

Perfect Zero-Knowledge PCPs for $\#P$

Tom Gur *

Jack O'Connor[†]

Nicholas Spooner[‡]

Abstract

We construct perfect zero-knowledge probabilistically checkable proofs (PZK-PCPs) for every language in $\#P$. This is the first construction of a PZK-PCP for any language outside BPP. Furthermore, unlike previous constructions of (statistical) zero-knowledge PCPs, our construction simultaneously achieves non-adaptivity and zero knowledge against arbitrary (adaptive) polynomial-time malicious verifiers.

Our construction consists of a novel masked sumcheck PCP, which uses the combinatorial nullstellensatz to obtain antisymmetric structure within the hypercube and randomness outside of it. To prove zero knowledge, we introduce the notion of locally simulatable encodings: randomised encodings in which every local view of the encoding can be efficiently sampled given a local view of the message. We show that the code arising from the sumcheck protocol (the Reed–Muller code augmented with subcube sums) admits a locally simulatable encoding. This reduces the algebraic problem of simulating our masked sumcheck to a combinatorial property of antisymmetric functions.

*University of Cambridge. Email: tom.gur@cl.cam.ac.uk.

[†]University of Cambridge. Email: jo496@cam.ac.uk.

[‡]University of Warwick/NYU. Email: nicholas.spooner@warwick.ac.uk.

Contents

1	Introduction	3
1.1	Our results	4
1.2	Acknowledgements	5
2	Techniques	5
2.1	Structure versus randomness	6
2.2	Combinatorial structure of antisymmetric functions	7
2.3	Local simulation of random low-degree extensions	9
2.4	Local simulation of subcube sums of LDEs	12
3	Preliminaries	13
3.1	Coding theory	13
3.2	Linear-algebraic claims	14
3.3	The combinatorial nullstellensatz	15
4	Locally simulatable encodings	15
4.1	Linear randomised encoding functions	16
5	Constraint location for random low-degree extensions	18
5.1	A constraint detector for random low-degree extensions of the zero function	19
5.2	The decision problem	21
5.3	Search-to-decision reduction	22
6	Constraint location for subcube sums of random low-degree extensions	24
6.1	Local characterisation of ΣRM	24
6.2	An ℓ -constraint locator for ΣRM	28
7	Constraint location for $\Sigma\text{AntiSym}$	30
7.1	Properties of $\Sigma\text{AntiSym}$	31
7.2	Algorithms for $\Sigma\text{AntiSym}$	34
7.3	Proof of Theorem 7.2	35
8	ZKPCP for $\#\text{P}$	36
8.1	Zero-Knowledge Probabilistically Checkable Proofs (ZK-PCPs)	36
8.2	Our Construction	37
A	Deferred proofs	40
A.1	Proof of Claim 5.10	40
A.2	Proof of Lemma 5.11	40
A.3	Proof of Claim 5.15	41
A.4	Proof of Lemma 5.17	42
A.5	Proof of Claim 6.11	42
	References	44

1 Introduction

The notion of a *zero-knowledge (ZK) proof*, a proof that conveys no information except the truth of a statement, is one of the most influential ideas in cryptography and complexity theory of the past four decades.

Zero knowledge was originally defined by Goldwasser, Micali and Rackoff [GMR89] in the context of interactive proofs (IPs). The deep and beautiful insight in this work is that it is possible to rigorously *prove* that an interaction does not convey information, by exhibiting an efficient algorithm called the *simulator* which can generate the distribution of protocol transcripts without interacting with the prover. In that work, the authors identify three different notions of zero knowledge, depending on the quality of the simulation. These are: (1) *perfect* zero knowledge (PZK), where the simulator’s distribution is *identical* to the real distribution of transcripts; (2) *statistical* zero knowledge (SZK), where the distributions are inverse-exponentially close; and (3) *computational* zero knowledge (CZK), where the distributions are computationally indistinguishable.

This hierarchy led naturally to the study of the complexity classes PZK, SZK, CZK of languages admitting ZK interactive proofs. Seminal results show that PZK contains interesting “hard” languages, including quadratic residuosity and nonresiduosity [GMR89] and graph isomorphism and nonisomorphism [GMW91]. Despite this, the structure of the class PZK, and its relation to other complexity classes, remains poorly understood. In light of this difficulty, the study of zero knowledge has followed two main routes: (1) studying the “relaxed” notions of SZK and CZK, and (2) studying zero knowledge in other models.

The former line of work has proved highly fruitful. A seminal result of [GMW91] showed that $CZK = IP = PSPACE$, which launched the cryptographic study of ZK proofs, yielding a plethora of theoretical and practical results [Tha22]. The study of SZK has revealed that this class has a rich structure, and deep connections within complexity theory (cf. [Vad99]).

The present work lies along the second route, also hugely influential. The seminal work of Ben-Or, Goldwasser, Kilian, and Wigderson [BGKW88] introduced the model of multi-prover interactive proofs (MIP) in order to achieve perfect zero-knowledge without any computational assumptions [LS95; DFKNS92]. This in turn inspired some of the most important models and results in contemporary theoretical computer science, including entangled-prover interactive proofs (MIP*), interactive oracle proofs (IOPs), and most notably, probabilistically checkable proofs (PCPs). Indeed, the celebrated PCP theorem [ALMSS92; AS03; Din07] is widely recognised as one of the most important achievements of modern complexity theory [AAV13].

Zero-knowledge PCPs. Recall that a PCP is a proof which can be verified, with high probability, by a verifier that only makes a small number of queries (even $O(1)$) to the proof (cf. [AB09]). A zero-knowledge PCP is a *randomised* proof which, in addition to being locally checkable, satisfies a zero knowledge condition: the view of any efficient (malicious) verifier can be efficiently simulated (see [Wei22] for a survey). Similarly to the IP case, we can distinguish PCPs with perfect, statistical and computational zero knowledge. Note that there appears to be no formal relationship between ZK-IPs and ZK-PCPs: standard transformations from IP to PCP spoil zero knowledge.

The first *statistical* zero-knowledge PCPs appeared in the seminal work of Kilian, Petrank and Tardos [KPT97]. Later works [IMS12; IW14] simplified this construction, and extended it to PCPs of proximity. These constructions operate via a 2-step compilation of (non-zero-knowledge) PCP: first the PCP is endowed with honest-verifier statistical zero-knowledge (HVSZK), and then the latter property is boosted into full statistical zero knowledge by “forcing” the malicious verifier’s query pattern to be similarly distributed to that of the honest verifier, using an object called a locking scheme. Kilian, Petrank and Tardos use this transformation to prove that NEXP has SZK-PCPs (with a polynomial-time verifier) that are zero knowledge against all polynomial-time malicious verifiers; i.e., $SZK\text{-}PCP[\text{poly}, \text{poly}] = \text{NEXP}$.

Unfortunately, locking schemes have two inherent drawbacks: they require the honest PCP verifier to be

adaptive and they cannot achieve perfect zero knowledge. Another line of work, motivated by cryptographic applications, focuses on obtaining SZK-PCPs for NP with a *non-adaptive* honest verifier from leakage resilience. These results come with caveats, achieving either a weaker notion of zero knowledge known as *witness indistinguishability* [IWY16], or simulation against adversaries making only quadratically many more queries than the honest verifier [HVW22].

The complexity landscape of perfect zero knowledge is subtle. As discussed above, we know that $\text{PZK-MIP} = \text{MIP} = \text{NEXP}$ [BGKW88], where MIP is the class of languages with a multi-prover interactive proof. The quantum analogue of this result, $\text{PZK-MIP}^* = \text{MIP}^* = \text{RE}$, is also known to hold [CFG22; GSY19]. We know a similar result for interactive PCPs (IPCPs) [KR08], an interactive generalisation of PCPs (and special case of IOPs): $\text{PZK-IPCP} = \text{IPCP} = \text{NEXP}$ [CFS17]. For IPs, the picture is very different. We know that $\text{PZK} \subseteq \text{SZK} \subseteq \text{AM} \cap \text{coAM}$ [For87; AH91], and so it is unlikely that $\text{PZK} = \text{IP}$ ($= \text{PSPACE}$), or even $\text{NP} \subseteq \text{PZK}$. It is unknown whether $\text{PZK} = \text{SZK}$; indeed, it was recently shown that there is an oracle relative to which this equality does not hold [BCHTV20].

In contrast, *nothing at all* is known about the class PZK-PCP except for the trivial inclusion $\text{BPP} \subseteq \text{PZK-PCP}$. In particular, the following key question in the study of perfect zero knowledge remains open:

*Do perfect zero-knowledge PCPs exist for **any** non-trivial language?*

1.1 Our results

In this work, we give strong positive answer to this question, showing that there exist perfect zero-knowledge PCPs (PZK-PCPs) for all of $\#\text{P}$.

Theorem 1 (Informally stated, see Theorem 8.1). $\#\text{P} \subseteq \text{PZK-PCP}[\text{poly}, \text{poly}]$.

We prove Theorem 1 by constructing a PZK-PCP for the $\#\text{P}$ -complete language

$$\#\text{SAT} = \{(\Phi, N) : \Phi \text{ is a CNF, } \sum_{x \in \{0,1\}^n} \Phi(x) = N\}.$$

This is the first construction of a PZK-PCP for a language (believed to be) outside BPP. Furthermore, unlike previous constructions of zero-knowledge PCPs, our construction simultaneously achieves *non-adaptivity* for the honest verifier and zero knowledge against arbitrary (adaptive) polynomial-time malicious verifiers. We stress that Theorem 1 is unconditional and does not rely on any cryptographic assumptions.

Remark 1.1. As with its IP counterpart PZK, PZK-PCP is a class of *decision* problems, and so the inclusion $\#\text{P} \subseteq \text{PZK-PCP}$ refers to the decision version of $\#\text{P}$ (for which $\#\text{SAT}$ is complete). This class contains UP and coNP (in a natural way) but may not contain NP. Toda’s theorem, that $\text{PH} \subseteq \text{P}^{\#\text{P}}$, does not directly imply that $\text{PH} \subseteq \text{PZK-PCP}$, as the $\#\text{P}$ oracle in this inclusion is a function oracle.

On the way to proving Theorem 1, we solve the following general algebraic-algorithmic problem, which we consider to be of independent interest.

Problem 1 (Local simulation of low-degree extensions). *Given oracle access to a function $f : \{0, 1\}^m \rightarrow \mathbb{F}$, where \mathbb{F} is a finite field, efficiently simulate oracle access to a uniformly random **degree- d extension** of f ; that is, a function \tilde{f} drawn uniformly at random from the set of m -variate polynomials over \mathbb{F} of individual degree d that agree with f on $\{0, 1\}^m$.*

By “efficient” we mean polynomial in m , d and $\log |\mathbb{F}|$. Chen et al. [CCGOS23] give a *query-efficient* simulator for $d \geq 2$, based on observations of Aaronson and Wigderson [AW09]. (For $d = 1$, there is a query lower bound of 2^m [JKRS09].) In this work we give a *computationally* efficient simulator for $d \geq 2$.

Theorem 2 (informal, see Theorem 5.4). *There is an efficient algorithm solving Problem 1 for $d \geq 2$.*

1.2 Acknowledgements

The authors thank Alessandro Chiesa, Hendrik Waldner and Arantxa Zapico for helpful discussions. TG is supported by UKRI Future Leaders Fellowship MR/S031545/1, EPSRC New Horizons Grant EP/X018180/1, and EPSRC RoaRQ Grant EP/W032635/1.

2 Techniques

We start by outlining how to construct (non-ZK) PCPs for $\#\text{SAT}$ via the sumcheck protocol [LFKN92]. Recall that the prover and verifier in the sumcheck protocol have oracle access to an m -variate polynomial P (the arithmetisation of Φ) of individual degree d over a finite field \mathbb{F} , and the prover wishes to convince the verifier that $\sum_{\vec{a} \in \{0,1\}^m} P(\vec{a}) = \gamma$ for some $\gamma \in \mathbb{F}$. This is achieved via a $2m$ -message protocol, in which for every $i \in [m]$,

- the $2i$ -th message, sent by the verifier, is a uniformly random challenge $c_i \in \mathbb{F}$; and
- the $(2i - 1)$ -th message, sent by the prover, is the univariate polynomial

$$g_i(X) = \sum_{\vec{a} \in \{0,1\}^{m-i}} P(c_1, \dots, c_{i-1}, X, a_{i+1}, \dots, a_m).$$

The verifier accepts if (a) $g_1(0) + g_1(1) = \gamma$; (b) for each $1 \leq i \leq m - 1$, $g_i(c_i) = g_{i+1}(0) + g_{i+1}(1)$; and (c) $g_m(c_m) = P(c_1, \dots, c_m)$. This protocol is complete and sound for sufficiently large \mathbb{F} . Moreover, we can “unroll” the interactive sumcheck protocol into an (exponentially large) PCP by writing down the prover’s answers to each possible (sub-)sequence of challenges.

Unfortunately, the sumcheck PCP is *clearly not zero knowledge*: even the prover’s first message g_1 reveals information about P which is $\#\text{P}$ -hard to compute. In this overview, we will explain how to modify the sumcheck PCP in order to achieve perfect zero knowledge.

Zero knowledge IOPs. Our approach to constructing a PZK-PCP for sumcheck is inspired by the [BCFGRS17] construction of a perfect zero knowledge sumcheck in the *interactive oracle proof* (IOP) model. The IOP model is an interactive generalisation of PCPs, in which the prover and verifier interact across multiple rounds, with the prover sending a PCP oracle in each round.

In their IOP, the prover first sends the evaluation table of a uniformly random m -variate polynomial R such that $\sum_{\vec{a} \in \{0,1\}^m} R(\vec{a}) = 0$, which will be used as a mask. Next, to ensure soundness in case R does not sum to zero, the verifier sends a challenge α . Finally, the prover sends a sumcheck PCP for the statement $\sum_{\vec{a} \in \{0,1\}^m} \alpha P(\vec{a}) + R(\vec{a}) = \alpha\gamma$.

Intuitively, this protocol does not leak much information about P , because $\alpha P + R$ is a *uniformly random* polynomial that sums to $\alpha\gamma$. [BCFGRS17] shows, using techniques from algebraic complexity theory, that this IOP is indeed perfect zero knowledge; we will make use of these techniques later in our construction.

Unfortunately, the interaction in this protocol is crucial in order to balance soundness and zero knowledge. Indeed, one could imagine “unrolling” this IOP into a PCP in the same way as for sumcheck itself: writing down, for each $\alpha \in \mathbb{F}$, a sumcheck PCP Π_α for $\alpha P + R$. This preserves soundness, but completely breaks zero knowledge: since the sumcheck PCP is a linear function of the underlying polynomial, given any *two* $\Pi_\alpha, \Pi_{\alpha'}$ for distinct α, α' , we can recover the sumcheck PCP for P as $(\Pi_\alpha - \Pi_{\alpha'}) / (\alpha - \alpha')$.

On the other hand, if we have the prover send only *one* Π_α , then soundness is lost, as the prover can easily cheat by sending R that does not sum to zero. We could attempt to fix the soundness issue by having the prover

additionally *prove*, via another sumcheck PCP Π_R , that $\sum_{\vec{a}} R(\vec{a}) = 0$. While this would restore soundness, we once again lose zero knowledge, as the sumcheck PCP for P can be recovered as $(\Pi_\alpha - \Pi_R)/\alpha$.

2.1 Structure versus randomness

From the above discussion, we see that the central obstacle to obtaining a zero knowledge PCP is that the prover can break soundness by sending some R with $\sum_{\vec{a}} R(\vec{a}) \neq 0$, and known methods for detecting this strategy (without interaction) break zero knowledge.

Our approach is to *prevent* this malicious prover strategy by choosing a polynomial R so that $\sum_{\vec{a}} R(\vec{a}) = 0$ by *definition*. Of course, we must take care to ensure that R still hides certain information about P . But what kind of information? We observe that, without loss of generality, we can view the sumcheck PCP for $P + R$ as an oracle Π such that

$$\Pi(c_1, \dots, c_i) = \sum_{\vec{a} \in \{0,1\}^{m-i}} (P + R)(c_1, \dots, c_i, a_{i+1}, \dots, a_m),$$

for all $i \in [m]$ and $c_1, \dots, c_i \in \mathbb{F}$. Computing the “subcube sum” $\sum_{\vec{a} \in \{0,1\}^{m-i}} P(c_1, \dots, c_i, a_{i+1}, \dots, a_m)$ requires 2^{m-i} queries to P . Hence, in order to hope to simulate, we must ensure that R hides such subcube sums for any $1 \leq i < m - O(\log m)$. So, at a minimum, it should be the case that for all such i , and all $c_1, \dots, c_i \in \mathbb{F}$, the subcube sum $\sum_{\vec{a} \in \{0,1\}^{m-i}} R(c_1, \dots, c_i, a_{i+1}, \dots, a_m)$ is marginally uniform.

In short: (1) for soundness, we need that the *full* sum of the masking polynomial R over the hypercube $\{0, 1\}^n$ is fixed, and (2) for zero knowledge, other than sum over the hypercube, we would like R to be distributed as random as possible.

We start by ensuring that any *partial* subcube sum of R is distributed uniformly at random. Our key observation here is that while the full sum is *invariant* under any reordering of the variables X_1, \dots, X_m , any nontrivial subcube sum ($i < m$) is *not*. This leads us to the following choice for the masking polynomial:

$$R(X_1, \dots, X_m) := Q(X_1, X_2, \dots, X_m) - Q(X_m, X_{m-1}, \dots, X_1),$$

where Q is a uniformly random m -variate polynomial of individual degree d .

Clearly, by the permutation invariance of the complete sum, $\sum_{\vec{a} \in \{0,1\}^m} R(\vec{a}) = 0$. On the other hand, consider the special case of a subcube sum of R on the hypercube, i.e., for $(c_1, \dots, c_i) \in \{0, 1\}^i$:

$$\sum_{\vec{a} \in \{0,1\}^{m-i}} Q(c_1, \dots, c_i, a_{i+1}, \dots, a_m) - \sum_{\vec{a} \in \{0,1\}^{m-i}} Q(a_m, \dots, a_{i+1}, c_i, \dots, c_1).$$

This expression can only be identically zero for all Q if $\{c_1\} \times \dots \times \{c_i\} \times \{0, 1\}^{m-i} = \{0, 1\}^{m-i} \times \{c_i\} \times \dots \times \{c_1\}$, which is clearly only possible if $i \in \{0, m\}$. By linearity, if this subcube sum is not identically zero for all Q , then it is uniform when Q is uniform.

Still, compared to the IOP case, the masking polynomial R has extra structure: it satisfies an “antisymmetry” $R(\alpha_1, \dots, \alpha_m) = -R(\alpha_m, \dots, \alpha_1)$ for all $(\alpha_1, \dots, \alpha_m) \in \mathbb{F}^m$. This structure means that $P + R$ is no longer uniformly random and may leak nontrivial information about P .

To mitigate this, we modify R to make it “as random as possible” while preserving antisymmetry on $\{0, 1\}^m$ and low-degree structure over \mathbb{F}^m . Our final choice of R is as follows:

$$R(\vec{X}) = Q(\vec{X}) - Q(\vec{X}_{\text{rev}}) + \sum_{i=1}^m X_i(1 - X_i)T_i(\vec{X}), \quad (1)$$

where Q is as before, $\vec{X}_{\text{rev}} := (X_m, \dots, X_1)$, and each T_i is a uniformly random m -variate polynomial where $\deg_{X_i}(T_i) = d - 2$ and $\deg_{X_j}(T_i) = d$ for $j \neq i$. What is this extra term? It was observed by [CCGOS23], via the combinatorial nullstellensatz [Alo99], that $Z(\vec{X}) := \sum_{i=1}^m X_i(1 - X_i)T_i(\vec{X})$ is a uniformly random polynomial subject to the condition that, for all $\vec{a} \in \{0, 1\}^m$, $Z(\vec{a}) = 0$. Adding this to R retains the antisymmetric structure on the hypercube, but has the effect of “masking out” any structure that appears outside of $\{0, 1\}^m$; in particular, R is no longer antisymmetric over all of \mathbb{F}^m .

Key challenge: local simulation of combinatorial and algebraic structure. Showing the completeness and soundness of the PCP obtained via the foregoing approach is straightforward. Hence, it remains to prove that the perfect zero-knowledge condition holds.

Starting with the trivial case, note that our observation about subcube sums in the discussion above provides a simple strategy for simulating any *single* query (c_1, \dots, c_i) within the hypercube $\{0, 1\}^i$ to Π . Alas, this clearly does not suffice, as even the honest sumcheck verifier makes multiple queries to the proof, for $(c_1, \dots, c_i) \notin \{0, 1\}^i$. Hence, we must simulate multiple queries beyond the hypercube.

Indeed, the simulation of multiple queries to the low-degree extension of the hypercube is where the key challenge arises. More specifically, the difficulty is that our masking polynomial R has both *algebraic* structure arising from the polynomial degree bound and *combinatorial* structure arising from the antisymmetric reordering of the variables. In order to simulate, we must not only understand both types of structure, but crucially, also *how they interact*.

In the remainder of this proof overview, we explain how we design an efficient simulator that can simulate responses to any number of queries to Π (the sumcheck PCP for $P + R$), over the entire space $\mathbb{F}^{\leq m}$. Towards this end, we introduce the notion of *locally-simulatable encodings*.

Loosely speaking, locally-simulatable encodings are randomised encodings in which every local view of the encoding can be efficiently sampled from a local view of the message. More precisely, we say that a randomised encoding function $\text{ENC}: (H \rightarrow \mathbb{F}) \rightarrow (D \rightarrow \mathbb{F})$ is locally simulatable if there is an algorithm \mathcal{S} which, given oracle access to a message $m: H \rightarrow \mathbb{F}$ and a set $S \subseteq D$, samples from the distribution of $\text{ENC}(m)|_S \subseteq \mathbb{F}^S$ in time polynomial in $|S|$ (which may be much less than $|H|$).

The notion of locally-simulatable encodings is tightly connected to ZKPCPs: we can view the mapping from instances to corresponding PCPs as a randomised encoding. A local simulator for this encoding is a zero knowledge simulator for the PCP system.¹ In the following sections, we will provide an overview of the proofs that the combinatorial antisymmetric structure of our masking, the algebraic structure of the Reed-Muller code, and their augmentations with subcube sums that arise in the sumcheck protocol admit locally simulatable encodings.

2.2 Combinatorial structure of antisymmetric functions

As discussed above, the response to any *single* $(c_1, \dots, c_i) \in \{0, 1\}^i$ to Π can be perfectly simulated. In this section, we explain how to simulate any number of queries to Π , provided that those queries lie in the set $\{0, 1\}^{\leq m} := \cup_{i=1}^m \{0, 1\}^i$. Note that evaluated over the hypercube, the last term of Eq. 1 cancels, hence we can think of the masking polynomials as $R(\vec{X}) = Q(\vec{X}) - Q(\vec{X}_{\text{rev}})$, where $\vec{X}_{\text{rev}} := (X_m, \dots, X_1)$.

Observe that, when restricted to $\{0, 1\}^m$, the distribution of R is well structured; namely, it is a uniformly random element of the vector space AntiSym of “antisymmetric functions”², i.e., functions $f: \{0, 1\}^m \rightarrow \mathbb{F}$

¹To handle malicious verifiers that make adaptive queries, we will actually require a stronger notion of local simulation that allows for conditional sampling; see Section 4 for details.

²This is terminology that we introduce, which is distinct from the usual notion of antisymmetry for polynomials.

such that for all $\vec{c} \in \{0, 1\}^m$, $f(\vec{c}) = -f(\vec{c}_{\text{rev}})$, where $\vec{c}_{\text{rev}} = (c_m, \dots, c_1)$ is the reverse of \vec{c} . We can view the restriction of Π to $\{0, 1\}^{\leq m}$ as the output of a randomised encoding function $\text{ENC}_{\Sigma\text{AntiSym}} : (\{0, 1\}^m \rightarrow \mathbb{F}) \rightarrow (\{0, 1\}^m \rightarrow \mathbb{F})$, applied to the restriction of P to $\{0, 1\}^m$. The function $\text{ENC}_{\Sigma\text{AntiSym}}$ has the following description.

$\text{ENC}_{\Sigma\text{AntiSym}}(f)$:

1. Choose a uniformly random $r \in \text{AntiSym}$.
2. Output the function $\Sigma[f + r]$, given by

$$\Sigma[f + r](c_1, \dots, c_i) := \sum_{\vec{a} \in \{0, 1\}^{m-i}} (f + r)(c_1, \dots, c_i, a_{i+1}, \dots, a_m), \quad (2)$$

for all $i \in [m]$; i.e., the function $f + r$, augmented with all of its subcube sums.

