S_i is in $\{0, 1\}^{\log T + O(\log \log T + \log n)}$, which gives the bound above. \square

The above Theorem 7.7 is our Lemma 7.6 with $m = 1$ except for one difference: the neighborhood size η is $\text{poly}(n + \log T(n))$ in Theorem 7.7 whereas it is $\text{polylog}(T(n))$ in our Lemma 7.6. In our proof below, we decrease the neighborhood size from $\text{poly}(n + \log T(n))$ to $\text{polylog}(T(n))$, and generalize to arbitrary $m > 1$.

We first focus on reducing η for the case $m = 1$, via a padding argument; at the end of the proof we explain how to handle the general case.

For $z \in \{0, 1\}^\ell$ let $\bar{z} = \sum_{i=0}^{\ell-1} z_i 2^i$ be the integer whose binary representation is z . Let $\mathcal{R}_{\text{pad}} := \{(\mathbf{x}_{\text{pad}}, (\mathbf{x}, \mathbf{w})) \mid |\mathbf{x}| = \bar{\mathbf{x}}_{\text{pad}}, (\mathbf{x}, \mathbf{w}) \in \mathcal{R}\}$. Consider the nondeterministic machine M_{pad} that on input \mathbf{x}_{pad} of length n_{pad} and witness $\mathbf{w}_{\text{pad}} = (\mathbf{x}, \mathbf{w})$ accepts if and only if $|\mathbf{x}| = \bar{\mathbf{x}}_{\text{pad}}$ and $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$; M_{pad} decides \mathcal{R}_{pad} in time $T_{\text{pad}}(n_{\text{pad}}) = T(2^{n_{\text{pad}}}) + O(2^{n_{\text{pad}}})$. Applying Theorem 7.7 to \mathcal{R}_{pad} proves the theorem (again, for $m = 1$) for the relation \mathcal{R}_{pad} , because Theorem 7.7 gives us a relation $\mathcal{R}'_{\text{pad}} \subseteq \mathcal{R}_{\text{GAP}}$ where all parameters are as required; for instance $\eta_{\text{pad}}(n_{\text{pad}}) := \text{poly}(\lceil \log n_{\text{pad}} \rceil + \log T(n_{\text{pad}})) = \text{polylog}(T(n_{\text{pad}}))$ because $T(n_{\text{pad}}) \geq n_{\text{pad}}$. However, $\text{Lan}(\mathcal{R}_{\text{pad}})$ does not equal $\text{Lan}(\mathcal{R})$. We now address this issue.

We use \mathcal{R}_{pad} to obtain a relation \mathcal{R}_{fin} with $\text{Lan}(\mathcal{R}_{\text{fin}}) = \text{Lan}(\mathcal{R})$ and the aforementioned efficiency parameters of \mathcal{R}_{pad} , as described next; the output of our reduction was a relation $\mathcal{R}'_{\text{fin}} \subseteq \mathcal{R}_{\text{GAP}}$ capturing \mathcal{R}_{fin} as a group-reserving algebraic problem. The instance space of \mathcal{R}_{fin} is the same as that of \mathcal{R} and witnesses \mathbf{w}_{fin} are of the form (\mathbf{x}, \mathbf{w}) just like for \mathcal{R}_{pad} . We define:

$$\mathcal{R}_{\text{fin}} := \{(\mathbf{x}_{\text{fin}}, \mathbf{w}_{\text{fin}}) \mid \exists \mathbf{w} : \mathbf{w}_{\text{fin}} = (\mathbf{x}_{\text{fin}}, \mathbf{w}), (|\mathbf{x}_{\text{fin}}|, (\mathbf{x}_{\text{fin}}, \mathbf{w})) \in \mathcal{R}_{\text{pad}}\}.$$

The crux of the proof is reducing \mathcal{R}_{fin} to a group-preserving algebraic problem so that it retains the efficiency properties achieved by $\mathcal{R}'_{\text{pad}}$. We now describe our construction, turning to its efficiency analysis afterwards.

