Electronic Colloquium on Computational Complexity, Report No. 150 (2026)

Private PCPs from Product Expansion

Mitali Bafna* Nikhil Vyas†

Abstract

The quantum analogue of the PCP theorem for QMA remains wide open. A central obstacle is the local indistinguishability of quantum codes: every sufficiently small view of an encoded witness is independent of the witness, seemingly preventing a local verifier from distinguishing YES from NO instances. One approach to this apparent paradox [ABN24] is to encode the witness in a quantum code and compute on it fault-tolerantly, successively reducing the encoding length until the answer is revealed. We study a classical relaxation of this approach based on the equivalent view of quantum codes as private randomized encodings from multiparty computation. Specifically, we ask whether one can construct circuits that compute on a private encoding of an NP witness, such that every sufficiently small fractional[1] view of the honest computation transcript is independent of the witness (i.e. has quantum distance) and the computation remains correct despite a small fraction of adversarial bit-flip or X-errors in every layer.

We construct such circuits and use the classical Cook–Levin theorem to obtain *private PCPs* for NP: the prescribed PCP encoding is randomized, and for each instance every view below the privacy threshold has the same distribution for all assignments, whether satisfying or not. For Circuit-SAT instances of size n, we get $\sqrt{n}$-query private PCPs of size $O(n \log n)$ over alphabet size $\mathrm{poly}(n)$, fractional privacy $\Omega(1/\log n)$ and constant soundness gap. Furthermore, conditional on a high-dimensional *product expansion* conjecture for Reed–Solomon codes, our PCPs have length $n^{1+o(1)}$, use $n^{o(1)}$ queries, and have fractional privacy and soundness gap $n^{-o(1)}$.

Our construction is based on a new family of small-alphabet quantum codes which have near-linear rate, sparse X-checks that enable local testability for X-errors, support for multiplication (or transversal CCZ gates on the full logical space) and near-linear quantum distance using product expansion. The codes are obtained using tensor products of Reed–Solomon codes, and our key innovation is to choose the evaluation domains as multiplicative subgroups of pairwise coprime orders of $\mathbb{F}_q^\star$. A proof obtained by ChatGPT 5.6 Sol establishes constant 2-dimensional product expansion whenever both rates are bounded away from one, crossing the tight sum-of-rates-below-one barrier in the product expansion theorem of Polishchuk and Spielman [PS94]. We conjecture the analogous statement in higher dimensions; a candidate proof generated by ChatGPT is available for this statement (link), but the authors have not verified its correctness.

*University of Washington. Email: mitalib@cs.washington.edu

†Email: vyasnikhil96@gmail.com

[1]The fraction is under an appropriate non-uniform measure on the transcript.

ISSN 1433-8092

Contents

1 Introduction

The classical PCP theorem [ALM+98, AS98, Din07] is one of the most important results in computer science, with many applications, especially to hardness of approximation. It says that the n-bit witness of any NP problem such as 3SAT can be encoded into a constant-alphabet proof that can be verified probabilistically using only a constant number of queries: in the YES case the verifier always accepts, whereas in the NO case it rejects with constant probability. The analogous statement for QMA, the quantum analogue of NP, is wide open. In particular, no construction is known even with $n^{o(1)}$ queries, an $n^{o(1)}$-bit alphabet, and soundness gap $n^{-o(1)}$, let alone even a *polylog-quantum-PCP* with each $n^{o(1)}$ replaced with $\mathsf{polylog}n$.

A classical PCP typically works by encoding the NP witness in an error-correcting code, so that whether it is a YES or NO witness can be determined with good probability by examining only a few random locations. A major obstacle to quantizing this paradigm is that quantum codes have *local indistinguishability*: if the distance is d, then every view of fewer than d coordinates is independent of the encoded witness; we refer to this as d-privacy. Thus one cannot directly use local constraints to distinguish encodings of YES and NO witnesses as in a classical PCP proof.

One proposal of [ABN24] to get around this paradox is to build a circuit which takes in a witness encoded in a quantum code, computes on it while always keeping it encoded, but finally extracts a YES/NO answer by successively reducing the encoding length until the one-bit is revealed. Furthermore, the quantum code and the circuit should correctly compute the output even if each layer undergoes a small fraction of adversarial faults. Thus, although any small view across a layer reveals nothing about the witness, the final output in a $\mathsf{polylog}n$-depth and fanin circuit can depend on every bit. At the same time, the encoding can retain its fractional privacy at every layer as its length shrinks; since the adversary can only corrupt a small fraction of each layer, once the encoding has shrunk to a single bit, it cannot corrupt the final bit. They show a quantum analogue of the Cook–Levin theorem which given such an adversarially robust circuit with $\mathsf{polylog}n$ fanin and $1/\mathsf{polylog}n$-robustness to errors, produces a polylog-quantum-PCP with a weaker soundness guarantee called "combinatorial soundness"; they conjecture that their reduction also gives the standard soundness guarantee (known as "energy").

We propose a relaxed classical version of this question: can one build fault-tolerant circuits for verifying an NP witness, while keeping the witness encoded in a quantum code, but tolerance only against classical (bit-flip) adversarial errors? Fortunately, quantum codes have an equivalent classical view as secret-sharing schemes or *private codes*[2], first studied in multiparty computation (MPC) [BOGW88, CCD88]: whereas an ordinary error-correcting code assigns one string in Σ^N to each message x, a private code assigns a set of strings $\mathrm{Enc}(x) \subset \Sigma^N$, such that the marginal on any small fraction of coordinates of the uniform distribution over $\mathrm{Enc}(x)$ is independent of x. Thus, given a function f, the goal is to build an encoding function $\mathrm{Enc}(\cdot)$ and a *private* and *robust* circuit that computes directly on this encoding. We call the circuit δ-robust if while starting from an input that is δ-close to some word in $\mathrm{Enc}(x)$, an execution with at most a δ-fraction of faulty gates in each layer outputs $f(x)$. We call it ν-private if every ν-fractional view of any layer (including the input layer), or spread jointly across the honest computation transcript, is independent of x.

A major motivation for studying this question is to isolate, in a simple classical model, what we view as a central challenge to building quantum PCPs,

Question 1.1. *Can we compute* $\mathrm{poly}(n)$*-sized circuits on* n *bits using* $n^{o(1)}$*-fanin and* $\mathrm{poly}(n)$*-sized circuits over alphabet size* $\mathrm{poly}(n)$ *that are* $n^{-o(1)}$*-robust and* $n^{-o(1)}$*-private, ideally with* $n^{o(1)} = \mathsf{polylog}n$*?*

[2]We elaborate upon this equivalence in Section 3; see also [RZW26] where they call these zero-knowledge codes. Although these are the same objects we use the term private codes while referring to the distributional view of quantum codes.

Although this notion is the same as perfect information-theoretic privacy underlying multiparty computation, the focus on *fractional* privacy here makes the question significantly harder. In ABN's proposal, the fractional tolerance against quantum errors is at most the fractional privacy, analogous to how in classical PCPs, the soundness gap is at most the fractional distance of the encoding. All the prior works in MPC and zero-knowledge-PCPs [KPT97, IMS12, GOS25], encode the input bit-by-bit – this increases the absolute privacy, but does not increase the privacy as a fraction of the encoding length. Our parameter regime therefore forces us to encode the witness globally and compute robustly on this global encoding. In fact, it is not even clear whether one can compute simple functions like the SUM and AND of n bits in this model.

Analogous to [ABN24], given such a circuit, we can use the classical Cook-Levin theorem to compile it into a *private PCP*: the transcript of the circuit gives us an honest encoding of any assignment w into $\Pi(w)$ that is randomized and the privacy of the circuit gives that any view of relative size less than ν (in the measure induced by the verifier's queries) is independent of the assignment w (see Definition 8.2), whereas the robustness of the circuit implies the gap of the PCP. This is stronger than the zero-knowledge requirement, which concerns the privacy of witnesses only in the YES case.[3] There is also a simple restriction on the encoding length. Write the randomized encoder as $\mathrm{Enc}(x) = D(x; r)$, where D is deterministic once the random seed r is fixed. If r contains $|r|$ symbols, then it can be shown that the privacy threshold of $\mathrm{Enc}(\cdot)$ is at most $|r|+1$, and hence its fractional privacy is at most $(|r|+1)/N$. Since $|r| < \mathrm{poly}(n)$, to get fractional privacy $n^{-o(1)} = |r|^{-o(1)}$, the blowup incurred by $D(\cdot)$ must be almost-linear in its input size. So one must aim to construct almost-linear-size PCPs; classically we know of only two such constructions in this gap regime [BS08, BKK+16].

Our Results: Conditional on a certain *product expansion* conjecture about Reed–Solomon codes (Conjecture 1.4), our main result is the following private PCP construction,

Theorem 1.2 (Informal)**.** *For $k = 2$ unconditionally, and for every integer $k > 2$ assume Conjecture 1.4 and let $\rho = \rho(k)$ be the product expansion constant therein. For every integer $n > k$, Circuit-SAT on instances of size n has a private PCP with size $N = n \log n \exp(O(k \log k))$, over alphabet size $\mathrm{poly}(n)$ and query size $O(kN^{1/k})$. It has perfect completeness, soundness gap $\exp(-O(k))$, and relative privacy at least $\frac{\rho^{\mathrm{poly}(k)}}{\log(n)}$.*

Unconditionally, we obtain $\sqrt{n}$-query private PCPs over an alphabet of size $\mathrm{poly}(n)$, with privacy $\Omega(1/\log n)$ and robustness $\Omega(1)$. This already gives a nontrivial PCP construction analogous to the classical construction of [BKK+16] (with alphabet size $O(1)$). Furthermore, if the conjecture holds, we can let $k = k(n)$ grow sufficiently slowly so that all dependencies on k and $\rho(k)$ are $n^{o(1)}$, giving proof length $n^{1+o(1)}$, $n^{o(1)}$ queries, and privacy and soundness gap $n^{-o(1)}$. The circuit underlying our PCP construction has the same parameters. This also implies perfect-zero-knowledge PCPs with the same asymptotic parameters although we do not obtain an efficient simulator.

The technical heart of our construction is a new private code with linear distance and privacy, sublinear-sized parity checks that allow for local testability, and support for multiplication: the coordinate-wise product of representatives of $\mathrm{Enc}(x)$ and $\mathrm{Enc}(y)$ gives an encoding (in a slightly larger code) of the coordinate-wise multiplication of x and y. We then show how these codes allow us to compute functions like SUM and AND of n symbols robustly and privately. Finally, we adapt the 'arithmetization' procedure used in the almost-linear-size PCP constructions of [BS08, BKK+16] to show the NP-hardness of an appropriately chosen algebraic problem that is compatible with our codes. Computing this problem robustly using our

[3] Perfect zero-knowledge is usually stated using a simulator that, given only x, reproduces the view of a bounded-query verifier. Our definition directly compares the distributions of small views and does not require a simulator.

codes, gives the desired private PCPs for NP. Although we begin from a classical perspective, the equivalence between private and quantum codes yields the first quantum codes satisfying analogous guarantees which are of considerable interest for fault-tolerant quantum computation.

New Quantum Codes with Transversal CCZ Gates from Product Expansion: Our codes are based on tensor products of Reed–Solomon (RS) codes over a prime field $\mathbb{F}_q$ of polynomial size. The tensor product structure naturally gives local testability but the challenge is to design these codes so that they are private in a regime that retains the multiplication property of RS codes.

The privacy of tensor product codes boils down to product expansion, which is a form of high-dimensional expansion that was introduced by [PK22] in their construction of the first good quantum LDPC codes. To describe the two-dimensional case, let $C_1 \subseteq \mathbb{F}_q^{S_1}$ and $C_2 \subseteq \mathbb{F}_q^{S_2}$ be linear codes, where $n_i = |S_i|$, and view a word on $S_1 \times S_2$ as a matrix. For a matrix M, let $\mathrm{supp}(M)$ be the set of its nonzero entries, and let $|M|_{\mathrm{col}}$ and $|M|_{\mathrm{row}}$ denote the number of its nonzero columns and rows, respectively.

Definition 1.3 (Two-dimensional product expansion)**.** *The pair (C_1, C_2) is ρ-product-expanding if every matrix M that is the sum of a matrix whose columns belong to C_1 and a matrix whose rows belong to C_2 admits such a decomposition $M = M_1 + M_2$ satisfying*

$$|\operatorname{supp}(M)| \geqslant \rho\big(n_1|M_1|_{\mathrm{col}} + n_2|M_2|_{\mathrm{row}}\big).$$

It turns out that 2d-product expansion of (C_1, C_2) is equivalent to robust testability of the tensor code $C_1 \otimes C_2$. Polishchuk and Spielman [PS94] proved that any pair of RS codes is product expanding with $\rho = \Omega(1)$ when the *sum* of their rates is *less than one*. This result played a crucial role in the almost-linear size classical PCP constructions of [PS94, BS08].

Their sum-of-rates condition is tight, and ultimately a barrier to obtaining linear privacy and multiplication support from tensor product of RS codes (see Section 2.1 for details). We cross this barrier by choosing the evaluation domains of the pair of RS codes carefully: we prove that any pair of RS codes, evaluated on *multiplicative subgroups of coprime orders* of $\mathbb{F}_q$, is product-expanding as long as both rates are bounded away from one (Theorem 2.1).

The higher-dimensional version of product expansion can be defined analogously (Definition 4.1). For $k > 2$, [KP25] showed that k-tuples of random codes have constant product expansion, but over exponential sized fields $q = 2^{\Theta(N)}$. Our conjecture stated below would instead give a structured polynomial-alphabet analogue useful for multiplication,

Conjecture 1.4 (Informal)**.** *Let q be a prime, $k \geqslant 2$ and $\mathcal{C} = (C_1, \ldots, C_k)$ be a k-tuple of Reed–Solomon codes evaluated on multiplicative subgroups of $\mathbb{F}_q^\star$ of pairwise coprime orders, with $\mathrm{rate}(C_i) < 1 - \varepsilon$. Then $\mathcal{C}$ is ρ-product-expanding with ρ being only a function of ε and k.*

We have a candidate AI-generated proof of this conjecture which gives $\rho = \varepsilon^{\exp(k)}$, but it has not been verified by the authors.

Our product expansion conjecture yields private codes, and by the equivalence, also quantum codes of rate and distance $\rho(k)^k$ over an alphabet of size $\mathrm{poly}(n)$ (Corollary 4.5). The X-check family has $O(kN^{1/k})$-weight generators and the codes support transversal CCZ gates (the quantum analogue of multiplication) on the full logical space. Using the subsystem-product framework of [GG25], the same construction gives quantum 'subsystem' codes with both $O(kN^{1/k})$-weight X- and Z-gauge generators. To our knowledge, we did not know of quantum codes matching our unconditional result even for $k = 2$ – the closest such result is in the work of [GG25] that used the product expansion result of [PS94] to show such codes with transversal CCZ gates but only over a subset of the logical space.

Towards Quantum Fault-Tolerant Circuits: Our question relaxes the ABN question in two ways: (i) we compute on private codes, equivalently codes with quantum distance, but require robustness only against classical bit-flip (X) errors; and (ii) we restrict attention to classical verification circuits. We view the first relaxation as the more fundamental one. Indeed, secret-sharing-based MPC techniques naturally extend to the fault-tolerant implementation of quantum gates [AB97, Smi01], suggesting that our scheme could extend to the quantum setting. The code underlying our construction already has quantum distance and sparse X checks which allow for local testability for X errors, but its Z-basis checks are dense and hence do not provide a local test for Z-errors. An immediate open question is whether one can extend our codes to be locally testable in both bases. Combined with our computation scheme, one could then get adversarially-robust quantum computation, and finally the quantum Cook–Levin theorem of [ABN24] would then give a quantum PCP with the same parameters and the weaker guarantee of combinatorial soundness.

2 Proof Overview

Throughout this section, q is prime. We first focus on the case $k = 2$ in Theorem 1.2. This case already contains the main ideas and gives square-root fanin circuits with linear privacy and tolerance to faults. The same strategy extends to higher dimensions.

2.1 Private Codes with Local Testability and Multiplication Support

We first consider a message with one symbol $s \in \mathbb{F}_q$. A natural starting point is Reed–Solomon-based secret sharing, as used in secure multiparty computation (MPC) [BOGW88]. Choose a uniformly random polynomial p of degree at most d, subject to $p(0) = s$, and output its evaluations on $\mathbb{F}_q^\star = \mathbb{F}_q \setminus \{0\}$. Any d evaluations are independent of s, while the code has relative distance $1 - \frac{d}{q-1}$. Thus privacy and distance are both linear when d is a suitable constant fraction of q. The obstacle appears when we ask for robust recovery with small-fanin gates. Every nonzero dual check reads $\Omega(q)$ symbols, so even checking proximity to the code appears to require linear fanin.

Local Testability from Tensor Product of Reed–Solomon Codes A standard way around this obstacle is to pass to bivariate polynomials, or equivalently to tensor the encoding. We choose

$$\mathrm{Enc}(s) \sim \mathrm{Unif}\left\{ \big(p(x,y)\big)_{(x,y)\in(\mathbb{F}_q^\star)^2} : \deg_X p, \deg_Y p \leqslant d,\ p(0,0) = s \right\}.$$

The encoding has length $N = (q-1)^2$ and relative distance $(1 - \frac{d}{q-1})^2$. Each row and column is a degree-d Reed–Solomon codeword, so the row–column test uses only $q - 1 = \sqrt{N}$ inputs per gate. We can recover s with the same fanin: decode each x-line and evaluate it at 0, obtaining the values $p(0, y)$, and then interpolate $p(0, Y)$ and evaluate it at 0. This gives us a depth-2, $\sqrt{N}$-fanin robust circuit for recovery.

The Diagonal Attack on Privacy in the Low-Degree Regime: This solves the local checking and recovery problem, but introduces a new privacy attack. If $2d < q - 1$, then $r(x) = p(x, x)$ has degree less than $q - 1$. Its values on the diagonal therefore determine $r(0) = p(0, 0) = s$. So only $q - 1 = \sqrt{N}$ entries, a $o(1)$-fraction of the grid, suffice to recover s.

Privacy in High-Degree Regime from Product Expansion: This suggests moving to the high-degree regime. Put $m = q - 1$, and take $(1/2 + \varepsilon)m < d < (1 - \varepsilon)m$. The diagonal attack no longer applies; in fact, the product expansion theorem of [PS94] rules out every other sparse recovery rule as well. The proof of this implication is essentially contained in [GG25, Lemma 4.9].

To see why, consider the subcode $\mathcal{C}_0$ of bivariate evaluation tables satisfying $p(0,0) = 0$. Let V be the univariate Reed–Solomon code of degree at most d, and let $K \subset V$ consist of the words obtained from polynomials f satisfying $f(0) = 0$. Then $\mathcal{C}_0 = K \otimes V + V \otimes K$, while the full bivariate code is $\mathcal{C} = V \otimes V$.

Suppose the view on $A \subseteq (\mathbb{F}_q^\star)^2$ depends on $p(0,0)$. By linearity of the encoding, there is a coefficient matrix M, supported on A, such that

$$p(0,0) = \sum_{(x,y)\in A} M_{x,y} p(x,y)$$

for every p of individual degree at most d. Equivalently, $M \in \mathcal{C}_0^\perp \setminus \mathcal{C}^\perp$.

Set $U = V^\perp$ and $W = K^\perp$. The words in U are ordinary parity checks for V, while those in $W \setminus U$ recover a nonzero multiple of $f(0)$. Moreover, W has distance $d+1$, and $\dim U + \dim W = 2m - 2d - 1 < (1 - \Omega(\varepsilon))m$. Since U and W are Reed–Solomon codes up to coordinate rescaling, the Polishchuk–Spielman theorem gives $\rho = \rho(\varepsilon) > 0$ product expansion for (W, U).

Since $M \perp K \otimes V$, product expansion (Definition 1.3) gives a decomposition $M = M_{\mathrm{col}} + M_{\mathrm{row}}$, where the columns of M_{col} lie in W, the rows of M_{row} lie in U, and $\mathrm{wt}(M)$ is at least ρm times the total number of nonzero columns and rows in these two parts. The row part lies in $\mathcal{C}^\perp$, so $M_{\mathrm{col}} \in \mathcal{C}_0^\perp \setminus \mathcal{C}^\perp$. Contracting its columns against a suitable word in V produces a word in $W \setminus U$ supported on the nonzero columns of M_{col}. Hence at least $d + 1$ columns are nonzero, and

$$|A| \geqslant \mathrm{wt}(M) \geqslant \rho m(d+1) = \Omega(q^2) = \Omega(N).$$

Thus tensoring gives $\sqrt{N}$-sized tests, while product expansion restores linear privacy; the code also has constant relative distance.