In this perspective, our task is now to provide a local simulator for $\text{ENC}_{\Sigma\text{AntiSym}}$. That is, for a query set $S \subseteq \{0, 1\}^{\leq m}$, we want to efficiently simulate $\Sigma[f + r]|_S$ for a uniformly random $r \in \Sigma\text{AntiSym}$, given oracle access to f .

Let $\Sigma\text{AntiSym}|_S$ denote the vector space $\{\Sigma[r]|_S : r \in \text{AntiSym}\}$, and let B be a basis for its dual code $(\Sigma\text{AntiSym}|_S)^\perp$; that is, $v \in \Sigma\text{AntiSym}|_S$ if and only if $Bv = 0$. Then, by linearity, $\Sigma[f + r]|_S$ is distributed as a uniformly random vector $w : S \rightarrow \mathbb{F}$ such that $Bw = B(\Sigma[f]|_S)$.

With this in mind, our local simulator \mathcal{S} could work as follows.

$\mathcal{S}^f(S)$:

1. Compute a basis B for $(\Sigma\text{AntiSym}|_S)^\perp$.
2. Compute $y = B(\Sigma[f]|_S)$ by querying f as necessary.
3. Output a random w such that $Bw = y$.

Correctness is straightforward. The only potential issue is efficiency: evaluating $\Sigma[f]$ at a point in S may require exponentially many queries to f . Note, however, that if the column of B corresponding to $\vec{c} := (c_1, \dots, c_i)$ is all zero, then we do not need to know $\Sigma[f](c_1, \dots, c_i)$ to compute $B(\Sigma[f]|_S)$. Hence, for efficiency, it would suffice to (efficiently) find a basis B for $(\Sigma\text{AntiSym}|_S)^\perp$ such that $\sum_{\vec{c} \in \text{supp}(B)} 2^{m-\ell(\vec{c})} = \text{poly}(|S|)$, where $\text{supp}(B)$ is the set of nonzero columns of B and $\ell(\vec{c}) = i$ for $c \in \{0, 1\}^i$.

Does such a basis exist? Not exactly: for example, if $S = \{(0), (1)\}$, then $\text{supp}(B) = \{(0), (1)\}$, since $\Sigma[r](0) + \Sigma[r](1) = \sum_{\vec{a} \in \{0, 1\}^{m-1}} r(0, \vec{a}) + \sum_{\vec{a} \in \{0, 1\}^{m-1}} r(1, \vec{a}) = 0$ for all $r \in \text{AntiSym}$. However, this counterexample does not actually cause a problem for zero knowledge, since for this basis B , it holds that $B(\Sigma[P]|_S) = \sum_{\vec{a} \in \{0, 1\}^m} P(\vec{a}) = \gamma$, which is part of the input to the problem. Our key technical result in this section is that all possible counterexamples are essentially of this form.

Lemma 1. *There is a polynomial p such that, for any prefix-free³ $S \subseteq \{0, 1\}^{\leq m}$, there exists a basis B of $(\Sigma\text{AntiSym}|_S)^\perp$ where each row b_i of B is a 0-1 vector, and letting $T(b_i) := \{(\vec{c}, \vec{a}) : b_i(\vec{c}) = 1, \vec{a} \in \{0, 1\}^{m-|\vec{c}|}\}$, either*

$$|T(b_i)| \leq p(|S|) \quad \text{or} \quad |T(b_i)| \geq 2^m - p(|S|).$$

Moreover, B can be efficiently computed from S .

³A set $S \subseteq \{0, 1\}^{\leq m}$ is *prefix-free* if for any $\vec{c}_1, \vec{c}_2 \in S$, if \vec{c}_1 is a prefix of \vec{c}_2 then $\vec{c}_1 = \vec{c}_2$.

$$\begin{pmatrix} 0 & a_{12} & a_{13} & a_{14} & a_{15} \\ -a_{12} & 0 & a_{23} & a_{24} & a_{25} \\ -a_{13} & -a_{23} & 0 & a_{34} & a_{35} \\ -a_{14} & -a_{24} & -a_{34} & 0 & a_{45} \\ -a_{15} & -a_{25} & -a_{35} & -a_{45} & 0 \end{pmatrix}$$

Figure 1: A general antisymmetric matrix ($n = 5$), with an element of the dual highlighted ($r = 2, t = 7$).

For the purposes of this overview, we will assume that S is indeed prefix-free; in the full proof we show that this holds without loss of generality.

To gain some intuition for this result, let us consider the 2-dimensional case of antisymmetric functions $A: [n]^2 \rightarrow \mathbb{F}$; i.e., $n \times n$ antisymmetric matrices A over \mathbb{F} . Suppose that $X \subseteq [n] \times [n]$ is a set consisting of r full rows and t individual entries (so $|X| = rn + t$) such that for all antisymmetric A , it holds that $\sum_{(i,j) \in X} a_{ij} = 0$. The latter condition implies that the indicator matrix for X must be symmetric; one possible choice of X is highlighted in Fig. 1. It is straightforward to show that this symmetry implies $t \geq r \cdot (n - r)$, from which we obtain the pair of inequalities

$$r \leq \frac{1}{2} \cdot n \left(1 - \sqrt{1 - \frac{4t}{n^2}} \right) \approx \frac{t}{n} + \frac{t^2}{n^3} \quad \text{or} \quad r \geq \frac{1}{2} \cdot n \left(1 + \sqrt{1 - \frac{4t}{n^2}} \right) \approx n - \frac{t}{n} - \frac{t^2}{n^3}.$$

This in turn yields appropriate bounds on $|X|$: $|X| \lesssim 2t + t^2/n^2$ or $|X| \gtrsim n^2 - t^2/n^2$. The proof of Lemma 1 generalises this approach to higher dimensions.

Lemma 1 suggests the following modified local simulator, which we additionally provide with the total sum γ :

$$\mathcal{S}_{\text{AntiSym}}^f(S; \gamma):$$

1. Compute a basis B for $(\Sigma \text{AntiSym}|_S)^\perp$ as in Lemma 1.
2. For each row b_i of B :
 - (a) if $|T(b_i)| \leq 2^{m-1}$, compute $y_i = \sum_{\vec{x} \in T(b_i)} f(\vec{x})$;
 - (b) otherwise, compute $y_i = \gamma - \sum_{\vec{x} \in \{0,1\}^m \setminus T(b_i)} f(\vec{x})$.
3. Output a random w such that $Bw = y$.

Clearly, the number of queries to f required is $\min(|T(b_i)|, 2^m - |T(b_i)|)$, which is $\text{poly}(|S|)$ by Lemma 1; hence the overall running time of the simulator is $\text{poly}(|S|)$.

2.3 Local simulation of random low-degree extensions

In the previous section we addressed the *combinatorial* problem of simulating queries to Π that lie in the set $\{0, 1\}^{\leq m}$, which exhibits antisymmetric structure. In this section we consider the *algebraic* problem of simulating queries to Π that also lie outside of the hypercube (i.e., general point queries), which exhibits pseudorandom structure. We then bring these two parts together in Section 2.4.

Recall that our choice of R is

$$R(\vec{X}) = Q(\vec{X}) - Q(\vec{X}_{\text{rev}}) + \sum_{i=1}^m X_i(1 - X_i)T_i(\vec{X}),$$

where Q is as before, $\vec{X}_{\text{rev}} := (X_m, \dots, X_1)$, and each T_i is a uniformly random m -variate polynomial where $\deg_{X_i}(T_i) = d - 2$ and $\deg_{X_j}(T_i) = d$ for $j \neq i$.

Denote by $\text{RM}[\mathbb{F}, m, d] \subseteq \mathbb{F}^m \rightarrow \mathbb{F}$ the *Reed–Muller* code of evaluations of m -variate polynomials of individual degree d over \mathbb{F} . For a function $f: \{0, 1\}^m \rightarrow \mathbb{F}$, denote by $\text{LD}[f, d] := \{\hat{f} \in \text{RM}[\mathbb{F}, m, d] : \hat{f}(x) = f(x) \forall x \in \{0, 1\}^m\}$ the set of *degree- d extensions* of f . With the above modification, we can describe the sumcheck PCP for $P + R$, restricted to \mathbb{F}^m , as the output of the following randomised encoding:

$\text{ENC}_{\text{RMAntiSym}}(P)$:

1. Sample a random antisymmetric function $f: \{0, 1\}^m \rightarrow \mathbb{F}$.
2. Sample R uniformly at random from $\text{LD}[f, d]$.
3. Output $P + R$.

Our problem now becomes to design a local simulator for $\text{ENC}_{\text{RMAntiSym}}$.

We will do this by solving a much more general problem: *we prove that random low-degree extensions are locally simulatable*. More precisely, let $\text{ENC}_{\text{RM}}^d(f)$ be the encoding function that outputs a uniformly random element of $\text{LD}[f, d]$. We prove that ENC_{RM}^d has a local simulator $\mathcal{S}_{\text{RM}}^d$ for any $d \geq 2$. We can then easily obtain a simulator for $\text{ENC}_{\text{RMAntiSym}}$ by composing \mathcal{S}_{RM} with the local simulator for AntiSym described in the previous section.

Our starting point. We will build on the “stateful emulator” that was constructed in [CCGOS23] (for a particular oracle model that they introduced), which can be conceptualised as an *inefficient* local simulator for ENC_{RM} . To describe it, we first introduce some notation.

For $w \in \{0, 1\}^m$, we denote by δ_w the unique multilinear polynomial satisfying $\delta_w(w) = 1$ and $\delta_w(x) = 0$ for all $x \in \{0, 1\}^m \setminus \{w\}$. For a set $S \subseteq \mathbb{F}^m$, we say that w is *S -good* if there exists an m -variate polynomial q_w of individual degree at most d such that (i) $q_w(x) = 0$ for every $x \in \{0, 1\}^m \setminus \{w\}$; (ii) $q_w(z) = 0$ for every $z \in S$; and (iii) $q_w(w) = 1$. We say that w is *S -bad* if it is not S -good. Intuitively, when \hat{f} is a uniformly random low-degree extension of f , $\hat{f}|_S$ only conveys information about $f(w)$ for S -bad points w . For example, any point in $\{0, 1\}^m \cap S$ is trivially S -bad. Less trivially, if S consists of sufficiently many points on a curve passing through $w \in \{0, 1\}^m$, then w may be S -bad even if $\{0, 1\}^m \cap S$ is empty.

The simulator of Chen et al. works roughly as follows:⁴

$\mathcal{S}_{\text{RM}}^f(S)$:

1. Compute the set of S -bad points W .
2. Query $f(w)$ at all $w \in W$, sample a uniformly random polynomial P such that $P(w) = f(w)$ for all $w \in W$, and output $P|_S$.

Chen et al. prove that \mathcal{S}_{RM} is correct, and moreover that it is *query-efficient*, provided that $d \geq 2$: the number of queries it makes to f (equal to $|W|$) is at most $|S|$ (this follows from [AW09, Lemma 4.3]).

However, they leave open the question of whether \mathcal{S}_{RM} can be made *computationally* efficient. We note that Step 2 can be achieved efficiently using an algorithm of [BCFGRS17]. The issue is in Step 1; namely, it is not clear how to compute the set W from S . Indeed, at first glance this appears to be hopeless: naively, computing whether a single w is S -bad requires solving an exponentially large linear system, and there are

⁴Their analysis requires sampling P in a different (and more complicated) way.

2^m possible choices for w . Nonetheless, we will present an efficient algorithm which, given S , outputs a list of all S -bad points.⁵

Step 1: solving the decision problem. Our first step is to consider the *decision* variant of this problem: given a set S , does there *exist* an S -bad w ? Our key insight here is that we can relate the existence of an S -bad w to the dimension of the vector space $\text{LD}[Z, d]|_S$, where Z is the constant zero function. We show that all w are S -good if and only if $\text{LD}[Z, d]|_S = \text{RM}[\mathbb{F}, m, d]|_S$. Indeed, if every w is S -good, then for any degree- d polynomial P , the polynomial $P - \sum_{w \in \{0,1\}^m} P(w) \cdot q_w$ is a degree- d extension of Z that agrees with P on S . On the other hand, if $\text{LD}[Z, d]|_S = \text{RM}[\mathbb{F}, m, d]|_S$, then for every w , there is a $\hat{Z}_w \in \text{LD}[Z, d]$ that agrees with δ_w on S ; hence we can take $q_w = \delta_w - \hat{Z}_w$, and so w is S -good.

Hence to solve the decision problem it suffices to compute (the dimensions of) $\text{RM}[\mathbb{F}, m, d]|_S$ and $\text{LD}[Z, d]|_S$. The former can be achieved efficiently (in time $\text{poly}(|S|, m, d, \log |\mathbb{F}|)$) using the *succinct constraint detector* for Reed–Muller discovered by [BCFGRS17].

To compute the latter, we build on ideas introduced in [CCGOS23]. As discussed above, in that work the authors observe that to sample a vector $v \leftarrow \text{LD}[Z, d]|_S$, it suffices to (lazily) sample random polynomials $T_i, i \in [m]$, of the appropriate degrees, and then set $v(\vec{\alpha}) := \sum_{i=1}^m \alpha_i (1 - \alpha_i) T_i(\vec{\alpha})$ for each $\vec{\alpha} \in S$. Similarly, we show how to compute a basis for $\text{LD}[Z, d]|_S$ by combining bases for the subspaces $\mathcal{P}_i := \{T_i|_S \in \text{RM}[\mathbb{F}, d, m]|_S : \deg_{X_i}(T_i) = d - 2\}$, for $i \in [m]$. These bases can also be computed using the succinct constraint detector for Reed–Muller; see Section 5.

Step 2: search-to-decision reduction. Next, we will build upon the techniques we developed in Step 1 to solve the original search problem. A natural strategy is to employ a binary search: for each $b \in \{0, 1\}$, test if $\{b\} \times \{0, 1\}^{m-1}$ contains any S -bad points, and recurse if it does. Since the number of S -bad points is bounded by $|S|$, the recursion will terminate quickly. It therefore suffices to give an efficient algorithm that, given a set A of the form $\{a_1\} \times \cdots \times \{a_i\} \times \{0, 1\}^{m-i}$, determines whether A contains any S -bad point.

For $A \subseteq \{0, 1\}^m$, let $\text{LD}[Z_A, d]|_S := \{P \in \text{RM}[\mathbb{F}, m, d] : P(\vec{x}) = 0 \forall \vec{x} \in A\}$. Our algorithm for computing a basis for $\text{LD}[Z, d]|_S$ can be straightforwardly extended to compute a basis for $\text{LD}[Z_A, d]|_S$. We can also extend the reasoning from Step 1 to show that if every $w \in A$ is S -good, then $\text{LD}[Z_A, d]|_S = \text{RM}[\mathbb{F}, m, d]|_S$. Unfortunately, the converse does not hold: it may be that A contains S -bad points but $\text{LD}[Z_A, d]|_S = \text{RM}[\mathbb{F}, m, d]|_S$. This can happen, for example, if S contains points on a curve passing through both $w \in A$ and $w' \in \{0, 1\}^m \setminus A$. In particular, the argument from Step 1 fails: for $\hat{Z}_w \in \text{LD}[Z_A, d]$, $\delta_w - \hat{Z}_w$ is not necessarily zero on $\{0, 1\}^m \setminus A$.

To obtain our final algorithm, we will instead exploit the algebraic relationship between $\text{RM}[\mathbb{F}, m, d]$ and its subcode $\text{RM}[\mathbb{F}, m, d - 1]$ to derive a *sufficient* condition for every $w \in A$ being S -good. Observe that if $\text{LD}[Z_A, d - 1]|_S = \text{RM}[\mathbb{F}, m, d - 1]|_S$, then for every $w \in A$ there is a $\hat{Z}_w \in \text{LD}[Z_A, d - 1]$ such that $p_w := \delta_w - \hat{Z}_w$ satisfies $p_w(w) = 1$ and $p_w(\vec{\alpha}) = 0$ for all $\vec{\alpha} \in S$. It follows that all $w \in A$ are S -good, as we can set $q_w = p_w \cdot \delta_w$.

Thus, our efficient test for whether A contains any S -bad point is as follows: compute bases for $\text{LD}[Z_A, d - 1]|_S$ and $\text{RM}[\mathbb{F}, m, d - 1]|_S$, and output “no” if they are of the same dimension; otherwise output “maybe”. By the above discussion, this test does not have any false negatives; however, there may be false positives. Provided there are not too many false positives, this does not cause a problem (we can either include them in W or perform an extra test to filter them out). We bound the number of false positives by noting that if the test says “maybe”, this in fact means that there exists an “ S -bad” point in A with respect to

⁵Chen et al. follow a different route: roughly, they show that because their f is *extremely sparse*, it is very unlikely that $f(w) \neq 0$ for any $w \in W \setminus S$, and so for simulation one can pretend that f is zero on $\{0, 1\}^m \setminus S$. This does not work for us for two reasons: first, a random antisymmetric function is not sparse with overwhelming probability, and second, this strategy introduces a small statistical error in simulation, which would spoil perfect zero knowledge.

polynomials of degree $d - 1$. By [AW09], provided $d \geq 2$, there are at most $|S|$ such points.

2.4 Local simulation of subcube sums of LDEs

So far, we have built a simulator that can answer arbitrary queries to our sumcheck PCP Π , provided those queries lie within the set $\{0, 1\}^{\leq m} \cup \mathbb{F}^m$, handling the combinatorial structure on the hypercube and the algebraic, pseudorandom algebraic structure outside of it. In this final part of the overview, we outline how this simulator can be extended to handle all queries; i.e., queries in the set $\mathbb{F}^{\leq m}$.

To do this, it suffices to show that the code ΣRM admits a locally-simulatable encoding, where $\Sigma\text{RM}[\mathbb{F}, m, d] = \{\Sigma[P] : P \in \text{RM}[\mathbb{F}, m, d]\}$ and $\Sigma[P]$ is defined as in Eq. 2 (where (c_1, \dots, c_i) now ranges over \mathbb{F}^i). Specifically, we show that the following encoding function is locally simulatable:

$$\underline{\text{ENC}_{\Sigma\text{RM}}^d(\Sigma[f])} :$$

1. Sample \hat{f} uniformly at random from $\text{LD}[f, d]$.
2. Output $\Sigma[\hat{f}]$.

Note that the *message* (input to $\text{ENC}_{\Sigma\text{RM}}$) is $\Sigma[f]$, even though $\text{ENC}_{\Sigma\text{RM}}$ operates only on f . This is necessary for local simulation, since individual locations in $\Sigma[\hat{f}]$ depend on partial sums of f which cannot be computed from few queries to f itself.

We would like to follow the strategy from the previous section: given a set $S \subseteq \mathbb{F}^{\leq m}$ on which we want to simulate $\text{ENC}_{\Sigma\text{RM}}(\Sigma[f])$, compute the set of all “bad” points $W \subseteq \{0, 1\}^{\leq m}$, i.e., $w \in W$ if $\text{ENC}_{\Sigma\text{RM}}(\Sigma[f])|_S$ conveys information about $\Sigma[f](w)$. Unfortunately, unlike in the “plain” Reed–Muller case, it is not clear how to even bound the *size* of W , let alone compute it.

The issue here is that a single evaluation of $\Sigma[\hat{f}](\vec{\alpha})$ for $\vec{\alpha} \in \mathbb{F}^i$ depends on $\hat{f}(\vec{\beta})$ for every point $\vec{\beta} \in \vec{\alpha} \times \{0, 1\}^{m-i}$. Naively applying the lemma of [AW09] to this set of points yields a set $W \subseteq \{0, 1\}^m$ of size 2^{m-i} . To bound $|W|$ by a polynomial, we must therefore crucially make use of the fact that making a few queries to $\Sigma[\hat{f}]$ can reveal only a few (possibly large) *partial sums* of \hat{f} — which we can hope to deduce from a small number of queries to $\Sigma[f]$.

To achieve this we will give a *decomposition* of ΣRM as a sequence of RM codes on different numbers of variables, “tied together” with constraints that enforce summation structure. In more detail, let $T \subseteq \mathbb{F}^{\leq m}$ be a set with the following special structure, which we call “closed”: if $(s_1, \dots, s_i) \in S$, then its “parent” (s_1, \dots, s_{i-1}) and its “siblings” $(s_1, \dots, s_{i-1}, 0)$ and $(s_1, \dots, s_{i-1}, 1)$ are also in S . The set $\{0, 1\}^{\leq m}$ is closed; it is also easy to see that any set T can be made closed by adding at most $3m|T|$ points. We prove the following theorem about the structure of restrictions of ΣRM to closed sets T :

Theorem 2.1 (Informally stated, see Theorem 6.5). *For any closed set T , any constraint $z \in (\Sigma\text{RM}[\mathbb{F}, m, d]|_T)^\perp$ lies in the span of the following two types of constraint:*

- summation constraints, which ensure that $\Sigma[\hat{f}](t_1, \dots, t_i) = \Sigma[\hat{f}](t_1, \dots, t_i, 0) + \Sigma[\hat{f}](t_1, \dots, t_i, 1)$ for $(t_1, \dots, t_i) \in T$, $i < m$; and
- low-degree constraints, which are given by, for each i , the i -variate Reed–Muller constraints on the set $T_i := T \cap \mathbb{F}^i$; i.e., $(\text{RM}[\mathbb{F}, i, d]|_{T_i})^\perp$.

If we take $T = \{0, 1\}^{\leq m} \cup S$, then $T_i = \{0, 1\}^i \cup (S \cap \mathbb{F}^i)$ (assuming S is closed). Using this decomposition, we show that we can take as the “bad” set $W := \bigcup_{i=1}^m W_i$, where W_i is the set of $(S \cap \mathbb{F}^i)$ -bad points in $\text{RM}[\mathbb{F}, i, d]$. Each W_i can be computed efficiently using the algorithm described in Section 2.3.

We conclude by giving some brief intuition on how we prove Theorem 2.1. We show that any constraint z supported on $T_i \cup T_{i+1} \cup \dots \cup T_m$ can be “flattened” into a constraint supported on T_i only. Such constraints belong to $(\text{RM}[\mathbb{F}, i, d]|_{T_i})^\perp$. We then use a dimension-counting argument to show that these, augmented with the summation constraints, span $(\Sigma\text{RM}[\mathbb{F}, m, d]|_T)^\perp$.

3 Preliminaries

For a vector $\vec{a} \in \Sigma^i$, we denote by $\ell(\vec{a})$ the *number of entries* in \vec{a} , i.e. $\ell(\vec{a}) = i$. Throughout, \mathbb{F} is a finite field. For two vectors $x, y \in \mathbb{F}^n$, we denote the dot product by $x \cdot y := \sum_{i=1}^n x_i y_i$. For sets A, B , we denote by $A \sqcup B$ the *disjoint union* of A and B .

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

Polynomials. For $d \in \mathbb{N}$, we write $\mathbb{F}^{\leq d}[X_1, \dots, X_m]$ for the ring of polynomials in m variables over \mathbb{F} of individual degree d . For $\vec{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$, we use $\mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ to denote the ring of polynomials in m variables over \mathbb{F} of degree d_i in X_i for each $i \in [m]$.

Low-degree extensions. Given a function $f: A_1 \times \dots \times A_m \rightarrow \mathbb{F}$, $A_1, \dots, A_m \subseteq \mathbb{F}$, we denote the set of all low-degree extensions of f by $\text{LD}_{\vec{d}}[f] := \{p \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m] : p(x) = f(x) \forall x \in A_1 \times \dots \times A_m\}$.

Lagrange polynomials. Let \mathbb{F} be a field, and let $m \in \mathbb{N}$. Given $S_1, \dots, S_m \subseteq \mathbb{F}$ and a point $w \in S_1 \times \dots \times S_m$, we define the *Lagrange interpolating polynomial for \mathcal{S} and w* to be

$$L_{\mathcal{S}, w}(x) := \prod_{i=1}^m \prod_{z \in S_i \setminus \{w_i\}} \frac{x_i - z_i}{w_i - z_i},$$

where denote $\mathcal{S} := S_1 \times \dots \times S_m$. $L_{\mathcal{S}, w}$ is of individual degree $|S_i| - 1$ in its i -th variable.

Fact 3.1. The set of Lagrange interpolators $\{L_{\mathcal{S}, w}\}_{w \in \mathcal{S}}$ form a basis for the space $\mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ where $\vec{d} := (|S_1| - 1, \dots, |S_m| - 1)$.