We first apply Theorem 7.7 on \mathcal{R}_{pad} obtaining a group-preserving algebraic problem $\mathcal{R}'_{\text{pad}}$ and the corresponding instance, witness and extractor maps. Instances of $\mathcal{R}'_{\text{pad}}$ are of the form $\mathbf{x}'_{\text{pad}} = (1^{n_{\text{pad}}}, H, Q, \vec{N})$ and witnesses of

$\mathcal{R}'_{\text{pad}}$ are of the form $w'_{\text{pad}} = (A)$, where A is an assignment polynomial. Property (3) of Theorem 7.7 implies that for every A that satisfies $x'_{\text{pad}} := \text{inst}(x)$, the sequence $w_{\text{pad}} := (A(\alpha_1), \dots, A(\alpha_{n+T}))$ belongs to $\{0, 1\}^{n+T}$ and $(x_{\text{pad}}, w_{\text{pad}}) \in \mathcal{R}_{\text{pad}}$. It follows that the sequence $x' := (A(\alpha_1), \dots, A(\alpha_n))$ belongs to $\text{Lan}(\mathcal{R})$. However, as explained above $\text{Lan}(\mathcal{R}) \neq \text{Lan}(\mathcal{R}_{\text{pad}})$: while $|x| = |x'|$, the x' does not necessarily equal x . We now explain how to modify \mathcal{R}_{pad} so it encodes a concrete instance x , not just fixes its length $|x|$.

We define an auxiliary instance x_{aux} of \mathcal{R}_{GAP} to be $x_{\text{aux}} := (1^{n_{\text{pad}}}, H_{\text{aux}}, Q_{\text{aux}}, \vec{N}_{\text{aux}})$, where we choose $H_{\text{aux}}, Q_{\text{aux}}, \vec{N}_{\text{aux}}$ as follows:

- $H_{\text{aux}} := \beta + H$, where $\beta \in F \setminus H$ (so H_{aux} is a β -coset of H);
- $\eta_{\text{aux}} := 1$ and $N_{\text{aux},1}(X) := X - \beta$; and
- $Q_{\text{aux}}(X, Y_1) := \sum_{i=1}^n Z_{\beta+(H \setminus \{\alpha_i\})}(X) \cdot (Y_1 - x_i)$.

It can be verified that A satisfies x_{aux} if and only if x' defined above equals x : indeed, when evaluating $(Q_{\text{aux}} \circ A \circ \vec{N}_{\text{aux}})(\alpha_i)$ all terms vanish except for $Z_{\beta+(H \setminus \{\alpha_i\})}(\alpha_i) \cdot (x_i - x'_i)$. We now use the fact that if set S consists of all points of an affine subspace, then the polynomial $Z_S(X)$ has at most $\dim(S) + 1$ non-zero terms [Ore33]. Therefore, each summand of Q_{aux} can be computed by an arithmetic circuit of size $O(\dim(H)^2)$ because each $\beta + (H \setminus \{\alpha_i\})$ can be partitioned in $\dim(H)$ affine subspaces as $\beta + (H \setminus \{\alpha_i\}) = \bigcup_{i=0}^{\dim(H)-1} (\beta_i + H_i)$, where the affine subspace $\beta_i + H_i$ has dimension i .¹

We combine x_{aux} and x'_{pad} to obtain $x'_{\text{fin}} := (1^{n_{\text{pad}}}, H_{\text{fin}}, Q_{\text{fin}}, \vec{N}_{\text{fin}})$ as follows:

- $H_{\text{fin}} := H \cup (\beta + H)$;
- $\eta_{\text{fin}} := \eta + 1$ and $\vec{N}_{\text{fin}} := \vec{N} \cup \vec{N}_{\text{aux}} = \{\vec{N}_1, \dots, \vec{N}_\eta, \vec{N}_{\eta+1} := (X - \beta)\}$; and
- $Q_{\text{fin}}(X, Y_1, \dots, Y_{\eta+1}) := Z_{\beta+H}(X) \cdot Q(X, Y_1, \dots, Y_\eta) + Z_H(X) \cdot Q_{\text{aux}}(X, Y_{\eta+1})$.

It can also be verified that A satisfies x_{aux} if and only if A encodes instance x' and witness w' with $(x', w') \in \mathcal{R}$ and $x = x'$. We now verify that all parameters of $\mathcal{R}'_{\text{fin}}$ match those in the statement of Lemma 7.6.