Privacy in Low-Degree Regime from Punctured Reed-Solomon Codes: Suppose, given the encodings P, P' of two symbols s, s', we want to get an encoding of $s \cdot s'$. The obvious way to implement this would be to multiply P and P' entrywise, since $(PP')(0) = P(0)P'(0) = ss'$, thus we will get a valid encoding of ss'. But we run into an immediate issue. Privacy needs the regime $2d > q$, in which case we can no longer recover our information from PP', so multiplication support is in the regime where $2d < q$.

Our main idea here is to come up with an appropriate evaluation domain $\Omega \subset \mathbb{F}_q \times \mathbb{F}_q$, so that the resulting tensor code has linear distance even in the regime $2d \ll q$. First observe that the reason we wanted $2d < q$ is because of a 'diagonal' attack, where to break privacy an adversary can look at the values of P on the diagonal $\{(x, x) : x \in \mathbb{F}_q^\star\}$ to recover $P(0,0)$. One attempt to avoid such attacks could be to take different evaluation sets for the rows and columns and in fact we want these to be so different that all such 'low-degree' attacks are ruled out. Inspired by the use of multiplicative subgroups in [BS08], we let the evaluation sets S_1, S_2 be multiplicative subgroups of $\mathbb{F}_q^\star$, but of pairwise coprime orders. That is, suppose that two coprime numbers n_1, n_2 divide $q-1$, then, given a generator ω of $\mathbb{F}_q^\star$, let $S_i = \{\omega^{\frac{j(q-1)}{n_1}} : j \in [n_i]\}$. This choice allows us to cross the sum-of-rates less than 1 condition of [PS94],

Theorem 2.1 (2d Product expansion)**.** *Fix $\varepsilon > 0$. There exists $\rho = \rho(\varepsilon) > 0$ such that the following holds. Let $S_i \leqslant \mathbb{F}_q^\times$ have coprime orders $n_i = |S_i|$ with $|S_1| \approx |S_2|$ and let $C_i = \mathrm{RS}(S_i, d_i)$, where $d_i \leqslant (1-\varepsilon)n_i$.*

Every matrix M that is a sum of C_1-columns and C_2-rows admits a decomposition $M = A + B$ such that

$$\mathrm{wt}(M) \geqslant \rho\big(n_1 \cdot \#\{\textit{nonzero columns of } A\} + n_2 \cdot \#\{\textit{nonzero rows of } B\}\big).$$

By the same proof as above, this new product expansion result implies linear privacy of encoding s via the tensor code $\mathrm{RS}(S_1, d_1') \otimes \mathrm{RS}(S_2, d_2')$, even in the regime where $d_i' = \varepsilon|S_i|$. Thus, this gives an encoding that has linear privacy and supports multiplication.

2.2 Private and Robust Circuits for SUM & AND

Having built codes for encoding one symbol, we extend the construction to a vector of symbols arranged as a hidden grid $D = H_1 \times H_2$. For $i \in \{1, 2\}$, let $S_i \leqslant \mathbb{F}_q^\times$ be multiplicative subgroups of coprime orders with $|S_1| \approx |S_2|$. Let $\Gamma_i \leqslant S_i$, and take $H_i = \zeta_i \Gamma_i$ to be a coset disjoint from S_i; it suffices to take any set H_i that is disjoint from S_i to get linear privacy, but we choose these cosets to get a code with symmetries that is useful to compute permutations of the message.

Write $D = H_1 \times H_2$, $\Omega = S_1 \times S_2$, $h_i = |H_i|$, and $s_i = |S_i|$. For a table $w \in \mathbb{F}_q^D$, choose P uniformly among the bivariate polynomials satisfying $P|_D = w$ and $\deg_{X_i} P \leqslant d_i$, where $d_i \geqslant h_i - 1$. The encoding consists of the evaluations of P on Ω and has length $N = |\Omega|$. For suitable constant-rate choices with $h_i = \Theta(d_i) = \Theta(s_i)$, where $d_i - h_i$ and $s_i - d_i$ are both linear in s_i, the same argument as above gives constant relative distance and $\Omega(N)$ privacy. We also choose $q > N$, and hence $q > |D|$.

Computing SUM: This encoding supports a simple robust computation of $\mathrm{SUM}(w) = \sum_{(a,b)\in D} w(a, b)$.

1. For every $y \in S_2$, decode the S_1-column through y and output $Q(y) = \sum_{a\in H_1} P(a, y)$. These are evaluations on S_2 of the degree-d_2 polynomial $Q(Y) = \sum_{a\in H_1} P(a, Y)$.
2. Decode Q from its evaluations on S_2 and output $\sum_{b\in H_2} Q(b) = \sum_{(a,b)\in D} w(a, b)$.

Since $|S_1| \approx |S_2|$, every gate has fan-in $O(\sqrt{N})$. The main robustness intuition is that changing the decoded value of a line requires corrupting a constant fraction of that line. Since the lines are disjoint, a small fraction of corruptions and faulty gates produces only a small fraction of incorrect first-stage outputs. These remain within the decoding radius of the outer Reed–Solomon code, so the final gate still recovers the correct polynomial Q. Moreover, a uniformly random representation P induces a uniformly random representation $Q|_{S_2}$ of the partial-sum vector $\big(\sum_{a\in H_1} w(a, b)\big)_{b\in H_2}$, so the intermediate layer remains private.

Computing AND: Given $w \in \mathbb{F}_q^D$, define $\mathrm{AND}(w)$ to be 1 if and only if $w \equiv 0$. At a high level, multiplication support gives the following computation:

1. Using repeated squaring and multiplication, compute an encoding of $w^{\odot(q-1)}$ in $O(\log q)$ layers.
2. Compute $\sigma = \mathrm{SUM}\big(w^{\odot(q-1)}\big)$ using the circuit above and output $\mathbb{1}[\sigma = 0]$.

Indeed, $w(z)^{q-1}$ is 0 when $w(z) = 0$ and 1 otherwise. Thus σ is the number of nonzero entries of w, viewed as an element of $\mathbb{F}_q$. Since $q > |D|$, it vanishes exactly when $w \equiv 0$. The degree growth can be handled by interspersing the multiplications with degree reduction along rows and columns. We do not elaborate on this approach because our construction in the next section uses a more direct method based on auxiliary proofs.

Symmetries of Our Code: Multiplication by any element of Γ_i permutes both H_i and S_i. Hence, such a multiplicative shift in one coordinate permutes both the hidden grid D and the visible grid Ω, allowing the corresponding permutation of the message to be implemented by permuting the encoded coordinates.

2.3 Private PCPs for NP

We can now use the ingredients developed in the previous two subsections to build the verifier circuit needed for every NP problem. We first reduce Circuit-SAT to 3SAT and then to the following algebraic problem tailored to our private codes,

Definition 2.2 (Informal Zero-on-Subcube)**.** *A Zero-on-Subcube instance consists of a hidden grid $D = H_1 \times H_2 \subset \mathbb{F}_q \times \mathbb{F}_q$, a small collection of maps $\pi_1, \dots, \pi_t$, and a fixed low-degree polynomial C. Each π_i is a multiplicative shift in one coordinate that preserves D. Given a table A on D, define*

$$B_A(x, y) = C\big(x, y, A(\pi_1(x, y)), \dots, A(\pi_t(x, y))\big).$$

We say that the instance is satisfiable if there exists A for which $B_A(x, y) = 0$ for every $(x, y) \in D$.

Adapting the reduction of Ben-Sasson and Sudan [BS08] we can reduce 3SAT to the Zero-on-Subcube problem with $t = O(\log n)$.

We first note the two features of Zero-on-Subcube that make it compatible with our private codes. Firstly, every π_i is a multiplicative shift, and by our choices of H_i and S_i these shifts preserve both D and $\Omega = S_1 \times S_2$. Consequently, all the values needed to compute $B_A(x, y)$, for $(x, y) \in \Omega$, lie on the row and column through (x, y) in Ω. A gate can obtain them by decoding these two lines and therefore needs only $O(|S_i|) = O(\sqrt{N})$ inputs. Second, the same rule that defines B_A on the hidden grid also defines a pointwise low-degree representation of B_A on the disjoint visible grid Ω. Thus the circuit can compute a (slightly higher degree) encoding of B_A from $\mathrm{Enc}(A)$. In contrast, the usual implementation of sumcheck-based PCPs checks a zero claim using mixed-grid evaluations such as $H_1 \times S_2$ and $S_1 \times H_2$, which would break privacy.

Fix a table A and let $B = B_A$; we show how to compute whether A satisfies the zero-on-subcube instance. We start by computing an encoding of B from $\mathrm{Enc}(A)$. Then we want to check whether $\mathrm{Enc}(B)|_D$ is the zero table. We could do this using the AND circuit from the previous subsection. Instead, we give a simpler construction using auxiliary proofs. We use the Combinatorial Nullstellensatz which says that every low-degree polynomial Q vanishing on $D = H_1 \times H_2$ can be written as

$$Q(x, y) = Z_1(x)P_1(x, y) + Z_2(y)P_2(x, y), \qquad Z_i(z) = \prod_{h \in H_i} (z - h).$$

The polynomials P_1, P_2 form a certificate whose defining identity can be checked pointwise on Ω.

The simplest approach would therefore be to ask the prover for a certificate that B vanishes. This does not fit our private circuit model though: such a certificate exists only when A satisfies the instance, whereas the prescribed circuit transcript must exist for every A. We modify the construction as follows.

1. Starting from $\mathrm{Enc}(A)$, the circuit computes $\mathrm{Enc}(B)$ which lies in a slightly higher-degree code.
2. The prover supplies an encoding of the table $J(z) = \begin{cases} B(z)^{-1}, & B(z) \neq 0, \\ 0, & B(z) = 0. \end{cases}$

3. Using the multiplication primitive from Section 2.1, the circuit computes the encodings of the tables,

$$U = B \odot J \qquad \text{and} \qquad V = B \odot (1 - U).$$

4. The prover supplies a Nullstellensatz certificate (V_1, V_2) that V vanishes on D, and the circuit checks the identity $V = Z_1 V_1 + Z_2 V_2$ on Ω.

5. The circuit from the previous subsection computes $\sigma = \mathrm{SUM}(U)$ over D, and outputs $\mathbb{1}[\sigma = 0]$.

6. Every encoded table is checked using the row–column test from the previous subsections.

For the honest choice of J, the table U indicates exactly the violated constraints, while V vanishes everywhere. Thus a certificate for V exists for every A. Conversely, once the certificate enforces $V = 0$, this conclusion no longer depends on J being chosen honestly: where $B = 0$, we have $U = BJ = 0$, while where $B \neq 0$, the identity $V = B(1 - U) = 0$ forces $U = 1$. The final sum is therefore the number of violated constraints. Because the field is larger than the number of constraints, this sum is zero exactly when there are no violations.

Privacy: Privacy of each layer follows from the privacy of the underlying encoding.[4] Privacy of views that touch multiple layers turns out to be much more subtle though and boils down to showing privacy of a random nullstellensatz certificate. Using techniques from [DLV24, KP25], in Section 4.3 we first derive a robust-testability statement from product expansion and then use it to prove joint privacy in Sections 6.2 and 6.3.

Robustness: As in the previous subsections, every encoded table is checked by the row–column test. The circuit's tester chooses uniformly among these tests and the other checks in the construction. Tensor-code soundness bounds the distance of each tested table from its code in terms of the tester's rejection probability. Thus, if only a small fraction of gates are faulty and the tester's rejection probability is sufficiently small, each layer remains close to the intended low-degree table, and the constraint identity, vanishing certificate, and final SUM force the correct output. We give the formal argument in the robustness section.

Reducing the fan-in: We have described the two-dimensional construction. Replacing $D = H_1 \times H_2$ and $\Omega = S_1 \times S_2$ by $D = \prod_{i=1}^{k} H_i$ and $\Omega = \prod_{i=1}^{k} S_i$, and rows and columns by axis-parallel lines, gives depth $O(k)$ and fan-in $O(kN^{1/k})$. The privacy proof uses the corresponding higher-dimensional product-expansion statement.

Private PCP: The PCP proof is simply the transcript of this circuit, including its auxiliary words. With probability $1/3$ each, the verifier checks the final output bit, checks a random computation gate in a random layer, or runs the circuit tester. Honest correctness gives perfect completeness, while robustness gives soundness. The natural privacy measure first chooses a circuit layer and then a wire within that layer; because the layers have different sizes, the resulting PCP is lopsided. For the non-lopsided perfect-zero-knowledge variant we use a different proof: on a satisfying witness, B already vanishes on D, so the prover can supply a random Nullstellensatz certificate for B directly.

[4] The formal construction rerandomizes the representation of each layer after each algebraic operation, which we have suppressed here for simplicity.

2.4 2D-Product Expansion of Punctured Reed Solomon Codes

We give the proof overview of Theorem 2.1 in this section. We want to show that whenever $M = A + B$, where the columns of A lie in C_1 and the rows of B lie in C_2, there is another such decomposition whose weighted number of nonzero rows and columns is $O(\mathrm{wt}(M))$. A standard argument from [PS94] reduces this to proving a divisibility lemma.

Lemma 2.3 (Informal Lemma 4.13)**.** *Let $S_i \leqslant \mathbb{F}_q^\times$ have order n_i, with $\gcd(n_1, n_2) = 1$, and set*

$$h = \min\{n_1 - k_1, n_2 - k_2\}, \qquad N = n_1 n_2.$$

Suppose that $\deg_X E, \deg_Y E \leqslant t$, where $t \leqslant ch^3/N$ for a sufficiently small constant c, and that $\deg_X P < k_1 + t$ and $\deg_Y P < k_2 + t$. If

$$P(X, \beta) = E(X, \beta)A_\beta(X) \quad (\forall \beta \in S_2), \qquad P(\alpha, Y) = E(\alpha, Y)B_\alpha(Y) \quad (\forall \alpha \in S_1),$$

then $E \mid P$.

For simplicity, assume that E is irreducible and depends on both variables, and write $a = \deg_X E$ and $b = \deg_Y E$. Suppose toward a contradiction that $E \nmid P$.

Let Z_1 be the points on $E = 0$ whose x-coordinate lies in S_1. For each $\alpha \in S_1$, the polynomial $E(\alpha, Y)$ has b roots, counted with multiplicity, so $|Z_1| = n_1 b$. The row identities imply that P vanishes at all these points. Similarly, the column identities force P to vanish on a set Z_2 of size $n_2 a$, consisting of the points whose y-coordinate lies in S_2.

Let $\Gamma_E = |Z_1 \cap Z_2|$, counting multiplicities. Since $E \nmid P$, Bézout's theorem gives

$$|Z_1 \cup Z_2| \leqslant (k_1 + t)b + (k_2 + t)a.$$

Inclusion–exclusion therefore gives

$$n_1 b + n_2 a - \Gamma_E = |Z_1 \cup Z_2| \leqslant (k_1 + t)b + (k_2 + t)a,$$

and hence $\Gamma_E \geqslant (h - 2t)(a + b)$.

Since $S_i = \{z \in \mathbb{F}_q^\times : z^{n_i} = 1\}$, the points counted by Γ_E are exactly the common solutions to

$$E(x, y) = 0, \qquad x^{n_1} = 1, \qquad y^{n_2} = 1.$$

Set $u = x^{n_1}$ and $v = y^{n_2}$. If u and v are *multiplicatively independent* on $E = 0$, Corvaja–Zannier [CZ13] gives

$$\Gamma_E \leqslant O\Big(N^{1/3}(ab)^{2/3} + ab\Big).$$

Since $a, b \leqslant t \leqslant ch^3/N$, this is smaller than $(h - 2t)(a + b)$ for sufficiently small c, a contradiction.

Suppose instead that u and v are *multiplicatively dependent*, which means that on the curve $E = 0$, u, v satisfy the relation $u^r v^s = c$ for some coprime integers r, s; in particular each common solution has $c = 1$. This equation has a one-variable parametrization $u = w^s, v = w^{-r}$. Indeed, choose ℓ, m with $\ell s - mr = 1$ and set $w = u^\ell v^m$. Here w is a rational function on the one-dimensional curve $E = 0$. The common solutions $u = v = 1$ are exactly the solutions to $w = 1$, so their number, counted with multiplicity, is at most $\deg w$.

It remains to bound this degree. Since $u = w^s$ and $v = w^{-r}$,

$$\deg u = |s|\deg w, \qquad \deg v = |r|\deg w.$$

Moreover, $\deg x = b$, since for a generic α the equation $E(\alpha, Y) = 0$ has b solutions; similarly, $\deg y = a$. Therefore, $\deg u = n_1 b, \deg v = n_2 a$, thus $\deg w$ divides both $n_1 b$ and $n_2 a$, and

$$\Gamma_E \leqslant \deg w \leqslant \gcd(n_1 b, n_2 a) \leqslant ab,$$

where the last inequality uses $\gcd(n_1, n_2) = 1$. Since $a, b \leqslant t \ll h$, this again contradicts the lower bound. For example, on $E(X, Y) = XY - 1$ one may take $w = x$.

We conclude that $E \mid P$. The general case follows by applying the same argument to every irreducible factor of E. To formalize this argument, one passes to the 'projective closure' of $E = 0$ and then to its 'normalization' to justify the degree and multiplicity counts and to apply Corvaja–Zannier.

2.5 Discussion and Future Directions

Lopsidedness of Private PCPs. The private PCP obtained from the full circuit transcript uses a lopsided measure, giving comparable weight to circuit layers of different sizes. This arises from the stronger privacy guarantee, which compares all witnesses, whether or not they satisfy the instance: the final output bit distinguishes witnesses on which the circuit has different outputs. If privacy is required only among satisfying witnesses, the output is fixed, and the construction in Section 8 gives a non-lopsided perfect-zero-knowledge PCP with essentially the same parameters. The resulting simulator is not efficient. A quantum PCP based on the ABN proposal would suffer from the same lopsidedness. Thus to get a uniform quantum PCP, one would need to figure out how to get quantum gap while breaking the quantum encoding.

Improving Parameters to Polylogarithmic. Our private PCP construction is based on tensor-products of Reed-Solomon codes so that for all k we get fan-in to $O(kN^{1/k})$, but the current privacy and robustness bounds can at best be $\exp(-O(k))$. Thus using this construction we can hope to get a private-PCP with all parameters being $n^{o(1)} = 2^{\log n^{1-\Omega(1)}}$, instead of polylogarithmic. A possible alternative is to recursively apply the two-dimensional construction, following Ben-Sasson and Sudan [BS08]. The difficulty arises even in a single recursion step, where the proof exposes a bivariate polynomial on both a grid and a curve. We do not know how to prove privacy for their joint local views. Resolving this question could lead to the two-dimensional construction to be applied recursively like in [BS08], giving better parameters and also potentially removing the need for the high-dimensional product expansion conjecture.

Universal Encoding and Low-communication MPC. Our present encoding depends on the circuit being computed. This suffices for constructing a PCP for NP: we first reduce Circuit-SAT to the Zero-on-Subcube problem and tailor the encoding and robust circuit to the resulting instance. Ideally, one would instead construct a universal encoding, independent of the circuit, together with a compiler that transforms any circuit into a private and fault-tolerant circuit that computes directly on this global encoding.

Such a construction could also lead to new MPC protocols. Existing approaches based on bit-by-bit encodings require large communication. Computing directly on a single global encoding could potentially reduce this communication to sublinear.

Concurrent work. This work was done independently of the recent related works [ABG+26, BGV26].

Organization. Section 3 reviews the connection between quantum codes and secret sharing, as well as the necessary facts about Reed–Solomon codes. Section 4 develops the product-expansion framework and derives its consequences for relative dual distance and relative robust testability, needed for showing the privacy of our construction. Section 5 presents the formal circuit model we work with, and then constructs a private and robust circuit for Zero-on-Subcube. The privacy and robustness are proved in Sections 6 and 7. Finally, Section 8 constructs private and perfect-zero-knowledge PCPs for NP.

3 Preliminaries

Equivalence between Quantum Codes and Secret Sharing:

Definition 3.1 (CSS codes)**.** *Let $\mathcal{K} \subsetneq \mathcal{C} \subseteq \mathbb{F}_q^\Omega$. The associated CSS code has $C_X = \mathcal{C}$ and $C_Z^\perp = \mathcal{K}$. Its computational-basis codewords are the coset states*

$$|x + \mathcal{K}\rangle = \frac{1}{\sqrt{|\mathcal{K}|}} \sum_{z \in \mathcal{K}} |x + z\rangle\,, \qquad x + \mathcal{K} \in \mathcal{C}/\mathcal{K}.$$

Its Z-distance is

$$d_Z = \min\{\mathrm{wt}(M) : M \in C_Z \setminus C_X^\perp\} = \min\{\mathrm{wt}(M) : M \in \mathcal{K}^\perp \setminus \mathcal{C}^\perp\}. \tag{1}$$

By [RZW26, Corollary 3.2], sampling a uniform element of each coset of $\mathcal{K}$ in $\mathcal{C}$ has privacy threshold

$$d_Z - 1 = \min\{\mathrm{wt}(M) : M \in \mathcal{K}^\perp \setminus \mathcal{C}^\perp\} - 1. \tag{2}$$

Equivalently, every view of fewer than d_Z coordinates is independent of the encoded coset/message, and some view of d_Z coordinates is not. We will therefore only lower-bound $\min\{\mathrm{wt}(M) : M \in \mathcal{K}^\perp \setminus \mathcal{C}^\perp\}$ to bound the privacy of our randomized encoding.