Vanishing functions. Given a subset $S \subseteq \mathbb{F}$ we denote the function which is vanishing on S by $\mathcal{Z}_S(x) := \prod_{s \in S} (x - s)$.

Vector reversal. For a vector $\vec{x} = (x_1, x_2, \dots, x_n)$, we denote by \vec{x}_{rev} the *reversed vector* $(x_n, x_{n-1}, \dots, x_1)$.

3.1 Coding theory

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

Definition 3.2. Given a function $w: D \rightarrow \mathbb{F}$ and a subset $I \subseteq D$, we denote by $w|_I$ the *restriction of w to I* , that is the function $w|_I: I \rightarrow \mathbb{F}$ such that $w(x) = w|_I(x)$ for all $x \in I$.

Definition 3.3. Given a linear code $C \subseteq \mathbb{F}^D$ and a subset $I \subseteq D$, we denote by $C|_I$ the *code consisting of codewords from C restricted to I* , that is

$$C|_I = \{w|_I : w \in C\}.$$

Note that $C|_I$ is itself a linear code, as it is a linear subspace of \mathbb{F}^I .

Linear codes. Let \mathbb{F} be a finite field. A code $C: \mathbb{F}^k \rightarrow \mathbb{F}^\ell$ is *linear* if it is an \mathbb{F} -linear map from \mathbb{F}^k to \mathbb{F}^ℓ .

Reed-Muller codes and low-degree extensions. The Reed-Muller (RM) code is the code consisting of evaluations of multivariate low-degree polynomials over a finite field.

Definition 3.4. Given a field \mathbb{F} , a positive integer m , and a degree vector $\vec{d} = (d_1, \dots, d_m) \in \mathbb{N}^m$, we denote by $\text{RM}[\mathbb{F}, m, \vec{d}]$ the code of evaluations over \mathbb{F}^m of polynomials in $\mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$. For $d \in \mathbb{N}$, $\text{RM}[\mathbb{F}, m, d] := \text{RM}[\mathbb{F}, m, (d, \dots, d)]$.

Zero and sum codes. A *zero code* is a subcode of a linear code that is zero on a given subset.

Definition 3.5 (Zero codes). Given a $C \subseteq \mathbb{F}^D$, and $S \subseteq D$, we define the **zero code**

$$\mathbf{Z}_S(C) = \{w \in C : w|_S = \vec{0}\}.$$

A *sum code* is the extension of a linear code obtained by appending all sums of entries over a given subdomain.

Definition 3.6 (Sum codes). For a product set $\mathcal{A} := A_1 \times \dots \times A_m$, we define $\mathcal{A}_i := A_1 \times A_2 \times \dots \times A_i$ and $\mathcal{A}_{>i} := A_{i+1} \times \dots \times A_m$; $\mathcal{A}_0 := \{\perp\}$. We denote $\bar{\mathcal{A}} := \cup_{i=0}^m \mathcal{A}_i$.

Given a code $C \subseteq \mathbb{F}^D$ where $D = D_1 \times \dots \times D_m$ and $\mathcal{A} \subseteq D$, we define the **sum code** $\Sigma_{\mathcal{A}}C \subseteq \mathbb{F}^{\bar{D}}$ to be the code consisting of all of the subcube sums over \mathcal{A} of the codewords in C . That is,

$$\Sigma_{\mathcal{A}}C := \{\Sigma_{\mathcal{A}}[w] : w \in C\}$$

where $\Sigma_{\mathcal{A}}[w]: \mathbb{F}^D \rightarrow \mathbb{F}^{\bar{D}}$ is given by

$$\Sigma_{\mathcal{A}}[w](\vec{x}) := \sum_{\vec{y} \in \mathcal{A}_{>\ell(\vec{x})}} w(\vec{x}, \vec{y}).$$

If $D = \mathcal{A}$, we will typically omit \mathcal{A} from the notation (i.e., we write $\Sigma C, \Sigma[w]$).

3.2 Linear-algebraic claims

We will make use of the following simple linear-algebraic algorithms and facts.

Claim 3.7. Let $U \leq \mathbb{F}^n$ be a vector space. There is a polynomial-time algorithm which, given a basis for U , computes a basis for U^\perp .

Claim 3.8. Let $U \leq \mathbb{F}^n, V \leq \mathbb{F}^m$ be vector spaces. There is a polynomial-time algorithm which, given a linear map $M \in \mathbb{F}^{m \times n}$ such that $\{Mu : u \in U\} = V$ and basis $B \in \mathbb{F}^{n \times k}$ for U^\perp , outputs a basis for V^\perp .

Proof. Use Claim 3.7 to compute a basis $B' \in \mathbb{F}^{n \times (n-k)}$ for U . Then compute $A = MB' \in \mathbb{F}^{m \times (n-k)}$; we have that $V = \text{colspan}(A)$ by assumption on M . Compute a basis for V from A by Gaussian elimination, and from this a basis for V^\perp by Claim 3.7. \square

Claim 3.9. Let $T: V \rightarrow W$ be a linear map. If $v \sim \mathcal{U}(V)$, then $T(v) \sim \mathcal{U}(T(V))$.

Proof. We have that

$$\Pr_{v \leftarrow V}[T(v) = w] = \Pr_{v \leftarrow V}[v \in \ker(T) + w] = \frac{|\ker(T)|}{|V|} = \frac{1}{|T(V)|}.$$

\square

Definition 3.10. Let $(D, <)$ be an ordered domain. For a vector $v: D \rightarrow \mathbb{F}$, we define $\lambda(v) := \min\{i \in D : v(i) \neq 0\}$, where the minimum is taken with respect to $<$. A sequence of vectors $v_1, \dots, v_n: D \rightarrow \mathbb{F}$ is in echelon form if, for all $i < j$, $\lambda(v_i) < \lambda(v_j)$, and any all-zero vectors are at the end of the sequence.

It is a straightforward consequence of Gauss–Jordan elimination that every subspace of $(D \rightarrow \mathbb{F})$ has a basis in echelon form. Conversely, a sequence of vectors in echelon form (not containing the all-zero vector) forms a linearly independent set.

Definition 3.11. Let $(R, <_R)$ and $(C, <_C)$ be ordered domains and let $A \in \mathbb{F}^{R \times C}$. The **leading entry** of a row $w \in \mathbb{F}^C$ of A is $w(\lambda(w))$, i.e., the first non-zero entry in that row, with respect to the ordering $<_C$.

Definition 3.12. Let $(R, <_R)$ and $(C, <_C)$ be ordered domains. A matrix $A \in \mathbb{F}^{R \times C}$ is said to be in **reduced row echelon form** if (i) the columns of A , listed in increasing order by their indices with respect to $<_C$ are in row echelon form (with respect to $<_D$); (ii) all leading entries are 1; and (iii) all entries in the same column as a leading entry are 0.

Every matrix can be converted to reduced row echelon form via Gauss-Jordan elimination, and moreover, the resulting matrix is unique. We will often represent systems of linear equations in their matrix form. We refer to variables with a leading entry in their column of the reduced row echelon form matrix as *leading variables*, and the remaining variables as *free variables*.

3.3 The combinatorial nullstellensatz

The combinatorial nullstellensatz, due to Alon [Alo99], is a powerful tool in combinatorial number theory. It is stated as follows.

Lemma 3.13. Let \mathbb{F} be an arbitrary field, and let f be a polynomial in $\mathbb{F}[X_1, \dots, X_n]$. Let S_1, \dots, S_n be nonempty subsets of \mathbb{F} and define $\mathcal{Z}_{S_i}(x_i) = \prod_{s \in S_i} (x_i - s)$. If f vanishes over all the common zeros of $\mathcal{Z}_{S_1}, \dots, \mathcal{Z}_{S_n}$ (that is, if $f(s_1, \dots, s_n) = 0$ for all $s_i \in S_i$), then there are polynomials $h_1, \dots, h_n \in \mathbb{F}[X_1, \dots, X_n]$ so that $\deg(h_i) \leq \deg(f) - \deg(\mathcal{Z}_{S_i})$ so that

$$f = \sum_{i=1}^n h_i \mathcal{Z}_{S_i} .$$

4 Locally simulatable encodings

In this section we introduce the notion of locally simulatable encodings. First, we define randomised encoding functions, which are randomised mappings from a message space \mathcal{M} into a code \mathcal{C} .

Definition 4.1 (Randomised encoding function). Let Σ be an alphabet, S, D be sets, and $\mathcal{M} \subseteq \Sigma^S, \mathcal{C} \subseteq \Sigma^D$ be the message space and code space, respectively. A **randomised encoding function** is a random variable ENC taking values in $\mathcal{M} \rightarrow \mathcal{C}$.

Next, we define what it means for a randomised encoding function to be locally simulatable. At a high level, we require that the (conditional) distribution of any local view of the encoding can be efficiently simulated using only a local view of the message.

Definition 4.2 (Locally simulatable encoding). Let $\text{ENC}: \mathcal{M} \rightarrow \mathcal{C}$ be a randomised encoding function, let $t: \mathbb{N} \rightarrow \mathbb{N}$ and let $\ell: \mathbb{N} \rightarrow \mathbb{N}$. We say that ENC is (t, ℓ) -**locally simulatable** if there is a probabilistic algorithm \mathcal{S} , called a **local simulator**, which receives oracle access to a message $m \in \mathcal{M}$, a query-answer set $T \in (D \times \Sigma)^n$ and a new query $\alpha \in D$, runs in time $t(n)$ and makes at most $\ell(n)$ queries to m , satisfying:

$$\Pr[\mathcal{S}^m(T, \alpha) = \beta] = \Pr[\text{ENC}(m)_\alpha = \beta \mid \text{ENC}(m)_x = y \forall (x, y) \in T],$$

for all $m \in \mathcal{M}, \alpha \in D, \beta \in \Sigma$, and for all T such that there exists $c \in \text{supp}(\text{ENC}_m)$ with $c_x = y$ for all $(x, y) \in T$.

To give a sense of the definition, we consider two examples.

Example 4.3. The randomised encoding function which maps $m \in \mathbb{F}^S$, $S \subseteq \mathbb{F}$, to the evaluation of a random univariate polynomial p over \mathbb{F} of degree $2(|S| - 1)$ such that $p(x) = m(x)$ for all $x \in S$, is $(\text{poly}(n, d, \log |\mathbb{F}|), n)$ -locally simulatable. Indeed, for $I \subseteq \mathbb{F}$, $|I| \leq |S| - 1$, $p|_I$ is distributed as $(m(x))_{x \in I \cap S} \parallel (\mathcal{U}(\mathbb{F}))_{x \in I \setminus S}$. For $|I| \geq |S|$, the simulator \mathcal{S} can simply fully determine m and directly conditionally sample p .

Example 4.4. The encoding function which maps $m \in \mathbb{F}^S$, $S \subseteq \mathbb{F}$, to the evaluation of the *unique* univariate polynomial p over \mathbb{F} of degree $|S| - 1$ such that $p(x) = m(x)$ for all $x \in S$, is *not* (t, ℓ) -locally simulatable unless $\ell(n) \geq |S|$ for all $n > 0$. Indeed, determining $p(x)$ for $x \in \mathbb{F} \setminus S$ requires knowledge of the entire message m .

4.1 Linear randomised encoding functions

In this work we will focus on a special family of randomised encoding functions, for which the randomised encoding is given by appending a random vector to the message, and then applying a linear function.

Definition 4.5 (Linear randomised encoding function). We say that $\text{ENC}: \mathcal{M} \rightarrow \mathcal{C}$ is a **linear randomised encoding function** if $\Sigma = \mathbb{F}$ for some finite field \mathbb{F} , and $\text{ENC}(m)$ is distributed as $\text{ENC}(m; \mu)$ where $\mu \sim \mathcal{U}(\mathbb{F}^r)$ and $\text{ENC}: \mathcal{M} \times \mathbb{F}^r \rightarrow \mathcal{C}$ is a (deterministic) linear function (of both message and randomness).

That is, we use ENC to refer to both the fixed encoding function and the resulting random variable.

Linear randomised encodings have certain structural properties that make them easier to work with. One important property of linear randomised encodings is that the distribution of $\text{ENC}(m)|_I$ can be characterised as the solution space of an affine system. An ℓ -constraint locator is an algorithm which outputs a description of this system, defined as follows.

Definition 4.6. Let $\text{ENC}: \mathcal{M} \rightarrow \mathcal{C}$ be a linear randomised encoding function, and let $\ell: \mathbb{N} \rightarrow \mathbb{N}$. We say that an algorithm CL is an **ℓ -constraint locator for ENC** if, for every subdomain $I \subseteq D$, it satisfies

$$\text{CL}(I) = (R, Z),$$

where $R \subseteq S$ and $Z \in \mathbb{F}^{k \times (R \sqcup I)}$ are such that $|R| = \ell(|I|)$, and for every message $m \in \mathcal{M}$ and $\vec{\beta} \in \mathbb{F}^I$, $(m|_R, \vec{\beta})^T \in \ker(Z) \subseteq \mathbb{F}^{R \sqcup I}$ if and only if there exists randomness $\mu \in \mathbb{F}^r$ such that:

$$\text{ENC}(m; \mu)|_I = \vec{\beta}.$$

Our notion of an ℓ -constraint locator generalises the notion of a “constraint detector” introduced in [BCFGRS17]. A constraint detector for a linear code $C \subseteq \mathbb{F}^D$ takes as input a subdomain $I \subseteq D$ and outputs a matrix representing all linear relations that hold between codewords in $C|_I$. Setting the randomised encoding function for our ℓ -constraint locator to be the uniform distribution over all codewords in C , we recover the definition of a constraint detector.

Definition 4.7 (Constraint detector). *Let $C \subseteq \mathbb{F}^D$ be a linear code. A **constraint detector for C** is a 0-constraint locator for $\text{ENC}_C: \emptyset \rightarrow C$, the linear randomised encoding function given by $\text{ENC}_C \sim \mathcal{U}(C)$.*

Given an efficient constraint locator, it is straightforward to construct a local simulator.

Claim 4.8. *Let $\text{ENC}: \mathcal{M} \rightarrow \mathcal{C}$ be a linear randomised encoding with an ℓ -constraint locator running in time t . Then ENC is (t, ℓ) -locally simulatable.*

Construction 4.9. A local simulator \mathcal{S} for a linear randomised encoding function $\text{ENC}: \mathbb{F}^S \rightarrow \mathbb{F}^D$ given an ℓ -constraint locator CL for ENC .

$\mathcal{S}_{\text{ENC}}^m(T, \alpha)$:

1. Set $(R, Z) = \text{CL}(\text{supp}(T) \cup \{\alpha\})$.
2. For each $\gamma \in R$, query the message and set $m_\gamma := m(\gamma)$.
3. If there exists a row $z \in \mathbb{F}^{R \cup \text{supp}(T) \cup \{\alpha\}}$ of the matrix Z with $z(\alpha) \neq 0$, output $\beta := -\frac{1}{z(\alpha)} \sum_{x \in \text{supp}(T)} y_x z(x) + \sum_{x \in R} m_x z(x)$, where $(x, y_x) \in T$.
4. Otherwise, output $\beta \leftarrow \mathbb{F}$.

Proof. First we introduce some notation. Denote $p_\beta := \Pr[\text{ENC}(m)_\alpha = \beta \mid \text{ENC}(m)_x = y \forall (x, y) \in T]$. Denote $I := \text{supp}(T) \cup \{\alpha\}$ and denote $\kappa := -\frac{1}{z(\alpha)} \sum_{x \in \text{supp}(T)} y_x z(x) + \sum_{x \in R} m_x z(x)$. Define the vector $\vec{y} \in \mathbb{F}^{\text{supp}(T)}$ to be the vector consisting of values y_x for each $(x, y_x) \in T$.

We will break the proof into two parts, corresponding to the two possibilities for the output of the construction in Step 3 and Step 4. First, we show that if there exists a row $z \in \mathbb{F}^{R \cup \text{supp}(T) \cup \{\alpha\}}$ of the matrix Z with $z(\alpha) \neq 0$ (as in Step 3), then $p_\beta = 1$ if $\beta = \kappa$ and $p_\beta = 0$ otherwise. If there exists such a row, then the vector $(\beta, \vec{y}, m|_R) \in \ker(Z)$ if and only if $\beta = \kappa$. By the correctness of CL , $(\beta, \vec{y}, m|_R) \in \ker(Z)$ if and only if there exists randomness μ such that $\text{ENC}(m; \mu)|_I = (\beta, \vec{y})$. So for all randomness $\mu \in \mathbb{F}^r$, $\text{ENC}(m; \mu)|_I = (\beta, \vec{y})$ if and only if $\beta = \kappa$; hence $p_\beta = 1$ if $\beta = \kappa$ and $p_\beta = 0$ otherwise.

Second, we show that if there does not exist a row z with $z(\alpha) \neq 0$ (as in Step 4), then $p_\beta = \frac{1}{|\mathbb{F}|}$, for all $\beta \in \mathbb{F}$. For each $\beta \in \mathbb{F}$, consider the set of all possible choices of randomness which agree with $T \cup (\alpha, \beta)$, that is $Q_\beta := \{\mu \in \mathbb{F}^r : \text{ENC}(m, \mu)|_{\{\alpha\} \cup \text{supp}(T)} = (\beta, \vec{y})\}$. We will show that $|Q_\beta| = |Q_{\beta'}$ for each $\beta, \beta' \in \mathbb{F}$ by giving a bijection, which implies that $p_\beta = \frac{1}{|\mathbb{F}|}$, since μ is chosen uniformly at random.

In this case, $(\beta, \vec{y}, m|_R) \in \ker(Z)$ for all $\beta \in \mathbb{F}$. Then, by the correctness of CL , for all $\beta \in \mathbb{F}$, there exists randomness $\mu \in \mathbb{F}^r$ such that $\text{ENC}(m, \mu)|_I = (\beta, \vec{y})$, that is, Q_β is non-empty for each $\beta \in \mathbb{F}$. Let $\mu_\beta \in Q_\beta$ and let $\mu_{\beta'} \in Q_{\beta'}$. It is easy to verify that the map $M: Q_\beta \rightarrow Q_{\beta'}$ defined by $M(\mu) := \mu + (\mu_{\beta'} - \mu_\beta)$ is bijective, completing the proof. \square

We prove that ℓ -constraint locators satisfy a useful composition property. That is, given constraint locators for encodings ENC_{in} and ENC_{out} , we can construct a constraint locator for the composed encoding $\text{ENC}_{\text{out}} \circ \text{ENC}_{\text{in}}$.

Claim 4.10. Let $\text{ENC}_{in}: \mathbb{F}^S \rightarrow \mathbb{F}^{D_{in}}$ and $\text{ENC}_{out}: \mathbb{F}^{D_{in}} \rightarrow \mathbb{F}^{D_{out}}$ admit ℓ_{in} -constraint location and ℓ_{out} -constraint location respectively. Then $\text{ENC}_{out} \circ \text{ENC}_{in}(M; \mu_{in}, \mu_{out}) := \text{ENC}_{out}(\text{ENC}_{in}(m; \mu_{in}); \mu_{out})$ admits $(\ell_{in} \circ \ell_{out})$ -constraint location.

Construction 4.11. A constraint locator CL_* for ENC_* , which receives oracle access to the message m and is given constraint locators CL_{in} and CL_{out} for ENC_{in} and ENC_{out} respectively. It receives as input a subdomain $I \subseteq D_{out}$.

$\text{CL}_*(I)$:

1. Compute $(R_{out}, Z_{out}) := \text{CL}_{out}(I)$.
2. Compute $(R_{in}, Z_{in}) := \text{CL}_{in}(R_{out})$.
3. Define Z_* to the $(k_{in} + k_{out}) \times (|I| + |R_{out}| + |R_{in}|)$ matrix defined as

$$Z_* := \begin{pmatrix} Z_{out|I} & Z_{out|R_{out}} & \mathbf{0} \\ \mathbf{0} & Z_{in|R_{out}} & Z_{in|R_{in}} \end{pmatrix},$$

where k_{in} and k_{out} are the number of rows in Z_{in} and Z_{out} respectively.

4. Using Gaussian elimination, compute a basis \mathcal{B} for the space $V := \ker(Z_*)$.
5. Compute a basis \mathcal{B}' for the space $(V|_{I \cup R_{in}})^\perp$, using Claim 3.7 and Claim 3.8. Define Z to be the matrix whose rows are the elements of \mathcal{B}' .
6. Output (R_{in}, Z) .

Proof. Let $m \in \mathbb{F}^S$, and denote $w_m := \text{ENC}_{in}(m; \mu_{in})$ and $c_w := \text{ENC}_{out}(w; \mu_{out})$.

By the correctness of CL_{out} , it holds that $\vec{\alpha} \in \mathbb{F}^I$ satisfies $(\vec{\alpha}, w|_{R_{out}}) \in \ker(Z_{out})$ if and only if there exists randomness μ_{out} such that for every message $w \in \mathbb{F}^{D_{in}}$, $c_w|_I = \vec{\alpha}$. Similarly, we have that $\vec{\beta} \in \mathbb{F}^{R_{out}}$ satisfies $(\vec{\beta}, m|_{R_{in}}) \in \ker(Z_{in})$ if and only if there exists randomness μ_{in} such that for every message $m \in \mathbb{F}^S$, $w_m|_{R_{out}} = \vec{\beta}$.

Then by definition of Z_* we must have that $(\vec{\alpha}, \vec{\beta}, m|_{R_{in}}) \in \ker(Z_*)$ if and only if there exists μ_{out}, μ_{in} such that for all messages $m \in \mathbb{F}^S$, it holds that $w_m|_{R_{out}} = \vec{\beta}$ and $c_{w_m}|_I = \vec{\alpha}$. Finally, as Z is defined to be the matrix whose rows are the elements of a basis for $(V|_{I \cup R_{in}})^\perp = (\ker(Z_*|_{I \cup R_{in}}))^\perp$, we have that $(\vec{\alpha}, m|_{R_{in}}) \in \ker(Z)$ if and only if there exists randomness μ_{in}, μ_{out} , for which $\text{ENC}_*(m; \mu_{in}, \mu_{out})|_I = \vec{\alpha}$. \square

Note that this claim does not seem to generalise to permit composition of arbitrary locally simulatable encodings. Indeed, here we are implicitly using the fact that a constraint locator allows efficient sampling of a random input to ENC_{out} consistent with (the restriction to I of) a given output. For general locally simulatable encodings, such an algorithm may not exist.

5 Constraint location for random low-degree extensions

In this section we construct an efficient constraint locator for random low-degree extensions of a given function. In Section 5.1, we construct an efficient constraint *detector* (cf. Definition 4.7) for polynomials which are zero on a given product set \mathcal{S} , i.e., for the code $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])$. In Section 5.2, we use this constraint detector to construct a “decision” constraint locator for random low-degree extensions, which accepts on inputs which are constrained, and rejects on inputs which are unconstrained. In Section 5.3, we efficiently reduce the task of (search) constraint location to the decision version.

In order to state the main theorem of this section, we first define a randomised encoding that maps a function f to a random extension of specified degree.

Definition 5.1. Let $\text{ENC}_{\text{RM}(\vec{d}, \mathcal{S})}$ to be the randomised encoding function that maps $f: \mathcal{S} \rightarrow \mathbb{F}$, where $\mathcal{S} := S_1 \times \cdots \times S_m \subseteq \mathbb{F}^m$, $d_i \geq 2(|S_i| - 1)$ for all $i \in [m]$, to a random degree- \vec{d} extension of f , i.e., a uniformly random element of $\text{LD}_{\vec{d}}[f]$.

Note the restriction that $|d_i| \geq 2(|S_i| - 1)$; this is tight in the case that $|S_1| = \cdots = |S_m| = 2$. We prove the following theorem.

Theorem 5.2. There is an n -constraint locator for $\text{ENC}_{\text{RM}(\vec{d}, \mathcal{S})}$ running in time $\text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$.

As a corollary, we obtain a local simulator for $\text{ENC}_{\text{RM}(\vec{d}, \mathcal{S})}$: that is, we obtain an algorithm that can efficiently simulate a random low-degree extension of f given oracle access to f .

Corollary 5.3. $\text{ENC}_{\text{RM}(\vec{d}, \mathcal{S})}$ is (t, n) -locally simulatable for $t(n) = \text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$.

5.1 A constraint detector for random low-degree extensions of the zero function

We construct an efficient constraint detector for $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])$, where \mathcal{S} is a product set, and $\mathbf{Z}_{\mathcal{S}}(\cdot)$ denotes the subcode consisting of codewords which are zero on \mathcal{S} . We shall employ the following result, which is proved in [BCFGRS17].