All parameters are closely tied to those of $\mathcal{R}'_{\text{pad}}$ and are established as follows. F is inherited from $\mathcal{R}'_{\text{pad}}$ and is $\mathbb{F}_{2^{\log T + O(\log \log T) + s}}$. The vanishing set has size $h_{\text{fin}} = \tilde{O}(T)$, which matches the theorem statement as $m = 1$. Number of neighborhood functions is $\text{polylog}(T)$ and total degree of Y in Q_{fin} is $d = O(1)$, both as required. Since $\beta + H$ is an affine space and H is a linear space we see that Q_{fin} can be computed by a circuit of size $\sigma_{\text{fin}} = \sigma_{\text{pad}} + \text{polylog}(T)$ because $\dim(H) = O(\log T(n))$ (again, using the theorem of [Ore33]). Finally, H_{fin} is an additive group (it is union of a linear space space and its shift and $\text{char}(F) = 2$) and it can be verified that each $\vec{N}_{\text{fin},j}$ satisfies $\vec{N}_{\text{fin},j}(H) \subseteq H^*$, therefore $x'_{\text{fin}} \in \mathcal{R}'_{\text{pad}}$ as claimed.

We now prove the case $m = m(n) > 1$. First, we invoke the reduction described above on \mathcal{R} , next we modify its output $\mathcal{R}'_{\text{fin}}$ to achieve the required parameters m and h , as follows. Suppose that $\mathcal{R}'_{\text{fin}}$ is defined over field F_{fin} with $|F_{\text{fin}}| = 2^r$ and the vanishing set H_{fin} has size h_{fin} . The theorem statement guarantees that $h^m = \Omega(T \log T)$ and therefore $r \leq h \cdot m$. We first discuss the case that $r = h \cdot m$. Notice that the additive group $\mathbb{F}_{2^r}^+$ is isomorphic to $(\mathbb{F}_{2^h}^+)^m$, therefore $H_{\text{fin}}^+ \subseteq \mathbb{F}_{2^r}^+$ is isomorphic to some H^m and we are done. If $r < h \cdot m$, then we just increase the dimension of the field F_{fin} and modify Q_{fin} to vanish whenever any variable X_i is set in the padded coordinates. No other parameters are changed, except for the gain of an additive m term in σ to account for the additional constraints. \square

7.5 Combining the two reductions

By combining Lemma 7.6 and Lemma 7.5, we obtain the following theorem, which gives the reduction claimed at the beginning of this section.

Theorem 7.9 (NTIME $\rightarrow \mathcal{R}_{\text{RLA}}$). *For every $T, t, s: \mathbb{N} \rightarrow \mathbb{N}$, $\delta, \epsilon: \mathbb{N} \rightarrow (0, 1)$ with $2^{s(n)} \geq \Omega(\frac{t}{(1-\delta)T \text{polylog}(T)})$, and $\mathcal{R} \in \text{NTIME}(T)$ there exist a relation \mathcal{R}' and algorithms $\text{inst}, \text{wit}_1, \text{wit}_2$ satisfying the following conditions:*

- **EFFICIENT REDUCTION.** *For every instance x , letting $x' := \text{inst}(x)$:*
 - *for every witness w , if $(x, w) \in \mathcal{R}$ then $(x', \text{wit}_1(x, w)) \in \mathcal{R}'$;*

¹For example, one can express α_i in the basis of H and define $\beta_i + H_{\dim(H)-k}$ to be the affine subspace where first $k-1$ basis vectors are fixed with coefficients equal to that of α_i , k -th basis vector is fixed with the opposite coefficient, and the last $\dim(H) - k$ basis vectors remain unfixed.

– for every witness w' , if $(x', w') \in \mathcal{R}'$ then $(x, \text{wit}_2(x, w')) \in \mathcal{R}$.

Moreover, inst runs in time $\text{poly}(n + \log(\frac{T(n)+t(n)}{1-\delta(n)}))$ and $\text{wit}_1, \text{wit}_2$ run in time $\text{poly}(n) \cdot \tilde{O}(\frac{T(n)+t(n)}{1-\delta(n)})$.