Reed Solomon Codes: For a finite set $S \subseteq \mathbb{F}_q$ and an integer a with $0 \leqslant a \leqslant |S|$, define $\mathrm{RS}(S, a) = \{(f(\alpha))_{\alpha \in S} : f \in \mathbb{F}_q[X],\ \deg f < a\}$.

Fact 3.2 (Dual of a Reed–Solomon code)**.** *Let $S \subseteq \mathbb{F}_q$ have size n, and let $0 \leqslant a \leqslant n$. For each $\alpha \in S$, define $\lambda_\alpha = \left(\prod_{\beta \in S \setminus \{\alpha\}} (\alpha - \beta)\right)^{-1}$. Then*

$$\mathrm{RS}(S, a)^\perp = \mathrm{diag}((\lambda_\alpha)_{\alpha \in S})\, \mathrm{RS}(S, n - a).$$

In particular, the dual is a generalized Reed–Solomon code of dimension $n - a$ and, when $a < n$, minimum distance $a + 1$. More generally, for every invertible diagonal matrix D, $(D\,\mathrm{RS}(S, a))^\perp = D^{-1}\,\mathrm{RS}(S, a)^\perp$.

Fact 3.3 (Vanishing on a Cartesian grid [Alo99])**.** *Let $H_1, \ldots, H_k \subseteq \mathbb{F}_q$ be finite sets, put*

$$D = \prod_{i=1}^{k} H_i, \qquad Z_i(X_i) = \prod_{a \in H_i} (X_i - a),$$

and write $h_i = |H_i|$. For a degree profile $\mathbf{a} = (a_1, \ldots, a_k)$ satisfying $a_i \geqslant h_i$ for every i, let

$$\mathcal{Q}_i(\mathbf{a}) = \{Q \in \mathbb{F}_q[X_1, \ldots, X_k] : \deg_{X_i} Q \leqslant a_i - h_i,\ \deg_{X_j} Q \leqslant a_j \text{ for } j \neq i\}.$$

A polynomial $R \in \mathbb{F}_q[X_1, \ldots, X_k]$ with degree profile $\mathbf{a}$ vanishes on D if and only if $R = \sum_{i=1}^k Z_i Q_i$ where each $Q_i \in \mathcal{Q}_i(\mathbf{a})$.

4 Product Expansion and Implications to Quantum Codes

In this section we show that product expansion implies the privacy of tensor-product based codes, and also a robust testability type of statement which will be needed to prove the joint privacy of our circuit. We start with the general setup.

General Setup: Fix $k \geqslant 2$, finite sets $S_1, \dots, S_k$, and write

$$s_i = |S_i|, \qquad \Omega = \prod_{i=1}^{k} S_i, \qquad N = |\Omega| = \prod_{i=1}^{k} s_i.$$

An i-axis line is obtained by fixing every coordinate except the ith. For a code $C_i \subseteq \mathbb{F}_q^{S_i}$, let

$$L_i(C_i) = \mathbb{F}_q^{S_1} \otimes \cdots \otimes \mathbb{F}_q^{S_{i-1}} \otimes C_i \otimes \mathbb{F}_q^{S_{i+1}} \otimes \cdots \otimes \mathbb{F}_q^{S_k}.$$

For $M_i \in L_i(C_i)$, let $\ell_i(M_i)$ be the number of nonzero i-axis lines. We use the same notation on a subproduct of Ω; the ambient coordinate set will always be clear.

Definition 4.1 (Product expansion)**.** *A tuple $(C_i)_{i\in I}$, where $I \subseteq [k]$, is ρ-product-expanding if every $M \in \sum_{i\in I} L_i(C_i)$ has a decomposition $M = \sum_{i\in I} M_i$, with $M_i \in L_i(C_i)$, such that*

$$\mathrm{wt}(M) \geqslant \rho \sum_{i\in I} s_i \ell_i(M_i).$$

We now fix nested codes

$$0 \neq K_i \subsetneq V_i \subseteq \mathbb{F}_q^{S_i}, \qquad U_i = V_i^{\perp} \subsetneq W_i = K_i^{\perp}.$$

Define

$$\mathcal{C} = \bigotimes_{i=1}^{k} V_i, \qquad E_i = K_i \otimes \bigotimes_{\substack{j=1 \\ j\neq i}}^{k} V_j, \qquad \mathcal{K} = \sum_{i=1}^{k} E_i.^{5}$$

Put

$$\mathcal{B} = \mathcal{C}^{\perp}, \qquad \mathcal{D}_i = E_i^{\perp}, \qquad \mathcal{D} = \bigcap_{i=1}^{k} \mathcal{D}_i = \mathcal{K}^{\perp}.$$

The standard tensor-product dual identities give

$$\mathcal{B} = \sum_{i=1}^{k} L_i(U_i), \qquad \mathcal{D}_i = L_i(W_i) + \sum_{\substack{j=1 \\ j\neq i}}^{k} L_j(U_j) = \mathcal{B} + L_i(W_i), \qquad \mathcal{D} = \mathcal{B} + \bigotimes_{i=1}^{k} W_i. \tag{3}$$

For a coordinate subtuple, we form the corresponding spaces from the restricted pairs (K_i, V_i) and use the same notation; no extra subtuple index will be needed. For the nested CSS code with

$$C_X = \mathcal{C}, \qquad C_Z^{\perp} = \mathcal{K}$$

we have $C_X^{\perp} = \mathcal{B}$, $C_Z = \mathcal{D}$, so the relative weight to be bounded is

$$d_Z = \min\{\mathrm{wt}(M) : M \in \mathcal{D} \setminus \mathcal{B}\}.$$

Sections 4.2 and 4.3 derive relative distance and robust testability (with respect to the spaces $\mathcal{D}_i$) from product expansion. We first state the corollaries of these general statements to our setting of punctured RS codes.

4.1 Punctured Reed–Solomon Codes

Let $H_i \subseteq \mathbb{F}_q$ have size $h_i \geqslant 1$, with $H_i \cap S_i = \varnothing$, and put $D = \prod_i H_i$. Fix an individual degree profile $\mathbf{a} = (a_1, \ldots, a_k)$ with $h_i \leqslant a_i < s_i$, and define

$$\begin{aligned} V_i &= \mathrm{RS}(S_i, a_i + 1), \\ K_i &= \{\mathrm{ev}_{S_i}(p) : \deg p \leqslant a_i,\ p|_{H_i} = 0\}, \\ U_i &= V_i^{\perp}, \qquad W_i = K_i^{\perp}. \end{aligned}$$

Let $\mathcal{P}_{\mathbf{a}}$ be the polynomials of individual degree at most $\mathbf{a}$, and let $\mathcal{Z}_{\mathbf{a}} \subseteq \mathcal{P}_{\mathbf{a}}$ be those that vanish on D. The polynomially defined CSS code fits the generic setup above:

$$C_X = \mathrm{ev}_\Omega(\mathcal{P}_{\mathbf{a}}) = \bigotimes_{i=1}^{k} V_i = \mathcal{C}, \qquad C_Z^{\perp} = \mathrm{ev}_\Omega(\mathcal{Z}_{\mathbf{a}}) = \sum_{i=1}^{k} K_i \otimes \bigotimes_{\substack{j=1 \\ j \neq i}}^{k} V_j = \mathcal{K}. \tag{4}$$

Indeed, the first identity follows by multivariate interpolation, and the second follows from Fact 3.3. Because evaluation on Ω is injective, $\mathcal{K}$ is the kernel of the map $\mathcal{C} \to \mathbb{F}_q^D$ sending $\mathrm{ev}_\Omega(P)$ to $P|_D$. This map is surjective by interpolation on D, since $a_i \geqslant h_i - 1$. Consequently, it identifies the cosets of $\mathcal{K}$ in $\mathcal{C}$ with tables on D.

Writing $Z_i(X) = \prod_{a \in H_i}(X - a)$, we have

$$K_i = \mathrm{diag}(Z_i|_{S_i})\,\mathrm{RS}(S_i, a_i - h_i + 1).$$

Hence $U_i \subseteq W_i$, and Fact 3.2 gives generalized Reed–Solomon codes, up to diagonal rescaling, with

$$\dim U_i = s_i - a_i - 1, \qquad \dim W_i = s_i - a_i + h_i - 1, \qquad d(W_i) = a_i - h_i + 2. \tag{5}$$

Conjecture 4.2 (Product Expansion of Punctured Reed–Solomon Codes)**.** *For every $\eta \in (0, 1)$, balance parameter $\beta \geqslant 1$, and integer $r \geqslant 2$, there is $\rho = \rho(r, \eta, \beta) \in (0, 1]$ such that the following holds. Let q be prime, and let $S_1, \ldots, S_r \leqslant \mathbb{F}_q^{\times}$ be multiplicative subgroups of pairwise coprime orders $s_i = |S_i|$ satisfying*

$$\beta^{-1} \leqslant \frac{s_i}{s_j} \leqslant \beta \qquad (i, j \in [r]).$$

For integers t_i with $1 \leqslant t_i \leqslant (1 - \eta)s_i$, and invertible diagonal matrices Λ_i indexed by S_i, the tuple

$$\big(\Lambda_i\,\mathrm{RS}(S_i, t_i)\big)_{i=1}^{r}$$

is ρ-product-expanding.

Note that, without loss of generality, we may assume that $\rho(r, \eta, \beta)$ is nonincreasing in r. Now to bound the privacy of our codes, it suffices to bound the relative dual distance (see (2) from Section 3),

Corollary 4.3 (RS relative dual distance)**.** *If $k > 2$, assume Conjecture 4.2. Suppose $a_i - h_i \geqslant \eta s_i$ and $a_i + 1 < s_i$ for every i, and suppose the sets S_i satisfy the hypotheses of Conjecture 4.2 with balance parameter β. Put $\rho = \rho(k, \eta, \beta)$. Then*

$$\mathrm{wt}(M) \geqslant \eta\rho^{k-1}N \qquad (M \in \mathcal{K}^{\perp} \setminus \mathcal{C}^{\perp}).$$

Proof. Equation (5) shows that every word in $W_i \setminus U_i$ has weight at least $d(W_i) \geqslant \eta s_i$. Moreover,

$$1 \leqslant \dim U_i \leqslant \dim W_i \leqslant (1-\eta)s_i.$$

Every tuple $(W_t, U_{t+1}, \dots, U_k)$ is therefore a tuple of nonzero generalized Reed–Solomon codes of rate at most $1-\eta$, with the required balance and coprimality. Every required subtuple of arity $r \leqslant k$ is $\rho(r,\eta,\beta)$-product-expanding: for $r = 2$, this follows from the two-dimensional theorem, and for $r > 2$, it follows from Conjecture 4.2. Since $\rho(r,\eta,\beta)$ is nonincreasing in r, every such subtuple is $\rho(k,\eta,\beta) = \rho$-product-expanding. The relative-distance claim follows from Lemma 4.6. □

The same verification in the preceding proof allows us to apply Corollary 4.10 to get,

Corollary 4.4 (RS relative robust testability)**.** *Under the hypotheses of Corollary 4.3, every $M \in \mathbb{F}_q^{\Omega}$ satisfies*

$$\mathit{dist}(M, \mathcal{K}^{\perp}) \leqslant C(\rho,k) \sum_{i=1}^{k} \mathit{dist}(M, \mathcal{D}_i), \qquad C(\rho,k) = \exp\bigl(O\bigl(k^3 + k^2 \log(1/\rho)\bigr)\bigr),$$

where $\rho = \rho(k,\eta,\beta)$.

Finally, we state the implications of our results to quantum codes (with the parameter setting from Lemma 5.4),

Corollary 4.5 (Quantum codes from punctured tensor Reed–Solomon codes)**.** *Assume Conjecture 4.2, and write $\rho(k) = \rho(k, 1/128, 4)$, decreasing this product-expansion constant if necessary so that $\rho(k) \leqslant 1/128$.*

For every $n \geqslant 1$ and every integer $k = k(n) \geqslant 2$, there are integers ℓ, N and a prime q satisfying

$$n < \ell < N < q \leqslant N^6, \qquad N \leqslant n \log n \exp(O(k \log k)), \qquad \frac{\ell}{N} \geqslant \exp(-O(k \log k)).$$

Specialize the construction in Section 4.1 to $a_i = d_i$, and put

$$\mathcal{C} = \bigotimes_{i=1}^{k} V_i, \qquad \mathcal{K} = \sum_{i=1}^{k} K_i \otimes \bigotimes_{j \neq i} V_j, \qquad \widetilde{\mathcal{K}} = \sum_{i=1}^{k} L_i(K_i).$$

Here an X-check means a parity check for C_X, and similarly for Z.

1. *The CSS code $C_X = \mathcal{C}$, $C_Z = \mathcal{K}^{\perp}$ encodes ℓ qudits. Its relative X-distance is at least $(3/4)^k$, and its relative Z-distance is at least $\rho(k)^k$. Its X-checks have weight at most $4N^{1/k}$, whereas every nonzero Z-check has weight at least $(3/4)^k N$.*
2. *The subsystem CSS code $C_X = \mathcal{C}$, $C_Z = \widetilde{\mathcal{K}}^{\perp}$ encodes ℓ protected qudits. Both dressed relative distances are at least $\rho(k)^k$, and both its X- and Z-gauge checks have weight at most $4N^{1/k}$.*

Moreover, the CSS code supports transversal CCZ on all ℓ logical qudits. The subsystem code supports the same transversal gate after fixing all gauge qudits to $|0\rangle$ in the Z basis (see Definition 3.33 in [GG25]).

Proof. Apply Lemma 5.4 and take $a_i = d_i$. By (4), restriction to D identifies $\mathcal{C}/\mathcal{K}$ with $\mathbb{F}_q^D$. Thus $\ell = |D| = \prod_i h_i$. In the construction of Lemma 5.4, $s_i = b_i h_i$ with $b_i \leqslant Ck\log(k+1)$, so $\ell/N = \prod_i b_i^{-1} \geqslant (Ck\log(k+1))^{-k}$. The remaining bounds on ℓ, N, and q are those in Lemma 5.4.

For the CSS code, the logical space is $\mathcal{C}/\mathcal{K}$, so its dimension is ℓ. By (21), $d_i + 1 \leqslant v_i + 1 \leqslant s_i/4$. Therefore

$$d_X \geqslant d(\mathcal{C}) = \prod_i (s_i - d_i) \geqslant (3/4)^k N.$$

By Corollary 4.3,

$$d_Z = \min\{\mathrm{wt}(c) : c \in \mathcal{K}^\perp \setminus \mathcal{C}^\perp\} \geqslant \frac{1}{128}\rho(k)^{k-1} N \geqslant \rho(k)^k N.$$

Equation (3) gives $\mathcal{C}^\perp = \sum_i L_i(V_i^\perp)$. Hence the X-check space has generators supported on individual axis lines, each of weight at most $\max_i s_i \leqslant 4N^{1/k}$. The Z-check space is $\mathcal{K}$. Since $\mathcal{K} \subseteq \mathcal{C}$, every nonzero Z-check has weight at least $d(\mathcal{C}) \geqslant (3/4)^k N$.

For the subsystem code, the standard tensor-quotient identity gives $\mathcal{C} \cap \widetilde{\mathcal{K}} = \mathcal{K}$ and $\widetilde{\mathcal{K}}^\perp = \bigotimes_i K_i^\perp$. Thus its protected logical space is again $\mathcal{C}/\mathcal{K}$. Equivalently, in the convention of [GG25], this is the subsystem product of $Q^i = (K_i^\perp, V_i)$, and [GG25, Proposition 3.23] gives dimension $\prod_i h_i = \ell$.

Its dressed Z-distance equals the ordinary Z-distance, since

$$\widetilde{\mathcal{K}}^\perp + \mathcal{C}^\perp = (\widetilde{\mathcal{K}} \cap \mathcal{C})^\perp = \mathcal{K}^\perp.$$

Thus Corollary 4.3 gives $d_Z^{\mathrm{sub}} \geqslant \rho(k)^k N$.

For the dressed X-distance, the parameter ρ_Z^i in [GG25, Theorem 4.3] is the product expansion of $(K_1, \ldots, K_{i-1}, V_i)$. For $i \geqslant 2$, every member of this tuple is a nonzero scaled Reed–Solomon code and

$$1 \leqslant \dim K_j \leqslant \dim V_j = d_j + 1 \leqslant v_j + 1 \leqslant s_j/4.$$

By Lemma 5.4, the evaluation domains have pairwise coprime, 4-balanced orders. Hence Conjecture 4.2 gives product expansion at least $\rho(i, 1/128, 4) \geqslant \rho(k)$. For $i = 1$, the product expansion of V_1 is its relative distance, which is at least $3/4$. Therefore [GG25, Lemma 4.9] gives $d_X^{\mathrm{sub}} \geqslant \rho(k)^k N$.

The two gauge-check spaces are $\mathcal{C}^\perp = \sum_i L_i(V_i^\perp)$ and $\widetilde{\mathcal{K}} = \sum_i L_i(K_i)$. Both have generators supported on individual axis lines and therefore of weight at most $4N^{1/k}$.

Finally, (21) in Lemma 5.4 gives $3d_i < s_i$. Evaluation on S_i is therefore injective for polynomials of degree at most $3d_i$, so there are coefficients $\lambda_{i,x} \in \mathbb{F}_q$, for $x \in S_i$, such that $\sum_{a \in H_i} f(a) = \sum_{x \in S_i} \lambda_{i,x} f(x)$ for every such polynomial f. Set $\lambda_x = \prod_i \lambda_{i,x_i}$ for $x = (x_1, \ldots, x_k) \in \Omega$. Then, for all $P_1, P_2, P_3 \in \mathcal{P}_{\mathbf{d}}$,

$$\sum_{x \in \Omega} \lambda_x P_1(x) P_2(x) P_3(x) = \sum_{a \in D} P_1(a) P_2(a) P_3(a).$$

Under the identification $\mathcal{C}/\mathcal{K} \cong \mathbb{F}_q^D$, this says that the physical transversal layer induces coordinatewise logical CCZ on all ℓ logical qudits, by Lemma 3.34 of [GG25].

For the subsystem code, $\mathcal{C} \cap \widetilde{\mathcal{K}} = \mathcal{K}$. Fixing the gauge qudits to their Z-basis zero state therefore gives the ordinary CSS codespace spanned by the states $|c + \mathcal{K}\rangle$, $c \in \mathcal{C}$. The same identity implements coordinatewise logical CCZ on all ℓ protected logical qudits in this gauge-fixed subspace(see Definition 3.33 and Lemma 3.34 in [GG25])

□

4.2 Relative Dual Distance from Product Expansion

Here we derive the dual distance from product expansion.

For $i \in [k]$, put

$$\delta_i = \frac{\min\{\mathrm{wt}(w) : w \in W_i \setminus U_i\}}{s_i}, \qquad \delta = \min_{i \in [k]} \delta_i.$$

Lemma 4.6 (Relative dual distance)**.** *Suppose that, for every $t < k$, the tuple $(W_t, U_{t+1}, \ldots, U_k)$ is ρ-product-expanding. Then*

$$\mathrm{wt}(M) \geqslant \delta \rho^{k-1} N \qquad (M \in \mathcal{K}^\perp \setminus \mathcal{C}^\perp).$$

Proof. For $t \in [k]$, define the suffix codes

$$\mathcal{C}_t = \bigotimes_{i=t}^{k} V_i.$$

Define $\mathcal{K}_t$ recursively by

$$\mathcal{K}_k = K_k, \qquad \mathcal{K}_t = K_t \otimes \mathcal{C}_{t+1} + V_t \otimes \mathcal{K}_{t+1} \quad (t < k).$$

Thus $\mathcal{C}_1 = \mathcal{C}$ and $\mathcal{K}_1 = \mathcal{K}$. We prove by backward induction on t that every $M \in \mathcal{K}_t^\perp \setminus \mathcal{C}_t^\perp$ has weight at least

$$\delta \rho^{k-t} \prod_{i=t}^{k} s_i. \tag{6}$$

For $t = k$, such an M belongs to $W_k \setminus U_k$, so the claim follows from the definition of δ.

Now let $t < k$ and take $M \in \mathcal{K}_t^\perp \setminus \mathcal{C}_t^\perp$. Since M is orthogonal to $K_t \otimes \mathcal{C}_{t+1} \subseteq \mathcal{K}_t$,

$$M \in L_t(W_t) + \sum_{i=t+1}^{k} L_i(U_i).$$

Product expansion gives a decomposition

$$M = M_t + \sum_{i=t+1}^{k} M_i, \qquad M_t \in L_t(W_t), \quad M_i \in L_i(U_i),$$

such that

$$\mathrm{wt}(M) \geqslant \rho \sum_{i=t}^{k} s_i \ell_i(M_i) \geqslant \rho s_t \ell_t(M_t). \tag{7}$$

The sum $\sum_{i=t+1}^{k} M_i$ lies in $\mathcal{C}_t^\perp \subseteq \mathcal{K}_t^\perp$. Hence $M_t \in \mathcal{K}_t^\perp$, and $M_t \notin \mathcal{C}_t^\perp$, since otherwise $M \in \mathcal{C}_t^\perp$.