Theorem 5.4. There is a constraint detector $\text{CD}_{\vec{d}}$ for $\text{RM}[\mathbb{F}, m, \vec{d}]$, running in time $\text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$.

We prove the following.

Lemma 5.5. Construction 5.6 is a constraint detector for $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])$, running in time $\text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$.

Construction 5.6. A constraint detector for $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])$, where $\mathcal{S} := S_1 \times \cdots \times S_m$, given a constraint detector $\text{CD}_{\vec{d}}$ for $\text{RM}[\mathbb{F}, m, \vec{d}]$. Receives as input a subdomain $I \subseteq \mathbb{F}^m$.

$\text{CD}_{\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])}(I)$:

1. For each $i \in [m]$, compute $Z_i := \text{CD}_{\vec{d}_i}(I)$ where $\vec{d}_i := (d_1, \dots, d_i - |S_i|, \dots, d_m)$.
2. Define the block diagonal matrix

$$Z' := \begin{pmatrix} Z_1 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & Z_2 & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & Z_m \end{pmatrix}.$$

3. Using Gaussian elimination, compute a basis b_1, \dots, b_k for $\ker(Z')$ and define \mathcal{B} to be the $k \times |I| \cdot m$ matrix whose i -th row is b_i .
4. Define the $|I| \cdot m \times |I|$ matrix A_I to be the matrix whose $((\vec{\alpha}, i), \vec{\gamma})$ -entry is $\mathcal{Z}_{S_i}(\alpha_i)$ if $\vec{\alpha} = \vec{\gamma}$ and 0 otherwise (where $\mathcal{Z}_{S_i}(x) := \prod_{s \in S_i} (x - s)$).
5. Compute $G := (\mathcal{B}A_I)^T$, an $|I| \times k$ matrix, which is a generator matrix for $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_I$.
6. Using G , compute a parity check matrix H for $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_I$, and output H .

Proof. In order to prove that Construction 5.6 is a constraint detector for $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])$, it suffices to show that $G := (\mathcal{B}A_I)^T$ is a valid generator matrix for $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_I$; i.e., $\text{Im}(G) = \mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_I$.

We show that $\text{Im}(G) \subseteq \mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_I$ using the fact that $\ker(Z') = \prod_{i=1}^m \text{RM}[\mathbb{F}, m, \vec{d}_i]|_I$ (Theorem 5.4) and some linear algebra. We prove the reverse inclusion via the combinatorial nullstellensatz (see Lemma 3.13).

Before we proceed, it will be very useful to write down an expression for an individual entry of G . For $\vec{\beta} \in I, \ell \in [k]$, the $(\vec{\beta}, \ell)$ -entry of G , which we will denote by $G_{\ell}(\vec{\beta})$, is given by the expression

$$G_{\ell}(\vec{\beta}) = \sum_{i=1}^m b_{\ell}(i, \vec{\beta}) \mathcal{Z}_{S_i}(\beta_i), \quad (3)$$

where b_{ℓ} is the ℓ -th basis element computed in Step 3 and $\mathcal{Z}_{S_i}(X) := \prod_{s \in S_i} (x - s)$. We now proceed with the proof.

Correctness. We first show $\text{Im}(G) \subseteq \mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_I$. Let $y \in \text{Im}(G)$. Then, for some scalars $a_1, \dots, a_k \in \mathbb{F}$, and for each $\vec{\beta} \in I$, we can write

$$y(\vec{\beta}) = \sum_{j=1}^k a_j G_j(\vec{\beta}) = \sum_{j=1}^k a_j \sum_{i=1}^m b_j(i, \vec{\beta}) \mathcal{Z}_{S_i}(\beta_i) = \sum_{i=1}^m \mathcal{Z}_{S_i}(\beta_i) \sum_{j=1}^k a_j b_j(i, \vec{\beta}),$$

where the second equality follows from Eq. 3. We also know that $b_j(i, \vec{X}) \in \text{RM}[\mathbb{F}, m, \vec{d}_i]|_I$ for each $j \in [k]$, and $i \in [m]$ because b_1, \dots, b_k is a basis for $\ker(Z') = \prod_{i=1}^m \text{RM}[\mathbb{F}, m, \vec{d}_i]|_I$. Then by the linearity of $\text{RM}[\mathbb{F}, m, \vec{d}_i]$, we have that for each $i \in [m]$, $\sum_{j=1}^k a_j b_j(i, \vec{X}) := R_i(\vec{X}) \in \text{RM}[\mathbb{F}, m, \vec{d}_i]|_I$. Substituting this into our expression for $y(\vec{\beta})$,

$$y(\vec{\beta}) = \sum_{i=1}^m R_i(\vec{\beta}) \mathcal{Z}_{S_i}(\beta_i).$$

Therefore $y \in \text{RM}[\mathbb{F}, m, \vec{d}]|_I$. As $\mathcal{Z}(\beta_i) = 0$ for all $\beta_i \in S_i$ and $i \in [m]$, we have that $y(\vec{\beta}) = 0$ for all $\vec{\beta} \in \mathcal{S}$, so $y \in \mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_I$.

Now we show $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_I \subseteq \text{Im}(G)$. Let $\tilde{Z} \in \mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_I$. Then by the combinatorial nullstellensatz (Lemma 3.13) there exists $\vec{R} := (R_1, \dots, R_m) \in \prod_{i=1}^m \text{RM}[\mathbb{F}, m, \vec{d}_i]|_I$ such that

$$\tilde{Z}(\vec{\beta}) = \sum_{i=1}^m \mathcal{Z}_{S_i}(\beta_i) R_i(\vec{\beta}),$$

for all $\beta \in I$. As $\{b_1, \dots, b_k\}$ is a basis for $\prod_{i=1}^m \text{RM}[\mathbb{F}, m, \vec{d}_i]|_I$, we can write

$$R_i(\vec{\beta}) = \sum_{j=1}^k a_j b_j(i, \vec{\beta}),$$

for some $a_j \in \mathbb{F}$. Then substituting into the previous expression for \tilde{Z} yields

$$\tilde{Z}(\vec{\beta}) = \sum_{i=1}^m \mathcal{Z}_{S_i}(\beta_i) \sum_{j=1}^k a_j b_j(i, \vec{\beta}) = \sum_{j=1}^k a_j \left(\sum_{i=1}^m b_j(i, \vec{\beta}) \mathcal{Z}_{S_i}(\beta_i) \right) = \sum_{j=1}^k a_j G_j(\vec{\beta}).$$

Efficiency. Construction 5.6 makes m calls to $\text{CD}_{\vec{d}_i}$, which runs in time $\text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$. In the remaining steps, the algorithm performs only Gaussian elimination and matrix multiplication of matrices whose sizes are polynomial in $|I|$ and m . \square

5.2 The decision problem

Before stating the main result of this section we define the notion of a constraint with respect to a linear code on a subset of its domain.

Definition 5.7. Let $C \subseteq \mathbb{F}^D$ be a linear code. A subset $I \subseteq D$ is **constrained with respect to C** if there exists a nonzero vector $z \in \mathbb{F}^I$ such that, for every codeword $w \in C$, $z \cdot w|_I = 0$ (equivalently, if there exists $z \neq 0 \in C^\perp$ with $\text{supp}(z) \subseteq I$); we refer to z as a **constraint with respect to C on I** . We say that I is **unconstrained with respect to C** if it is not constrained with respect to C .

We construct an algorithm `CheckConstraints` which, given as input a (large) product set \mathcal{S} and a polynomial-size set I , efficiently determines whether $I \cup \mathcal{S}$ is constrained with respect to the Reed-Muller code.

Lemma 5.8. Let $\mathcal{S} := S_1 \times \dots \times S_m$, $I \subseteq \mathbb{F}^m$ and $\vec{d} = (d_1, \dots, d_m)$ with $d_i \geq |S_i|$ for all $i \in [m]$. `CheckConstraints $_{\vec{d}}$` (I, \mathcal{S}) outputs **yes** if and only if $I \cup \mathcal{S}$ is constrained with respect to $\text{RM}[\mathbb{F}, m, \vec{d}]$, and runs in time $\text{poly}(\log |\mathbb{F}|, m, \max_i d_i, |I|)$.

Construction 5.9. An algorithm which efficiently determines whether a set $I \cup \mathcal{S}$ is constrained with respect to the Reed–Muller code, given an efficient constraint detector $\text{CD}_{\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])}(I)$ for $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])$.

`CheckConstraints $_{\vec{d}}$` (I, \mathcal{S}):

1. Compute $H = \text{CD}_{\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])}(I \setminus \mathcal{S})$ (see Construction 5.6).
2. If $H = 0$, output **no**, otherwise output **yes**.

To prove the correctness of Construction 5.9, we show an equivalence between constraint detection for $\mathbf{Z}_{\mathcal{S}}(\text{RM})$ and deciding whether $I \cup \mathcal{S}$ is constrained with respect to RM . In order to do so, we shall require the following two properties of linear codes, whose proofs we defer to Appendix A.1 and Appendix A.2.

Claim 5.10. Let $C \subseteq \mathbb{F}^D$ be a linear code and let $I \subseteq D$. Then I is unconstrained with respect to C if and only if for all $x \in I$ there exists a codeword $w_x \in C$ satisfying (i) $w_x(x) = 1$ and (ii) $w_x(y) = 0$ for all $y \in I \setminus \{x\}$.

Lemma 5.11. Let $C \subseteq \mathbb{F}^D$ be a linear code, and let $I, \mathcal{S} \subseteq D$. Suppose that \mathcal{S} is unconstrained with respect to C . Then $I \cup \mathcal{S}$ is constrained with respect to C if and only if $I \setminus \mathcal{S}$ is constrained with respect to $\mathbf{Z}_{\mathcal{S}}(C)$.

As a corollary, we show an equivalence between constraints with respect to the general Reed–Muller code and the Reed-Muller code fixed to be zero on a product set.

Corollary 5.12. Let $m \in \mathbb{N}$, $\vec{d} \in \mathbb{N}^m$, $I \subseteq \mathbb{F}^m$, let $S_1, \dots, S_m \subseteq \mathbb{F}$, be such that $|S_i| \leq d_i + 1$ for all $i \in [m]$, and denote $\mathcal{S} := S_1 \times \dots \times S_m$. Then $I \cup \mathcal{S}$ is constrained with respect to $\text{RM}[\mathbb{F}, m, \vec{d}]$ if and only if $I \setminus \mathcal{S}$ is constrained with respect to $\mathbf{Z}_{\mathcal{S}}(\text{RM}[\mathbb{F}, m, \vec{d}])$.

Proof. By Lemma 5.11, it suffices to show that \mathcal{S} is unconstrained with respect to $\text{RM}[\mathbb{F}, m, \vec{d}]$. Consider, for each $w \in \mathcal{S}$, the Lagrange polynomial $L_{\mathcal{S}, w}(\vec{X}) \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$. As $|S_i| \leq d_i + 1$ for all $i \in [m]$, the evaluation table of $L_{\mathcal{S}, w}(\vec{X})$ is a codeword in $\text{RM}[\mathbb{F}, m, \vec{d}]$, for each $w \in \mathcal{S}$. These codewords satisfy $L_{\mathcal{S}, w}(w) = 1$ and $L_{\mathcal{S}, w}(x) = 0$ for all $x \neq w \in \mathcal{S}$, so by Claim 5.10, \mathcal{S} is unconstrained with respect to $\text{RM}[\mathbb{F}, m, \vec{d}]$. \square

Proof of Lemma 5.8. The correctness follows from Corollary 5.12 and the correctness of Construction 5.6. The runtime follows from the fact that Construction 5.6 runs in time $\text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$. \square

5.3 Search-to-decision reduction

In this section we show a reduction from constraint location of random low-degree extensions (a search problem) to CheckConstraints. Before proving Theorem 5.2, we define what it means for a subset I of the domain D of a code to *determine* a point $x \in D$. We then prove some useful technical lemmas.

Definition 5.13. *Let $C \subseteq \mathbb{F}^D$ be a linear code. We say that $I \subseteq D$ **determines** $x \in D$ with respect to C if $x \in I$ or there exists a constraint z with respect to C on $I \cup \{x\}$ such that $z(x) \neq 0$.*

The following lemma is central to our reduction. It shows that points which are constrained with respect to the Reed–Muller code are determined with respect to another Reed–Muller code of lower degree. Thus, in our algorithm, it suffices to check the latter condition.

Lemma 5.14. *Let $m \in \mathbb{N}$, $\vec{d} := (d_1, \dots, d_m) \in \mathbb{N}^m$, $I \subseteq \mathbb{F}^m$, let $S_1, \dots, S_m \subseteq \mathbb{F}$ and denote $\mathcal{S} := S_1 \times \dots \times S_m$. If there exists a constraint $z: I \cup \mathcal{S} \rightarrow \mathbb{F}$ on $I \cup \mathcal{S}$ with respect to $\text{RM}[\mathbb{F}, m, \vec{d}]$ such that for some $w \in \mathcal{S}$, $z(w) \neq 0$, then w is determined by I with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$, where $\vec{d}' := (d_1 - (|S_1| - 1), \dots, d_m - (|S_m| - 1))$.*

Proof. We proceed via contrapositive: if $w \in \mathcal{S}$ is not determined by I with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$, then $I \cup \mathcal{S}$ is unconstrained with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$.

Let $w \in \mathcal{S}$ be not determined by I with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$. Then there exists $p(\vec{X}) \in \mathbb{F}^{\leq \vec{d}'}[X_1, \dots, X_m]$ such that (i) $p(w) = 1$ and (ii) $p(i) = 0$ for all $i \in I$. Now define $q(\vec{X}) := p \cdot L_{\mathcal{S}, w}(\vec{X}) \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$, where $L_{\mathcal{S}, w}$ denotes the Lagrange polynomial. Note $q(\vec{X})$ satisfies: (i) $q(w) = 1$, (ii) $q(x) = 0$ for all $x \in \mathcal{S} \setminus \{w\}$ and (iii) $q(i) = 0$ for all $i \in I$. By Claim 5.10 $I \cup \mathcal{S}$ is unconstrained with respect to $\text{RM}[\mathbb{F}, m, \vec{d}]$. \square

We shall also need the following claim which defines the notion of an *interpolating set* and gives an algorithm for reducing an arbitrary subdomain of a linear code to an interpolating set. We defer the proof of this claim to Appendix A.3.

Claim 5.15. *Let $C \subseteq \mathbb{F}^D$ be a linear code which has an efficient constraint detector. Given a subdomain $I \subseteq D$, there exists an efficient algorithm, which computes an **interpolating set for I with respect to C** : a set $I' \subseteq I$ such that: (i) I' is unconstrained with respect to C ; and (ii) for every $\mathcal{S} \subseteq D$ which is unconstrained with respect to C and satisfies $I \cap \mathcal{S} = \emptyset$, if there exists a constraint $z: I \cup \mathcal{S} \rightarrow \mathbb{F}$ such that $z(s) \neq 0$ for some $s \in \mathcal{S}$, then there exists a constraint $z': I' \cup \mathcal{S} \rightarrow \mathbb{F}$ such that $z'(s) \neq 0$.*

We will also require the following simple monotonicity property of CheckConstraints (see Construction 5.9).

Claim 5.16. *If $\text{CheckConstraints}(I, \mathcal{S}) = \text{yes}$ and $\mathcal{S} \subseteq \mathcal{S}'$, then $\text{CheckConstraints}(I, \mathcal{S}') = \text{yes}$.*

Proof. Clear from Lemma 5.8. \square

Before we give our reduction, we give a slight generalisation of Lemma 4.2 from [AW09]. Its proof is a relatively straightforward adaptation of the proof of that lemma and is deferred to Appendix A.4.

Lemma 5.17. *Let $m \in \mathbb{N}$, let $I \subseteq \mathbb{F}^m$, let $S_1, \dots, S_m \subseteq \mathbb{F}$ and denote $\mathcal{S} := S_1 \times \dots \times S_m$. Set $\vec{d} := (|S_1| - 1, \dots, |S_m| - 1)$. There exists a set $G_I \subseteq \mathcal{S}$ with $|G_I| \geq |\mathcal{S}| - |I|$, such that for all $g \in G_I$ there exists a polynomial $p_g(\vec{X}) \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ satisfying*

- (i) $p_g(x) = 0$ for all $x \in I \cup G_I \setminus \{g\}$ and
- (ii) $p_g(g) = 1$.

We are now ready to prove the main result of this section.

Proof of Theorem 5.2. We give a reduction from constraint location to $\text{CheckConstraints}_{\vec{d}}$.

Reduction 1 Reduction from constraint location for $\text{ENC}_{\text{RM}(\vec{d}, \mathcal{S})}$ to $\text{CheckConstraints}_{\vec{d}}$ and a constraint detector $\text{CD}_{\vec{d}}$ for $\text{RM}[\mathbb{F}, m, \vec{d} := (d_1, \dots, d_m)]$.

Set $\vec{d}' := (d_1 - (|S_1| - 1), \dots, d_m - (|S_m| - 1))$.

Search($I, t, M := (S_1, \dots, S_m)$):

if $M = \perp$ **then**

Return t .

else

Set $T = \emptyset$.

for all $s \in S_i$ **do**

if $\text{CheckConstraints}_{\vec{d}'}(I, t \times \{s\} \times S_{i+1} \times \dots \times S_m) = \text{yes}$ **then**

Add $\text{Search}(I, t \times \{s\}, (S_{i+1}, \dots, S_m))$ to T .

Return T .

Compute I' to be an interpolating set for I .

Compute $T := \text{Search}(I', \perp, (S_1, \dots, S_m))$.

Output $(T, \text{CD}_{\vec{d}}(I \cup T))$.

Correctness. First, we show that if $I \cup \mathcal{S}$ is unconstrained with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$ then Reduction 1 outputs \perp . In this case, as $T \subseteq \mathcal{S}$, $I \cup T$ is unconstrained so $\text{CD}_{\vec{d}}(I \cup T) = \perp$, as required.

Now we consider the case where $I \cup \mathcal{S}$ is constrained with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$. We show that if $s \in S_1 \times \dots \times S_m$ is constrained, then $s \in T$, where T is the output of Reduction 1. By inspection of Reduction 1 and Claim 5.16, we see that $s \in T$ if and only if $\text{CheckConstraints}_{\vec{d}'}(I', \{s_1\} \times \dots \times \{s_m\}) = \text{yes}$. If there exists a constraint $z: I \cup \mathcal{S} \rightarrow \mathbb{F}$ with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$ with $z(s) \neq 0$, then by Claim 5.15, there exists a constraint $z': I' \cup \mathcal{S}$ with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$ with $z'(s) \neq 0$. Further, by Lemma 5.14, s is determined by I' with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$.

Lastly, $\text{CD}_{\vec{d}}(I \cup T)$ will output a basis for the space of constraints on $I \cup T$, so the output is in the desired form.

Efficiency. We show that the reduction is efficient by bounding the number of calls made to $\text{CheckConstraints}_{\vec{d}'}$. For $i \in [m]$, we say a point $(s_1, \dots, s_i) \in S_1 \times \dots \times S_i$ is *i-accepting* if $\text{CheckConstraints}_{\vec{d}'}(I, \{s_1\} \times \dots \times \{s_i\} \times S_{i+1} \times \dots \times S_m) = \text{yes}$. Note that by correctness of CheckConstraints and the fact that I' is unconstrained, this means that a point $(s_1, \dots, s_i) \in S_1 \times \dots \times S_i$ is *i-accepting* if and only if there exists $(s_{i+1}, \dots, s_m) \in S_{i+1} \times \dots \times S_m$ such that (s_1, \dots, s_m) is involved in a non-zero constraint with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$. For all $i \in [m]$, let a_i denote the number of points which are *i-accepting*.

Now we count the number of calls our reduction makes to CheckConstraints in terms of a_1, \dots, a_{m-1} . Note that regardless of the value of I , Reduction 1 makes $|S_1|$ calls to $\text{CheckConstraints}_{\vec{d}'}$ at the beginning. Subsequently, for each $i \in [m-1]$, Reduction 1 makes $a_i |S_{i+1}|$ many calls to $\text{CheckConstraints}_{\vec{d}'}$. Overall, the total number of calls that Reduction 1 makes to $\text{CheckConstraints}_{\vec{d}'}$ is:

$$N = |S_1| + \sum_{i=1}^{m-1} a_i |S_{i+1}|.$$

Now we bound the size of a_i for each $i \in [m]$. As $d'_i \geq |S_i| - 1$ for all $i \in [m]$, Lemma 5.17 implies that there exists a set G_I with $|G_I| \geq |S| - |I|$, such that for all $g \in G_I$ there exists a polynomial $p_g(\vec{X}) \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ satisfying (i) $p_g(x) = 0$ for all $x \in I \cup G_I \setminus \{g\}$ and (ii) $p_g(g) = 1$. Thus, by Claim 5.10, at most $|I|$ many points in S are constrained by I with respect to $\text{RM}[\mathbb{F}, m, \vec{d}']$. In other words, $a_i \leq |I|$ for all $i \in [m-1]$. Therefore

$$N \leq |S_1| + |I| \sum_{i=1}^{m-1} |S_{i+1}|.$$

Thus, as each call to `CheckConstraints` takes time $\text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$, and $|S_i|$ is bounded by a polynomial in d_i for each i , we have that $t(n) = \text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$.

Finally, the fact that at most $|I|$ many points in S are constrained implies that $\ell(n) = n$. \square

6 Constraint location for subcube sums of random low-degree extensions

In this section we extend the efficient constraint locator of the previous section to support queries to subcube sums. We begin by defining the corresponding randomised encoding.

Definition 6.1. Let $\mathcal{A} := A_1 \times \dots \times A_m \subseteq \mathbb{F}^m$. We define $\text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})} : \Sigma_{\mathcal{A}}(\mathcal{A} \rightarrow \mathbb{F}) \rightarrow (\mathbb{F}^{\leq m} \rightarrow \mathbb{F})$ as the randomised encoding $\text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})}(\Sigma_{\mathcal{A}}[F]) \sim \mathcal{U}(\Sigma_{\mathcal{A}}[\text{LD}_{\vec{d}}[F]])$, where $d_i \geq |A_i| - 1$ for all $i \in [m]$.

That is, $\text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})}$ encodes a message $\Sigma_{\mathcal{A}}[F] \in \Sigma_{\mathcal{A}}(\mathcal{A} \rightarrow \mathbb{F})$ by first choosing a random degree- \vec{d} extension \hat{F} of F , and then outputting the word $\Sigma_{\mathcal{A}}[\hat{F}] : \mathbb{F}^{\leq m} \rightarrow \mathbb{F}$ obtained by augmenting \hat{F} with partial sums over \mathcal{A} (see Definition 3.6). Note that the input to $\text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})}$ includes partial sums of F ; these are not necessary to define the encoding but are essential to permit local simulation.

The main theorem of this section is the following.

Theorem 6.2. There is a (t, ℓ) -constraint locator for $\text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})}$, where $\ell(n) = nm(m(a+1)+1)^2$, for $a := \max_i |A_i|$ and $t(n) = \text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$.

We prove this theorem by giving, in Section 6.1, a local characterisation of ΣRM in terms of the plain Reed–Muller code. We give our construction based on this characterisation in Section 6.2.

6.1 Local characterisation of ΣRM

In this subsection we give a local characterisation of ΣRM in terms of the plain Reed–Muller code. We define the code ΣRM as follows.

Definition 6.3. Define $\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}] := \Sigma_{\mathcal{A}}\text{RM}[\mathbb{F}, m, \vec{d}] \subseteq (\mathbb{F}^{\leq m} \rightarrow \mathbb{F})$.

Our local characterisation will hold for certain “nice” subsets of $\mathbb{F}^{\leq m}$ which we call “ \mathcal{A} -closed”.

Definition 6.4. A set $S \subseteq \mathbb{F}^{\leq m}$ is \mathcal{A} -closed if the following two conditions hold:

(i) if $(s_1, \dots, s_t) \in S$, then $(s_1, \dots, s_{t-1}) \in S$; and

(ii) if $(s_1, \dots, s_\ell) \in S$, then for all $a_\ell \in A_\ell$, $(s_1, \dots, s_{\ell-1}, a_\ell) \in S$.

Given a set $X \subseteq \mathbb{F}^{\leq m}$, we refer to the smallest \mathcal{A} -closed set containing X as the \mathcal{A} -closure of X , i.e., \hat{X} is the \mathcal{A} -closure of X if \hat{X} is \mathcal{A} -closed, $X \subseteq \hat{X}$, and for all \mathcal{A} -closed $Y \supseteq X$, it holds that $\hat{X} \subseteq Y$.