- RANDOMIZABLE LINEAR ALGEBRAIC CSP. The relation \mathcal{R}' is a subset of

$$\mathcal{R}_{\text{RLA}} \left[\begin{array}{ll} \text{field} & f = \mathbb{F}_{2^{\log(T+t)+O(\log \log(T+t))+s}} \\ \text{block length} & \ell = \tilde{O}(\frac{T+t}{1-\delta}) \\ \text{rate} & \rho = 1 - \delta \\ \text{relative distance} & \delta \\ \text{map locality} & q = \text{polylog}(T) \\ \text{map efficiency} & c = \text{poly}(n + \log T) \\ \text{map sampling I} & \gamma = \text{polylog}(T) \cdot \epsilon \\ \text{map sampling II} & \epsilon \\ \text{randomizability} & t \\ \text{randomize time} & r = \tilde{O}(\frac{T+t}{1-\delta}) \end{array} \right].$$

- RS CODES OVER CHARACTERISTIC 2. $\mathcal{C}_{\mathcal{R}', \circ} \subseteq \mathcal{RS}_{\ell, \rho, \chi}^*$ and $\mathcal{C}_{\mathcal{R}', \bullet} \subseteq \mathcal{VRS}_{\ell, \rho, \chi}^*$ with $\chi := 2^s \cdot \text{polylog}(T)$ (see Section 2.4).

Proof of Theorem 7.9. First, we invoke Lemma 7.6 with h, m, T, s such that $m(n) = 1$ and $h(n) = O(T(n) \log T(n))$; this yields a relation $\mathcal{R}^{(1)}$ and algorithms $\text{inst}^{(1)}, \text{wit}_1^{(1)}, \text{wit}_2^{(1)}$ such that: (i) $\text{inst}^{(1)}, \text{wit}_1^{(1)}, \text{wit}_2^{(1)}$ provide a reduction from $\mathcal{R} \in \text{NTIME}(T)$ to $\mathcal{R}^{(1)}$, with $\text{inst}^{(1)}(x)$ running in time $\text{poly}(n + \log h(n) + m(n))$ and $\text{wit}_1^{(1)}(x, w), \text{wit}_2^{(1)}(x, w^{(1)})$ in time $\tilde{O}(T(n))$; and (ii) $\mathcal{R}^{(1)}$ is a subset of

$$\mathcal{R}_{\text{GAP}} \left[\begin{array}{ll} \text{field} & F = \mathbb{F}_{2^{\log T + O(\log \log T) + s}} \\ \text{size of subset} & h = O(T(n) \log T(n)) \\ \text{copies of subset} & m = 1 \\ \text{number of neighbors} & \eta = \text{polylog}(T) \\ \text{circuit total degree} & d = O(1) \\ \text{circuit size} & \sigma = \text{poly}(n + \log T) \\ + & \end{array} \right].$$

Next, we invoke Lemma 7.5 on $\mathcal{R}^{(1)}$, using δ, ϵ, t from the theorem statement. Note that the conditions of the theorem are satisfied as $|F| \geq \frac{(h+t)d}{1-\delta} + h \geq \hat{h}$. Therefore this yields a relation $\mathcal{R}^{(2)}$ and algorithms $\text{inst}^{(2)}, \text{wit}_1^{(2)}, \text{wit}_2^{(2)}$ such that: (i) $\text{inst}^{(2)}, \text{wit}_1^{(2)}, \text{wit}_2^{(2)}$ provide a reduction from $\mathcal{R}^{(1)}$ to $\mathcal{R}^{(2)}$, with $\text{inst}^{(2)}(x^{(1)})$ running in time $\text{poly}(|x^{(1)}|)$, $\text{wit}_1^{(2)}(x^{(1)}, w^{(1)})$ in time $\text{poly}(|x^{(1)}|) \cdot \tilde{O}(|w^{(1)}| \cdot \eta \cdot \sigma)$ and $\text{wit}_2^{(2)}(x^{(1)}, w^{(2)})$ in time $\text{poly}(|x^{(1)}|) \cdot \tilde{O}(|w^{(2)}|)$; and (ii) $\mathcal{R}^{(2)}$ is a subset of

$$\mathcal{R}_{\text{RLA}} \left[\begin{array}{ll} \text{field} & f = F \\ \text{block length} & \ell = O(\frac{h+t}{1-\delta}) \\ \text{rate} & \rho = 1 - \delta \\ \text{relative distance} & \delta \\ \text{map locality} & q = \eta \\ \text{map efficiency} & c = \sigma + \eta \\ \text{map sampling I} & \gamma = \eta \epsilon \\ \text{map sampling II} & \epsilon \\ \text{randomizability} & t \\ \text{randomize time} & r = \tilde{O}(\frac{h+t}{1-\delta}) \end{array} \right].$$

One can check that $\mathcal{R}^{(2)}$ achieves the parameters specified in the theorem statement.