Pure tensors span $\mathcal{C}_t = V_t \otimes \mathcal{C}_{t+1}$, so there are $v \in V_t$ and $P \in \mathcal{C}_{t+1}$ such that $\langle M_t, v \otimes P \rangle \neq 0$. Contract M_t with v:

$$M'(x_{t+1}, \ldots, x_k) = \langle M_t(\,\cdot\,, x_{t+1}, \ldots, x_k), v \rangle.$$

Then $\langle M', P \rangle \neq 0$, so $M' \notin \mathcal{C}_{t+1}^\perp$. For every $R \in \mathcal{K}_{t+1}$, the recursive definition gives $v \otimes R \in \mathcal{K}_t$, and hence

$$\langle M', R \rangle = \langle M_t, v \otimes R \rangle = 0.$$

Thus $M' \in \mathcal{K}_{t+1}^{\perp} \setminus \mathcal{C}_{t+1}^{\perp}$. Also, $\mathrm{wt}(M') \leqslant \ell_t(M_t)$, because a zero t-line of M_t contributes a zero coordinate to M'. The induction hypothesis and (7) now give

$$\ell_t(M_t) \geqslant \mathrm{wt}(M') \geqslant \delta\rho^{k-t-1} \prod_{i=t+1}^{k} s_i,$$

and therefore

$$\mathrm{wt}(M) \geqslant \rho s_t \ell_t(M_t) \geqslant \delta\rho^{k-t} \prod_{i=t}^{k} s_i.$$

This proves (6); the case $t = 1$ proves the lemma. □

4.3 Relative Robust Testability from Product Expansion

We first recall collective product expansion and then prove a collective relative-robust-testability statement by induction. All weights and distances below are unnormalized.

For a finite family $\mathbf{X} = (X^{(a)})_{a\in\mathcal{A}}$ on the same coordinate set, let

$$\mathrm{wt}_{\cup}(\mathbf{X}) = \left| \bigcup_{a\in\mathcal{A}} \mathrm{supp}(X^{(a)}) \right|.$$

For a linear space E, define

$$\mathrm{dist}_{\cup}(\mathbf{X}, E) = \min_{\mathbf{Y}\in E^{\mathcal{A}}} \mathrm{wt}_{\cup}(\mathbf{X} - \mathbf{Y}).$$

For a singleton family, these are the ordinary Hamming weight and distance. For a family $\mathbf{X}_i = (X_i^{(a)})_{a\in\mathcal{A}}$ in $L_i(C_i)$, let $\ell_{i,\cup}(\mathbf{X}_i)$ be the number of i-axis lines on which some $X_i^{(a)}$ is nonzero.

Definition 4.7 (Collective product expansion)**.** *The tuple $(C_1, \dots, C_r)$ is collectively P-product-expanding if every finite family $\mathbf{X} = (X^{(a)})_{a\in\mathcal{A}}$ in $\sum_{i\in[r]} L_i(C_i)$ has decompositions*

$$X^{(a)} = \sum_{i\in[r]} X_i^{(a)}, \qquad X_i^{(a)} \in L_i(C_i),$$

such that

$$\mathrm{wt}_{\cup}(\mathbf{X}) \geqslant P \sum_{i\in[r]} s_i \ell_{i,\cup}\big((X_i^{(a)})_{a\in\mathcal{A}}\big). \tag{CPE}$$

The implication from product expansion to collective product expansion is standard and was also proved in [DLV24, Lemma 5.13 and Proposition 5.14]. We use the sharper quantitative bound obtained from Kalachev and Panteleev [KP25, Lemmas 17 and 18], and include the short proof for completeness.

Lemma 4.8 (Collective PE from PE)**.** *Let $r \geqslant 2$, let $\rho \in (0, 1]$, and let $C_i \subsetneq \mathbb{F}_q^{S_i}$, $i \in [r]$, be proper linear codes. If $(C_1, \dots, C_r)$ is ρ-product-expanding, then it is collectively $P(\rho, r)$-product-expanding, where*

$$P(\rho, r) = \frac{\rho^r}{r(2^r + 1)^r}.$$

Proof. Fix a finite family $\mathbf{X} = (X^{(a)})_{a\in\mathcal{A}}$ in $\sum_{i=1}^{r} L_i(C_i)$, and let $A = \bigcup_{a\in\mathcal{A}} \operatorname{supp}(X^{(a)})$. Starting from A, repeatedly adjoin every i-axis line ℓ that meets the current set in at least ρs_i points, and let $B = [A]_\rho$ be the stabilized set. Thus every i-axis line is either contained in B or meets B in fewer than ρs_i points; such a set is called ρ-closed. Lemma 18 of [KP25] shows that for square grids,

$$|B| \leqslant \left(\frac{2^r+1}{\rho}\right)^r |A|.$$

Their proof extends directly to rectangular grids by replacing the size $n^{|J|}$ of a J-dimensional coordinate flat by $\prod_{j\in J} s_j$, and using the threshold ρs_i on an i-axis line.

By [KP25, Lemma 17], ordinary ρ-product expansion implies that every ρ-closed set has what they call "inner generation": for every $a \in \mathcal{A}$, there is a decomposition

$$X^{(a)} = \sum_{i=1}^{r} X_i^{(a)}, \qquad X_i^{(a)} \in L_i(C_i),$$

such that $X_i^{(a)}$ is a sum of C_i-codewords placed on i-axis lines contained in B.

Let $\mathbf{X}_i = (X_i^{(a)})_{a\in\mathcal{A}}$. Every line counted by $\ell_{i,\cup}(\mathbf{X}_i)$ is contained in B. Since distinct i-axis lines are disjoint and each has s_i points,

$$s_i \ell_{i,\cup}(\mathbf{X}_i) \leqslant |B|.$$

Thus,

$$\sum_{i=1}^{r} s_i \ell_{i,\cup}(\mathbf{X}_i) \leqslant r|B| \leqslant r\left(\frac{2^r+1}{\rho}\right)^r |A| = \frac{1}{P(\rho,r)} \mathrm{wt}_\cup(\mathbf{X}),$$

which is precisely collective $P(\rho, r)$-product expansion. □

Lemma 4.9 (Collective relative robust testability)**.** *Suppose every mixed subtuple $(C_i)_{i\in I}$, with $C_i \in \{U_i, W_i\}$ is collectively P-product-expanding. Then there are constants $K_r = K_r(P)$, $1 \leqslant r \leqslant k$, such that*

$$K_1 = 1, \qquad K_r = 1 + \frac{r-1}{P} K_{r-1}, \tag{CR}$$

and the following holds. For every family $\mathbf{M}$, indexed by a finite set $\mathcal{A}$, of words in $\bigotimes_{j=1}^{r} \mathbb{F}_q^{S_j}$,

$$\mathit{dist}_\cup\left(\mathbf{M}, \bigcap_{i=1}^{r} \mathcal{D}_i\right) \leqslant K_r \sum_{i=1}^{r} \mathit{dist}_\cup(\mathbf{M}, \mathcal{D}_i), \tag{8}$$

Our proof uses an induction on r, like the proof of [DLV24, Lemma 5.13 and Proposition 5.14], except that our inductive invariant is different. DLV fix a coboundary x and bound the union block support of a cochain $\widetilde{y}$ satisfying $\delta\widetilde{y} = x$ in terms of the union block support of x. For vector-valued coefficients, an entire face is counted once whenever any component on that face is nonzero. In their setting, PE supplies the top level, while Proposition 5.14 supplies the collective block-support control needed after peeling. Here we fix an ambient family $\mathbf{M}$ and control the union Hamming support of a correction into $\bigcap_i \mathcal{D}_i$. Thus the lemma adapts their dimension-peeling argument, but it is not a formal consequence of their coboundary-expansion theorem.

Proof. We induct on r. The case $r = 1$ is immediate, since there is only one local space.

The anchor and the goal. Suppose $r \geqslant 2$, and set $t_i = \mathrm{dist}_\cup(\mathbf{M}, \mathcal{D}_i)$. Choose $\mathbf{Y}_i \in \mathcal{D}_i^{\mathcal{A}}$ satisfying

$$\mathrm{wt}_\cup(\mathbf{M} - \mathbf{Y}_i) = t_i. \tag{9}$$

We take $\mathbf{Y}_r$ as the anchor. It suffices to construct $\mathbf{Q} \in (\bigcap_{i=1}^r \mathcal{D}_i)^{\mathcal{A}}$ and bound $\mathrm{wt}_\cup(\mathbf{Y}_r - \mathbf{Q})$, since

$$\begin{aligned} \mathrm{dist}_\cup\left(\mathbf{M}, \bigcap_{i=1}^r \mathcal{D}_i\right) &\leqslant \mathrm{wt}_\cup(\mathbf{M} - \mathbf{Q}) \\ &\leqslant t_r + \mathrm{wt}_\cup(\mathbf{Y}_r - \mathbf{Q}). \end{aligned} \tag{10}$$

Pairwise repairs of the anchor. For each $i < r$, we first construct a correction on few r-lines that places the corrected anchor in $\mathcal{D}_i \cap \mathcal{D}_r$. Since $U_j \subseteq W_j$,

$$\mathbf{Y}_i - \mathbf{Y}_r \in \left(L_i(W_i) + L_r(W_r) + \sum_{j \neq i,r} L_j(U_j)\right)^{\mathcal{A}}.$$

Let us now apply collective PE to the family $\mathbf{Y}_i - \mathbf{Y}_r$ with the mixed tuple having W_i and W_r in directions i, r, and U_j in all other directions. Let $\mathbf{X}_i \in L_r(W_r)^{\mathcal{A}}$ be the r-direction component of the resulting decomposition and let λ_i denote $\ell_{r,\cup}(\mathbf{X_i})$, the number of r-lines on which some member of $\mathbf{X}_i$ is nonzero. Then we get,

$$\begin{aligned} s_r \lambda_i &\leqslant P^{-1} \mathrm{wt}_\cup(\mathbf{Y}_i - \mathbf{Y}_r) \\ &\leqslant P^{-1}(t_i + t_r), \end{aligned} \tag{11}$$

where $s_r = |S_r|$ and the second inequality follows from subadditivity of $\mathrm{wt}_\cup$ and (9).

We have that,

$$\mathbf{Y}_r + \mathbf{X}_i \in (\mathcal{D}_i \cap \mathcal{D}_r)^{\mathcal{A}}. \tag{12}$$

Indeed, the remaining components give $\mathbf{Y}_i - \mathbf{Y}_r - \mathbf{X}_i \in \left(L_i(W_i) + \sum_{j \neq i,r} L_j(U_j)\right)^{\mathcal{A}}$ which implies $\mathbf{Y}_r + \mathbf{X}_i \in \mathcal{D}_i^{\mathcal{A}}$ and since $\mathbf{X}_i \in L_r(W_r)^{\mathcal{A}}$, we have $\mathbf{Y}_r + \mathbf{X}_i \in \mathcal{D}_r^{\mathcal{A}}$ putting the family in the required intersection.

Isolating the lower-dimensional word. We now split $\mathbf{Y}_r$ into a part that already lies in every $\mathcal{D}_i$ and a vector-valued word on the first $r-1$ directions that remains to be repaired. Write L_j^- for the lift in the first $r-1$ directions, and put $A^- = \bigotimes_{j<r} \mathbb{F}_q^{S_j}$, $\mathcal{B}^- = \sum_{j<r} L_j^-(U_j)$, and $\mathcal{D}_i^- = \mathcal{B}^- + L_i^-(W_i)$ for $i < r$. These are the spaces in the lemma for the first $r-1$ directions. Since $\mathbf{Y}_r \in \mathcal{D}_r^{\mathcal{A}}$ and $\mathcal{D}_r = (\mathcal{B}^- \otimes \mathbb{F}_q^{S_r}) + (A^- \otimes W_r)$, we may write

$$\mathbf{Y}_r = \mathbf{b} + \mathbf{w}, \qquad \mathbf{b} \in (\mathcal{B}^- \otimes \mathbb{F}_q^{S_r})^{\mathcal{A}}, \qquad \mathbf{w} \in (A^- \otimes W_r)^{\mathcal{A}}. \tag{13}$$

Choose a complement H_r such that $W_r = U_r \oplus H_r$, and a linear projection $p : \mathbb{F}_q^{S_r} \longrightarrow H_r$ onto H_r, whose restriction to W_r has kernel U_r. Let $\Pi = \mathrm{Id}_{A^-} \otimes p$, and set

$$\mathbf{g} = \mathbf{b} + (\mathbf{w} - \Pi\mathbf{w}), \qquad \mathbf{h} = \Pi\mathbf{w}.$$

Since $\mathbf{w} - \Pi\mathbf{w} \in (A^- \otimes U_r)^{\mathcal{A}}$, we have

$$\mathbf{Y}_r = \mathbf{g} + \mathbf{h}, \qquad \mathbf{g} \in \mathcal{B}^{\mathcal{A}} \subseteq \left(\bigcap_{i=1}^r \mathcal{D}_i\right)^{\mathcal{A}}.$$

Thus we keep $\mathbf{g}$ fixed and repair only $\mathbf{h}$.

The lower-dimensional distance bounds. The pairwise corrections now show that the same word $\mathbf{h}$ is close to every lower-dimensional space. Indeed, by (12) and (13), $\mathbf{w} + \mathbf{X}_i = \mathbf{Y}_r + \mathbf{X}_i - \mathbf{b} \in \mathcal{D}_i^{\mathcal{A}}$. Moreover, $\Pi(\mathcal{D}_i) \subseteq \mathcal{D}_i^- \otimes H_r$. Indeed, $\mathcal{D}_i = (\mathcal{D}_i^- \otimes \mathbb{F}_q^{S_r}) + (A^- \otimes U_r)$, and Π kills the second summand. Viewing $A^- \otimes H_r$ as the space of H_r-valued words on $\prod_{j<r} S_j$, we also have $\mathrm{wt}_\cup(\Pi\mathbf{X}_i) \leqslant \lambda_i$, since projection cannot make a zero r-line nonzero. Since $\mathbf{h} + \Pi\mathbf{X}_i \in (\mathcal{D}_i^- \otimes H_r)^{\mathcal{A}}$, it follows that

$$\mathrm{dist}_\cup(\mathbf{h}, \mathcal{D}_i^- \otimes H_r) \leqslant \lambda_i.$$

Applying induction to the H_r-valued family. We now rewrite this vector-valued problem as one scalar-family problem to which the collective induction hypothesis applies. Let $h_1, \ldots, h_d$ be a basis of H_r. For every $a \in \mathcal{A}$, write the vector $h^{(a)} \in \mathbf{h}$ as,

$$h^{(a)} = \sum_{s=1}^{d} m_s^{(a)} \otimes h_s, \qquad (\Pi\mathbf{X}_i)^{(a)} \quad = \sum_{s=1}^{d} x_{i,s}^{(a)} \otimes h_s,$$

where $m_s^{(a)}, x_{i,s}^{(a)} \in A^-$.

Membership in $\mathcal{D}_i^- \otimes H_r$ is coefficientwise: for $v_1, \ldots, v_d \in A^-$,

$$\sum_{s=1}^{d} v_s \otimes h_s \in \mathcal{D}_i^- \otimes H_r \quad \Longleftrightarrow \quad v_s \in \mathcal{D}_i^- \quad \text{for every } s.$$

Consequently, $m_s^{(a)} + x_{i,s}^{(a)} \in \mathcal{D}_i^-$ for every $a \in \mathcal{A}$ and $s \in [d]$. Thus the enlarged family $\left(m_s^{(a)}\right)_{(a,s)\in\mathcal{A}\times[d]}$ has distance at most λ_i from $\mathcal{D}_i^-$. Indeed,

$$\mathrm{dist}_\cup\left(\left(m_s^{(a)}\right)_{a,s}, \mathcal{D}_i^-\right) \leqslant \mathrm{wt}_\cup\left(\left(x_{i,s}^{(a)}\right)_{a,s}\right) = \mathrm{wt}_\cup(\Pi\mathbf{X}_i) \leqslant \lambda_i,$$

where the equality holds because an H_r-valued entry is nonzero exactly when at least one of its basis coefficients is nonzero.

Apply the collective induction hypothesis once to this enlarged family, with family index set $\mathcal{A} \times [d]$. It gives $q_s^{(a)} \in \bigcap_{i<r} \mathcal{D}_i^-$ for every $a \in \mathcal{A}$ and $s \in [d]$, such that

$$\mathrm{wt}_\cup\left(\left(m_s^{(a)} - q_s^{(a)}\right)_{a,s}\right) \leqslant K_{r-1} \sum_{i<r} \lambda_i.$$

For each $a \in \mathcal{A}$, define $\widetilde{h}^{(a)} = \sum_{s=1}^{d} q_s^{(a)} \otimes h_s$, and let $\widetilde{\mathbf{h}} = (\widetilde{h}^{(a)})_{a\in\mathcal{A}}$. Then

$$\widetilde{\mathbf{h}} \in \left(\left(\bigcap_{i<r} \mathcal{D}_i^-\right) \otimes H_r\right)^{\mathcal{A}}. \tag{14}$$

Reassembling the coefficients does not change their union support on $\prod_{j<r} S_j$. Therefore

$$\mathrm{wt}_\cup(\mathbf{h} - \widetilde{\mathbf{h}}) = \mathrm{wt}_\cup\left(\left(m_s^{(a)} - q_s^{(a)}\right)_{a,s}\right) \leqslant K_{r-1} \sum_{i<r} \lambda_i. \tag{15}$$

Here the union weight counts the r-lines on which some member of the H_r-valued difference is nonzero.

Lifting the common repair. We now expand the repaired vector-valued word back into r-lines and add it to the part $\mathbf{g}$ that already lies in every $\mathcal{D}_i$. Define

$$\mathbf{Q} = \mathbf{g} + \widetilde{\mathbf{h}} = \mathbf{b} + (\mathbf{w} - \Pi\mathbf{w}) + \widetilde{\mathbf{h}}. \tag{16}$$

For every $i < r$, (14) gives $\widetilde{\mathbf{h}} \in (\mathcal{D}_i^- \otimes H_r)^{\mathcal{A}} \subseteq \mathcal{D}_i^{\mathcal{A}}$. We also have $\widetilde{\mathbf{h}} \in (A^- \otimes H_r)^{\mathcal{A}} \subseteq L_r(W_r)^{\mathcal{A}} \subseteq \mathcal{D}_r^{\mathcal{A}}$. Since $\mathbf{g} \in \mathcal{B}^{\mathcal{A}}$, it follows that

$$\mathbf{Q} \in \left(\bigcap_{i=1}^{r} \mathcal{D}_i\right)^{\mathcal{A}}. \tag{17}$$

By (16), $\mathbf{Y}_r - \mathbf{Q} = \mathbf{h} - \widetilde{\mathbf{h}}$. Each point $u \in \prod_{j<r} S_j$ at which some member of $\mathbf{h} - \widetilde{\mathbf{h}}$ is nonzero corresponds to one r-line $\{u\} \times S_r$, containing s_r scalar coordinates. Thus, by (15) and (11),

$$\mathrm{wt}_{\cup}(\mathbf{Y}_r - \mathbf{Q}) \leqslant s_r K_{r-1} \sum_{i<r} \lambda_i \leqslant \frac{K_{r-1}}{P} \sum_{i<r} (t_i + t_r).$$

Combining this with (10) yields

$$\begin{aligned} \mathrm{dist}_{\cup}\left(\mathbf{M}, \bigcap_{i=1}^{r} \mathcal{D}_i\right) &\leqslant t_r + \frac{K_{r-1}}{P} \sum_{i<r} (t_i + t_r) \\ &\leqslant \left(1 + \frac{r-1}{P} K_{r-1}\right) \sum_{i=1}^{r} t_i. \end{aligned}$$

This is (8) with the recurrence (CR). □

Corollary 4.10 (Relative robust testability)**.** *Let $\rho \in (0, 1]$. Suppose every mixed subtuple $(C_i)_{i \in I}$, where $C_i \in \{U_i, W_i\}$ and $|I| \geqslant 2$, is ρ-product-expanding. Then every $M \in \mathbb{F}_q^{\Omega}$ satisfies*

$$\mathit{dist}(M, \mathcal{D}) \leqslant C(\rho, k) \sum_{i=1}^{k} \mathit{dist}(M, \mathcal{D}_i), \qquad C(\rho, k) = \exp\left(O\left(k^3 + k^2 \log(1/\rho)\right)\right).$$

Proof. For every mixed subtuple of arity r, where $2 \leqslant r \leqslant k$, Lemma 4.8 gives collective $P(\rho, r)$-product expansion. Since $P(\rho, r) \geqslant P(\rho, k)$, all the required mixed subtuples are collectively P-product-expanding for

$$P = P(\rho, k) = \frac{\rho^k}{k(2^k + 1)^k}.$$

Apply Lemma 4.9 to the singleton family $\{M\}$. Since $P \leqslant 1$, its recurrence gives

$$K_k \leqslant k! P^{-(k-1)} = k! \left(\frac{k(2^k + 1)^k}{\rho^k}\right)^{k-1} = \exp\left(O\left(k^3 + k^2 \log(1/\rho)\right)\right).$$

Taking $C(\rho, k) = K_k$ proves the claim. □

4.4 Proof of Two-Dimensional Product Expansion

In this section we prove product expansion for $r = 2$.