Globally the ΣRM code can be viewed as a collection of plain Reed–Muller codes on $1, 2, \dots, m$ variables, related by summation constraints. The following theorem asserts that this characterisation also holds locally: over any \mathcal{A} -closed subdomain S , $(\Sigma\text{RM}|_S)^\perp$ is spanned by summation constraints $\{z_{\vec{s}}\}_{\vec{s} \in S^*}$ and the local constraints on a collection of Reed–Muller codes. For technical reasons, we show this with respect to both the “plain” ΣRM code and the subcode $\mathbf{Z}_{\mathcal{A}}(\Sigma\text{RM})$ of encodings of the zero word.

Theorem 6.5. *Let $A_1, \dots, A_m \subseteq \mathbb{F}$, $\mathcal{A} := A_1 \times \dots \times A_m$ and let S be an \mathcal{A} -closed set. Then for $X \in \{\emptyset, \mathcal{A}\}$, we have*

$$(\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_S)^\perp = \text{span} \left(\{z_{\vec{s}}\}_{\vec{s} \in S^*} \cup \bigcup_{i=1}^m (\mathbf{Z}_{X_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)])|_{S_i})^\perp \right)$$

where $X_i := \{\vec{x} \in X : \ell(\vec{x}) = i\}$, $S_i := \{\vec{s} \in S : \ell(\vec{s}) = i\}$, $S^* := \{(s_1, \dots, s_{\ell-1}) \in S : \exists s_\ell \text{ s.t. } (s_1, \dots, s_{\ell-1}, s_\ell) \in S\}$, and

$$z_{\vec{s}}(\vec{t}) := \begin{cases} 1 & \text{if } \vec{t} = \vec{s}; \\ -1 & \text{if } \vec{t} = (\vec{s}, a) \text{ for some } a \in A_{\ell(\vec{s})+1}; \\ 0 & \text{otherwise.} \end{cases}$$

We will use the above notation for X_i , S_i , S^* and $z_{\vec{s}}(\vec{t})$ throughout the remainder of this section.

We prove Theorem 6.5 by a dimensionality argument. In particular, we show that any basis for $(\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_S)^\perp$ which is in echelon form can be mapped to a set of linearly independent vectors in the span of summation constraints and Reed–Muller constraints.

The key technical tool used here is a “flattening lemma” (Lemma 6.8). This lemma provides a means to map any constraint z on subcube sums of m -variate polynomials to a new constraint z' on subcube sums of $(m-1)$ -variate polynomials, while preserving the value of $\lambda(z)$. Crucially, in the case where $|\lambda(z)| = m-1$, all dependence of z' on subcube sums vanishes, and we obtain a constraint with respect to the plain Reed–Muller code.

By repeated applications of the flattening lemma we map any constraint over subcube sums to a Reed–Muller constraint; we then show by a counting argument that these “flattened” constraints, along with the summation constraints, span the dual code.

Proof of Theorem 6.5. Define $Z^* := \{z_{\vec{s}}\}_{\vec{s} \in S^*}$, and $Z_i := (\mathbf{Z}_{X_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)])|_{S_i})^\perp$ for $i \in [m]$. It follows easily from the definition of ΣRM that $Z_1, \dots, Z_m \subseteq (\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_S)^\perp$. As S is \mathcal{A} -closed we also have that $Z^* \subseteq (\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_S)^\perp$. Hence it suffices to show that $\dim(\text{span}(Z^* \cup \bigcup_{i=1}^m Z_i)) \geq \dim((\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_S)^\perp)$.

Let $\Delta := \dim((\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_S)^\perp)$. Fix some ordering $<$ on S so that (i) $\ell(\vec{s}) < \ell(\vec{t})$ then $\vec{s} < \vec{t}$, and (ii) if $\ell(\vec{s}) = \ell(\vec{t})$, $\vec{s} \notin S^*$ and $\vec{t} \in S^*$ then $\vec{s} < \vec{t}$. Let b_1, \dots, b_Δ be a basis for $(\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_S)^\perp$ which is in echelon form with respect to $<$.

We show that $\dim(\text{span}(Z^* \cup \bigcup_{i=1}^m Z_i)) \geq \Delta$ by exhibiting a sequence of linearly independent vectors $b'_1, \dots, b'_\Delta \in Z^* \cup \bigcup_{i=1}^m Z_i$. These vectors are defined as follows:

$$b'_i := \begin{cases} z_{\vec{\lambda}(b_i)} & \text{if } \lambda(b_i) \in S^*; \text{ and} \\ R(b_i) & \text{otherwise,} \end{cases}$$

where R is as guaranteed by Claim 6.6 below. To show that the b'_i are linearly independent, it suffices to show that they are in echelon form with respect to $<$. This follows since $\lambda(b'_i) = \lambda(b_i)$ for all $i \in [\Delta]$: if $\lambda(b_i) \in S^*$, then $\lambda(b'_i) = \lambda(z_{\lambda(b_i)}) = \lambda(b_i)$; and if $\lambda(b_i) \notin S^*$ then $\lambda(b'_i) = \lambda(R(b_i)) = \lambda(b_i)$ by Claim 6.6. \square

It remains to prove the following technical claim.

Claim 6.6. *Let $S \subseteq \mathbb{F}^{\leq m}$, $\mathcal{A} := A_1 \times \dots \times A_m$, $\vec{d} := (d_1, \dots, d_m)$ and $X \in \{\emptyset, \mathcal{A}\}$. If $d_i \geq |A_i| - 1$ for all $i \in [m]$, then there is a function $T: (S \rightarrow \mathbb{F}) \rightarrow \bigcup_{i=1}^m (S_i \rightarrow \mathbb{F})$ such that for all $z \in (\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}]|_S)^\perp)$, $T(z) \in (\mathbf{Z}_{X_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{S_i})^\perp)$ for $i = |\lambda(z)|$, and if $\lambda(z) \notin S^*$, $\lambda(T(z)) = \lambda(z)$.*

This claim will be a consequence of Lemma 6.8 (the “flattening lemma”), which provides a way to transform a constraint on an m -variate ΣRM code into a related constraint on an $(m-1)$ -variate ΣRM code.

Definition 6.7. *Let $S \subseteq \mathbb{F}^{\leq m}$ and $A_1 \times \dots \times A_m \subseteq \mathbb{F}^m$. For $i \in [m]$ and an element $a \in A_i$, we define the “flattening map” $Q_{i,a}: (S \rightarrow \mathbb{F}) \rightarrow (S \setminus S_i \rightarrow \mathbb{F})$ as follows, for $\vec{s} \in S \setminus S_i$:*

$$Q_{i,a}(z)(\vec{s}) := \begin{cases} z(\vec{s}) + \sum_{s_i \in S_i} z(\vec{s}, s_i) L_{A_i, a}(s_i) & \text{if } \vec{s} \in S_{i-1} \cap S^*, \text{ and} \\ z(\vec{s}) & \text{otherwise.} \end{cases}$$

Lemma 6.8 (Flattening lemma). *Let $S \subseteq \mathbb{F}^{\leq m}$, $\mathcal{A} := A_1 \times \dots \times A_m$, $\vec{d} := (d_1, \dots, d_m)$, $X \in \{\emptyset, \mathcal{A}\}$, and let $z \in (\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}]|_S)^\perp)$, where $d_m \geq |A_m| - 1$. For any point $a_m^* \in A_m$, the function $Q_{m, a_m^*}: (S \rightarrow \mathbb{F}) \rightarrow (S \setminus S_m \rightarrow \mathbb{F})$ (see Definition 6.7) satisfies the following properties:*

1. *If z satisfies $|\lambda(z)| < m - 1$, then $Q_{m, a_m^*}(z) \in (\mathbf{Z}_{X_{m-1}}(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}_{m-1}, (d_1, \dots, d_{m-1})]|_{S \setminus S_i})^\perp)$ and $\lambda(Q_{m, a_m^*}(z)) = \lambda(z)$.*
2. *If z satisfies $|\lambda(z)| = m - 1$, then $Q_{m, a_m^*}(z) \in (\mathbf{Z}_{X_{|\lambda(z)|}}(\text{RM}[\mathbb{F}, |\lambda(z)|, (d_1, \dots, d_{|\lambda(z)|})]|_{S_{|\lambda(z)|}})^\perp)$, and furthermore if $\lambda(z) \notin S^*$, then $\lambda(Q_{m, a_m^*}(z)) = \lambda(z)$.*

Proof. We will proceed by showing how a constraint on subcube sums of an m -variate polynomials induces a constraint on subcube sums of $(m-1)$ -variate polynomials.

Let $z \in (\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}]|_S)^\perp)$. Define the set $P_0 := \{p \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m] : p(\vec{x}) = 0 \forall \vec{x} \in X\}$. Then for all polynomials $p \in P_0$ it holds that

$$\sum_{\vec{s} \in S} z(\vec{s}) \sum_{\vec{a} \in \mathcal{A}_{>\ell(\vec{s})}} p(\vec{s}, \vec{a}) = 0. \quad (4)$$

Let $a_m^* \in A_m$ and define the set of $(m-1)$ -variate polynomials $Q_0 := \{q \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_{m-1}] : q(\vec{x}) = 0 \forall \vec{x} \in X_{m-1}\}$, where $\vec{d}' := (d_1, \dots, d_{m-1})$. Then for any $q \in Q_0$, the polynomial defined by

$q_{a_m^*}(x_1, \dots, x_m) := q(x_1, \dots, x_{m-1}) \cdot L_{A_m, a_m^*}(x_m)$ satisfies $q_{a_m^*}(\vec{x}) = 0$ for all $\vec{x} \in X$. So $q_{a_m^*} \in P_0$. Thus, substituting $q_{a_m^*}$ into Equation 4 implies that for all polynomials $q \in Q_0$,

$$\sum_{\vec{s} \in S \setminus (S_{m-1} \cup S_m)} z(\vec{s}) \sum_{\vec{a} \in \mathcal{A}'_{>\ell(\vec{s})}} q(\vec{s}, \vec{a}) + \sum_{\vec{s} \in S_{m-1}} z(\vec{s})q(\vec{s}) + \sum_{\vec{s} \in S_{m-1}} q(\vec{s}) \sum_{\substack{s_m \\ (\vec{s}, s_m) \in S_m}} z(\vec{s}, s_m)L_{A_m, a_m^*}(s_m) = 0,$$

where $\mathcal{A}' := \mathcal{A}_{m-1}$. Recall that for points $(s_1, \dots, s_{m-1}) \notin S^*$, there does not exist $s_m \in S_m$ such that $(s_1, \dots, s_{m-1}, s_m) \in S$. Thus we can rearrange this expression as:

$$\sum_{\vec{s} \in S \setminus (S_{m-1} \cup S_m)} z(\vec{s}) \sum_{\vec{a} \in \mathcal{A}'_{>\ell(\vec{s})}} q(\vec{s}, \vec{a}) + \sum_{\vec{s} \in S_{m-1} \setminus S^*} z(\vec{s})q(\vec{s}) + \sum_{\vec{s} \in S_{m-1} \cap S^*} q(\vec{s}) \left(z(\vec{s}) + \sum_{\substack{s_m \\ (\vec{s}, s_m) \in S_m}} z(\vec{s}, s_m)L_{A_m, a_m^*}(s_m) \right) = 0. \quad (5)$$

Notice that the coefficients of Eq. 5 are precisely the values of Q_{m, a_m^*} (Definition 6.7) at each point in $S \setminus S_m$.

This demonstrates that for all $z \in (\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_S)^\perp$, $Q_{m, a_m^*}(z) \in (\mathbf{Z}_{X_{m-1}}(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}', \vec{d}'])|_{S \setminus S_m})^\perp$.

Now we will deal with the two cases of the lemma statement separately.

Case (i) $|\lambda(z)| + 1 < m$. In this case we show that $\lambda(Q_{m, a_m^*}(z)) = \lambda(z)$.

As $|\lambda(z)| + 1 < m$, $\lambda(z) \in S \setminus (S_m \cup S_{m-1})$. But we have that $Q_{m, a_m^*}(z)(\vec{s}) = z(\vec{s})$ for all $\vec{s} \in S \setminus (S_m \cup S_{m-1})$, so in particular $Q_{m, a_m^*}(z)(\vec{s}) = z(\vec{s}) = 0$ for all $\vec{s} < \lambda(z)$, and $Q_{m, a_m^*}(z)(\lambda(z)) = z(\lambda(z)) \neq 0$. Thus $\lambda(Q_{m, a_m^*}(z)) = \lambda(z)$.

Case (ii) $|\lambda(z)| + 1 = m$. First, we show that $Q_{m, a_m^*}(z) \in (\mathbf{Z}_{X_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{S_i})^\perp$, where $i := |\lambda(z)|$. Second, we show that if $\lambda(z) \notin S^*$, then $\lambda(Q_{m, a_m^*}(z)) = \lambda(z)$.

For the first part, as $|\lambda(z)| = m - 1$, $z(\vec{s}) = 0$ for all $\vec{s} \in S \setminus (S_{m-1} \cup S_m)$. As a result, Equation 5 simplifies greatly and becomes:

$$\sum_{\vec{s} \in S_{|\lambda(z)|} \setminus S^*} z(\vec{s})q(\vec{s}) + \sum_{\vec{s} \in S_{|\lambda(z)|} \cap S^*} q(\vec{s}) \left(z(\vec{s}) + \sum_{\substack{s_m \\ (\vec{s}, s_m) \in S_m}} z(\vec{s}, s_m)L_{A_m, a_m^*}(s_m) \right) = 0. \quad (6)$$

Note that Equation 6 is a constraint with respect to the Reed-Muller code $(\mathbf{Z}_{X_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{S_i})^\perp$. So $Q_{m, a_m^*}(z) \in (\mathbf{Z}_{R_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{S_i})^\perp$.

For the second part, assume $\lambda(z) \notin S^*$. By the definition of $<$, if $\vec{s} < \lambda(z)$, then it must be that $\vec{s} \notin S^*$. However, if $\vec{s} \notin S^*$, then $Q_{m, a_m^*}(z)(\vec{s}) = z(\vec{s})$. In particular, for $\vec{s} < \lambda(z)$, $Q_{m, a_m^*}(z)(\vec{s}) = z(\vec{s}) = 0$, and $Q_{m, a_m^*}(z)(\lambda(z)) = z(\lambda(z)) \neq 0$, so $\lambda(Q_{m, a_m^*}(z)) = \lambda(z)$. \square

We now apply Lemma 6.8 to prove Claim 6.6, which will complete the proof of Theorem 6.5.

Proof of Claim 6.6. For each $i \in \{|\lambda(z)| + 1, \dots, m\}$ arbitrarily chose an element $a_i \in A_i$, and for notational convenience, denote by $Q_i := Q_{i, a_i}$ the flattening map (Definition 6.7). For each $z \in (\mathbf{Z}_X(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_S)^\perp$, define the value of $T(z) := (Q_{|\lambda(z)|+1} \circ \dots \circ Q_{m-1} \circ Q_m)(z)$.

By part 1 of Lemma 6.8, we see that $(Q_{|\lambda(z)|+2} \circ \dots \circ Q_m)(z) \in (\mathbf{Z}_{X_{i+1}}(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}_{i+1}, (d_1, \dots, d_{i+1})])|_{S \setminus S_m})^\perp$, with $\lambda((Q_{|\lambda(z)|+2} \circ \dots \circ Q_m)(z)) = \lambda(z)$. Combining this with part 2 of Lemma 6.8 we see that $T(z) := (Q_{|\lambda(z)|+1} \circ \dots \circ Q_m)(z) \in (\mathbf{Z}_{X_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{S_i})^\perp$, and provided that $\lambda(z) \notin S^*$, then $\lambda(T(z)) = \lambda(z)$. \square

6.2 An ℓ -constraint locator for ΣRM

In this section we construct an constraint locator for ΣRM , by combining the structure implied by Theorem 6.5 with the constraint locator we constructed in Section 5. We will often denote $\Sigma\text{RM} := \Sigma\text{RM}[\mathbb{F}, m, \vec{d}]$ when \mathbb{F}, m, \vec{d} are clear from context.

Construction 6.9. An ℓ -constraint locator for $\text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})}$. For each $i \in [m]$, let CL_i be an ℓ_i -constraint locator for $\text{ENC}_{\text{RM}((d_1, \dots, d_i), \mathcal{A}_i)}$ (cf. Theorem 5.2). Input: a subdomain $I \subseteq \mathbb{F}^m$.

$\text{CL}_\Sigma(I)$:

1. Compute the \mathcal{A} -closure \hat{I} of I .
2. For each $i \in [m]$, compute $(R_i, Z_i) := \text{CL}_i(\hat{I}_i)$, where $\hat{I}_i := \{\vec{x} \in \hat{I} : \ell(\vec{x}) = i\}$.
3. Define the set $R := \bigcup_{i=1}^m R_i$ and compute the \mathcal{A} -closure \hat{R} of R .
4. For each $i \in [m]$, define Z'_i to be the natural embedding of Z_i into the space of matrices $\mathbb{F}^{k_i \times \hat{R}_i \sqcup \hat{I}}$ (by appropriate padding with $|\hat{R}_i| - |R_i|$ columns of zeros), where k_i is the number of rows of Z_i .
5. Construct the matrix

$$Z' := \begin{pmatrix} Z'_1 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & Z'_2 & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & Z'_m \end{pmatrix}.$$

6. Construct the matrix Z^* whose $(\vec{s}, \vec{\gamma})$ -th entry, for $\vec{s} \in (\hat{R} \sqcup \hat{I})^*$, and $\vec{\gamma} \in \hat{R} \sqcup \hat{I}$, is defined by

$$z_{\vec{s}}(\vec{\gamma}) = \begin{cases} 1 & \text{if } \vec{\gamma} = \vec{s}, \\ -1 & \text{if } \vec{\gamma} = (\vec{s}, a) \text{ for some } a \in A_{\ell(\vec{\gamma})+1}, \\ 0 & \text{otherwise.} \end{cases}$$

7. Define Z to be the matrix obtained by vertically stacking Z' on top of Z^* .
8. Compute a basis B for the space $\{(\vec{r}, \vec{\beta}) \in \mathbb{F}^{\hat{R}} \times \mathbb{F}^{\hat{I}} : \exists \vec{\gamma} \in \mathbb{F}^{\hat{I} \setminus I}, Z(\vec{r}, \vec{\beta}, \vec{\alpha})^T = 0\}$.
9. Output (\hat{R}, B) .

Our proof of correctness for Construction 6.9 consists of two main steps. The first is showing that $\ker(B) = \Sigma\text{RM}|_{\hat{R} \sqcup \hat{I}}$, which we accomplish by appealing to Theorem 6.5. We then show that $(\vec{r}, \vec{\beta}) \in \Sigma\text{RM}|_{\hat{R} \sqcup \hat{I}}$ if and only if for all messages $m \in \mathcal{M}$ such that $m|_{\hat{R}} = \vec{r}$, we have $(m, \vec{\beta}) \in \Sigma\text{RM}|_{\hat{\mathcal{A}} \sqcup \hat{I}}$.

Claim 6.10. Let $I \subseteq \mathbb{F}^m$ and let $(\hat{R}, B) := \text{CL}_\Sigma(I)$ (cf. Construction 6.9), then $\ker(B) = \Sigma\text{RM}[\mathbb{F}, m, \vec{d}]|_{\hat{R} \sqcup \hat{I}}$.

Proof. It suffices to show that $\ker(Z) = \Sigma\text{RM}|_{\hat{R} \sqcup \hat{I}}$ (for Z as in Construction 6.9). This is because $(\vec{r}, \vec{\beta}) \in \ker(B)$ if and only if there exists $\vec{\gamma} \in \mathbb{F}^{\hat{I} \setminus I}$ such that $(\vec{r}, \vec{\beta}, \vec{\gamma}) \in \ker(Z)$. So $\ker(B) = \ker(Z)|_{\hat{R} \sqcup \hat{I}}$. If $\ker(Z) = \Sigma\text{RM}|_{\hat{R} \sqcup \hat{I}}$, then upon restricting both sides of the equivalence to $\hat{R} \sqcup \hat{I}$, we have that $\ker(B) = \Sigma\text{RM}|_{\hat{R} \sqcup \hat{I}}$, as required.

To see that $\ker(Z) = \Sigma\text{RM}|_{\hat{R} \sqcup \hat{I}}$, consider the following. By construction, the rows of Z consist of trivial constraints over $(\hat{I} \sqcup \hat{R})^*$ and Reed-Muller constraints on the set $R_i \sqcup \hat{I}$ padded with zeros so as to be supported on $\hat{R} \sqcup \hat{I}$. That is,

$$\text{rowspace}(Z) = \text{span} \left(\{z_{\vec{s}}\}_{\vec{s} \in (\hat{R} \sqcup \hat{I})^*} \cup \bigcup_{i=1}^m (\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{R_i \sqcup \hat{I}_i})^\perp \times \{0\}^{\hat{R}_i \setminus R_i} \right).$$

By the correctness of each CL_i , we know that for all $i \in [m]$, for all constraints $z \in (\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{\mathcal{A}_i \sqcup \hat{I}})^\perp$, $z(\vec{\alpha}) = 0$ for all $\vec{\alpha} \in \mathcal{A}_i \setminus R_i$. Therefore as $R_i \subseteq \hat{R}_i \subseteq \mathcal{A}_i$, we have that $(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{\hat{R}_i \sqcup \hat{I}})^\perp = (\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{R_i \sqcup \hat{I}})^\perp \times \{0\}^{\hat{R}_i \setminus R_i}$. Hence,

$$\text{rowspace}(Z) = \text{span} \left(\{z_{\vec{s}}\}_{\vec{s} \in (\hat{R} \sqcup \hat{I})^*} \cup \bigcup_{i=1}^m (\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{\hat{R}_i \sqcup \hat{I}})^\perp \right) = (\Sigma \text{RM}|_{\hat{R} \sqcup \hat{I}})^\perp,$$

where the final equality follows from Theorem 6.5 and the fact that $\hat{R} \sqcup \hat{I}$ is \mathcal{A} -closed. In particular, we have shown that Z is a parity check matrix for $\Sigma \text{RM}|_{\hat{R} \sqcup \hat{I}}$. In other words, $\ker(Z) = \Sigma \text{RM}|_{\hat{R} \sqcup \hat{I}}$. \square

Now we focus on the second part of our proof of correctness: showing that $(\vec{r}, \vec{\beta}) \in \Sigma \text{RM}|_{\hat{R} \sqcup \hat{I}}$ if and only if for all messages $m \in \mathcal{M}$ such that $m|_{\hat{R}} = \vec{r}$, we have $(m, \vec{\beta}) \in \Sigma \text{RM}|_{\mathcal{A} \sqcup \hat{I}}$. We accomplish this via the following property of linear codes, whose proof is deferred to Appendix A.5.

Claim 6.11. *Let $C \subseteq \mathbb{F}^D$ be a linear code, and let $U \subseteq V \subseteq D$, $I \subseteq D$ be such that $\mathbf{Z}_U(C)|_I = \mathbf{Z}_V(C)|_I$. Then $(u, x) \in C|_{U \sqcup I}$ if and only if for all $v \in C|_V$ such that $v|_U = u$, $(v, x) \in C|_{V \sqcup I}$.*

We prove that $\mathbf{Z}_{\mathcal{A}}(\Sigma \text{RM})|_{\hat{I}} = \mathbf{Z}_{\hat{R}}(\Sigma \text{RM})|_{\hat{I}}$ where \hat{I}, \hat{R} are as above. To do this, we first show an analogous statement for “plain” RM, and then lift that result to ΣRM via Theorem 6.5.

Claim 6.12. *Let $\mathcal{A} = A_1 \times \dots \times A_m \subseteq \mathbb{F}^m$ and let $S \subseteq \mathbb{F}^m$. It holds that $\mathbf{Z}_{\mathcal{A}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_S = \mathbf{Z}_R(\text{RM}[\mathbb{F}, m, \vec{d}])|_S$, where $(R, Z) := \text{CL}(S)$, for an ℓ -constraint locator CL for $\text{ENC}_{\text{RM}(\vec{d}, \mathcal{A})}$.*

Proof. As $R \subseteq \mathcal{A}$, it is clear that $\mathbf{Z}_{\mathcal{A}}(\text{RM}[\mathbb{F}, m, \vec{d}])|_S \subseteq \mathbf{Z}_R(\text{RM}[\mathbb{F}, m, \vec{d}])|_S$. We now argue the reverse containment. For brevity we write RM for $\text{RM}[\mathbb{F}, m, \vec{d}]$.