The desired reduction from \mathcal{R} to $\mathcal{R}^{(2)}$ is given by the algorithms $\text{inst}(x) := \text{inst}^{(2)}(\text{inst}^{(1)}(x))$, $\text{wit}_1(x, w) := \text{wit}_1^{(2)}(\text{inst}^{(1)}(x), \text{wit}_1^{(1)}(x, w))$, and $\text{wit}_2(x, w') := \text{wit}_2^{(2)}(x, \text{wit}_2^{(1)}(\text{inst}^{(1)}(x), w'))$. One can verify that inst runs in time $\text{poly}(n + \log(\frac{T(n)+t(n)}{1-\delta(n)}))$ and $\text{wit}_1, \text{wit}_2$ run in time $\text{poly}(n) \cdot \tilde{O}(\frac{T(n)+t(n)}{1-\delta(n)})$.

Finally, Lemma 7.6 implies that $\mathfrak{x}^{(1)} = (1^n, H, Q, \vec{N})$ where H is an additive subgroup of $F(n)$, which has characteristic 2. Hence, Lemma 7.5 implies that $\mathfrak{x}^{(2)} = (1^n, C_\circ, C_\bullet, g)$ where $C_\circ = \text{RS}[F(n), \hat{H}, \frac{h(n)+t(n)}{\hat{h}(n)}] \subseteq \text{RS}[F(n), \hat{H}, \delta]$ and $C_\bullet = \text{VRS}[F(n), \hat{H}, \frac{(h(n)+t(n))d(n)}{\hat{h}(n)}, H] \subseteq \text{VRS}[F(n), \hat{H}, \delta, H]$ where \hat{H} is a union of cosets of the subgroup H with cardinality $\hat{h}(n)$. Thus, $C_\circ \in \mathcal{RS}_{\ell, \rho, \chi}^*$ and $C_\bullet \in \mathcal{VRS}_{\ell, \rho, \chi}^*$, as claimed. \square

8 Proof of Theorem 4.1

Proof of Theorem 4.1. We combine Theorem 7.9 and Lemma 6.4 (and Theorem 2.5) so to obtain Theorem 4.1, as follows. Let \mathcal{R} be a relation in $\text{NTIME}(T)$; we construct a duplex PCP system for \mathcal{R} with the claimed parameters.

We first reduce **NTIME** to randomizable linear algebraic CSPs: invoke Theorem 7.9 on \mathcal{R} , with parameters $T, t_1, s_1, \epsilon_1, \delta_1$ (where $t_1, s_1, \epsilon_1, \delta_1$ are specified below), to obtain a relation \mathcal{R}' and algorithms $\text{inst}, \text{wit}_1, \text{wit}_2$ such that: (i) $\text{inst}, \text{wit}_1, \text{wit}_2$ provide a reduction from \mathcal{R} to \mathcal{R}' , with inst running in time $\text{poly}(n + \log(\frac{T+t_1}{1-\delta_1}))$ and $\text{wit}_1, \text{wit}_2$ in time $\tilde{O}(\frac{T+t_1}{1-\delta_1})$; and (ii) \mathcal{R}' is a subset of

$$\mathcal{R}_{\text{RLA}} \left[\begin{array}{ll} \text{field} & f_1 = \mathbb{F}_{2^{\log(T+t_1)+O(\log \log(T+t_1))+s_1}} \\ \text{block length} & \ell_1 = \tilde{O}(\frac{T+t_1}{1-\delta_1}) \\ \text{rate} & \rho_1 = 1 - \delta_1 \\ \text{relative distance} & \delta_1 \\ \text{map locality} & q_1 = \text{polylog}(T) \\ \text{map efficiency} & c_1 = \text{poly}(n + \log T) \\ \text{map sampling I} & \gamma_1 = \text{polylog}(T) \cdot \epsilon_1 \\ \text{map sampling II} & \epsilon_1 \\ \text{randomizability} & t_1 \\ \text{randomize time} & r_1 = \tilde{O}(\frac{T+t_1}{1-\delta_1}) \end{array} \right].$$

Above, we chose the parameters $t_1, s_1, \epsilon_1, \delta_1$ for Theorem 7.9 as follows: ϵ_1 such that $\gamma_1 = \text{polylog}(T) \cdot \epsilon_1 \leq \frac{2}{9}$, then $\delta_1 := 1 - \epsilon_1/4$, and $t_1 := k \cdot q_1 = k \cdot \text{polylog}(T)$; finally, s_1 is set as explained below.