Theorem 4.11 (Two-dimensional product expansion of subgroup Reed–Solomon codes)**.** *Fix $\varepsilon \in (0,1)$ and $K \geqslant 1$. There is a constant $\rho = \Omega(\varepsilon^6/K^3) \in (0,1]$ with the following property. Let q be prime, and let $S_1, S_2 \leqslant \mathbb{F}_q^\times$ be multiplicative subgroups of coprime orders $n_i = |S_i|$ satisfying $K^{-1} \leqslant \frac{n_1}{n_2} \leqslant K$. For integers k_i satisfying $1 \leqslant k_i \leqslant (1-\varepsilon)n_i$, let $C_i = \mathrm{RS}(S_i, k_i)$. Then (C_1, C_2) is ρ-product-expanding. The same conclusion holds after independently rescaling the coordinates of C_1 and C_2 by nonzero field elements.*

We first prove a bivariate divisibility lemma for coprime multiplicative subgroups using an algebraic geometry theorem from the work of [CZ13]. Once this lemma is established, the rest of the proof is along the lines of [PS94].

4.4.1 Algebraic Geometry Preliminaries

We recall the algebraic notation used in the divisibility lemma.

Let $\mathbb{K} = \overline{\mathbb{F}}_q$. The projective line $\mathbf{P}^1(\mathbb{K})$ is the affine line $\mathbb{K}$ together with one point ∞. Thus $\mathbf{P}^1 \times \mathbf{P}^1$ is obtained from $\mathbb{K}^2$ by allowing either coordinate to equal ∞. We use it only to keep track of intersections that escape to infinity.

For $F \in \mathbb{K}[X,Y]$, its *bidegree* is $(\deg_X F, \deg_Y F)$. Fixing $X = \alpha$ gives the $X = \alpha$ coordinate line; the restriction of F to this line is $F(\alpha, Y)$; fixing $Y = \beta$ gives the $Y = \beta$ coordinate line $F(X, \beta)$. In the table indexed by $S_1 \times S_2$, these are a row and a column, respectively.

Suppose that F is irreducible and has bidegree (a,b) with $a, b > 0$. The affine equation $F = 0$ defines a curve in $\mathbb{K}^2$. Its closure in $\mathbf{P}^1 \times \mathbf{P}^1$ adds the points at which one of the coordinates is infinite. An affine point is *singular* if both partial derivatives of F vanish there; otherwise it is nonsingular. A nonsingular point has a single local branch, whereas a singular point may have several branches meeting there.

We work on the *normalization* $\mathcal{X}_F$ of this closure. A nonsingular point has one point of $\mathcal{X}_F$ above it, while the points above a singular point correspond to its different local branches. Thus normalization lets us measure zeros and poles separately along each branch. It gives a smooth projective curve, as required by the Corvaja–Zannier theorem below. The notation $\mathbb{K}(\mathcal{X}_F)$ denotes the field of rational functions on this curve, and $x, y \in \mathbb{K}(\mathcal{X}_F)$ are its two coordinate functions. Thus a point $\nu \in \mathcal{X}_F$ above an affine point (α, β) satisfies $x(\nu) = \alpha$ and $y(\nu) = \beta$.

A generic line $X = \alpha$ meets $F = 0$ in b points, counted with multiplicity, and a generic line $Y = \beta$ meets it in a points. Here generic means all but finitely many choices of the fixed coordinate, and a repeated root is counted according to its multiplicity. Equivalently, $\deg(x) = b$ and $\deg(y) = a$, where the degree of a rational function is its total pole multiplicity.

For a point $\nu \in \mathcal{X}_F$ and a nonzero rational function f, the integer $\mathrm{ord}_\nu(f)$ is the multiplicity of the zero of f at ν; it is negative when f has a pole, and is zero when f is finite and nonzero there. Thus $-\mathrm{ord}_\nu(f)$ is the pole multiplicity when this order is negative. The total zero multiplicity equals the total pole multiplicity, and their common value is $\deg(f)$. In particular, if $H(X,Y)$ has bidegree at most (D_X, D_Y) and does not vanish identically on $\mathcal{X}_F$, then $H(x,y)$ has at most $D_X b + D_Y a$ zeros, counted with multiplicity. This is the form of Bézout's theorem used below. Specifically, in the proof of Lemma 4.13, the assumption $F \nmid P$ ensures that $P(x,y)$ does not vanish identically, and this bound gives at most $(k_1 + e_X)b + (k_2 + e_Y)a$ zeros.

If $\mathrm{char}(\mathbb{K}) \nmid n$, $\alpha^n = 1$, and $x(\nu) = \alpha$, then $T^n - 1$ has a simple root at α, so $\mathrm{ord}_\nu(1 - x^n) = \mathrm{ord}_\nu(x - \alpha)$. The common value can nevertheless exceed 1, because $x - \alpha$ may vanish to higher order along the branch represented by ν.

For positive integers n_1, n_2, define the weighted common-zero count

$$\Gamma_F = \sum_{\nu \in \mathcal{X}_F:\ x(\nu),y(\nu)\in \mathbb{K}^\times} \min\bigl\{\mathrm{ord}_\nu(1-x^{n_1}), \mathrm{ord}_\nu(1-y^{n_2})\bigr\}. \tag{18}$$

Here the subgroup grid consists of the pairs (α, β) satisfying $\alpha^{n_1} = \beta^{n_2} = 1$. The sum is over points of the normalization, so a singular affine grid point can contribute separately along each of its branches. At a nonsingular grid point, at least one of $x-\alpha$ and $y-\beta$ vanishes to order 1, so the point contributes exactly 1. Thus, if every grid point on $F = 0$ is nonsingular, Γ_F is simply the number of such points. Along a singular branch, the contribution can exceed 1 only if both coordinate differences vanish to order greater than 1; tangency to only one coordinate line does not increase the minimum. The minimum records the multiplicity common to the two zero sets, which is exactly the overlap subtracted when the zero counts from the rows $X = \alpha$ and columns $Y = \beta$ are combined below.

For a rational function u on $\mathcal{X}_F$, let du denote its formal differential. Since $\mathbb{K}$ is perfect, $du = 0$ exactly when u is a qth power in $\mathbb{K}(\mathcal{X}_F)$. We say that two rational functions u, v are *multiplicatively independent modulo constants* if $u^r v^s \in \mathbb{K}^\times$ for $r, s \in \mathbb{Z}$ only when $r = s = 0$.

We use the following specialization of the Corvaja–Zannier gcd estimate.

Theorem 4.12 (Specialized Corvaja–Zannier estimate)**.** *Let q be prime, let $F \in \mathbb{K}[X, Y]$ be irreducible of bidegree (a, b), where $a, b > 0$, and let x, y be the coordinate functions on $\mathcal{X}_F$. Let n_1, n_2 be positive integers and put $N = n_1 n_2$. Suppose that $N < q$, that x^{n_1} and y^{n_2} are multiplicatively independent modulo constants, and that $d(x^{n_1}), d(y^{n_2}) \neq 0$. Then, for an absolute constant C,*

$$\Gamma_F \leqslant \max\left\{3\sqrt[3]{2}\left(2Na^2b^2\right)^{1/3}, \frac{12Nab}{q}\right\} \leqslant C\bigl(N^{1/3}(ab)^{2/3} + ab\bigr). \tag{19}$$

Proof. Let $\mathcal{S}$ be the set of zeros and poles of x and y, let g be the genus of $\mathcal{X}_F$, and put $\chi = |\mathcal{S}| + 2g - 2$. Apply [CZ13, Theorem 2] to $u = x^{n_1}$ and $v = y^{n_2}$. It gives

$$\Gamma_F \leqslant \max\left\{3\sqrt[3]{2}\left(\deg(u)\deg(v)\chi\right)^{1/3}, \frac{12\deg(u)\deg(v)}{q}\right\}.$$

Here $\deg(u) = n_1 b$ and $\deg(v) = n_2 a$. Moreover, $g \leqslant (a-1)(b-1)$, while the zeros and poles of x and y have total support at most $2a + 2b$. Hence $\chi \leqslant 2ab$. Substituting these bounds gives the first inequality in (19); the second uses $N < q$. □

4.4.2 Bivariate divisibility on a coprime subgroup grid

We now prove a bivariate divisibility lemma analogous to [PS94, Lemma 8]. Their lemma has a condition of the form '$t + k < n/2$', which forces the rate of the component codes to be smaller than $1/2$. By moving to multiplicative subgroups of coprime order, we avoid such a condition, which eventually allows us to handle codes with rate close to 1.

Lemma 4.13 (Coprime-subgroup bivariate divisibility)**.** *There is an absolute constant $c_0 > 0$ such that the following holds. Let q be prime, let $S_i \leqslant \mathbb{F}_q^\times$ have coprime orders n_i, and fix integers $0 \leqslant k_i < n_i$. Put $N = n_1 n_2$ and $h = \min\{n_1 - k_1, n_2 - k_2\}$. Let $E, P \in \mathbb{F}_q[X, Y]$, with $E \neq 0$, and set $e_X = \deg_X E$, $e_Y = \deg_Y E$, and $t = \max\{e_X, e_Y\}$. Suppose that*

$$t \leqslant c_0 \frac{h^3}{N}, \qquad \deg_X P < k_1 + e_X, \qquad \deg_Y P < k_2 + e_Y,$$

and that for every $\alpha \in S_1$ and $\beta \in S_2$ there are polynomials $R_\alpha(Y)$ and $C_\beta(X)$ satisfying

$$P(\alpha, Y) = E(\alpha, Y)R_\alpha(Y), \qquad P(X, \beta) = E(X, \beta)C_\beta(X).$$

Then $E \mid P$. Consequently, $P = EQ$ for some $Q \in \mathbb{F}_q[X, Y]$ satisfying $\deg_X Q < k_1$ and $\deg_Y Q < k_2$.

Proof. We work first over $\mathbb{K}$. The claim is immediate if $P = 0$. Otherwise, let F be an irreducible factor of E that depends on both variables, and write $a = \deg_X F > 0$ and $b = \deg_Y F > 0$. Suppose for contradiction that $F \nmid P$, and use the notation from the preceding subsection.

The divisibilities after fixing one coordinate give a lower bound on the number of zeros of $P(x, y)$. Indeed, for each $\alpha \in S_1$, the polynomial $F(\alpha, Y)$ is nonzero and divides $P(\alpha, Y)$. We can therefore write $P(X, Y) = F(X, Y)H_\alpha(Y) + (X - \alpha)K_\alpha(X, Y)$. On $\mathcal{X}_F$, this gives $P(x, y) = (x - \alpha)K_\alpha(x, y)$. Hence, at every point above $x = \alpha$ at which y is finite, the order of vanishing of $P(x, y)$ is at least that of $x - \alpha$. The analogous statement holds above each $y = \beta$, for $\beta \in S_2$, when x is finite.

The zero divisors of $x^{n_1} - 1$ and $y^{n_2} - 1$ have degrees $n_1 b$ and $n_2 a$. Their common zeros are counted twice, and their weighted overlap is exactly Γ_F. The zeros of $x^{n_1} - 1$ that occur at $y = \infty$ have total multiplicity at most ab: there are at most a such points, and the multiplicity at each is at most b. The symmetric loss is also at most ab. Thus $P(x, y)$ has at least $n_1 b + n_2 a - \Gamma_F - 2ab$ zeros, whereas the degree bounds on P and Bézout's theorem give at most $(k_1 + e_X)b + (k_2 + e_Y)a$ zeros. Since $a \leqslant e_X \leqslant t$ and $b \leqslant e_Y \leqslant t$, the last two inequalities imply

$$\Gamma_F \geqslant (h - 2t)(a + b). \tag{20}$$

We next prove an upper bound. Since n_1 and n_2 are coprime subgroup orders, $N = n_1 n_2$ divides $q - 1$, and hence $N < q$. Suppose first that x^{n_1} and y^{n_2} are multiplicatively independent modulo constants. For a sufficiently small choice of c_0, we have $a, b \leqslant t < q$. If $dx = 0$, then x is a qth power in $\mathbb{K}(\mathcal{X}_F)$, which would make q divide $\deg(x) = b$. Thus $dx \neq 0$, and similarly $dy \neq 0$. Since $q \nmid n_1 n_2$, the differentials of x^{n_1} and y^{n_2} are also nonzero. We may therefore apply Theorem 4.12.

It remains to handle multiplicative dependence. A nontrivial relation between x^{n_1} and y^{n_2} gives, after clearing negative exponents, an equation $x^A y^B = \lambda$ or $x^A = \lambda y^B$ for some $A, B > 0$. Factoring this binomial over $\mathbb{K}$ and using the irreducibility of F shows that F is a primitive binomial $X^a Y^b = \lambda'$ or $X^a = \lambda' Y^b$, where $\gcd(a, b) = 1$. Its normalization has a coordinate z for which $x = \mu z^b$ and either $y = \nu z^a$ or $y = \nu z^{-a}$. Thus a common zero counted by Γ_F satisfies two equations of the form $z^{bn_1} = c_1$ and $z^{\pm an_2} = c_2$. The characteristic divides neither exponent because $a, b, n_1, n_2 < q$, so all these roots in $\mathbb{K}^\times$ are simple. Their number is at most $\gcd(bn_1, an_2) = \gcd(b, n_2)\gcd(a, n_1) \leqslant ab$. Hence $\Gamma_F \leqslant ab$ in the dependent case.

In both cases, $\Gamma_F \leqslant C\big(N^{1/3}(ab)^{2/3} + ab\big)$ for a possibly larger absolute constant C. Since $a, b \leqslant t$, we have $(ab)^{2/3} \leqslant \frac{1}{2}t^{1/3}(a + b)$ and $ab \leqslant \frac{1}{2}t(a + b)$. It follows that

$$\Gamma_F \leqslant C'\big(N^{1/3}t^{1/3} + t\big)(a + b).$$

Choosing the absolute constant c_0 sufficiently small makes this upper bound less than $h(a + b)/2$, while (20) is greater than $h(a + b)/2$. This is a contradiction, so $F \mid P$.

We have shown that every mixed irreducible factor of E divides P. Such a factor never becomes the zero polynomial after fixing X or Y. We may therefore cancel one copy of it from E, P, and every corresponding divisibility. The reduced polynomials satisfy the same degree bounds, with the degree of the cancelled factor subtracted on both sides. Repeating this argument shows that every mixed factor of E divides P with its full multiplicity.

After cancelling these factors, write the remaining part of E as $U(X)V(Y)$, and let P_0 be the corresponding quotient of P. For every $\alpha \in S_1$, we have $V(Y) \mid P_0(\alpha, Y)$. Divide P_0 by V in the variable Y. Each coefficient of the remainder is a polynomial in X of degree less than $k_1 + \deg U < n_1$ and vanishes on all of S_1. The remainder is therefore zero, so $V \mid P_0$. The symmetric argument gives $U \mid P_0$. Since U and V are relatively prime in $\mathbb{K}[X, Y]$, their product divides P_0. We conclude that $E \mid P$ over $\mathbb{K}$, and hence over $\mathbb{F}_q$.

Writing $P = EQ$, the degree bounds on Q follow by subtracting the degrees of E from those of P. $\square$

4.4.3 Proof of Theorem 4.11

We now complete the proof using the argument of Polishchuk and Spielman, who also deduce product expansion (equivalently, bivariate testing) from an analogous divisibility lemma [PS94, Lemma 8 and Theorem 9]. We give the proof for completeness.

Proof of Theorem 4.11. Put $N = n_1 n_2$, $m = \min\{n_1, n_2\}$, and $h = \min\{n_1 - k_1, n_2 - k_2\}$. Then $h \geqslant \varepsilon m$ and $N \leqslant K m^2$. Decrease the constant c_0 in Lemma 4.13 so that $c_0 \leqslant 1/8$, and set $\tau = c_0 h^3/N$ and $\rho = c_0^2 \varepsilon^6/(8K^3)$. Since $h^2 \leqslant N$, we have $\tau \leqslant h/8$.

Fix $M = A + B$, where every column of A lies in C_1 and every row of B lies in C_2, and write $w = |M|$. If $w = 0$, use the zero decomposition. Suppose first that $0 < w < \tau^2/4$, and let $t = \lfloor \sqrt{w} \rfloor$. There is a nonzero polynomial E of bidegree at most (t, t) that vanishes on $\operatorname{supp}(M)$, since the space of such polynomials has dimension $(t+1)^2 > w$.

Identify A and B with their grid interpolants. Since $EM = 0$ on the grid, let P be the polynomial of individual degrees less than (n_1, n_2) representing the common table $EA = -EB$. Writing $e_X = \deg_X E$ and $e_Y = \deg_Y E$, reduction modulo $Y^{n_2} - 1$ and $X^{n_1} - 1$, respectively, gives

$$\deg_X P < k_1 + e_X, \qquad \deg_Y P < k_2 + e_Y.$$

On each column $Y = \beta$, the polynomials $P(X, \beta)$ and $E(X, \beta)A(X, \beta)$ agree on S_1 and have degree less than n_1, so they are equal. The same argument on each row $X = \alpha$ gives

$$P(X, \beta) = E(X, \beta)A(X, \beta), \qquad P(\alpha, Y) = -E(\alpha, Y)B(\alpha, Y)$$

for all $\alpha \in S_1$ and $\beta \in S_2$. Since $t < \tau = c_0 h^3/N$, Lemma 4.13 gives $P = EQ$, where $\deg_X Q < k_1$ and $\deg_Y Q < k_2$. Hence the evaluation of Q lies in $C_1 \otimes C_2$, and $M = (A - Q) + (B + Q)$.

Let J be the set of nonzero columns of $A - Q$, and let I be the set of nonzero rows of $B + Q$. For $\beta \in J$, the column identity above and $P = EQ$ imply that $E(X, \beta) \equiv 0$. Hence $Y - \beta$ divides E, so there are at most e_Y such values of β. Thus $|J| \leqslant e_Y \leqslant t$. The row identity similarly gives $|I| \leqslant e_X \leqslant t$.

By the Reed–Solomon distance bound, every nonzero column of $A - Q$ and nonzero row of $B + Q$ has weight at least $h + 1$. Outside the rows in I, the table $B + Q$ vanishes, so $M = A - Q$. Therefore each column in J contributes at least $h + 1 - |I|$ nonzero entries to M. The symmetric row count gives

$$w \geqslant |J|(h + 1 - |I|), \qquad w \geqslant |I|(h + 1 - |J|).$$

Since $|I|, |J| \leqslant t < h/4$, this gives $|I|, |J| \leqslant 2w/h$. Consequently,

$$w \geqslant \frac{h}{2(n_1 + n_2)}\big(n_1|J| + n_2|I|\big) \geqslant \rho\big(n_1|J| + n_2|I|\big),$$

where $h/(2(n_1 + n_2)) \geqslant \varepsilon/(4K) \geqslant \rho$.

It remains to consider $w \geqslant \tau^2/4$. Keep the original decomposition $M = A + B$. Since A has at most n_2 nonzero columns and B has at most n_1 nonzero rows, $n_1|A|_{\mathrm{col}} + n_2|B|_{\mathrm{row}} \leqslant 2N$. Hence

$$w \geqslant \frac{\tau^2}{8N}\big(n_1|A|_{\mathrm{col}} + n_2|B|_{\mathrm{row}}\big) \geqslant \rho\big(n_1|A|_{\mathrm{col}} + n_2|B|_{\mathrm{row}}\big),$$

because $\tau^2/(8N) = c_0^2 h^6/(8N^3) \geqslant c_0^2\varepsilon^6/(8K^3) = \rho$.