Let $w \in \text{RM}$ such that $w(r) = 0$ for all $r \in R$. Then $w|_S$ is a general element of $\mathbf{Z}_R(\text{RM})|_S$. We will show that there exists $w' \in \text{RM}$ which agrees with w on S , but is zero on \mathcal{A} , i.e., $w'|_S = w|_S$ and $w'|_{\mathcal{A}} = \vec{0}$.

By the correctness of CL, all constraints on $\mathcal{A} \sqcup S$ with respect to RM are supported only on $R \sqcup S$; in particular, they have no support on the set $\mathcal{A} \setminus R$. For $\alpha \in \mathcal{A} \setminus R$, define $w_\alpha: \mathcal{A} \sqcup S \rightarrow \mathbb{F}$ by $w_\alpha(\alpha) := 1$ and $w_\alpha(\beta) := 0$ for all $\beta \in \mathcal{A} \sqcup S \setminus \{\alpha\}$. Then since w_α is only nonzero in $\mathcal{A} \setminus R$, $w_\alpha \in \text{RM}|_{\mathcal{A} \sqcup S}$, and so $w_\alpha = w'_\alpha|_{\mathcal{A} \sqcup S}$ for $w'_\alpha \in \text{RM}$. Consider the codeword defined by

$$w' := w - \sum_{\alpha \in \mathcal{A} \setminus R} w(\alpha) w'_\alpha.$$

This choice of w' satisfies $w' \in \mathbf{Z}_{\mathcal{A}}(\text{RM})$, and $w'|_S = w|_S$. Thus $w|_S \in \mathbf{Z}_{\mathcal{A}}(\text{RM})|_S$. \square

Corollary 6.13. *For \hat{I}, R as in Construction 6.9, $(\mathbf{Z}_{\mathcal{A}}(\Sigma \text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_{\hat{I}})^\perp = (\mathbf{Z}_{\hat{R}}(\Sigma \text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_{\hat{I}})^\perp$.*

Proof. It is clear that $(\mathbf{Z}_{\hat{R}}(\Sigma \text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_{\hat{I}})^\perp \subseteq (\mathbf{Z}_{\mathcal{A}}(\Sigma \text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}])|_{\hat{I}})^\perp$, so we show the reverse inclusion.

For R_i, Z_i, \hat{I}_i as in Construction 6.9, we have that

$$\begin{aligned}
(\mathbf{Z}_{\mathcal{A}}(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}]|_{\hat{I}})^\perp &= \text{span} \left(\{z_{\vec{s}}\}_{\vec{s} \in \hat{I}^*} \cup \bigcup_{i=1}^m (\mathbf{Z}_{\mathcal{A}_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{\hat{I}_i})^\perp) \right) \\
&= \text{span} \left(\{z_{\vec{s}}\}_{\vec{s} \in \hat{I}^*} \cup \bigcup_{i=1}^m (\mathbf{Z}_{R_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{\hat{I}_i})^\perp) \right) \\
&\subseteq \text{span} \left(\{z_{\vec{s}}\}_{\vec{s} \in \hat{I}^*} \cup \bigcup_{i=1}^m (\mathbf{Z}_{\hat{R}_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{\hat{I}_i})^\perp) \right) \\
&\subseteq (\mathbf{Z}_{\hat{R}}(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}]|_{\hat{I}})^\perp.
\end{aligned}$$

The first equality is a consequence of Theorem 6.5. The second equality is a consequence of Claim 6.12. The penultimate inclusion is true as $R_i \subseteq \hat{R}_i$ for each $i \in [m]$. The final inclusion follows from the fact that trivially $(\mathbf{Z}_{\hat{R}_i}(\text{RM}[\mathbb{F}, i, (d_1, \dots, d_i)]|_{\hat{I}_i})^\perp \subseteq (\mathbf{Z}_{\hat{R}}(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}]|_{\hat{I}})^\perp$ for each $i \in [m]$ and as \hat{R} is \mathcal{A} -closed, we also have that $\{z_{\vec{s}}\}_{\vec{s} \in \hat{I}^*} \subseteq (\mathbf{Z}_{\hat{R}}(\Sigma\text{RM}[\mathbb{F}, \mathcal{A}, \vec{d}]|_{\hat{I}})^\perp$. \square

Proof of Theorem 6.2. Correctness. By Claim 6.10, we have that $\ker(B) = \Sigma\text{RM}|_{\hat{R} \sqcup I}$. Then by Claim 6.11, setting $U := \hat{R}$ and $V := \bar{\mathcal{A}}$, and Corollary 6.13, noting that $\text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})}$ is a systematic encoding (that is $\mathcal{M} = \Sigma\text{RM}|_{\bar{\mathcal{A}}}$), it follows that $(\vec{r}, \vec{\beta}) \in \Sigma\text{RM}|_{\hat{R} \sqcup \hat{I}}$ if and only if for all messages $m \in \mathcal{M}$ such that $m|_{\hat{R}} = \vec{r}$, we have $(m, \vec{\beta}) \in \Sigma\text{RM}|_{\bar{\mathcal{A}} \sqcup \hat{I}}$. Thus, taking the restriction of both sides, $(\vec{r}, \vec{\beta}|_I) \in \Sigma\text{RM}|_{\hat{R} \sqcup I}$ if and only if for all messages $m \in \mathcal{M}$ such that $m|_{\hat{R}} = \vec{r}$, we have $(m, \vec{\beta}|_I) \in \Sigma\text{RM}|_{\bar{\mathcal{A}} \sqcup I}$. This is sufficient since, by definition of the encoding, $\vec{\beta} \in \text{supp}(\text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})}(m)|_I)$ if and only if $(m, \vec{\beta}) \in \Sigma\text{RM}|_{\bar{\mathcal{A}} \sqcup I}$.

Efficiency. By Claim 6.14 (below), $|\hat{R}| \leq (\sum_{i=1}^m |A_i| + m + 1)|R|$ and $|\hat{I}| \leq (\sum_{i=1}^m |A_i| + m + 1)|I|$. By Theorem 5.4, $|R_i| \leq |\hat{I}| \leq (\sum_{i=1}^m |A_i| + m + 1)|I|$. As $|R| \leq \sum_{i=1}^m |R_i| \leq m(\sum_{i=1}^m |A_i| + m + 1)|I|$, we have $|\hat{R}| \leq m(\sum_{i=1}^m |A_i| + m + 1)^2|I|$.

The runtime follows from Claim 6.14 and Theorem 5.2, as CL_i runs in time $\text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$. \square

Claim 6.14. *Let $I \subseteq \mathbb{F}^{\leq m}$ and $\mathcal{A} = A_1 \times \dots \times A_m$. Then there exists a set $\hat{I} \subseteq \mathbb{F}^{\leq m}$ such that $I \subseteq \hat{I}$ and \hat{I} is \mathcal{A} -closed, satisfying $|\hat{I}| \leq (\sum_{i=1}^m |A_i| + m + 1)|I|$. Moreover, \hat{I} can be computed in time $\text{poly}(m, |I|, |A_1|, \dots, |A_m|)$.*

Proof. We construct \hat{I} as follows. Let $I_i := \{x \in I : |x| = i\}$. Initialise $\hat{I} := I$. For each $(x_1, \dots, x_m) \in I_m$, we add the point (x_1, \dots, x_{m-1}) and for each $a_m \in A_m$ we add the point $(x_1, \dots, x_{m-1}, a_m)$ to \hat{I} . Then repeat this process for $I_{m-1}, I_{m-2}, \dots, I_1$. By construction, \hat{I} is \mathcal{A} -closed.

Now we argue about the size of \hat{I} . Note that for each point in I at most $\sum_{i=1}^m |A_i| + m$ points are added to \hat{I} . Thus it holds that $|\hat{I}| \leq (\sum_{i=1}^m |A_i| + m + 1)|I|$. The efficiency follows from this bound. \square

7 Constraint location for $\Sigma\text{AntiSym}$

In this section we define the code AntiSym of antisymmetric functions, and provide a constraint locator for an encoding function related to its sum code $\Sigma\text{AntiSym}$.

Definition 7.1. Let $\mathcal{A} := A_1 \times \cdots \times A_m$ be such that $A_i = A_{m-i+1}$ for all $1 \leq i \leq m$. We define

$$\text{AntiSym}[\mathcal{A}] := \{w \in \mathbb{F}^{\mathcal{A}} : \forall \vec{x} \in \mathcal{A}, w(\vec{x}) + w(\vec{x}_{\text{rev}}) = 0\}.$$

The code $\Sigma\text{AntiSym}$ is the sum code of AntiSym (see Definition 3.6).

The main theorem of this section is the following, which gives an efficient constraint locator for $\Sigma\text{AntiSym}$ with polynomial locality.

Theorem 7.2. Let \mathcal{A} be as above. Let $\mathcal{M} := \{F: \mathcal{A} \cup \{\perp\} \rightarrow \mathbb{F} \mid F(\perp) = \sum_{\vec{a} \in \mathcal{A}} F(\vec{a})\}$. Let $\text{ENC}_{\Sigma\text{AntiSym}}: \mathcal{M} \rightarrow \mathbb{F}^{\bar{\mathcal{A}}}$ be the following randomised encoding:

$$\text{ENC}_{\Sigma\text{AntiSym}}(F) := \Sigma[F|_{\mathcal{A}} + G],$$

where G is a codeword sampled uniformly at random from $\text{AntiSym}[\mathcal{A}]$.

$\text{ENC}_{\Sigma\text{AntiSym}}$ has an ℓ -constraint locator for $\ell(n) = n^2 m^2 + 1 + O(n^4 m^4 / |\mathcal{A}|)$, running in time $\text{poly}(a, m, n)$, where $a := \sum_i |A_i|$.

In Section 7.1 we analyse the combinatorial structure of $\Sigma\text{AntiSym}|_I^\perp$, obtaining a bound on the ‘‘complexity’’ of constraints in terms of the size of I . In Section 7.2, we show how this bound leads to an efficient constraint locator for $\text{ENC}_{\Sigma\text{AntiSym}}$.

7.1 Properties of $\Sigma\text{AntiSym}$

We start by introducing some useful notation.

Definition 7.3. We say $\vec{a} \in \bar{\mathcal{A}}$ is a prefix of $\vec{b} \in \bar{\mathcal{A}}$ if $\ell(\vec{a}) \leq \ell(\vec{b})$ and $(a_1, \dots, a_{\ell(\vec{a})}) = (b_1, \dots, b_{\ell(\vec{a})})$. A subset $J \subseteq \bar{\mathcal{A}}$ is **prefix-free** if for all $\vec{a}, \vec{b} \in J$, if \vec{a} is a prefix of \vec{b} then $\vec{a} = \vec{b}$.

We associate with an element $\vec{a} \in \bar{\mathcal{A}}$ the set $\vec{a} \times \mathcal{A}_{>\ell(\vec{a})}$; we denote this set also by \vec{a} . In particular, $|\vec{a}| = |\mathcal{A}_{>\ell(\vec{a})}| = \prod_{i=\ell(\vec{a})+1}^m |A_i|$. (See Definition 3.6.)

It is straightforward to see that $\vec{b} \subseteq \vec{a}$ if and only if \vec{a} is a prefix of \vec{b} . The following simple proposition gives a useful alternative characterisation of AntiSym .

Proposition 7.4. $w \in \text{AntiSym}[\mathcal{A}]$ if and only if there exists $f \in \mathbb{F}^{\mathcal{A}}$ such that $w(\vec{x}) = f(\vec{x}) - f(\vec{x}_{\text{rev}})$.

Proof. For the ‘‘if’’ direction, note that $w(\vec{x}) + w(\vec{x}_{\text{rev}}) = f(\vec{x}) - f(\vec{x}_{\text{rev}}) + f(\vec{x}_{\text{rev}}) - f(\vec{x}) = 0$ for all $\vec{x} \in \mathcal{A}$. For the ‘‘only if’’ direction, let $w \in \text{AntiSym}[\mathcal{A}]$, and consider some strict total ordering $<$ on \mathcal{A} . Define

$$f(\vec{x}) := \begin{cases} w(\vec{x}) & \text{if } \vec{x} < \vec{x}_{\text{rev}}, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Then for all \vec{x} ,

$$f(\vec{x}) - f(\vec{x}_{\text{rev}}) = \begin{cases} w(\vec{x}) & \text{if } \vec{x} < \vec{x}_{\text{rev}}, \text{ and} \\ -w(\vec{x}_{\text{rev}}) & \text{if } \vec{x}_{\text{rev}} < \vec{x} \end{cases} = w(\vec{x}). \quad \square$$

Definition 7.5. For $\vec{a} = (a_1, \dots, a_\ell) \in \bar{\mathcal{A}}$, we define the **reverse set** $\vec{a}_{\text{rev}} := \bar{\mathcal{A}}_{m-\ell} \times (a_\ell, \dots, a_1)$. For a set $H \subseteq \bar{\mathcal{A}}$, we define $H_{\text{rev}} := \{\vec{a}_{\text{rev}} : \vec{a} \in H\}$.

We say that H is **symmetric** if $\cup H = \cup H_{\text{rev}}$, where $\cup H := \cup_{\vec{h} \in H} \vec{h}$. We say that H is **minimal symmetric** if H is symmetric and for all nonempty $H' \subsetneq H$, H' is not symmetric.

Proposition 7.6. *If $H \subseteq \bar{\mathcal{A}}$ is symmetric and prefix-free, then $H = \cup_{i=1}^k H_i$ for $H_1, \dots, H_k \subseteq H$ where each H_i is minimal symmetric and $H_i \cap H_j = \emptyset$ for $i \neq j$.*

Proof. We proceed by induction on the size of H . Indeed, suppose that the statement is true for all sets $H' \subseteq \bar{\mathcal{A}}$ of size less than n , and let $H \subseteq \bar{\mathcal{A}}$ be a set of size n . If H is minimal symmetric, then the statement holds trivially, so suppose not. Then there is a nonempty symmetric $H' \subsetneq H$. Since H is prefix-free, $\cup(H \setminus H') = (\cup H) \setminus (\cup H') = (\cup H_{\text{rev}}) \setminus (\cup H'_{\text{rev}}) = \cup(H \setminus H')_{\text{rev}}$; hence $H \setminus H'$ is also symmetric. Since H' and $H \setminus H'$ are strictly smaller than H , we can apply the inductive hypothesis to complete the proof. \square

The following lemma shows that, for G prefix-free, $\Sigma\text{AntiSym}|_G^\perp$ has a basis of constraints of the form “ $\sum_{\vec{h} \in H} w(\vec{h}) = 0$ ”, taken over all minimal symmetric subsets $H \subseteq G$.

Lemma 7.7. *Let $G \subseteq \bar{\mathcal{A}}$ be prefix-free. For $H \subseteq G$, let $1_H: G \rightarrow \mathbb{F}$ be the indicator vector for H , i.e., $1_H(\vec{h}) = 1$ if $\vec{h} \in H$ and 0 otherwise. A basis for $\Sigma\text{AntiSym}[\mathcal{A}]|_G^\perp$ is $\{1_H\}_{H \in \mathcal{H}}$ where \mathcal{H} is the set of all minimal symmetric subsets of G .*

Proof. Let $B := \{1_H\}_{H \in \mathcal{H}}$. First, note that Proposition 7.6 implies that for all $H \neq H' \in \mathcal{H}$, $H \cap H' = \emptyset$. Hence B is a linearly independent set. It is also straightforward to see that $B \subseteq \Sigma\text{AntiSym}[\mathcal{A}]|_G^\perp$. It remains to show that $\text{span}(B) = \Sigma\text{AntiSym}[\mathcal{A}]|_G^\perp$.

Suppose that $z \in \Sigma\text{AntiSym}[\mathcal{A}]|_G^\perp$; equivalently, by Proposition 7.4, for all $f: \mathcal{A} \rightarrow \mathbb{F}$,

$$\sum_{\vec{a} \in G} z(\vec{a}) \sum_{\vec{x} \in \vec{a}} f(\vec{x}) = \sum_{\vec{a} \in G} z(\vec{a}) \sum_{\vec{x} \in \vec{a}} f(\vec{x}_{\text{rev}}).$$

Let $S_\gamma := \{\vec{a} \in G : z(\vec{a}) = \gamma\}$. Then we can write

$$\sum_{\gamma \in \mathbb{F}} \gamma \sum_{\vec{a} \in S_\gamma} \sum_{\vec{x} \in \vec{a}} f(\vec{x}) = \sum_{\gamma \in \mathbb{F}} \gamma \sum_{\vec{a} \in S_\gamma} \sum_{\vec{x} \in \vec{a}} f(\vec{x}_{\text{rev}}).$$

Fix $\gamma^* \in \mathbb{F}$, $\vec{a}^* \in S_{\gamma^*}$, and let $\vec{x}^* \in \vec{a}^*$. Let $f = 1_{\vec{x}^*}$ be the indicator function for \vec{x}^* . It follows that

$$\gamma^* = \sum_{\gamma \in \mathbb{F}} \gamma \sum_{\vec{a} \in S_\gamma} \sum_{\vec{x} \in \vec{a}} 1_{\vec{x}^*}(\vec{x}_{\text{rev}}) = \gamma'$$

where $\gamma' \in \mathbb{F}$ is unique such that $\vec{x}_{\text{rev}}^* \in \cup S_{\gamma'}$. Thus $\vec{x}_{\text{rev}}^* \in \cup S_{\gamma^*}$; equivalently, $\vec{x}^* \in \cup (S_{\gamma^*})_{\text{rev}}$. It follows that for all γ , $\cup S_\gamma \subseteq \cup (S_\gamma)_{\text{rev}}$; a similar argument establishes that in fact S_γ is symmetric. By Proposition 7.6, S_γ is a union of minimal symmetric sets. It follows that $z \in \text{span}(B)$. \square

Next, we prove the main technical lemma of this section.

Lemma 7.8. *Let $H, G \subseteq \bar{\mathcal{A}}$ be prefix-free with $|H| \cdot |G| \leq |\mathcal{A}|/4$. Suppose that $\cup H = \cup G_{\text{rev}}$. Then*

$$|\cup H| \geq \frac{K}{2}(1 + \sqrt{1 - 4|H| \cdot |G|/K}) \quad \text{or} \quad |\cup H| \leq \frac{K}{2}(1 - \sqrt{1 - 4|H| \cdot |G|/K}).$$

In particular,

$$\min(|\cup H|, |\mathcal{A}| - |\cup H|) \leq |H| \cdot |G| + O(|H|^2|G|^2/|\mathcal{A}|).$$

Taking $G = H$, we obtain strong bounds on the support size of symmetric sets: if H is symmetric then $\cup H$ must be of size either $\text{poly}(|H|)$ or $|\mathcal{A}| - \text{poly}(|H|)$.

Corollary 7.9. *If $H \subseteq \bar{\mathcal{A}}$ is symmetric and prefix-free, then*

$$\min(|\cup H|, |\mathcal{A}| - |\cup H|) \leq |H|^2 + O(|H|^4/|\mathcal{A}|).$$

Before proving Lemma 7.8, we state and prove a simple but crucial claim about the size of the intersection of sets \vec{a} and \vec{b}_{rev} for general $\vec{a}, \vec{b} \in \bar{\mathcal{A}}$.

Claim 7.10. *For any $\vec{a}, \vec{b} \in \bar{\mathcal{A}}$,*

$$|\vec{a} \cap \vec{b}_{\text{rev}}| \leq \begin{cases} \prod_{j=\ell(\vec{a})+1}^{m-\ell(\vec{b})} |A_j| & \text{if } \ell(\vec{a}) + \ell(\vec{b}) < m \\ 1 & \text{otherwise.} \end{cases}$$

Proof. Recall that, if $\ell(\vec{a}) + \ell(\vec{b}) < m$,

$$\begin{aligned} \vec{a} &= \{a_1\} \times \cdots \times \{a_{\ell(\vec{a})}\} \times A_{\ell(\vec{a})+1} \times \cdots \times A_{m-\ell(\vec{b})} \times \cdots \times A_m \\ \vec{b}_{\text{rev}} &= A_1 \times \cdots \times A_{\ell(\vec{a})} \times \cdots \times A_{m-\ell(\vec{b})} \times \{b_{\ell(\vec{b})}\} \times \cdots \times \{b_1\} \end{aligned}$$

and so

$$\vec{a} \cap \vec{b}_{\text{rev}} = \{a_1\} \times \cdots \times \{a_{\ell(\vec{a})}\} \times A_{\ell(\vec{a})+1} \times \cdots \times A_{m-\ell(\vec{b})} \times \{b_{\ell(\vec{b})}\} \times \cdots \times \{b_1\}.$$

On the other hand, if $\ell(\vec{a}) + \ell(\vec{b}) \geq m$ then

$$\vec{a} \cap \vec{b}_{\text{rev}} = \begin{cases} \{(a_1, \dots, a_{\ell(\vec{a})}, b_{m-\ell(\vec{a})}, \dots, b_1)\} & \text{if } a_i = b_{m-i+1} \text{ for all } m+1-\ell(\vec{b}) \leq i \leq \ell(\vec{a}); \\ \emptyset & \text{otherwise.} \end{cases} \quad \square$$

Proof of Lemma 7.8. Since $\cup H = \cup G_{\text{rev}}$, $(\cup H) \cap (\cup G_{\text{rev}}) = \cup H = \cup G_{\text{rev}}$. Since H, G are prefix-free,

$$|\cup H| = |(\cup H) \cap (\cup G_{\text{rev}})| = \sum_{\vec{a} \in H} \sum_{\vec{b} \in G} |\vec{a} \cap \vec{b}_{\text{rev}}| \leq |H| \cdot |G| + \sum_{\substack{\vec{a} \in H, \vec{b} \in G \\ \ell(\vec{a}) + \ell(\vec{b}) < m}} \prod_{j=\ell(\vec{a})+1}^{m-\ell(\vec{b})} |A_j| =: |H| \cdot |G| + N.$$

Now observe that

$$\begin{aligned} |\cup H|^2 &= |\cup H| \cdot |\cup G_{\text{rev}}| = \left(\sum_{\vec{a} \in H} \prod_{j=\ell(\vec{a})+1}^n |A_j| \right) \cdot \left(\sum_{\vec{b} \in G} \prod_{j=1}^{n-\ell(\vec{b})} |A_j| \right) \\ &= \sum_{\vec{a} \in H} \sum_{\vec{b} \in G} \prod_{j=\ell(\vec{a})+1}^n |A_j| \cdot \prod_{j=1}^{m-\ell(\vec{b})} |A_j| \\ &\geq |\mathcal{A}| \cdot \sum_{\substack{\vec{a} \in H, \vec{b} \in G \\ \ell(\vec{a}) + \ell(\vec{b}) < m}} \prod_{j=\ell(\vec{a})+1}^{m-\ell(\vec{b})} |A_j| = |\mathcal{A}| \cdot N. \end{aligned}$$

Hence, writing $t := |H| \cdot |G|$, $|\mathcal{A}| \cdot N \leq (t + N)^2$. Completing the square to solve for N , we obtain $(N + t - |\mathcal{A}|/2)^2 \geq \frac{|\mathcal{A}|^2}{4}(1 - 4t/|\mathcal{A}|)$. Hence either

$$N \geq \frac{|\mathcal{A}|}{2}(1 + \sqrt{1 - 4t/|\mathcal{A}|}) - t \quad \text{or} \quad N \leq \frac{|\mathcal{A}|}{2}(1 - \sqrt{1 - 4t/|\mathcal{A}|}) - t$$

The first part of the lemma follows by recalling that $|\cup H| = N + t$. The second part then follows since $\sqrt{1 - 4t/|\mathcal{A}|} = 1 - 2t/|\mathcal{A}| - O(t^2/|\mathcal{A}|^2)$. \square

7.2 Algorithms for $\Sigma\text{AntiSym}$

In this section, we present two useful algorithms for working with $\Sigma\text{AntiSym}$. The first (Lemma 7.11) takes as input an arbitrary set $I \subseteq \bar{\mathcal{A}}$ and outputs a prefix-free set $G \subseteq \bar{\mathcal{A}}$ of a similar size with the property that any $\vec{a} \in I$ can be obtained as the union of sets in G . The second (Lemma 7.13) takes as input a prefix-free set G and outputs the set of minimal symmetric subsets of G (i.e., a basis for $\Sigma\text{AntiSym}[\mathcal{A}]|_G^\perp$).