Next we obtain PCPP systems for the relations corresponding to codes appearing in instances of \mathcal{R}' . Theorem 7.9 guarantees that $\mathcal{C}_{\mathcal{R}', \circ} \subseteq \mathcal{RS}_{\ell_1, \rho_1, \chi}^*$ and $\mathcal{C}_{\mathcal{R}', \bullet} \subseteq \mathcal{VRS}_{\ell_1, \rho_1, \chi}^*$ with $\chi := 2^{s_1} \cdot \text{polylog}(T)$. We now invoke Theorem 2.5, choosing $\lambda := 2$, $\chi = 2^{s_1} \cdot \text{polylog}(T)$, $\varrho := \rho_1$; as for s , we choose it simultaneously with s_1 (which so far was unspecified) so that the field f_1 for \mathcal{R}' and the field \mathfrak{a}_2 for the PCPPs are equal; it suffices for them to be $\tilde{O}(\log \log(T + t_1))$ (and, since $k \leq T$ by hypothesis, we know that $t_1 = \tilde{O}(T)$ so that the condition on s_1 from Theorem 7.9 holds). Thus, the term χ in tp_2 and tv_2 drops off, and we get that:

$$\text{Rel}(\mathcal{C}_{\mathcal{R}', \circ}), \text{Rel}(\mathcal{C}_{\mathcal{R}', \bullet}) \in \text{PCPP} \left[\begin{array}{ll} \text{answer alphabet} & \mathfrak{a}_2 = \mathbb{F}_{2^{s+\log \ell_1}} \\ \text{proof length} & \ell_2 = \tilde{O}(\ell_1) \\ \text{query complexity} & \mathfrak{q}_2 = \text{polylog}(\ell_1) \\ \text{distance measure} & \Delta_2 = \Delta_{\mathfrak{a}}^{\text{Ham}} \\ \text{proximity parameter} & \mathfrak{d}_2 = \rho_1/2 \\ \text{soundness error} & \mathfrak{e}_2 = 1/4 \\ \text{prover time} & \text{tp}_2 = \text{poly}(s) \cdot \tilde{O}(\ell_1) \\ \text{verifier time} & \text{tv}_2 = \text{poly}(s + \log \ell_1) \\ \text{non-adaptive queries} & \text{na}_2 = \text{Yes} \end{array} \right].$$

Finally we invoke Theorem 6.4 for \mathcal{R}' to obtain a duplex PCP system for \mathcal{R}' , supplying the PCPPs we just obtained from Theorem 2.5. Note that our choices satisfy the hypothesis of Theorem 6.4 because the two fields match, ℓ_1, c_1, r_1 are polynomially bounded, and $\delta_1 \geq \frac{17}{18}$, we also have $\epsilon_1 < \min\{\frac{\delta_1}{2}, \delta_1 - \gamma_1\}$ and $\mathfrak{d}_2 \leq \epsilon_1/4$ (since we chose $\gamma_1, \epsilon_1 \leq \frac{2}{9}$). This establishes our claim that:

$$\mathcal{R} \in \text{DPCP}_{\text{pzk}} \left[\begin{array}{llll} \text{answer alphabet} & \mathfrak{a} = & & \mathbb{F}_{2^{\log(T+t_1)+O(\log \log(T+t_1))}} \\ \text{proof length} & \ell = & 2\ell_2(\ell_1) + 6\ell_1 & = \tilde{O}(T + t_1) \\ \text{query complexity} & \mathfrak{q} = & 2\mathfrak{q}_2(\ell_1) + \mathfrak{q}_1 + 7 & = \text{polylog}(T) \\ \text{soundness error} & \mathfrak{e} = & & 1/2 \\ \text{prover time} & \text{tp} = & \text{inst} + \text{wit}_1 + (2\text{tp}_2(\ell_1) + (c_1 + 5)\ell_1 + r_1) & = \text{poly}(n) \cdot \tilde{O}(T + k) \\ \text{verifier time} & \text{tv} = & \text{inst} + (2\text{tv}_2(\ell_1) + c_1 + \log \ell_1) & = \text{poly}(n + \log(T + k)) \\ \text{knowledge bound} & \mathfrak{k} & & \\ \text{non-adaptive queries} & \text{na} = & \text{na}_2 & = \text{Yes} \end{array} \right].$$