Finally, for coordinate-rescaled codes, apply the unscaled result to $D_1^{-1} M D_2^{-1}$ and rescale the resulting decomposition back. This preserves Hamming weight and the sets of nonzero rows and columns, so the same ρ works. □

5 Private and Robust Circuit for the Zero-on-Subcube Problem

We start by defining the zero-on-subcube problem,

Definition 5.1 (k-dimensional Zero-on-Subcube Problem)**.** *Fix $k \geqslant 2$ and a prime q. For every $i \in [k]$, let $\Gamma_i = \langle\omega_i\rangle \leqslant \mathbb{F}_q^\times$ have order h_i, choose $\zeta_i \in \mathbb{F}_q^\times$, and put $H_i = \zeta_i\Gamma_i$, $D = \prod_{i=1}^k H_i$. For $a \in \{0, \dots, h_i - 1\}$, define the one-coordinate shift $\sigma_{i,a}(x_1, \dots, x_k) = (x_1, \dots, \omega_i^a x_i, \dots, x_k)$. A k-dimensional zero-on-subcube instance over D is $I = (\mathbb{F}_q, D, \Psi, C)$, where $\Psi = (\psi_1, \dots, \psi_t)$ is a labeled list whose entries are the identity or one-coordinate shifts $\sigma_{i,a} : D \to D$, and*

$$C(\mathbf{X}, Y_1, \dots, Y_t) \in \mathbb{F}_q[X_1, \dots, X_k, Y_1, \dots, Y_t]$$

satisfies $\deg_{X_i} C < h_i$ for every i and $\deg_{\mathbf{Y}} C \leqslant 3$.

We use the same formulas to extend the maps in Ψ to $\mathbb{F}_q^k$. For any polynomial $P \in \mathbb{F}_q[X_1, \dots, X_k]$, define its constraint polynomial

$$B_P(\mathbf{X}) = C\big(\mathbf{X}, P(\psi_1(\mathbf{X})), \dots, P(\psi_t(\mathbf{X}))\big).$$

For a table $A : D \to \mathbb{F}_q$, let $\widetilde{A}$ be its unique interpolating polynomial satisfying $\deg_{X_i}\widetilde{A} < h_i$ for every i, and write $B_A := B_{\widetilde{A}}$.

We say that A satisfies I if $B_A|_D \equiv 0$, and that I is satisfiable if some table A satisfies I. Then define the zero-on-subcube function $f_I(A)$ as, $f_I(A) = 1$ if A satisfies I, and let $f_I(A) = 0$ otherwise.

Note that if $P|_D = A$, then, since every map in Ψ preserves D, $B_P|_D = B_A|_D$.

Next we define the circuit model we use formally.

5.1 Circuit Model

For words u, v on the same finite coordinate set, let $\mathrm{dist}(u, v)$ denote their Hamming distance. For a set of words $\mathcal{S}$, put $\mathrm{dist}(u, \mathcal{S}) = \min_{c\in\mathcal{S}} \mathrm{dist}(u, c)$.

Definition 5.2 (Randomized classical encoding)**.** *A randomized classical encoding is a map $\mathrm{Enc} : \mathbb{F}_q^m \times \mathcal{R} \longrightarrow \mathbb{F}_q^{N_0}$. For $x \in \mathbb{F}_q^m$, let $\mathrm{Enc}(x) = \{\mathrm{Enc}(x; r) : r \in \mathcal{R}\}$, and let C_x be the distribution of $\mathrm{Enc}(x; r)$ for uniform r. Thus C_x is the secret-sharing distribution for x. The underlying classical code is $C_{\mathrm{in}} = \bigcup_{x\in\mathbb{F}_q^m} \mathrm{Enc}(x)$. We assume that this code supports a randomized local membership test Test. On input a word W, the test uses fresh randomness and accepts or rejects, with $\Pr[\mathsf{Test}(W)$ rejects$] = 0 \iff W \in C_{\mathrm{in}}$.*

As noted in the proof overview, our construction for the zero-on-subcube is significantly simpler if one allows for auxiliary proofs supplied by an external prover (even though these are not necessary) in the circuit model. This is compatible with our final application to PCPs, so we use this stronger model,

Definition 5.3 (Private & Robust Circuit with Auxiliary Proofs)**.** *Let $f : \mathbb{F}_q^m \to \{0,1\}$. A private, robust circuit for f with respect to a randomized encoding $(\mathrm{Enc}(\cdot), \mathsf{Test})$ is a pair $(\mathcal{C}, \mathcal{T})$, where $\mathcal{C}$ is an L-layer computation circuit and $\mathcal{T}$ is a randomized local tester.*

The first layer of the circuit takes in inputs of length denoted by N_0. The ℓ^{th} layer of the circuit receives $(W_{\ell-1}, \Pi_\ell, R_\ell)$, where $W_{\ell-1}$ is the current state, Π_ℓ consists of auxiliary words supplied by an external prover, and R_ℓ consists of fresh random words independent of x and of one another. The construction specifies their honest distributions. The computation gates in layer ℓ produce W_ℓ, and each gate has fanin at most Δ_{fan}, which we call the fanin of the circuit.

The tester $\mathcal{T}$ has oracle access to the input and computation-output wires of all layers. It uses fresh randomness, independent of the circuit words and of the faulty gates, and accepts or rejects. For every fixing of its randomness, it queries at most Δ_{fan} wires. Rejection probabilities below are over this randomness for a fixed execution. The final state $W_L \in \{0,1\}$ is the circuit output.

1. ***Honest Correctness:*** *For every x, if $W_0 \sim C_x$, there exists a choice of distributions for the auxiliary words so that $\mathcal{T}$ accepts with probability one and $W_L = f(x)$.*
2. ***Privacy:*** *In an honest execution, where the auxiliary and random words follow their prescribed distributions,*
   (a) *We say the circuit is δ-**layerwise-private** if the values on any set containing fewer than a δ-fraction of the full input block $(W_{\ell-1}, \Pi_\ell, R_\ell)$ of any layer are independent of x.*
   (b) *Let $\mathsf{Tr}_{\mathcal{C}}(x)$ denote the collection of all input, auxiliary, random-input, and computation-output wires. Let μ be any probability measure on the coordinates of this transcript. We say that $\mathcal{C}$ is ν-**jointly private** under μ if, for every $x, x' \in \mathbb{F}_q^m$ and every set T of transcript coordinates with $\mu(T) < \nu$,*
   $$\mathsf{Tr}_{\mathcal{C}}(x)|_T \stackrel{\mathrm{d}}{=} \mathsf{Tr}_{\mathcal{C}}(x')|_T.$$
3. ***Robustness:*** *An ϵ_{lay}-faulty execution is one in which the adversary changes the outputs of fewer than an ϵ_{lay}-fraction of the computation gates in each layer. All remaining computation gates evaluate correctly on their actual inputs. The circuit has robustness parameters $(\gamma, \epsilon_{\mathrm{in}}, \epsilon_{\mathrm{lay}})$ if,*
   (a) *For arbitrary auxiliary and random words, if $\Pr[\mathcal{T}$ rejects$] \leqslant \gamma$, then dist$(W_0, C_{\mathrm{in}}) \leqslant \epsilon_{\mathrm{in}} N_0$.*
   (b) *For every x satisfying dist$(W_0, \mathrm{Enc}(x)) \leqslant \epsilon_{\mathrm{in}} N_0$, on any $\varepsilon_{\mathrm{lay}}$-faulty execution of the circuit, if $\Pr[\mathcal{T}$ rejects$] \leqslant \gamma$, then $W_L = f(x)$.*

5.2 Parameters and encodings

We use two Cartesian grids. The computation itself takes place on the constraint grid $D = \prod_{i=1}^k H_i$, whose points index the assignment and the constraints. A table $T : D \to \mathbb{F}_q$ represents the witness. To privately encode it, we sample uniformly a low-degree polynomial P subject to $P|_D = T$ and pass its evaluations on the larger grid $\Omega = \prod_{i=1}^k S_i$. Thus every table is represented by a randomized tensor Reed–Solomon codeword. Different choices of P represent the same table on D, and adding a fresh random polynomial that vanishes on D rerandomizes the representation without changing that table. The circuit maintains such

representations while it evaluates the constraints, forms the indicator of a violated constraint, and sums this indicator over D.

The parameters serve different purposes. The algebraic operations determine the degrees of the intermediate polynomials. The larger grid Ω must then be large enough relative to D to keep the corresponding codes at constant relative distance; this causes most of the increase in size. Privacy comes from the degree slack left after fixing a polynomial on D, together with our product expansion conjecture, whose hypotheses require the side lengths of Ω to be pairwise coprime and balanced. The BSS hypercube embedding first determines how large D must be. The computation then determines the degree profiles, after which we choose the larger side lengths needed for distance and privacy and realize both grids over a common field. Write $h_i = |H_i|$ and $s_i = |S_i|$.

Parameter goals and degree profiles. After the instance size n is fixed, we may choose the dimension $k = k(n)$. We take the hidden side lengths h_i to be comparable primes of size at least $(n \log(n+1))^{1/k}$, so that the BSS hypercube embeds into the constraint grid. We choose each evaluation side length to be $s_i = b_i h_i$, where $b_i = \Theta(k \log k)$, and give the input degree a constant fraction of s_i of slack beyond h_i. Specifically, we set $d_i = h_i + \left\lceil \frac{s_i}{128} \right\rceil$. The circuit operations then force the degree profiles $e_i = 3d_i + h_i - 1, u_i = 4d_i + h_i - 1, v_i = 7d_i + 2h_i - 2$. Here u_i is the degree of $P_B P_J$, while v_i is the degree of $P_B(1 - P_B P_J)$. The choices below make the s_i's pairwise coprime and 4-balanced. They also give every component code used by the circuit constant relative distance and leave an absolute constant fraction of privacy slack.

Lemma 5.4 (Parameter selection)**.** *There is an absolute constant $C \geqslant 1$ with the following property. Let $n \geqslant 1$, and let $k = k(n) \geqslant 2$ be any integer function of n. Then there are a prime q and grids $D = \prod_{i=1}^{k} H_i, \Omega = \prod_{i=1}^{k} S_i$, with sizes $h_i = |H_i|$, $s_i = |S_i|$, and $N = |\Omega| = \prod_{i=1}^{k} s_i$, satisfying the following,*

1. Degree margins. *For $\eta = 1/128$ and $v_i = 7d_i + 2h_i - 2$, for every $i \in [k]$,*

$$d_i - h_i \geqslant \eta s_i, \qquad v_i + 1 \leqslant s_i/4. \tag{21}$$

2. Domain sizes: *For $\lambda_k = k \log(k+1)$,*

$$n < |D| < N \leqslant C^k \max\{n \log(n+1)\lambda_k^k, \lambda_k^{2k}\}, \qquad \frac{1}{4} \leqslant \frac{s_i}{s_j} \leqslant 4 \quad (i, j \in [k]).$$

   Moreover, $2^{\sum_i \lfloor \log_2 h_i \rfloor} \geqslant cn \log(n+1)$, for the absolute constant in Lemma 8.1, thus D satisfies the hypotheses therein.

3. Field and subgroup choices. *The h_i are distinct primes, and the s_i are pairwise coprime. We have $N \mid q - 1$ and $N < q \leqslant N^6$. For every i, $S_i \leqslant \mathbb{F}_q^\times$ is a subgroup of order s_i, and there exist a subgroup $\Gamma_i \leqslant S_i$ of order h_i and an element $\zeta_i \in \mathbb{F}_q^\times$ such that $H_i = \zeta_i \Gamma_i$. Moreover, $H_i \cap S_i = \varnothing$, and both H_i and S_i are invariant under multiplication by Γ_i.*

Proof. Put $\lambda_k = k \log(k+1)$. By the prime number theorem, there is an absolute constant A_0 such that, for every integer $t \geqslant 1$ and every real $x \geqslant A_0 t \log(t+1)$, the interval $(x, 2x)$ contains at least t primes. Fix a sufficiently large absolute constant $A \geqslant A_0$, increasing it as needed below, and set

$$B = A\lambda_k, \qquad L = \max\{A(n \log(n+1))^{1/k}, 4B\}.$$

The scale B makes the b_i's sufficiently large for the degree margins. The first term in L ensures the hypercube-embedding condition in (2), while $L \geqslant 4B$ separates the two prime intervals. Both B and L satisfy the preceding prime-counting bound with $t = k$, thus we can choose distinct primes $b_1, \ldots, b_k \in (B, 2B)$ and $h_1, \ldots, h_k \in (L, 2L)$.

(1) Degree margins. Set $s_i = b_i h_i$ and $d_i = h_i + \left\lceil \frac{s_i}{128} \right\rceil$. It is clear that $d_i - h_i \geqslant \frac{s_i}{128}$ and since $b_i \geqslant 128$,

$$v_i + 1 = 9h_i + 7\left\lceil \frac{s_i}{128} \right\rceil - 1 \leqslant \frac{s_i}{8} + 6 \leqslant \frac{s_i}{4}.$$

(2) Domain sizes. For all $i, j \in [k]$, $1/4 < \frac{s_i}{s_j} < 4$. Since $b_i > 1$, $N > |D| = \prod_i h_i$. Moreover,

$$|D| > L^k \geqslant A^k n \log(n+1) > n$$

for sufficiently large A. Also, $s_i < 4BL = O\left(\lambda_k \max\{(n\log(n+1))^{1/k}, \lambda_k\}\right)$, thus, for some absolute constant C,

$$N = \prod_i s_i \leqslant C^k \max\{n\log(n+1)\lambda_k^k, \lambda_k^{2k}\}.$$

Finally, since $2^{\lfloor \log_2 h_i \rfloor} > h_i/2$ for every i,

$$2^{\sum_i \lfloor \log_2 h_i \rfloor} > \prod_i \frac{h_i}{2} > \left(\frac{L}{2}\right)^k \geqslant \left(\frac{A}{2}\right)^k n\log(n+1).$$

For sufficiently large A, the last expression is a sufficiently large absolute-constant multiple of $n\log(n+1)$, as required.

(3) Field and subgroup choices. Since $L \geqslant 4B$, all $2k$ primes are distinct. It follows immediately that the s_i's are pairwise coprime.

By Xylouris's form of Linnik's theorem [Xyl09], every sufficiently large M admits a prime $q \equiv 1 \pmod M$ with $q \leqslant M^6$. Since $N > (BL)^k$, our choice of A ensures that N exceeds this absolute threshold. Linnik's theorem therefore gives a prime $q \equiv 1 \pmod N$ with $N < q \leqslant N^6$. Since $\mathbb{F}_q^\times$ is cyclic, it has subgroups S_i of orders s_i and subgroups $\Gamma_i \leqslant S_i$ of orders h_i. Read indices cyclically modulo k, and choose $\zeta_i \in S_{i+1}$ of order b_{i+1}. Pairwise coprimality gives $S_i \cap S_{i+1} = \{1\}$, so $\zeta_i \notin S_i$. Thus $H_i = \zeta_i \Gamma_i$ is disjoint from S_i. The invariance properties follow from the definitions.

Exhaustive search with deterministic primality testing finds the primes in the two intervals and a prime $q \equiv 1 \pmod N$ below N^6. It also constructs the required subgroups in time polynomial in k and N. □

Fix a choice from Lemma 5.4. For a degree vector $\mathbf{a} = (a_1, \ldots, a_k) \in \mathbb{Z}_{\geqslant 0}^k$, write $\mathcal{P}_\mathbf{a} = \{P \in \mathbb{F}_q[X_1, \ldots, X_k] : \deg_{X_i} P \leqslant a_i \text{ for every } i\}$. We also write $\mathcal{P}_{a_1,\ldots,a_k}$ for this space and interpret vector inequalities and sums coordinatewise. Define the tensor code $\mathcal{C}_\mathbf{a} = \mathrm{ev}_\Omega(\mathcal{P}_\mathbf{a}) = \bigotimes_{i=1}^k \mathrm{RS}(S_i, a_i + 1)$.

A representation is simply a low-degree extension from D to Ω. More precisely, for a table $T : D \to \mathbb{F}_q$, let $\mathcal{R}_\mathbf{a}(T) = \{P \in \mathcal{P}_\mathbf{a} : P|_D = T\}$. A degree-$\mathbf{a}$ representation of T is $\mathrm{ev}_\Omega(P)$ for $P \in \mathcal{R}_\mathbf{a}(T)$. It is *uniform* if P is sampled uniformly from this affine space. Coordinatewise interpolation shows that the restriction map is surjective when $a_i \geqslant h_i - 1$ for every i. Thus $\mathrm{Enc}(A) = \{\mathrm{ev}_\Omega(P) : P \in \mathcal{R}_\mathbf{d}(A)\}$, and C_A is the distribution of $\mathrm{ev}_\Omega(P)$ for $P \sim \mathrm{Unif}(\mathcal{R}_\mathbf{d}(A))$. Let $\mathcal{Z}_\mathbf{a} := \mathcal{R}_\mathbf{a}(0) = \{R \in \mathcal{P}_\mathbf{a} : R|_D = 0\}$. For every $P_0 \in \mathcal{R}_\mathbf{a}(T)$,

$$\mathcal{R}_\mathbf{a}(T) = P_0 + \mathcal{Z}_\mathbf{a}. \tag{22}$$

Thus adding an independent uniform element of $\mathcal{Z}_{\mathbf{a}}$ to any degree-$\mathbf{a}$ representation of T produces a fresh uniform representation of the same table.

Let $Z_i(X_i) = \prod_{a \in H_i}(X_i - a)$. For a degree profile $\mathbf{a}$ with $a_i \geqslant h_i$, define

$$\Phi : \bigoplus_{i=1}^{k} \mathcal{Q}_i(\mathbf{a}) \longrightarrow \mathcal{Z}_{\mathbf{a}}, \qquad \Phi(Q_1, \ldots, Q_k) = \sum_{i=1}^{k} Z_i Q_i. \tag{23}$$

By Fact 3.3, $\Phi_{\mathbf{a}}$ is surjective, in fact, each polynomial $P \in \mathcal{Z}_a$ has the same number of preimages, denoted by $\Phi^{-1}(P)$.

5.3 The Circuit Construction for Zero-on-Subcube

The circuit computes the constraint values $B = B_A|_D$, converts them to the indicator $U = \mathbf{1}[B \neq 0]$, and sums U over D. It uses two primitives: $\mathsf{SampleMask}_{\mathbf{a}}$ samples a uniform polynomial in $\mathcal{Z}_{\mathbf{a}}$, and $\mathsf{Test}_{\mathbf{a}}$ tests a uniformly random axis line of a word.

Primitive $\mathsf{Test}_{\mathbf{a}}$.

Input. An index $j \in [k]$, a degree profile $\mathbf{a} = (a_j, \ldots, a_k)$, and a word W on $\prod_{i=j}^{k} S_i$.

Procedure. Choose $i \in \{j, \ldots, k\}$ uniformly, and then choose an i-axis line ℓ uniformly. Reject if $W|_\ell \notin \mathrm{RS}(S_i, a_i + 1)$, and otherwise accept.

Primitive $\mathsf{SampleMask}_{\mathbf{a}}$.

Input. Words $W_1, \ldots, W_k$ on Ω, supplied on random-input wires. Their prescribed honest distribution is $W_i = \mathrm{ev}_\Omega(R_i)$, where the R_i's are independent and R_i is uniform in $\mathcal{Q}_i(\mathbf{a})$.

Procedure. At every $x \in \Omega$, form $R(x) = \sum_{i=1}^{k} Z_i(x_i) W_i(x)$.

Output. Each data gate that uses $R(x)$ evaluates it directly from $W_1(x), \ldots, W_k(x)$ (there is no actual output).

We invoke $\mathsf{SampleMask}_{\mathbf{a}}$ as a black box: an invocation to get a mask R places its random-input words and the data gates that use R in the same layer. All mask calls named in one step and that step's data gates run in parallel.

Whenever a gate decodes a component-code word, it uses the unique codeword at distance less than half the minimum distance of the component code, if one exists; otherwise, it returns a fixed default codeword.

Circuit Construction.

1. **Input.** The initial state is the evaluation table P_A.

2. **Constraint evaluation.** Invoke $\mathsf{SampleMask}_{\mathbf{e}}$ to get R_B. In parallel, the data gate at $x \in \Omega$ decodes the k axis lines of P_A through x and computes

$$P_B(x) = C\big(x, (P_A(\psi(x)))_{\psi\in\Psi}\big) + R_B(x). \tag{24}$$

   For an identity query it uses the value from the first decoded line; for a shift in coordinate i it uses the corresponding value from the decoded i-line.

3. **Candidate inverse (auxiliary input).** For $B = B_A|_D$, the honest prover sets

$$J(z) = \begin{cases} B(z)^{-1}, & B(z) \neq 0, \\ 0, & B(z) = 0, \end{cases} \qquad z \in D, \tag{25}$$

   and independently samples $P_J \sim \mathrm{Unif}(\mathcal{R}_{\mathbf{d}}(J))$.

4. **Indicator.** Independently invoke $\mathsf{SampleMask}_{\mathbf{u}}$ and $\mathsf{SampleMask}_{\mathbf{v}}$ to get R_U and R_V, respectively. One layer has data gates that compute

$$P_U = P_B P_J + R_U, \tag{26}$$
$$P_V = P_B(1 - P_B P_J) + R_V. \tag{27}$$

5. **Vanishing certificate.** The prover supplies auxiliary words $(Q_1, \ldots, Q_k)$ on Ω. In the honest execution, this tuple is chosen uniformly at random from $\Phi^{-1}(P_V)$. For later use, define

$$E(x) = P_V(x) - \sum_{i=1}^{k} Z_i(x_i) Q_i(x) \tag{28}$$

6. **Aggregation.** Let $P^{(0)} = P_U$ and define $P^{(1)}, \ldots, P^{(k-1)}$ recursively by

$$P^{(i)}(x_{i+1}, \ldots, x_k) = \sum_{a\in H_i} \widehat{P}^{(i-1)}_{x'}(a), \tag{29}$$

   where $x' = (x_{i+1}, \ldots, x_k)$ and $\widehat{P}^{(i-1)}_{x'}$ is the polynomial decoded from the corresponding S_i-line. For $i = 1, \ldots, k-1$, one layer computes $P^{(i)}$ by (29).