Lemma 7.11. *There is a polynomial-time algorithm PrefixFree which, given as input a set $I \subseteq \bar{\mathcal{A}}$, outputs a prefix-free set G of size at most $|I| \cdot m$ and a list $(\Lambda_{\vec{a}} \subseteq G)_{\vec{a} \in I}$ such that for each $\vec{a} \in I$, $\cup \Lambda_{\vec{a}} = \vec{a}$.*

Proof. For $\vec{a} \in \bar{\mathcal{A}}$ and $G \subseteq \bar{\mathcal{A}}$, define $N_G(\vec{a}) := \{\vec{b} \in G : \vec{b} \subsetneq \vec{a}\}$. The algorithm PrefixFree operates as follows.

1. For each $\vec{a} \in I$, set $\Lambda_{\vec{a}} := \{\vec{a}\}$. Set $G := I$.
2. Let $\vec{a}^* = (a_1, \dots, a_{\ell^*})$ be a shortest element in G (i.e., ℓ^* is minimal) for which $N_G(\vec{a}^*)$ is nonempty. If there is no such term, output G .
3. Let $\vec{a}' = (a_1, \dots, a_{\ell'})$ be a shortest element in $N_G(\vec{a}^*)$. (Note $\ell' > \ell^*$.)
4. Remove \vec{a}^* from G .
5. For each $j \in \{\ell^* + 1, \dots, \ell'\}$, and for each $b_j \in A_j \setminus \{a_j\}$, add $\vec{\alpha}_j := (a_1, \dots, a_{j-1}, b_j)$ to G .
6. For each $\vec{a} \in I$ such that $\vec{a}^* \in \Lambda_{\vec{a}}$, remove \vec{a}^* from $\Lambda_{\vec{a}}$ and add $\{\vec{\alpha}_{\ell^*+1}, \dots, \vec{\alpha}_{\ell'}, \vec{a}'\}$.
7. Go to Step 2.

Clearly if the algorithm terminates then the correctness condition is satisfied, and so it remains to bound the number of iterations. Denote by Δ_i the value of $|\cup_{\vec{a} \in G} N_G(\vec{a})|$ at the beginning of the i -th iteration. Clearly Δ_1 cannot be larger than $|I|$, and if $\Delta_t = 0$ then the algorithm terminates at the beginning of the t -th iteration. We show that Δ is a progress measure for the above algorithm.

Claim 7.12. *If $\Delta_i > 0$ then $\Delta_{i+1} < \Delta_i$.*

Proof. Let \vec{a}' be the term chosen in Step 3 of the i -th iteration. We show that at the termination of the i -th iteration, \vec{a}' has been removed from $\cup_{\vec{a} \in G} N_G(\vec{a})$, and no element has been added. By choice of \vec{a}' , $\vec{a}' \in N_G(\vec{a}^*)$, and $\vec{a}' \notin N_G(\vec{a}'')$ for any $\vec{a}'' \neq \vec{a}^*$, since that would mean $\vec{a}' \subset \vec{a}'' \subset \vec{a}^*$. Moreover, for every element $\vec{\alpha}_j$ added to G in Step 5, $\vec{a}' \notin N_G(\vec{\alpha}_j)$ since $\vec{\alpha}_j$ is disjoint from \vec{a}' by construction, $N_G(\vec{\alpha}_j) \subset N_G(\vec{a})$ since $\vec{\alpha}_j \subset \vec{a}$, and $\vec{\alpha}_j \notin \cup_{\vec{a} \in G} N_G(\vec{a})$ by choice of \vec{a}^*, \vec{a}' . \square

It follows that the number of iterations is at most $|I|$; since each iteration clearly runs in polynomial time, the algorithm runs in polynomial time. The bound on $|G|$ is obtained by noting that, in each iteration, the size of G increases by at most $m - 1$. \square

Lemma 7.13. *There is a polynomial-time algorithm SymSets which, given as input a prefix-free set $G \subseteq \bar{\mathcal{A}}$, outputs the set \mathcal{H} of all minimal symmetric subsets H of G .*

Proof. The algorithm SymSets operates as follows:

1. Construct a graph $\Gamma = (G, E)$ where $(\vec{a}, \vec{b}) \in E$ if and only if $\vec{a} \cap \vec{b}_{\text{rev}} \neq \emptyset$.

2. Compute $C(\Gamma)$, the set of connected components of Γ .
3. Output the set $\{H \in C(\Gamma) : \cup H = \cup H_{\text{rev}}\}$.

It is straightforward to construct Γ and compute $C(\Gamma)$ in polynomial time. To determine whether $\cup H = \cup H_{\text{rev}}$, it suffices to compute $K := |(\cup H) \cap (\cup H_{\text{rev}})|$ and then check whether $|\cup H| = |\cup H_{\text{rev}}| = K$; this can also be achieved in polynomial time since H is prefix-free. The correctness of the algorithm is a consequence of the following claim.

Claim 7.14. *A set H is minimal symmetric if and only if $H \in C(\Gamma)$ and $\cup H = \cup H_{\text{rev}}$.*

Proof. First, suppose that H is minimal symmetric. Then it remains to show that H is a connected component in Γ . First, suppose that there is some edge $(\vec{a}, \vec{b}) \in E$ with $a \in H$ and $b \notin H$. Since J is prefix-free, $\vec{b} \cap (\cup H) = \emptyset$. However, since $\vec{a}_{\text{rev}} \in H_{\text{rev}}$ and $\vec{a}_{\text{rev}} \cap \vec{b} \neq \emptyset$, we have that $\vec{b} \cap (\cup H_{\text{rev}}) \neq \emptyset$, which is a contradiction since $\cup H = \cup H_{\text{rev}}$.

Hence H must be a disjoint union of k connected components H_1, \dots, H_k . Now consider two cases: either $\cup H_1 = \cup (H_1)_{\text{rev}}$, or not. If $\cup H_1 \neq \cup (H_1)_{\text{rev}}$, take $\vec{x} \in \cup H_1 \setminus \cup (H_1)_{\text{rev}}$ (so $\vec{x} \in \mathcal{A}$). Then $\vec{x}_{\text{rev}} \notin \cup H_i$ for all $i \in \{2, \dots, k\}$, otherwise there would be an edge from H_1 to H_i . Hence $\vec{x} \notin \cup H_{\text{rev}}$; but then $\cup H \neq \cup H_{\text{rev}}$, which contradicts that H is symmetric. Hence $\cup H_1 = \cup (H_1)_{\text{rev}}$; so since H is minimal, $H_1 = H$, and $k = 1$.

Now suppose conversely that H is a connected component in Γ with $\cup H = \cup H_{\text{rev}}$, and let $H' \subseteq H$ be minimal symmetric. By the above, H' is a connected component in Γ contained in H , whence $H' = H$. \square

This completes the proof. \square

7.3 Proof of Theorem 7.2

We construct a constraint locator $\text{CL}_{\Sigma\text{AntiSym}}$ for $\text{ENC}_{\Sigma\text{AntiSym}}$.

Construction 7.15. $\text{CL}_{\Sigma\text{AntiSym}}(I)$:

1. Compute $(G, (\Lambda_{\vec{a}})_{\vec{a} \in I}) := \text{PrefixFree}(I)$.
2. Compute $\mathcal{H} := \text{SymSets}(G)$. Let $\mathcal{H}_0 := \{H \in \mathcal{H} : |\cup H| \leq |\mathcal{A}|/2\}$, and $\mathcal{H}_1 := \{H \in \mathcal{H} : |\cup H| > |\mathcal{A}|/2\}$. Let $R := \{\perp\} \cup \bigcup_{H \in \mathcal{H}_0} (\cup H) \cup \bigcup_{H \in \mathcal{H}_1} (\mathcal{A} \setminus \cup H)$.
3. Let $B := (1_H)_{H \in \mathcal{H}} \in \mathbb{F}^{|\mathcal{H}| \times G}$, $C := ((1_{\cup H})_{H \in \mathcal{H}_0}, (1_{\perp} - 1_{\mathcal{A} \setminus \cup H})_{H \in \mathcal{H}_1}) \in \mathbb{F}^{|\mathcal{H}| \times R}$, $M := (1_{\Lambda_{\vec{a}}})_{\vec{a} \in I} \in \mathbb{F}^{I \times G}$, and $\mathbf{I} \in \mathbb{F}^{I \times I}$ be the identity matrix. Construct the block matrix

$$Y = \begin{pmatrix} C & -B & 0 \\ 0 & M & -\mathbf{I} \end{pmatrix} \in \mathbb{F}^{(|\mathcal{H}| + |I|) \times (R \sqcup G \sqcup I)}.$$

4. Output a basis Z for the space $\{(\vec{\alpha}, \vec{\gamma}) \in \mathbb{F}^R \times \mathbb{F}^I : \exists \vec{\beta}, Y(\vec{\alpha}, \vec{\beta}, \vec{\gamma})^T = 0\}$.

By Corollary 7.9, $|R| \leq 1 + \sum_{H \in \mathcal{H}} |H|^2 + O(|H|^4/|\mathcal{A}|) \leq 1 + |G|^2 + O(|G|^4/|\mathcal{A}|)$. The bound on $\ell(n)$ follows since $|G| \leq m \cdot |I|$. Efficiency follows from the efficiency of PrefixFree and SymSets .

It remains to show that for all $f \in \mathcal{M}$, $\vec{\gamma} \in \mathbb{F}^I$, it holds that $\vec{\gamma} \in \text{supp}(\text{ENC}_{\Sigma\text{AntiSym}}(f)|_I)$ if and only if $(f|_R, \vec{\gamma})^T \in \ker(Z)$. Note that, by Lemma 7.7, $\vec{\beta} \in \text{supp}(\text{ENC}_{\Sigma\text{AntiSym}}(f)|_G)$ if and only if

$B\vec{\beta} = B(\Sigma[f]|_G)$. For all f and all $H \in \mathcal{H}_0$, $1_H \cdot \Sigma[f]|_G = 1_{\cup H} \cdot f|_R$. Similarly, for all $H \in \mathcal{H}_1$, $1_H \cdot \Sigma[f]|_G = (1_{\perp} - 1_{A \cup H}) \cdot f|_R$. Hence $B(\Sigma[f]|_G) = C(f|_R)$.

By the guarantee of Lemma 7.11, we have that $\cup \Lambda_{\vec{a}} = \vec{a}$ for all $\vec{a} \in I$, and so for all functions $\Phi: \mathcal{A} \rightarrow \mathbb{F}$, $M(\Sigma[\Phi]|_G) = \Sigma[\Phi]|_I$. Hence $\vec{\gamma} \in \text{supp}(\text{ENC}_{\Sigma\text{AntiSym}}(f)|_I)$ if and only if there exists $\vec{\beta} \in \text{supp}(\text{ENC}_{\Sigma\text{AntiSym}}(f)|_G)$ such that $\vec{\gamma} = M\vec{\beta}$.

It follows that $\vec{\gamma} \in \text{supp}(\text{ENC}_{\Sigma\text{AntiSym}}(f)|_I)$ if and only if there exists $\vec{\beta} \in \text{supp}(\text{ENC}_{\Sigma\text{AntiSym}}(f)|_G)$ such that $B\vec{\beta} = C(f|_R)$ and $\vec{\gamma} = M\vec{\beta}$. By construction, this is equivalent to the existence of $\vec{\beta} \in \mathbb{F}^G$ such that $(f|_R, \vec{\beta}, \vec{\gamma})^T \in \ker(Y)$, and such a $\vec{\beta}$ exists if and only if $(f|_R, \vec{\gamma})^T \in \ker(Z)$.

8 ZKPCP for #P

In this section we prove our main theorem, that there exists a PZK-PCP for (decision) #SAT. First, we recall the definition:

$$\#\text{SAT} = \{(\Phi, N) : \Phi \text{ is a CNF, } \sum_{x \in \{0,1\}^n} \Phi(x) = N\}.$$

Theorem 8.1. *There exists a PZK-PCP for #SAT. The honest verifier is non-adaptive, and the zero-knowledge property holds against arbitrary (adaptive) polynomial-time malicious verifiers.*

To formally obtain Theorem 1, we must define what we mean by #P as a class of languages (typically, #P is treated as a class of function problems). One natural way is by reference to the number of accepting paths of a Turing machine, i.e.: a pair language $\mathcal{L} \subseteq (\{0,1\}^* \times \mathbb{N})$ is in #P if and only if there is a polynomial-time nondeterministic Turing machine M such that $x \in \mathcal{L}$ if and only if $M(x)$ has N accepting paths. #SAT is complete for #P under this definition.

One can also adopt a more permissive definition of #P as the set of all languages that have a deterministic polynomial-time reduction to #SAT. This permits, for example, the inclusion $\text{coNP} \subseteq \#\text{P}$, via reduction from UNSAT: $\Phi \mapsto (\Phi, 0)$. Theorem 1 follows from Theorem 8.1 also with respect to this broader definition.

Remark 8.2 (Sampling field elements). Our PZK-PCP construction relies on sampling uniformly random field elements. This can be an issue if we demand perfect simulation, as typically PPT algorithms are given a tape of uniformly random *bits*, from which it is not possible to perfectly sample random field elements in strict polynomial time if the field has characteristic different from 2. There are three ways to resolve this: (1) provide the simulator with uniformly random field elements on its random tape; (2) use rejection sampling, so that the simulator runs in expected polynomial time, or; (3) restrict attention to $\oplus P$, which can be verified by sumcheck over fields of characteristic 2.

8.1 Zero-Knowledge Probabilistically Checkable Proofs (ZK-PCPs)

In this section we formally define the notion of a (perfect) zero-knowledge probabilistically-checkable proof.

Definition 8.3 (PCP). *A probabilistically checkable proof (PCP) for a language \mathcal{L} consists of a prover \mathcal{P} and a PPT verifier \mathcal{V} such that the following holds.*

1. **Completeness.** *For every $x \in \mathcal{L}$,*

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

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

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

Definition 8.4 (View). For a PCP $(\mathcal{P}, \mathcal{V})$ and a (possibly malicious) verifier \mathcal{V}^* , $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x)$ denotes the view of \mathcal{V}^* with input x and oracle access to $\pi \leftarrow \mathcal{P}(x)$. That is, $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x)$ comprises \mathcal{V}^* 's random coins and all answers to \mathcal{V}^* 's queries to π .

Definition 8.5 (Perfect-Zero-Knowledge PCP). We say that a PCP system $(\mathcal{P}, \mathcal{V})$ for a language \mathcal{L} is **perfect zero-knowledge** (a PZK-PCP) if there exists an (expected) PPT algorithm Sim (the simulator), such that for every (possibly malicious) adaptive polynomial-time verifier \mathcal{V}^* , and for every $x \in \mathcal{L}$, $\text{Sim}^{\mathcal{V}^*}(x)$ is distributed identically to $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x)$.

8.2 Our Construction

We define the following randomised encoding function.

Definition 8.6. The linear randomised encoding $\text{ENC}_{\text{PCP}}: \mathbb{F}^{\mathcal{A}} \rightarrow \mathbb{F}^{\mathbb{F}^{\leq m}}$ is given by $\text{ENC}_{\text{PCP}}(m) := \text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})} \circ \text{ENC}_{\Sigma\text{AntiSym}}(m)$, for $\vec{d} \in \mathbb{F}^m$ and $\mathcal{A} \subseteq \mathbb{F}^m$ such that $d_i \geq 2(|A_i| - 1)$ for all $i \in [m]$.

Corollary 8.7. The encoding function ENC_{PCP} is (t, ℓ) -locally simulatable, for $\ell(n) = \ell_{\Sigma\text{RM}}(\ell_{\text{AntiSym}}(n))$ and $t(n) = \text{poly}(\log |\mathbb{F}|, m, \max_i d_i, n)$, where $\ell_{\Sigma\text{RM}}(n) := nm(m(a+1)+1)^2$, for $a := \max_i |A_i|$ and $\ell_{\text{AntiSym}}(n) := n^2m^2 + 1 + O(n^4m^4/|\mathcal{A}|)$.

Proof. By Theorem 6.2, there is an $\ell_{\Sigma\text{RM}}$ -constraint locator for $\text{ENC}_{\Sigma\text{RM}(\vec{d}, \mathcal{A})}$. By Theorem 7.2 there is an ℓ_{AntiSym} -constraint locator for $\text{ENC}_{\Sigma\text{AntiSym}}$. Thus Claim 4.10 and Claim 4.8 imply the result. \square

Theorem 8.8 (Low individual degree test [GS06; GR15]). Let $m \in \mathbb{N}$, $\vec{d} \in \mathbb{N}^m$ be such that $\sum_i d_i < |\mathbb{F}|/10$, $\varepsilon \in (0, 1/10)$ and $\delta \in [0, 1]$. There exists an efficient test that, given oracle access to a function $P: \mathbb{F}^m \rightarrow \mathbb{F}$, makes $O(md \cdot \text{poly}(1/\varepsilon) \cdot \log(1/\delta))$ queries to P , and:

- if $P \in \text{RM}[\mathbb{F}, m, \vec{d}]$ then the test accepts with probability 1;
- if P is ε -far from $\text{RM}[\mathbb{F}, m, \vec{d}]$ then the test accepts with probability at most δ .

Construction 8.9. A PZK-PCPP for sumcheck. Both parties receive the common input $(\mathbb{F}, m, d, H, \gamma)$, and oracle access to the evaluation table of $F: \mathbb{F}^m \rightarrow \mathbb{F}$. The proof proceeds as follows.

Proof:

1. Sample a polynomial $Q \leftarrow \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ uniformly at random. Compute the full evaluation table π_Q of Q .
2. For each $i \in [m]$, sample $T_i \leftarrow \mathbb{F}^{\leq \vec{d}_i}[X_1, \dots, X_m]$, where $\vec{d}_i = (d, \dots, d - |H|, \dots, d)$ is the vector which takes the value d in every coordinate except for the i -th location which takes the value $d - |H|$. Compute the full evaluation table π_{T_i} of T_i .
3. Define $R(\vec{X}) := Q(\vec{X}) - Q_{\text{rev}}(\vec{X}) + \sum_{i=1}^m Z_H(X_i)T_i(\vec{X})$, where $Z_H := \prod_{a \in H} (X - a)$, and $Q_{\text{rev}}(\vec{X}) := Q(\vec{X}_{\text{rev}})$.
4. Compute $\pi_\Sigma := \Sigma_{H^m}[F + R] \in \mathbb{F}^{\leq m}$.
5. Output $\pi := (\pi_\Sigma, \pi_Q, \pi_{T_1}, \dots, \pi_{T_m})$.

Verifier:

1. Run the sumcheck verifier on π , yielding a claim “ $(F + R)(\vec{\alpha}) = \beta$ ” for a uniformly random $\vec{\alpha} \in \mathbb{F}^m$. Query $F(\vec{\alpha})$. Evaluate $R(\vec{\alpha})$ by querying $\pi_{T_1}, \dots, \pi_{T_m}$ at $\vec{\alpha}$, and π_Q at $\vec{\alpha}, \vec{\alpha}_{\text{rev}}$. If $F(\vec{\alpha}) + R(\vec{\alpha}) \neq \beta$, reject.
2. Perform a low individual degree test (Theorem 8.8) on each of $\pi_{T_1}, \dots, \pi_{T_m}, \pi_Q$, with $\varepsilon = \frac{1}{4(m+2)}$, $\delta = 1/2$, and reject if any test rejects.
3. If none of the above tests reject, then accept.

Let S be the local simulator guaranteed by Corollary 8.7, and let CD be the t -constraint detector for RM guaranteed by Theorem 5.4.

Construction 8.10 (Simulator). A simulator for Construction 8.9. It receives as input $(\mathbb{F}, m, d, H, \gamma, \kappa)$, and oracle access to $F \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$.

$\text{Sim}^F(T_\Sigma, T_Q, T_{T_1}, \dots, T_{T_m}, (\mathcal{O}, \vec{\alpha}))$:

1. Sample $\beta \leftarrow S^F(T_\Sigma, \vec{\alpha})$, and store $(\vec{\alpha}, \beta)$ in T_Σ .
2. If $\vec{\alpha} \in \mathbb{F}^{\leq m}$ is a query to π_Σ , return β .
3. Otherwise, for each $i \in [m]$, compute $Z_{T_i} := \text{CD}_{\vec{d}_i}(\text{supp}(T_{T_i}) \cup \{\vec{\alpha}\})$ and $Z_Q := \text{CD}_{\vec{d}}(\text{supp}(T_Q) \cup \{\vec{\alpha}, \vec{\alpha}_{\text{rev}}\})$.
4. Let $z_{\Sigma, \vec{\alpha}} := (1, -1, Z_H(\alpha_1), \dots, Z_H(\alpha_m))$. For each $i \in [m]$, let $\gamma_i, \gamma_Q, \gamma_{Q_{\text{rev}}} \in \mathbb{F}$ be uniformly random solutions to the linear system:

$$\begin{aligned} z_{\Sigma, \vec{\alpha}} \cdot (\gamma_Q, \gamma_{Q_{\text{rev}}}, \gamma_1, \dots, \gamma_m) &= \beta - F(\vec{\alpha}) \\ Z_Q(T_Q, \gamma_Q, \gamma_{Q_{\text{rev}}})^T &= 0 \\ \forall i \in [m], Z_{T_i}(T_{T_i}, \gamma_i)^T &= 0 \end{aligned}$$

where we view each T as a vector in $\mathbb{F}^{\text{supp}(T)}$.

5. For each i , add $(\vec{\alpha}, \gamma_i)$ to T_{T_i} . Add $(\vec{\alpha}, \gamma_Q)$ and $(\vec{\alpha}_{\text{rev}}, \gamma_{Q_{\text{rev}}})$ to T_Q .
6. Return $T_{\mathcal{O}}(\vec{\alpha})$.

Lemma 8.11. *Construction 8.9 is a perfect zero-knowledge PCP of proximity for sumcheck. That is, if $md < |\mathbb{F}|/10$, and $d \geq |H| + 1$, then*

- **Completeness:** if $F \in \text{RM}[\mathbb{F}, m, d]$ and $\sum_{\vec{a} \in H^m} F(\vec{a}) = \gamma$, then the verifier accepts π with probability 1;
- **Soundness:** if $F \in \text{RM}[\mathbb{F}, m, d]$ and $\sum_{\vec{a} \in H^m} F(\vec{a}) \neq \gamma$, then for any proof π^* , the probability that the verifier accepts π^* is at most $1/2$;
- **Perfect zero knowledge:** there exists an algorithm Sim running in time $\text{poly}(\log |\mathbb{F}|, m, d, |H|)$ such that, for any (malicious) PPT verifier \mathcal{V}^* , $\text{Sim}^{F, \mathcal{V}^*}(\mathbb{F}, m, d, H, \gamma)$ outputs the distribution $\text{View}_{\mathcal{V}^*, \mathcal{P}}(x)$.

The proof is of length $O(\log |\mathbb{F}| \cdot m \cdot |\mathbb{F}|^m)$.

Proof. Completeness is straightforward; we prove soundness and zero knowledge.

Soundness. Suppose that F is such that $\sum_{\vec{a} \in H^m} F(\vec{a}) = \tilde{\gamma} \neq \gamma$, and let $\pi^* = (\pi_\Sigma, \pi_Q, \pi_{T_1}, \dots, \pi_{T_m})$. If any of $\pi_Q, \pi_{T_1}, \dots, \pi_{T_m}$ are more than ε -far from their corresponding RM codes, then the verifier rejects with probability $1/2$ as required. Otherwise, there exist polynomials $Q \in \mathbb{F}^{\leq d}[X_1, \dots, X_m]$ and $T_i \in \mathbb{F}^{\leq \vec{d}_i}[X_1, \dots, X_m]$ for each $i \in [m]$ whose evaluations over \mathbb{F}^m are each ε -close to their corresponding codewords in π^* .

Let $\hat{R} := Q - Q_{\text{rev}} + \sum_{i=1}^m X_i(1 - X_i)T_i$. Then $\hat{R} \in \mathbb{F}^{\leq d}[X_1, \dots, X_m]$ and, by construction, $\sum_{\vec{a} \in H^m} F(\vec{a}) + \hat{R}(\vec{a}) = \tilde{\gamma}$. By the soundness of the (standard) sumcheck PCPP, with probability $1 - \frac{md}{|\mathbb{F}|} \geq 3/4$ the verifier's claim " $(F + \hat{R})(\vec{\alpha}) = \beta$ " is false. Since $\vec{\alpha}$ is uniformly random in \mathbb{F}^m , and the evaluation of R queries each T_i at $\vec{\alpha}$ and Q at $\vec{\alpha}, \vec{\alpha}_{\text{rev}}$, the probability that $R(\vec{\alpha})$ as computed by the verifier is equal to $\hat{R}(\vec{\alpha})$ is at least $3/4$. The stated soundness follows by a union bound.