The precise expression for soundness error is $e := \max\{1 - \delta_1 + \gamma_1 + \epsilon_1, (1 - |f_1|^{-1}) \cdot \max\{\mathbf{e}_2, \epsilon_1/4\} + |f_1|^{-1}\}$, but it is upper bounded by $\frac{1}{2}$, because for us $1 - \delta_1 + \gamma_1 + \epsilon_1 \leq \frac{1}{2}$, $\max\{\mathbf{e}_2, \epsilon_1/4\} = \frac{1}{4}$ and $|f_1| \geq 4$. \square

Acknowledgments

We thank Yuval Ishai and Mor Weiss for helpful discussions. The research leading to these results has received funding from: the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 240258; the Israeli Science Foundation (grant 1501/14); and the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370.

References

- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998. Preliminary version in FOCS '92.
- [Alo99] Noga Alon. Combinatorial Nullstellensatz. *Combinatorics, Probability and Computing*, 8(1-2):7–29, Jan 1999.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: a new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998. Preliminary version in FOCS '92.
- [BCGT13a] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. Fast reductions from RAMs to delegatable succinct constraint satisfaction problems. In *Proceedings of the 4th Innovations in Theoretical Computer Science Conference*, ITCS '13, 2013.
- [BCGT13b] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. On the concrete efficiency of probabilistically-checkable proofs. In *Proceedings of the 45th ACM Symposium on the Theory of Computing*, STOC '13, 2013.
- [BCI⁺13] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. Succinct non-interactive arguments via linear interactive proofs. In *Proceedings of the 10th Theory of Cryptography Conference*, TCC '13, pages 315–333, 2013.
- [BFL91] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.
- [BFLS91] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, STOC '91, pages 21–32, 1991.
- [BG93] Mihir Bellare and Oded Goldreich. On defining proofs of knowledge. In *Proceedings of the 12th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '92, pages 390–420, 1993.
- [BGG⁺88] Michael Ben-Or, Oded Goldreich, Shafi Goldwasser, Johan Håstad, Joe Kilian, Silvio Micali, and Phillip Rogaway. Everything provable is provable in zero-knowledge. In *Proceedings of the 8th Annual International Cryptology Conference*, CRYPTO '88, pages 37–56, 1988.
- [BGH⁺04] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. Robust PCPs of proximity, shorter PCPs and applications to coding. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, STOC '04, pages 1–10, 2004.
- [BGH⁺05] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. Short PCPs verifiable in polylogarithmic time. In *Proceedings of the 20th Annual IEEE Conference on Computational Complexity*, CCC '05, pages 120–134, 2005.
- [BGH⁺06] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. Robust PCPs of proximity, shorter PCPs, and applications to coding. *SIAM Journal on Computing*, 36(4):889–974, 2006. Preliminary versions of this paper have appeared in Proceedings of the 36th ACM Symposium on Theory of Computing and in Electronic Colloquium on Computational Complexity.
- [BKK⁺13] Eli Ben-Sasson, Yohay Kaplan, Swastik Kopparty, Or Meir, and Henning Stichtenoth. Constant rate PCPs for Circuit-SAT with sublinear query complexity. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '13, pages 320–329, 2013.
- [BM88] László Babai and Shlomo Moran. Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity class. *Journal of Computer and System Sciences*, 36(2):254–276, 1988.
- [BOGKW88] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: how to remove intractability assumptions. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, STOC '88, pages 113–131, 1988.
- [BS08] Eli Ben-Sasson and Madhu Sudan. Short PCPs with polylog query complexity. *SIAM Journal on Computing*, 38(2):551–607, 2008. Preliminary version appeared in STOC '05.
- [BV14] Eli Ben-Sasson and Emanuele Viola. Short PCPs with projection queries. In *Proceedings of the 41st International Colloquium on Automata, Languages and Programming*, ICALP '14, pages 163–173, 2014.
- [DFK⁺92] Cynthia Dwork, Uriel Feige, Joe Kilian, Moni Naor, and Shmuel Safra. Low communication 2-prover zero-knowledge proofs for NP. In *Proceedings of the 12th Annual International Cryptology Conference*, CRYPTO '92, pages 215–227, 1992.
- [Din07] Irit Dinur. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3):12, 2007.