7. **Output.** The final layer has one data gate that computes

$$\nu = \sum_{a\in H_k} \widehat{P}^{(k-1)}(a) \tag{30}$$

   and outputs $\mathbf{1}[\nu = 0]$.

We pair the circuit above with the following randomized local tester,

Randomized Tester $\mathcal{T}$.

The tester chooses uniformly among the following $F = 2k + 8$ tests and runs the selected test:

1. Four tests: apply $\mathsf{Test}_{\mathbf{d}}$ to P_A, $\mathsf{Test}_{\mathbf{e}}$ to P_B, $\mathsf{Test}_{\mathbf{d}}$ to P_J, and $\mathsf{Test}_{\mathbf{v}}$ to P_V, respectively.
2. k tests: for each $i \in [k]$, apply $\mathsf{Test}_{(u_i,\dots,u_k)}$ to $P^{(i-1)}$, where $P^{(0)} = P_U$.
3. Three tests: one for each of the three mask calls to $\mathsf{SampleMask}_{\mathbf{a}}$, which in turn randomly runs one of k tests.
4. k tests: for each $i \in [k]$, apply $\mathsf{Test}_{(v_1,\dots,v_{i-1},v_i-h_i,v_{i+1},\dots,v_k)}$ to Q_i.
5. One test chooses $x \in \Omega$ uniformly and rejects if $E(x)$ from (28) is not equal to 0.

5.4 Guarantees of the Construction

Theorem 5.5 (Private, robust circuit for zero-on-subcube)**.** *Assume the setting of Lemma 5.4 and Conjecture 4.2 with product expansion constant being $\rho = \rho(k, 1/128, 4)$. Let $I = (\mathbb{F}_q, D, \Psi, C)$ be a k-dimensional zero-on-subcube instance over these parameters. Then f_I admits a private and robust circuit in the sense of Definition 5.3 with the following properties,*

1. *Honest correctness.*
2. *Layerwise privacy at least $\Omega(\rho^{k-1}/k)$ and joint privacy under the measure μ from Section 6.3 at least $\exp\big(-O(k^3 + k^2 \log(1/\rho))\big)/|\Psi|$.*
3. *Robustness with all parameters being at least $\exp(-O(k))$.*
4. *The circuit has encoding length N, fan-in $O(kN^{1/k})$, depth $O(k)$, and width and gate count $O(k^2 N)$. The tester makes $O(N^{1/k} + k)$ queries.*

Its description can be constructed in time polynomial in k, N, $\mathrm{size}(I)$, and $\log q$.

Proof. Honest correctness. Every axis-line check in $\mathcal{T}$ accepts with probability one in the honest execution. Each $\mathsf{SampleMask}_{\mathbf{a}}$ call evaluates a degree-$\mathbf{a}$ polynomial that vanishes on D.

The shift maps preserve individual degree, and C has degree less than h_i in X_i and query degree at most three. Hence

$$B_{P_A}|_D = B_A|_D, \qquad \deg_{X_i} B_{P_A} \leqslant 3d_i + h_i - 1 = e_i. \tag{31}$$

Thus $P_B = B_{P_A} + R_B$ has degree at most $\mathbf{e}$ and represents $B = B_A|_D$.

The degrees of P_U and P_V are at most $\mathbf{e}+\mathbf{d} = \mathbf{u}$ and $2\mathbf{e}+\mathbf{d} = \mathbf{v}$, respectively. Since R_B, R_U, R_V vanish on D, the honest inverse (25) and equations (26)– (27) give $U = BJ = \mathbf{1}[B \neq 0], V = B(1 - BJ) = 0$. In particular, $P_V \in \mathcal{Z}_{\mathbf{v}}$. By Fact 3.3, suitable certificate polynomials $Q_i \in \mathcal{Q}_i(\mathbf{v})$ exist, and $E(x) = 0$ for every $x \in \Omega$. Hence the certificate check accepts with probability one. Thus every constituent check of $\mathcal{T}$ accepts with probability one.

Successive aggregation computes

$$\nu = \sum_{z \in D} U(z) = \big|\{z \in D : B(z) \neq 0\}\big| \cdot 1_{\mathbb{F}_q}.$$

Because $q > |D|$, this is zero exactly when $B_A|_D \equiv 0$. The final output is therefore $f_I(A)$.

Privacy. Put $\eta = 1/128$ and

$$\tau_k = \eta\rho^{k-1}.$$

By Lemma 6.1, the circuit has layerwise privacy $\tau_k/(2k+2)$. Let $C(\rho,k) \geqslant 1$ be the constant from Corollary 4.4, and put

$$\nu_k = \frac{\tau_k}{3(2k+8)(|\Psi|+1)C(\rho,k)}.$$

By Lemma 6.3, the circuit is ν_k-jointly private under μ. The full-dimensional part of the construction has $O(1)$ layers, and each of the $k-1$ aggregation stages has $O(1)$ layers, so $L = O(k)$. Moreover, Corollary 4.4 gives $C(\rho,k) = \exp\big(O(k^3 + k^2\log(1/\rho))\big)$. Therefore

$$\nu_k \geqslant \frac{1}{|\Psi|+1}\exp\left(-O\left(k^3 + k^2\log(1/\rho)\right)\right),$$

as claimed.

Robustness. Lemma 7.1 supplies parameters $\epsilon_{\mathrm{in},k}, \epsilon_{\mathrm{lay},k}, \gamma_k$, each at least $\exp(-O(k))$, and proves the two robustness guarantees in Definition 5.3.

Circuit parameters. Put $\beta = 4$, as in Lemma 5.4, and let $s_{\max} = \max_i s_i$. The side-ratio bound gives

$$N \geqslant s_{\max}^k/\beta^{k-1}, \qquad \sum_i s_i \leqslant k\beta^{(k-1)/k}N^{1/k} < 4kN^{1/k}.$$

A constraint gate reads the k axis lines through its coordinate and one mask value. Every other gate reads at most one axis line or $k+2$ field values, proving the fan-in bound. The layer count above gives depth $O(k)$.

At most $O(k)$ full-grid words occur in any layer. Each constituent check of $\mathcal{T}$ queries either one axis line or at most $k+1$ field values. Thus its query complexity is $O(N^{1/k}+k)$, within the stated fan-in bound. The construction algorithms give the stated running-time bound. □

6 Privacy of the Circuit Construction

6.1 Layerwise Privacy of the Circuit

In this section we show that each input layer of an honest execution of the circuit is private. The proof is quite straightforward given the privacy of the underlying private code from Corollary 4.3 which shows that a sufficiently small view of a uniform representation is independent of the encoded message.

Lemma 6.1 (Layerwise privacy)**.** *Assume the setting of Lemma 5.4. If $k > 2$, assume Conjecture 4.2 and let $\rho = \rho(k,\eta,\beta)$ for $\eta = 1/128$ and $\beta = 4$, be the relevant product-expansion constant.*

In an honest execution of the circuit in Section 5.3, the values on any set containing fewer than a δ_k-fraction of the full input block $(W_{\ell-1}, \Pi_\ell, R_\ell)$ of any layer have a distribution independent of A, that is, the circuit is $\frac{\eta\rho^{k-1}}{2k+2}$-layerwise-private.

Proof. Fix an assignment A. All tables constructed from A, such as B, J, U, V, are then fixed; the randomness lies only in their representations and in the random zero representations used by the circuit.

The parameter setting in Lemma 5.4, point (1), implies that for $\eta = 1/128$ and for each i, $d_i - h_i \geqslant \eta s_i$ hence each degree profile using in the circuit construction, satisfies the hypotheses of Corollary 4.3 and has

relative privacy at least, $\tau_k = \eta\rho^{k-1}$ (by the equivalence between relative dual distance and privacy in (2) from Section 3). Thus, for any represented table, the distribution of a view containing at most τ_k-fraction of the coordinates is independent of that table.

Every layer input consists of equally sized words and one can check that each layer contains at most $2k + 2$ words. We will prove that, in each layer, every joint view containing fewer than a τ_k-fraction of one word is independent of A, thus giving us that touching at most $\tau_k/(2k + 2)$-fraction of the layer reveals no information about A.

Constraint-evaluation layer. The word P_A is a uniform representation of A. Consequently, the part of the view lying in P_A has a distribution independent of A. The component words defining R_B are sampled independently of (A, P_A). Their contribution to the view is therefore independent of both A and the viewed coordinates of P_A. Hence, the entire joint view is independent of A.

Indicator layer. Let $B = B_A|_D$. The constraint-evaluation layer computes

$$P_B = C\left(\mathbf{X}, \left(P_A(\psi(\mathbf{X}))\right)_{\psi\in\Psi}\right) + R_B.$$

The first summand represents B, and R_B is an independent uniform zero representation. Therefore, P_B is a uniform representation of B, whose distribution for fixed B does not depend on the particular value of P_A.

Let J be the table defined by $J(x) = B(x)^{-1}$ when $B(x) \neq 0$, and $J(x) = 0$ otherwise. The word P_J is sampled independently as a uniform representation of J. Thus, the words P_B and P_J are independent. The viewed coordinates in each word occupy fewer than a τ_k-fraction of that word, so code privacy makes the first distribution independent of B and the second independent of J. Their joint distribution is therefore fixed and independent of A.

Finally, the component words defining R_U and R_V are sampled independently of (A, P_B, P_J). Adding any coordinates from those words to the view preserves independence from A. Hence, the entire indicator-layer view is independent of A.

Vanishing-certificate layer. Define $U = BJ$ and $V = B(1 - BJ)$. By the definition of J, we have $V = 0$. As above, the independent uniform zero representations R_U and R_V make P_U and P_V independent uniform representations of U and V, respectively. In particular, P_V is a uniform zero representation and is independent of both A and P_U.

Given P_V, the prover samples $Q_1, \ldots, Q_k$ using fresh randomness according to the vanishing-certificate procedure. It follows that $(P_V, Q_1, \ldots, Q_k)$ is independent of (A, P_U). The part of the view lying in P_U is independent of A by code privacy, while all remaining viewed coordinates belong to a tuple independent of both A and P_U. Hence, the complete vanishing-certificate view is independent of A.

Aggregation and final layers. Put $P^{(0)} = P_U$. For every i, each coordinate of $P^{(i)}$ is obtained from the corresponding $S_1 \times \cdots \times S_i$ slice of P_U, and distinct coordinates of $P^{(i)}$ correspond to disjoint slices. Since

$$|P^{(i)}| = \frac{|P_U|}{|S_1| \cdots |S_i|},$$

a given fraction of the coordinates of $P^{(i)}$ is determined by the same fraction of the coordinates of P_U. In particular, a view containing fewer than a τ_k-fraction of $P^{(i)}$ is a deterministic function of fewer than a

τ_k-fraction of P_U. The latter view is independent of A by the privacy of P_U, and therefore so is the former. This applies to every aggregation layer, including the final-layer input $P^{(k-1)}$.

We have shown that every layer input is private against every joint view containing fewer than a τ_k-fraction of one word. Since each input contains at most $2k+2$ equally sized words, the circuit is layerwise private with $\frac{\tau_k}{2k+2} = \frac{\eta\rho^{k-1}}{2k+2}$. □

6.2 Joint Privacy of Nullstellensatz Certificates

For each message $m \in \mathbb{F}_q^D$, let $G(m) \in \mathcal{Z}_{\mathbf{a}}$ be a possibly random polynomial whose distribution may depend on m. In particular, $G(m)|_D = 0$. Conditioned on $G(m)$, sample

$$\mathrm{NS}(G(m)) = (Q_1, \dots, Q_k)$$

uniformly among the degree-bounded tuples satisfying $G(m) = \sum_{i=1}^k Z_i Q_i$.

Suppose that every sufficiently small local view of $G(m)$ is independent of m. We show that a small joint view of $G(m)$ and its certificate is also independent of m. Thus adjoining a uniformly random Nullstellensatz certificate reveals no more than a somewhat larger local view of $G(m)$.

Certificate spaces and relative robust testability. Continue with the degree profile $\mathbf{a}$ and the spaces $V_i, K_i, U_i, W_i, E_i, \mathcal{D}_i, \mathcal{D}$ from Section 4. We identify polynomials with their evaluation vectors on Ω. If $R_i := Z_i Q_i$, then $R_i \in E_i$, and the certificate identity becomes $\sum_{i=1}^k R_i = G$. We use relative robust testability in the form of Corollary 4.4.

Since evaluation is injective under the standing degree bounds, and Z_i is nonzero on S_i, multiplication by $Z_i|_{S_i}$ identifies the evaluation code of the ith quotient space with E_i. Consequently, conditioned on $G = g$, a uniform certificate induces the uniform distribution on the affine space

$$\mathcal{F}(g) = \left\{ (R_1, \dots, R_k) \in \bigoplus_{i=1}^k E_i : \sum_{i=1}^k R_i = g \right\}. \tag{32}$$

Fourier characters. Fix two messages m, m' and a queried view, and denote the resulting random vectors by X_m and $X_{m'}$. Our goal is to prove $X_m \overset{\mathrm{d}}{=} X_{m'}$.

Fix a nontrivial additive character $\psi : \mathbb{F}_q \to \mathbb{C}^\times$. For $\theta \in \mathbb{F}_q^T$, define

$$\chi_\theta(y) = \psi\left(\sum_{x \in T} \theta(x) y(x) \right).$$

As θ varies, these are all the characters of $\mathbb{F}_q^T$. Fourier inversion therefore shows that X_m and $X_{m'}$ have the same distribution if and only if $\mathbb{E}[\chi_\theta(X_m)] = \mathbb{E}[\chi_\theta(X_{m'})]$ for every θ. For example, over $\mathbb{F}_2^n$ the characters are

$$\chi_J(y) = (-1)^{\sum_{j \in J} y_j} = \prod_{j \in J} (-1)^{y_j}, \qquad J \subseteq [n].$$

Thus equality of distributions amounts to equality of all mixed moments of the sign variables $(-1)^{y_j}$, or equivalently of all multilinear moments of the bits y_j.

Lemma 6.2 (Joint privacy of Nullstellensatz certificates)**.** *Assume the hypotheses of Corollary 4.4 for the degree profile* $\mathbf{a}$*, and let* $C(\rho, k)$ *be the constant from that corollary. Suppose that, for every* $T \subseteq \Omega$ *with* $|T| \leqslant s$*, the distribution of* $G(m)|_T$ *is the same for every message* m*. Let* $T_G, T_1, \dots, T_k \subseteq \Omega$ *satisfy*

$$|T_G| + C(\rho, k) \sum_{i=1}^{k} |T_i| \leqslant s.$$

Then the distribution of

$$(G(m)|_{T_G}, Q_1|_{T_1}, \dots, Q_k|_{T_k})$$

is the same for every message m*.*

Proof. Fix a Fourier character of the queried view. Let $\mu, \lambda_1, \dots, \lambda_k \in \mathbb{F}_q^{\Omega}$ be its coefficient vectors, extended by zero outside $T_G, T_1, \dots, T_k$, respectively. We must show that

$$\mathbb{E}\left[\psi\left(\langle \mu, G(m)\rangle + \sum_{i=1}^{k} \langle \lambda_i, Q_i\rangle\right)\right] \tag{33}$$

is independent of m, where $\langle f, g\rangle = \sum_{x\in\Omega} f(x)g(x)$. The expectation is over the randomness of both $G(m)$ and the certificate.

For $x = (x_1, \dots, x_k) \in \Omega$, define

$$\alpha_i(x) = \frac{\lambda_i(x)}{Z_i(x_i)}.$$

The denominator is nonzero because $S_i \cap H_i = \varnothing$. Hence

$$\langle \lambda_i, Q_i\rangle = \langle \alpha_i, R_i\rangle, \qquad \mathrm{wt}(\alpha_i) = \mathrm{wt}(\lambda_i) \leqslant |T_i|.$$

Fix $G = g$ and one tuple $(R_1^0, \dots, R_k^0) \in \mathcal{F}(g)$. Every tuple in $\mathcal{F}(g)$ is uniquely of the form

$$(R_1, \dots, R_k) = (R_1^0, \dots, R_k^0) + (\Delta_1, \dots, \Delta_k),$$

where $(\Delta_1, \dots, \Delta_k)$ is uniform in

$$\mathcal{N} = \left\{(\Delta_1, \dots, \Delta_k) \in \bigoplus_{i=1}^{k} E_i : \sum_{i=1}^{k} \Delta_i = 0\right\}.$$

Therefore

$$\mathbb{E}\left[\psi\left(\sum_{i=1}^{k} \langle \alpha_i, R_i\rangle\right) \,\middle|\, G = g\right] = \psi\left(\sum_{i=1}^{k} \langle \alpha_i, R_i^0\rangle\right) \mathop{\mathbb{E}}_{\Delta\in\mathcal{N}} \psi\left(\sum_{i=1}^{k} \langle \alpha_i, \Delta_i\rangle\right).$$

The last average is zero unless

$$\sum_{i=1}^{k} \langle \alpha_i, \Delta_i\rangle = 0 \qquad \text{for every } \Delta \in \mathcal{N}. \tag{34}$$

Indeed, if the expression on the left is a nonzero $\mathbb{F}_q$-linear map on $\mathcal{N}$, its image is $\mathbb{F}_q$, and the average of the nontrivial character ψ over $\mathbb{F}_q$ is zero.

Suppose (34) holds. We use it to construct a single vector $L_0 \in \mathbb{F}_q^\Omega$ satisfying

$$L_0 - \alpha_i \in \mathcal{D}_i \qquad \text{for every } i;$$

we can then apply relative robust testability to L_0. Let $E = \sum_i E_i$. For $Y \in E$, choose $Y_i \in E_i$ such that $Y = \sum_i Y_i$, and define

$$\ell(Y) = \sum_i \langle \alpha_i, Y_i \rangle.$$

This definition does not depend on the chosen decomposition. Indeed, if $Y = \sum_i Y_i = \sum_i Y_i'$, then $(Y_i - Y_i')_i \in \mathcal{N}$, and hence (34) gives $\sum_i \langle \alpha_i, Y_i \rangle = \sum_i \langle \alpha_i, Y_i' \rangle$. Thus $\ell : E \to \mathbb{F}_q$ is a well-defined linear map. Every such map can be written as a dot product with some $L_0 \in \mathbb{F}_q^\Omega$, so choose L_0 such that

$$\ell(Y) = \langle L_0, Y \rangle \qquad \text{for every } Y \in E.$$

For any $Y_i \in E_i$, the definition of ℓ gives $\ell(Y_i) = \langle \alpha_i, Y_i \rangle$. Therefore

$$\langle L_0 - \alpha_i, Y_i \rangle = 0 \qquad \text{for every } Y_i \in E_i,$$

and hence $L_0 - \alpha_i \in E_i^\perp = \mathcal{D}_i$, as required.

By Corollary 4.4, choose $d \in \mathcal{D}$ nearest to L_0, and set $L = L_0 - d$. Then

$$\mathrm{wt}(L) \leqslant C(\rho, k) \sum_{i=1}^{k} \mathrm{wt}(\alpha_i) \leqslant C(\rho, k) \sum_{i=1}^{k} |T_i|. \tag{35}$$

Moreover, $d \in \mathcal{D} \subseteq \mathcal{D}_i$, so

$$L - \alpha_i \in \mathcal{D}_i \qquad \text{for every } i. \tag{36}$$

The vector L depends only on the fixed character and the spaces $\mathcal{D}_i$, not on g or m.

Since $\mathcal{D}_i = E_i^\perp$, every valid certificate satisfies

$$\begin{aligned} \sum_{i=1}^{k} \langle \lambda_i, Q_i \rangle &= \sum_{i=1}^{k} \langle \alpha_i, R_i \rangle \\ &= \sum_{i=1}^{k} \langle L, R_i \rangle = \left\langle L, \sum_{i=1}^{k} R_i \right\rangle = \langle L, G \rangle. \end{aligned} \tag{37}$$

Condition (34) depends only on the character, not on g. If it fails, the conditional expectation in (34) is zero for every g. If it holds, (37) shows that, conditioned on $G = g$, the character in (33) equals $\psi\big(\langle \mu + L, g \rangle\big)$. Thus (33) is either zero or equals

$$\mathbb{E}\left[\psi\big(\langle \mu + L, G(m) \rangle\big)\right].$$

Finally,

$$\mathrm{wt}(\mu + L) \leqslant |T_G| + C(\rho, k) \sum_{i=1}^{k} |T_i| \leqslant s.$$

The distribution of $G(m)$ on the support of $\mu + L$ is independent of m, so this expectation is independent of m. Every Fourier-character expectation of the queried view is therefore independent of m, and Fourier inversion proves the lemma. □

6.3 Joint Privacy of Circuit Construction

Layerwise privacy does not by itself give privacy for a view spread across several layers, even with a constant-factor loss depending on k. We therefore prove joint privacy under an appropriate measure that is induced by the verifier in our private-PCP construction.