Zero knowledge. Fix $\pi_\Sigma \in \text{supp}(\text{ENC}_{\text{PCP}}(F|_{H^m}))$, and consider the following distribution $D(\pi_\Sigma)$ on $(\pi_Q, \pi_{T_1}, \dots, \pi_{T_m})$:

Sample $Q \leftarrow \mathbb{F}^{\leq d}[X_1, \dots, X_m]$, $T_i \leftarrow \mathbb{F}^{\leq \vec{d}_i}[X_1, \dots, X_m]$ for $i \in [m]$, such that $\pi_\Sigma(\vec{\alpha}) = F(\vec{\alpha}) + Q(\vec{\alpha}) - Q(\vec{\alpha}_{\text{rev}}) + \sum_{i=1}^m \alpha_i(1 - \alpha_i)T_i(\vec{\alpha})$ for all $\vec{\alpha} \in \mathbb{F}^m$. Let $\pi_Q, \pi_{T_1}, \dots, \pi_{T_m}$ be the corresponding evaluation tables.

By construction, $D(\pi_\Sigma)$ is the distribution of $(\pi_Q, \pi_{T_1}, \dots, \pi_{T_m})$ produced by the sumcheck prover, conditioned on a fixed π_Σ . For any $S \subseteq \mathbb{F}^m$,

$$\Pr_{D(\pi_\Sigma)} \left[\begin{array}{l} \pi_Q|_{S \cup S_{\text{rev}}} = T_Q \\ \forall i, \pi_{T_i}|_S = T_{T_i} \end{array} \right] = \Pr \left[\begin{array}{l} Q|_{S \cup S_{\text{rev}}} = T_Q \\ \forall i, T_i|_S = T_{T_i} \end{array} \middle| \begin{array}{l} Q \leftarrow \mathbb{F}^{\leq d}[X_1, \dots, X_m] \\ T_i \leftarrow \mathbb{F}^{\leq \vec{d}_i}[X_1, \dots, X_m] \\ z_{\Sigma, \vec{\alpha}} \cdot (Q(\vec{\alpha}), Q(\vec{\alpha}_{\text{rev}}), T_1(\vec{\alpha}), \dots, T_m(\vec{\alpha})) = \\ \pi_\Sigma(\vec{\alpha}) - F(\vec{\alpha}) \forall \vec{\alpha} \in S \end{array} \right].$$

This implies, by the definition of a t -constraint detector, that under $D(\pi_\Sigma)$, $\pi_Q|_{S \cup S_{\text{rev}}}$ and $\pi_{T_i}|_S$, $i \in [m]$, are distributed as a uniformly random solution to the linear system \mathcal{L}_S given by

$$\begin{aligned} Z_Q \cdot \pi_Q|_{S \cup S_{\text{rev}}} &= 0 \\ \forall i \in m, Z_{T_i} \cdot \pi_{T_i}|_S &= 0 \\ Z_\Sigma(\pi_Q|_{S \cup S_{\text{rev}}}, \pi_{T_i}|_S) &= \pi_\Sigma|_S - F|_S \end{aligned}$$

where $Z_Q := \text{CD}_d(S \cup S_{\text{rev}})$, $Z_T := \text{CD}_{\vec{d}_i}(S)$, and $Z_\Sigma := (z_{\Sigma, \vec{\alpha}})_{\vec{\alpha} \in S}$.

Suppose, by induction, that $T_Q: S \cup S_{\text{rev}} \rightarrow \mathbb{F}$, $T_{T_i}: S \rightarrow \mathbb{F}$, $i \in [m]$, are distributed as a uniformly random solution to \mathcal{L}_S . Let $U := S \cup \{\vec{\alpha}\}$. By construction, $T_Q \cup \{(\vec{\alpha}, \gamma_Q), (\vec{\alpha}_{\text{rev}}, \gamma_{Q_{\text{rev}}})\}$ and $T_{T_i} \cup \{(\vec{\alpha}, \gamma_{T_i})\}$, $i \in [m]$, are distributed as a uniformly random solution to \mathcal{L}_U .

It remains to show that, in the real world, π_Σ is distributed as $\text{ENC}_{\text{PCP}}(F|_{H^m})$; i.e., a uniformly random element w of ΣRM such that $w(\vec{a}) = F(\vec{a}) + A(\vec{a})$ for all $\vec{a} \in H^m$, for a uniformly random element A of $\Sigma\text{AntiSym}$. By Proposition 7.4, for random Q , $(Q - Q_{\text{rev}})|_{H^m}$ is distributed as A since H^m is a (subset of an) interpolating set.

By the combinatorial nullstellensatz [Alo99; CCGOS23], the polynomial $Z = \sum_{i=1}^m Z_H T_i$ is distributed as a uniformly random element of $\mathbf{Z}_{H^m}(\text{RM}[\mathbb{F}, m, d])$. Observe that, for any degree- d extension \hat{f} of a function $f: H^m \rightarrow \mathbb{F}$, $\hat{f} + Z$ is a uniformly random degree- d extension of f . Hence, $F + R = F + Q - Q_{\text{rev}} + Z$ is distributed as a uniformly random degree- d extension of $(F + A)|_{H^m}$. Therefore $\pi_\Sigma = \Sigma_{H^m}[F + R]$ is distributed as w . \square

A Deferred proofs

A.1 Proof of Claim 5.10

Claim 5.10. *Let $C \subseteq \mathbb{F}^D$ be a linear code and let $I \subseteq D$. Then I is unconstrained with respect to C if and only if for all $x \in I$ there exists a codeword $w_x \in C$ satisfying (i) $w_x(x) = 1$ and (ii) $w_x(y) = 0$ for all $y \in I \setminus \{x\}$.*

Proof. First we show the forward implication. Assume I is unconstrained with respect to C . By definition, this means that any vector $z \in \mathbb{F}^I$ satisfying $z \cdot w|_I = 0$ for all $w \in C$ must be identically zero. In other words, $(C|_I)^\perp = \{\vec{0}\}$. However, we know that $\dim(C|_I) + \dim((C|_I)^\perp) = |I|$. As $\dim((C|_I)^\perp) = 0$, it must be that $\dim(C|_I) = |I|$. Therefore $C|_I$ is an $|I|$ -dimensional subspace of \mathbb{F}^I , meaning $C|_I$ is isomorphic to \mathbb{F}^I . Thus, for each $x \in I$ there exists a codeword $w_x|_I \in C|_I$ such that $w_x|_I(x) = 1$ and $w_x|_I(y) = 0$ for all $y \in I \setminus \{x\}$.

For the reverse implication, let $z \in \mathbb{F}^I$ be a constraint satisfying $z \cdot w|_I = 0$ for all $w \in C$. We will show that z is identically zero. By assumption, for each $x \in I$, there exists a codeword $w_x \in C$ satisfying $w_x(x) = 1$ and $w_x(y) = 0$ for all $y \in I \setminus \{x\}$. Substituting w_x into our constraint equation for each $x \in I$ give $z(x) = z \cdot w_x|_I = 0$ for each $x \in I$, that is, z is identically zero. In other words, I is unconstrained with respect to C . \square

A.2 Proof of Lemma 5.11

Lemma 5.11. *Let $C \subseteq \mathbb{F}^D$ be a linear code, and let $I, \mathcal{S} \subseteq D$. Suppose that \mathcal{S} is unconstrained with respect to C . Then $I \cup \mathcal{S}$ is constrained with respect to C if and only if $I \setminus \mathcal{S}$ is constrained with respect to $\mathbf{Z}_\mathcal{S}(C)$.*

Proof. First we show the forward implication. If $I \cup \mathcal{S}$ is constrained with respect to C then there exists a non-zero constraint $z \in \mathbb{F}^{I \cup \mathcal{S}}$ such that for every $w \in C$ we have

$$\sum_{x \in I \cup \mathcal{S}} z(x)w(x) = 0 .$$

As $\mathbf{Z}_\mathcal{S}(C)$ is a subcode of C , we have that all $w' \in \mathbf{Z}_\mathcal{S}(C)$ satisfy

$$\sum_{x \in I \cup \mathcal{S}} z(x)w'(x) = \sum_{x \in I} z(x)w'(x) = 0 ,$$

where we have used the fact that $w'(x) = 0$ for all $x \in \mathcal{S}$. Lastly, there must exist $x \in I \setminus \mathcal{S}$ for which $z(x) \neq 0$, for if not, then as z is non-zero, $z|_\mathcal{S}$ is a non-zero constraint on \mathcal{S} , contradicting the hypothesis. Hence, I is constrained with respect to $\mathbf{Z}_\mathcal{S}(C)$.

For the reverse implication, we will employ the contrapositive: if $I \cup S$ is unconstrained with respect to C , then $I \setminus S$ is unconstrained with respect to $\mathbf{Z}_S(C)$. If $I \cup S$ is unconstrained with respect to C , then, by the forward implication of Claim 5.10, for each $x \in I \cup S$ there exists $w_x \in C$ satisfying (i) $w_x(x) = 1$; and (ii) $w_x(y) = 0$ for all $y \in I \cup S \setminus \{x\}$. In particular, for every $x \in I \setminus S$, $w_x(y) = 0$ for all $y \in S$, so we have that for all $x \in I \setminus S$, $w_x \in \mathbf{Z}_S(C)$. Then the reverse implication of Claim 5.10 implies that $I \setminus S$ is unconstrained with respect to $\mathbf{Z}_S(C)$. \square

A.3 Proof of Claim 5.15

Claim 5.15. *Let $C \subseteq \mathbb{F}^D$ be a linear code which has an efficient constraint detector. Given a subdomain $I \subseteq D$, there exists an efficient algorithm, which computes an **interpolating set for I with respect to C** : a set $I' \subseteq I$ such that: (i) I' is unconstrained with respect to C ; and (ii) for every $S \subseteq D$ which is unconstrained with respect to C and satisfies $I \cap S = \emptyset$, if there exists a constraint $z: I \cup S \rightarrow \mathbb{F}$ such that $z(s) \neq 0$ for some $s \in S$, then there exists a constraint $z': I' \cup S \rightarrow \mathbb{F}$ such that $z'(s) \neq 0$.*

First we give the construction, as follows.

Construction A.1. An algorithm to compute the interpolating set for a linear code C given a constraint detector CD for C . Receives as input a subdomain $I \subseteq \mathbb{F}^m$.

Interpolate(I):

1. Compute $Z := \text{CD}(I)$.
2. Perform Gaussian elimination on Z to obtain a matrix Z' which is in reduced row echelon form.
3. Output the set $I' := \{x \in I : x \text{ is a free variable in } Z'\}$.

The efficiency is clear by construction. To see that this construction is correct, we will show that I' is unconstrained and each $x \in I \setminus I'$ is determined by I' . We then show that the second property in Claim 5.15 is a consequence of these two properties.

To see that I' is unconstrained, note that as each element $x \in I'$ is a free variable, for each choice of values $\vec{u} \in \mathbb{F}^{I'}$ for the variables in I' there will be some setting $\vec{v} \in \mathbb{F}^{I \setminus I'}$ of the remaining leading variables in $I \setminus I'$ such that $(\vec{u}, \vec{v}) \in \ker(Z')$. By the correctness of CD, $(\vec{u}, \vec{v}) \in \ker(Z')$ if and only if $(\vec{u}, \vec{v}) \in C|_I$. We have just demonstrated that $C|_{I'} = \mathbb{F}^{I'}$, in other words, I' is unconstrained.

For the second part, as each $x \in I \setminus I'$ is a leading variable, the row of Z' in which it is leading forms a constraint $z: \mathbb{F}^I \rightarrow \mathbb{F}$ with $z(x) = 1$, and $z(x') = 0$ for all $x' \in I \setminus I'$ (as Z' is in reduced row echelon form). Thus the restriction of this constraint to $I' \cup \{x\}$ is a constraint $z: \mathbb{F}^{I' \cup \{x\}} \rightarrow \mathbb{F}$ with $z(x) = 1$, so x is determined by I' .

For the third part, suppose there exists a constraint $z: I \cup S$ with respect to C such that $z(s) \neq 0$ for some $s \in S$. Denote $I^* := \{x \in I \setminus I' : z(x) \neq 0\}$. We assume that I^* is non-empty, for if it were not, then the restriction of z to $I' \cup S$ would suffice. Then we can write

$$\sum_{x \in I'} z(x)w(x) + \sum_{x^* \in I^*} z(x^*)w(x^*) + \sum_{x \in S} z(x)w(x) = 0 \quad (7)$$

for all $w \in C$. As we showed in the second part, each $x^* \in I^*$ is determined by I' . In particular, for each $x^* \in I^*$, we can write

$$w(x^*) = \sum_{x \in I'} a_{x,x^*} w(x),$$

for some scalars $a_{x,x^*} \in \mathbb{F}$. Substituting this expression into Eq. 7 we obtain

$$\sum_{x \in I'} z(x)w(x) + \sum_{x^* \in I^*} z(x^*) \left(\sum_{x \in I} a_{x,x^*} w(x) \right) + \sum_{x \in \mathcal{S}} z(x)w(x) = 0,$$

which simplifies to

$$\sum_{x \in I'} \left(z(x) + \sum_{x^* \in I^*} a_{x,x^*} z(x^*) \right) w(x) + \sum_{x \in \mathcal{S}} z(x)w(x) = 0.$$

Thus z induces a constraint $z' : I' \cup \mathcal{S} \rightarrow \mathbb{F}$ defined as follows:

$$z(x) := \begin{cases} z(x) + \sum_{x^* \in I^*} a_{x,x^*} z(x^*) & \text{if } x \in I', \\ z(x) & \text{otherwise.} \end{cases}$$

Observe that $z(s) = z'(s)$ for all $s \in \mathcal{S}$.

A.4 Proof of Lemma 5.17

Lemma 5.17. *Let $m \in \mathbb{N}$, let $I \subseteq \mathbb{F}^m$, let $S_1, \dots, S_m \subseteq \mathbb{F}$ and denote $\mathcal{S} := S_1 \times \dots \times S_m$. Set $\vec{d} := (|S_1| - 1, \dots, |S_m| - 1)$. There exists a set $G_I \subseteq \mathcal{S}$ with $|G_I| \geq |\mathcal{S}| - |I|$, such that for all $g \in G_I$ there exists a polynomial $p_g(\vec{X}) \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ satisfying*

- (i) $p_g(x) = 0$ for all $x \in I \cup G_I \setminus \{g\}$ and
- (ii) $p_g(g) = 1$.

Proof. By Fact 3.1, we can write any polynomial $p(\vec{X}) \in \mathbb{F}^{\leq \vec{d}}[X_1, \dots, X_m]$ as

$$p(\vec{X}) = \sum_{w \in \mathcal{S}} a_w L_{\mathcal{S},w}(\vec{X}),$$

for some scalars $a_w \in \mathbb{F}$. Then the constraints $p(x) = 0$ for all $x \in I$ can be written as an $|I| \times |\mathcal{S}|$ matrix A of linear equations over \mathbb{F} , relating the $|\mathcal{S}|$ many coefficients a_w . By applying Gaussian elimination we can uniquely convert A to a matrix B which is in reduced row echelon form. As $\text{rank}(B) = \text{rank}(A) \leq |I|$, B will have at most $|I|$ many ‘‘leading ones’’. Set G_I to be the set of all $w \in \mathcal{S}$ such that there is no leading one in the w -th column of B , i.e., for each $g \in G_I$, the coefficient a_g is a free variable with respect to the linear system. Thus, in particular, for any $g \in G_I$, we can set $a_g = 1$ and $a_{g'} = 0$ for all $g' \in G_I \setminus \{g\}$ and obtain a solution to the linear system, yielding the result. \square

A.5 Proof of Claim 6.11

Claim 6.11. *Let $C \subseteq \mathbb{F}^D$ be a linear code, and let $U \subseteq V \subseteq D$, $I \subseteq D$ be such that $\mathbf{Z}_U(C)|_I = \mathbf{Z}_V(C)|_I$. Then $(u, x) \in C|_{U \sqcup I}$ if and only if for all $v \in C|_V$ such that $v|_U = u$, $(v, x) \in C|_{V \sqcup I}$.*

Proof. The reverse implication is straightforward: if $(v, x) \in C|_{V \sqcup I}$, then $(v|_U, x) \in C|_{U \sqcup I}$ by taking the restriction to $U \sqcup I$.

For the forward implication, let $(u, x) \in C|_{U \sqcup I}$. Then, as $U \subseteq V$, trivially there exists $v \in C|_V$ such that $v|_U = u$ and $(v, x) \in C|_{V \sqcup I}$. Let $v' \in C|_V$ be an arbitrary message which agrees with v when restricted

to U , that is, $v'|_U = u$. Then for some $x' \in C|_I$, $(v', x') \in C|_{V \sqcup I}$. Restricting this vector to $U \sqcup I$, we have $(u, x') \in C|_{U \sqcup I}$. By linearity of $C|_{U \sqcup I}$, we have that $(u, x) - (u, x') = (\vec{0}_U, x - x') \in C|_{U \sqcup I}$, where $\vec{0}_U \in C|_U$ is the zero codeword. If we can show that $(\vec{0}_V, x - x') \in C|_{V \sqcup I}$, where $\vec{0}_V \in C|_V$ is the zero codeword, we will be done. Clearly $(u, x') \in \mathbf{Z}_U(C)|_I$, and by Corollary 6.13, we have that $\mathbf{Z}_U(C) = \mathbf{Z}_V(C)$, so $(u, x') \in \mathbf{Z}_V(C)|_I$. In other words, $(\vec{0}_V, x - x') \in C|_{V \sqcup I}$. Hence, by linearity again, $(v', x') + (\vec{0}_V, x - x') = (v', x) \in C|_{V \sqcup I}$, as required. \square

References

- [AAV13] D. Aharonov, I. Arad, and T. Vidick. “Guest column: the quantum PCP conjecture”. In: *Acm sigact news* 44.2 (2013), pp. 47–79.
- [AB09] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. 1st. New York, NY, USA: Cambridge University Press, 2009. ISBN: 0521424267, 9780521424264.
- [AH91] W. Aiello and J. Håstad. “Statistical Zero-Knowledge Languages can be Recognized in Two Rounds”. In: *Journal of Computer and System Sciences* 42.3 (1991). Preliminary version appeared in FOCS ’87., pp. 327–345.
- [ALMSS92] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. “Proof verification and hardness of approximation problems”. In: *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*. 1992, pp. 14–23.
- [Al099] N. Alon. “Combinatorial Nullstellensatz”. In: *Combinatorics, Probability and Computing* 8 (1999), pp. 7–29.
- [AS03] S. Arora and M. Sudan. “Improved Low-Degree Testing and its Applications”. In: *Combinatorica* 23.3 (2003). Preliminary version appeared in STOC ’97., pp. 365–426.
- [AW09] S. Aaronson and A. Wigderson. “Algebrization: A New Barrier in Complexity Theory”. In: *ACM Transactions on Computation Theory* 1.1 (2009), 2:1–2:54.
- [BCFGRS17] E. Ben-Sasson, A. Chiesa, M. A. Forbes, A. Gabizon, M. Riabzev, and N. Spooner. “Zero Knowledge Protocols from Succinct Constraint Detection”. In: *Proceedings of the 15th Theory of Cryptography Conference*. TCC ’17. 2017, pp. 172–206.
- [BCHTV20] A. Bouland, L. Chen, D. Holden, J. Thaler, and P. N. Vasudevan. “On the Power of Statistical Zero Knowledge”. In: *SIAM J. Comput.* 49.4 (2020).
- [BGKW88] M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson. “Multi-prover interactive proofs: how to remove intractability assumptions”. In: *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*. STOC ’88. 1988, pp. 113–131.
- [CCGOS23] M. Chen, A. Chiesa, T. Gur, J. O’Connor, and N. Spooner. “Proof-carrying data from arithmetized random oracles”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. EUROCRYPT ’23. 2023, pp. 379–404.
- [CFGs22] A. Chiesa, M. A. Forbes, T. Gur, and N. Spooner. “Spatial Isolation Implies Zero Knowledge Even in a Quantum World”. In: *Journal of the ACM* 69.2 (2022), pp. 1–44.
- [CFS17] A. Chiesa, M. A. Forbes, and N. Spooner. *A Zero Knowledge Sumcheck and its Applications*. Cryptology ePrint Archive, Report 2017/305. 2017.
- [DFKNS92] C. Dwork, U. Feige, J. Kilian, M. Naor, and S. Safra. “Low Communication 2-Prover Zero-Knowledge Proofs for NP”. In: *Proceedings of the 11th Annual International Cryptology Conference*. CRYPTO ’92. 1992, pp. 215–227.
- [Din07] I. Dinur. “The PCP theorem by gap amplification”. In: *Journal of the ACM* 54.3 (2007), p. 12.
- [For87] L. Fortnow. “The Complexity of Perfect Zero-Knowledge (Extended Abstract)”. In: *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*. STOC ’87. 1987, pp. 204–209.

- [GMR89] S. Goldwasser, S. Micali, and C. Rackoff. “The knowledge complexity of interactive proof systems”. In: *SIAM Journal on Computing* 18.1 (1989). Preliminary version appeared in STOC ’85., pp. 186–208.
- [GMW91] O. Goldreich, S. Micali, and A. Wigderson. “Proofs that Yield Nothing But Their Validity or All Languages in NP Have Zero-Knowledge Proof Systems”. In: *Journal of the ACM* 38.3 (1991). Preliminary version appeared in FOCS ’86., pp. 691–729.
- [GR15] T. Gur and R. D. Rothblum. “Non-Interactive Proofs of Proximity”. In: *Proceedings of the 6th Innovations in Theoretical Computer Science Conference*. ITCS ’15. 2015, pp. 133–142.
- [GS06] O. Goldreich and M. Sudan. “Locally testable codes and PCPs of almost-linear length”. In: *Journal of the ACM* 53 (4 2006). Preliminary version in STOC ’02., pp. 558–655.
- [GSY19] A. B. Grilo, W. Slofstra, and H. Yuen. “Perfect Zero Knowledge for Quantum Multiprover Interactive Proofs”. In: *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019*. IEEE Computer Society, 2019, pp. 611–635.
- [HVW22] C. Hazay, M. Venkatasubramanian, and M. Weiss. “ZK-PCPs from Leakage-Resilient Secret Sharing”. In: *J. Cryptol.* 35.4 (2022), p. 23.
- [IMS12] Y. Ishai, M. Mahmoody, and A. Sahai. “On Efficient Zero-Knowledge PCPs”. In: *Proceedings of the 9th Theory of Cryptography Conference on Theory of Cryptography*. TCC ’12. 2012, pp. 151–168.
- [IW14] Y. Ishai and M. Weiss. “Probabilistically Checkable Proofs of Proximity with Zero-Knowledge”. In: *Proceedings of the 11th Theory of Cryptography Conference*. TCC ’14. 2014, pp. 121–145.
- [IWY16] Y. Ishai, M. Weiss, and G. Yang. “Making the Best of a Leaky Situation: Zero-Knowledge PCPs from Leakage-Resilient Circuits”. In: *Proceedings of the 13th Theory of Cryptography Conference*. TCC ’16-A. 2016, pp. 3–32.
- [JKRS09] A. Juma, V. Kabanets, C. Rackoff, and A. Shpilka. “The Black-Box Query Complexity of Polynomial Summation”. In: *Computational Complexity* 18.1 (2009), pp. 59–79.
- [KPT97] J. Kilian, E. Petrank, and G. Tardos. “Probabilistically checkable proofs with zero knowledge”. In: *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*. STOC ’97. 1997, pp. 496–505.
- [KR08] Y. Kalai and R. Raz. “Interactive PCP”. In: *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*. ICALP ’08. 2008, pp. 536–547.
- [LFKN92] C. Lund, L. Fortnow, H. J. Karloff, and N. Nisan. “Algebraic Methods for Interactive Proof Systems”. In: *Journal of the ACM* 39.4 (1992), pp. 859–868.
- [LS95] D. Lapidot and A. Shamir. “A One-Round, Two-Prover, Zero-Knowledge Protocol for NP”. In: *Combinatorica* 15.2 (1995), pp. 204–214.
- [Tha22] J. Thaler. “Proofs, arguments, and zero-knowledge”. In: *Foundations and Trends® in Privacy and Security* 4.2–4 (2022), pp. 117–660.
- [Vad99] S. P. Vadhan. “A study of statistical zero-knowledge proofs”. PhD thesis. Massachusetts Institute of Technology, 1999.
- [Wei22] M. Weiss. “Shielding Probabilistically Checkable Proofs: Zero-Knowledge PCPs from Leakage Resilience”. In: *Entropy* 24.7 (2022), p. 970.