- [DR04] Irit Dinur and Omer Reingold. Assignment testers: Towards a combinatorial proof of the PCP theorem. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, FOCS '04*, pages 155–164, 2004.
- [GIMS10] Vipul Goyal, Yuval Ishai, Mohammad Mahmoody, and Amit Sahai. Interactive locking, zero-knowledge PCPs, and unconditional cryptography. In *Proceedings of the 30th Annual International Cryptology Conference, CRYPTO '10*, pages 173–190, 2010.
- [GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, STOC '85*, pages 291–304, 1985.
- [HS00] Prahladh Harsha and Madhu Sudan. Small PCPs with low query complexity. *Computational Complexity*, 9(3–4):157–201, Dec 2000. Preliminary version in STACS '91.
- [IKOS09] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge proofs from secure multiparty computation. *SIAM J. Comput.*, 39(3):1121–1152, 2009.
- [IMS12] Yuval Ishai, Mohammad Mahmoody, and Amit Sahai. On efficient zero-knowledge PCPs. In *Proceedings of the 9th Theory of Cryptography Conference on Theory of Cryptography, TCC '12*, pages 151–168, 2012.
- [IMSX15] Yuval Ishai, Mohammad Mahmoody, Amit Sahai, and David Xiao. On zero-knowledge PCPs: Limitations, simplifications, and applications, 2015. Available at <http://www.cs.virginia.edu/~mohammad/files/papers/ZKPCPs-Full.pdf>.
- [IW14] Yuval Ishai and Mor Weiss. Probabilistically checkable proofs of proximity with zero-knowledge. In *Proceedings of the 11th Theory of Cryptography Conference, TCC '14*, pages 121–145, 2014.
- [IY87] Russell Impagliazzo and Moti Yung. Direct minimum-knowledge computations. In *Proceedings of the 7th Annual International Cryptology Conference, CRYPTO '87*, pages 40–51, 1987.
- [KPT97] Joe Kilian, Erez Petrank, and Gábor Tardos. Probabilistically checkable proofs with zero knowledge. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing, STOC '97*, pages 496–505, 1997.
- [KR08] Yael Tauman Kalai and Ran Raz. Interactive PCP. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming, ICALP '08*, pages 536–547, 2008.
- [LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Nisan Noam. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992.
- [LS95] Dror Lapidot and Adi Shamir. A one-round, two-prover, zero-knowledge protocol for NP. *Combinatorica*, 15(2):204–214, 1995.
- [Mie08] Thilo Mie. Polylogarithmic two-round argument systems. *Journal of Mathematical Cryptology*, 2(4):343–363, 2008.
- [MX13] Mohammad Mahmoody and David Xiao. Languages with efficient zero-knowledge pcps are in SZK. In *Proceedings of the 10th Theory of Cryptography Conference, TCC '13*, pages 297–314, 2013.
- [Ore33] Oystein Ore. On a special class of polynomials. *Transactions of the American Mathematical Society*, 35(3):559–584, 1933.
- [OW93] Rafail Ostrovsky and Avi Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *Proceedings of the 2nd Israel Symposium on Theory of Computing Systems, ISTCS '93*, pages 3–17, 1993.
- [PS94] Alexander Polishchuk and Daniel A. Spielman. Nearly-linear size holographic proofs. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing, STOC '94*, pages 194–203, 1994.
- [Sha92] Adi Shamir. $IP = PSPACE$. *Journal of the ACM*, 39(4):869–877, 1992.
- [Spi95] Daniel Spielman. *Computationally Efficient Error-Correcting Codes and Holographic Proofs*. PhD thesis, Massachusetts Institute of Technology, 1995.
- [Sze99] Mario Szegedy. Many-valued logics and holographic proofs. In *Proceedings of the 26th International Colloquium on Automata, Languages and Programming, ICALP '99*, pages 676–686, 1999.