First, let $\mathcal{V}_{\mathcal{C}}$ be the following query procedure. With probability $1/3$, it queries the final output. With probability $1/3$, it chooses one of the L layers uniformly, then chooses a computation gate in that layer uniformly and queries its inputs and output. With probability $1/3$, it runs the circuit tester $\mathcal{T}$. For a random choice r of $\mathcal{V}_{\mathcal{C}}$, let $Q_{\mathcal{V}}(r)$ be the nonempty set of transcript coordinates queried, and let $p(r)$ be the probability of r. Then let μ be the measure induced by $\mathcal{V}_{\mathcal{C}}$ on the transcript: run $\mathcal{V}_{\mathcal{C}}$ and then choose a uniformly random coordinate among those queried.

Lemma 6.3 (Joint privacy)**.** *Continue with the setting and notation of Lemma 6.1. Let $C(\rho, k) \geqslant 1$ be the constant from Corollary 4.4. For $\eta = 1/128$, put*

$$\tau_k = \eta\rho^{k-1}, \qquad \nu_k = \frac{\tau_k}{3(2k+8)(|\Psi|+1)C(\rho,k)}.$$

Then the circuit from Section 5.3 is ν_k-jointly private under the measure μ defined above.

Proof. We first condition on all random-input words used to form R_B, R_U, R_V. Their values are independent of A, and after conditioning they determine fixed polynomials that vanish on D. We suppress these fixed terms and use the dependency bookkeeping of the circuit with the three mask-sampling calls omitted. The viewed mask coordinates are fixed under this conditioning. We average over the masks at the end.

For every assignment, $P_V \in \mathcal{Z}_{\mathbf{v}}$. Conditioned on P_V, the prover samples $(Q_1, \dots, Q_k)$ uniformly from $\Phi_{\mathbf{v}}^{-1}(P_V)$ using fresh randomness.

Let T be a set of transcript coordinates with $\mu(T) < \nu_k$. Note that the final output has μ-measure at least $1/3 > \nu_k$, so T cannot contain it.

Let Y consist of all its remaining coordinates other than the mask coordinates and those of $P_V, Q_1, \dots, Q_k$. None of the coordinates in Y depends on the fresh certificate randomness. Let T_V and T_i be the viewed coordinates of P_V and Q_i, respectively, and put

$$s = |T_V| + C(\rho, k)\sum_{i=1}^{k} |T_i|.$$

We first show that, for every $S \subseteq \Omega$ with $|S| \leqslant s$, the joint distribution of $(Y, P_V|_S)$ is independent of A.

For the measure bookkeeping, let

$$\mathcal{W} = \{P_A, P_B, P_J, P_V, P^{(0)}, \dots, P^{(k-1)}, Q_1, \dots, Q_k\}, \qquad P^{(0)} = P_U,$$

and write $T_W = T \cap W$. Every word $W \in \mathcal{W}$ has its own constituent axis-line test in $\mathcal{T}$. Conditional on this test, a random line together with the factor $1/|Q_{\mathcal{V}}(r)|$ assigns each coordinate of W weight exactly $1/|W|$. Since $\mathcal{T}$ chooses uniformly among $F = 2k + 8$ constituent tests,

$$\mu(T) \geqslant \frac{1}{3F}\sum_{W \in \mathcal{W}} \frac{|T_W|}{|W|}.$$

The other verifier tests only add nonnegative weight.

Fix $S \subseteq \Omega$ with $|S| \leqslant s$, and trace the coordinates of $(Y, P_V|_S)$ back to P_A and P_J. In an honest execution, a coordinate of P_B is determined by at most $|\Psi|$ coordinates of P_A, and a coordinate of P_U or P_V additionally uses one coordinate of P_J. A coordinate of $P^{(i)}$ is determined by a slice of P_U of size $N/|P^{(i)}|$, and these slices are disjoint for distinct coordinates. Replacing the certificate view by a view of P_V costs the factor $C(\rho, k)$ in the definition of s. Since $C(\rho, k) \geqslant 1$, $(Y, P_V|_S)$ is determined by sets $S_A, S_J \subseteq \Omega$ satisfying

$$\begin{aligned} |S_A| + |S_J| &\leqslant (|\Psi| + 1)C(\rho, k)N \sum_{W \in \mathcal{W}} \frac{|T_W|}{|W|} \\ &< 3F(|\Psi| + 1)C(\rho, k)N\nu_k = \tau_k N. \end{aligned}$$

The strict inequality follows from the preceding measure bound and $\mu(T) < \nu_k$; the equality uses $F = 2k+8$ and the definition of ν_k.

Conditioned on (A, J), the polynomials P_A and P_J are independent uniform representations of A and J. By Corollary 4.3 and (2), the joint law of $(P_A|_{S_A}, P_J|_{S_J})$ is independent of (A, J). Hence $(Y, P_V|_S)$ is independent of A, as claimed.

It remains to restore the certificate view. The law of Y is independent of A. Fix any y in its common support. For every $S \subseteq \Omega$ with $|S| \leqslant s$, the conditional distribution of $P_V|_S$ given $Y = y$ is independent of A. Moreover, conditioned on P_V and $Y = y$, the certificate remains uniform in $\Phi_{\mathbf{v}}^{-1}(P_V)$. Applying Lemma 6.2 to $G = P_V$ and its certificate $(Q_1, \ldots, Q_k)$, conditioned on $Y = y$, shows that

$$(P_V|_{T_V}, Q_1|_{T_1}, \ldots, Q_k|_{T_k} \,|\, Y = y)$$

has the same distribution for every assignment A. Combining this with the assignment-independent law of Y proves privacy of the entire view for every fixed realization of the random-input words. Averaging over their assignment-independent distribution proves the claim. □

7 Adversarial Robustness of Construction

Lemma 7.1 (Global robustness)**.** *Assume the setting of Lemma 5.4. There are positive parameters $\epsilon_{\mathrm{in},k}$, $\epsilon_{\mathrm{lay},k}$, and γ_k, each at least $\exp(-O(k))$, for which the circuit–tester pair in Section 5.3 satisfies the robustness clause of Definition 5.3. This holds for arbitrary prover words, arbitrary values on the random-input wires, and every $\epsilon_{\mathrm{lay},k}$-faulty execution. In particular:*

1. *if $\Pr[\mathcal{T}$ rejects$] \leqslant \gamma_k$, then the input is within $\epsilon_{\mathrm{in},k}N$ of C_{in}; and*
2. *if the input is within $\epsilon_{\mathrm{in},k}N$ of $\mathrm{Enc}(A)$, then every $\epsilon_{\mathrm{lay},k}$-faulty execution for which $\Pr[\mathcal{T}$ rejects$] \leqslant \gamma_k$ outputs $f_I(A)$.*

This section proves Lemma 7.1. Small axis-line-test rejection probability first gives nearby codewords. We compare the corrupted execution with a recursively defined reference execution, propagate this comparison through the certificate and aggregation steps, and recover the correct output.

7.1 Soundness of the Axis-Line Test

We use the following consequence of the tensor test. Its proof is standard and follows from [Vid15], but we include it for our setting of unequal component codes and side lengths in Appendix B.

Lemma 7.2 (Axis-line soundness)**.** *Let $r \geqslant 2$, and let the nonzero linear codes $C_i \subseteq \mathbb{F}_q^{S_i'}$, for $i \in [r]$, have relative distance at least Δ. Put*

$$\Omega' = \prod_{i=1}^{r} S_i', \qquad M = |\Omega'|, \qquad \mathcal{C} = \bigotimes_{i=1}^{r} C_i.$$

For a word $w \in \mathbb{F}_q^{\Omega'}$, let $b_i(w)$ be the fraction of i-axis lines ℓ for which $w|_\ell \notin C_i$. If

$$\tau_r(\Delta) = \frac{\Delta^{8r}}{10^6 r^{16}},$$

then

$$\Pr[\mathsf{Test}(w) \textit{ rejects}] = \frac{1}{r}\sum_{i=1}^{r} b_i(w) \geqslant \tau_r(\Delta)\frac{\textit{dist}(w,\mathcal{C})}{M}.$$

By Lemma 5.4, every component Reed–Solomon code used by an axis-line check of $\mathcal{T}$ has relative distance at least $3/4$. We therefore set

$$\Delta_* = \frac{3}{4}, \qquad \Delta_{\mathrm{ten}} = \Delta_*^k, \qquad \tau_k = \frac{\Delta_*^{8k}}{10^6 k^{16}}.$$

Every axis-line test used by $\mathcal{T}$ acts on an r-fold tensor product, for some $1 \leqslant r \leqslant k$, of these component codes. The corresponding tensor code therefore has relative distance at least

$$\Delta_*^r \geqslant \Delta_*^k = \Delta_{\mathrm{ten}}.$$

If $r \geqslant 2$, then $\tau_r(\Delta_*) \geqslant \tau_k$, so Lemma 7.2 applies. If $r = 1$, the unique axis line is the whole word, and the test rejects exactly when the word is not in the component code. Consequently, for every word w of length M to which $\mathcal{T}$ applies an axis-line test, if $\mathcal{C}$ denotes the corresponding tensor code, then

$$\Pr[\mathsf{Test}(w) \text{ rejects}] \geqslant \tau_k \frac{\mathrm{dist}(w,\mathcal{C})}{M}. \tag{38}$$

The next lemma is used after each computation step. Its first assertion gives a nearby codeword; its second identifies this codeword with a specified reference codeword whenever the latter is already sufficiently close.

Lemma 7.3 (Recovering the reference codeword)**.** *Let w be a word of length M to which $\mathcal{T}$ applies an axis-line test, and let $\mathcal{C}$ be the corresponding tensor code. If $\Pr[\mathsf{Test}(w)$ rejects$] \leqslant t$, then some $c' \in \mathcal{C}$ satisfies*

$$\textit{dist}(w,c') \leqslant \frac{tM}{\tau_k}.$$

Moreover, if $c \in \mathcal{C}$ satisfies $\textit{dist}(w,c) + \frac{tM}{\tau_k} < \Delta_{\mathrm{ten}} M$, then $c' = c$.

Proof. The first assertion follows from (38). Under the second assertion's hypothesis, the triangle inequality gives

$$\mathrm{dist}(c,c') < \Delta_{\mathrm{ten}} M.$$

Since $\mathcal{C}$ has relative distance at least Δ_{ten}, we must have $c = c'$. □

Lemma 7.4 (SampleMask proximity)**.** *Consider a mask-sampling call with parameter* $\mathbf{a}$ *and component input words* $W_1, \dots, W_k$ *on* Ω*. Let* p_{mask} *be the rejection probability of the constituent check that chooses* $i \in [k]$ *uniformly and applies the corresponding axis-line test to* W_i*. If* $p_{\mathrm{mask}} \leqslant t$*, then there are* $R_i^* \in \mathcal{Q}_i(\mathbf{a})$ *such that, for*

$$R^* = \sum_{i=1}^{k} Z_i R_i^* \in \mathcal{Z}_{\mathbf{a}} \qquad \textit{and} \qquad R_W(x) = \sum_{i=1}^{k} Z_i(x_i) W_i(x),$$

we have

$$\textit{dist}\big(R_W, \mathrm{ev}_\Omega(R^*)\big) \leqslant \frac{ktN}{\tau_k}.$$

Proof. Let p_i be the rejection probability of the axis-line test on W_i. Then $p_{\mathrm{mask}} = \frac{1}{k}\sum_{i=1}^{k} p_i$.

By Lemma 7.3, there are $R_i^* \in \mathcal{Q}_i(\mathbf{a})$ such that

$$\mathrm{dist}\big(W_i, \mathrm{ev}_\Omega(R_i^*)\big) \leqslant \frac{p_i N}{\tau_k}.$$

We can thus use a union bound to bound the size of the disagreement set between R_W and $\mathrm{ev}_\Omega(R^*)$,

$$\mathrm{dist}\big(R_W, \mathrm{ev}_\Omega(R^*)\big) \leqslant \frac{N}{\tau_k}\sum_{i=1}^{k} p_i = \frac{k p_{\mathrm{mask}} N}{\tau_k} \leqslant \frac{ktN}{\tau_k}.$$

□

7.2 Proof of Lemma 7.1

We prove that, if $\mathcal{T}$ has small rejection probability, every computed word remains close to a reference codeword. For a circuit-computed word, this reference is determined recursively from the references for the preceding step. For an arbitrary prover word, its test supplies a nearby codeword. For a mask-sampling primitive, its test shows that the mask values used by the computation gates are close to a polynomial that vanishes on D. We use these nearby words only as references in the analysis. We then propagate this comparison through the constraint, certificate, and aggregation steps to show that the circuit produces the correct output despite the faulty gates.

Proof of Lemma 7.1. Fix a sufficiently large absolute constant $C \geqslant 1$, then a sufficiently small absolute constant $c_0 > 0$, and set

$$\epsilon_* = \frac{c_0 \Delta_{\mathrm{ten}}}{k}, \qquad \gamma_k = \frac{c_0 \tau_k \epsilon_*}{k^2}. \tag{39}$$

Set

$$\epsilon_{\mathrm{in},k} = \epsilon_*, \qquad \epsilon_{\mathrm{lay},k} = \gamma_k. \tag{40}$$

The calculations below show that one choice of C, c_0 works for every $k \geqslant 2$.

The tester has $F = 2k+8 \leqslant 6k$ tests. For a fixed execution, if their rejection probabilities are $p_1, \dots, p_F$, then

$$\Pr[\mathcal{T} \text{ rejects}] = \frac{1}{F}\sum_{s=1}^{F} p_s. \tag{41}$$

Thus, if $\Pr[\mathcal{T} \text{ rejects}] \leqslant \gamma_k$, every test rejects with probability at most $t := F\gamma_k$. Moreover,

$$\frac{kt}{\tau_k} = c_0 \frac{F}{k} \epsilon_* \leqslant 6c_0\epsilon_* < \frac{\epsilon_*}{4} \tag{42}$$

for sufficiently small c_0. A direct calculation gives

$$\epsilon_{\mathrm{in},k} = \Theta\left(\frac{(3/4)^k}{k}\right), \qquad \epsilon_{\mathrm{lay},k}, \gamma_k = \Theta\left(\frac{(3/4)^{9k}}{k^{19}}\right).$$

In particular, all three parameters are at least $\exp(-O(k))$.

We prove the two assertions in Lemma 7.1 separately. First, every word in C_{in} is a degree-$\mathbf{d}$ evaluation, so $C_{\mathrm{in}} \subseteq \mathcal{C}_{\mathbf{d}}$. Conversely, if $c = \mathrm{ev}_\Omega(P) \in \mathcal{C}_{\mathbf{d}}$, then $P \in \mathcal{R}_{\mathbf{d}}(P|_D)$, so $c \in \mathrm{Enc}(P|_D)$. Thus $C_{\mathrm{in}} = \mathcal{C}_{\mathbf{d}}$.

If $\mathrm{dist}(W_0, C_{\mathrm{in}}) > \epsilon_* N$, then (38) shows that the test on W_0 rejects with probability greater than $\tau_k\epsilon_*$. Therefore

$$\Pr[\mathcal{T} \text{ rejects}] > \frac{\tau_k\epsilon_*}{F} > \gamma_k,$$

where the last inequality follows from $c_0 F/k^2 < 1$. This proves assertion (1).

For assertion (2), fix A such that $\mathrm{dist}(W_0, \mathrm{Enc}(A)) \leqslant \epsilon_* N$, and choose $\mathrm{ev}_\Omega(P_A^*) \in \mathrm{Enc}(A)$ within distance $\epsilon_* N$ of W_0. Consider an $\epsilon_{\mathrm{lay},k}$-faulty execution for which $\Pr[\mathcal{T} \text{ rejects}] \leqslant \gamma_k$.

For each nonfinal layer ℓ, let M_ℓ be the size of the grid supporting its computation outputs; for an aggregation layer, this is the shorter target grid. Let G_ℓ be the number of computation gates in that layer. Since $G_\ell \leqslant 2M_\ell$, the layer contains fewer than $2\gamma_k M_\ell$ faulty gates. The final layer has one computation gate. Since the fault bound is strict and $\gamma_k < 1$, this gate cannot be faulty.

Applying Lemma 7.3 to a tested word on a grid of size M gives a codeword within distance tM/τ_k. For an arbitrary prover word, we take this as its reference codeword. If the word already has a prescribed reference codeword at distance at most $Ck\epsilon_* M$, then the two codewords agree: indeed,

$$Ck\epsilon_* + \frac{t}{\tau_k} < \Delta_{\mathrm{ten}}$$

for sufficiently small c_0, because $k\epsilon_* = c_0\Delta_{\mathrm{ten}}$. Thus the test places the word within relative distance $t/\tau_k < \epsilon_*/4$ of its prescribed reference. This also applies to the initial word W_0.

For a mask call, Lemma 7.4 and (42) give a reference polynomial that vanishes on D whose evaluations differ from the mask values used by the computation gates in fewer than $\epsilon_* N/4$ coordinates.

We use one line-decoding calculation. If a word on a grid of size M differs from its reference codeword in at most $\epsilon_* M$ coordinates, any i-axis line decoded differently from the reference line contains at least $\Delta_*|S_i|/2$ errors. Hence there are at most

$$\frac{2\epsilon_* M}{\Delta_*|S_i|}$$

such lines. They affect an $O(\epsilon_*)$-fraction of the coordinates in constraint evaluation and an $O(\epsilon_*)$-fraction of the outputs in an aggregation.

Constraint and indicator layers. Write $w_A = W_0$, and let w_B, w_J, w_U, w_V denote the corresponding actual circuit words. The test on w_A identifies its nearby codeword with $\mathrm{ev}_\Omega(P_A^*)$, and the constraint-mask call supplies $R_B^* \in \mathcal{Z}_{\mathbf{e}}$. Set

$$P_B^* = B_{P_A^*} + R_B^*.$$

An incorrect constraint output can arise only from a bad decoded line in one of the k directions, an incorrect mask value, or a faulty gate. These affect at most $Ck\epsilon_* N$ coordinates in total. Thus w_B is within that distance of $\mathrm{ev}_\Omega(P_B^*)$. Its test identifies this prescribed reference and places w_B within relative distance $\epsilon_*/4$ of it. Since $R_B^*|_D = 0$ and $P_A^*|_D = A$, we have

$$B := P_B^*|_D = B_{P_A^*}|_D = B_A|_D.$$

The test on the arbitrary prover word w_J supplies a reference polynomial P_J^*; we assume no relation between P_J^* and the honest prover word. The two mask calls similarly supply $R_U^* \in \mathcal{Z}_{\mathbf{u}}$ and $R_V^* \in \mathcal{Z}_{\mathbf{v}}$. Define

$$P_U^* = P_B^* P_J^* + R_U^*, \qquad P_V^* = P_B^*(1 - P_B^* P_J^*) + R_V^*. \tag{43}$$

An indicator output can be incorrect only where one of its tested inputs or mask values is incorrect, or where its gate is faulty. Hence each of w_U, w_V is within $Ck\epsilon_* N$ of its prescribed reference codeword. Their tests identify these references and place the two words within relative distance $\epsilon_*/4$ of them.

Vanishing certificate. For every $i \in [k]$, the test on w_{Q_i} supplies a polynomial $Q_i^* \in \mathcal{Q}_i(\mathbf{v})$ whose evaluation is within relative distance $\epsilon_*/4$ of w_{Q_i}. Define

$$e(x) = w_V(x) - \sum_{i=1}^{k} Z_i(x_i) w_{Q_i}(x), \qquad E^* = P_V^* - \sum_{i=1}^{k} Z_i Q_i^*.$$

The polynomial E^* has degree at most $\mathbf{v}$, and e differs from its evaluation in fewer than $(k+1)\epsilon_* N/4$ coordinates. If $E^* \neq 0$, then its evaluation has weight at least $\Delta_{\mathrm{ten}} N$. Consequently,

$$\frac{\mathrm{wt}(e)}{N} > \Delta_{\mathrm{ten}} - \frac{(k+1)\epsilon_*}{4} > t$$

for sufficiently small c_0. The certificate test would therefore reject with probability greater than t, a contradiction. Hence $E^* = 0$.

Let J, U, V be the restrictions to D of P_J^*, P_U^*, P_V^*. Restricting $E^* = 0$ to D gives $V = 0$, while

$$U = BJ, \qquad V = B(1 - BJ) = B(1 - U).$$

Therefore

$$U(z) = \mathbf{1}[B(z) \neq 0] \qquad (z \in D). \tag{44}$$

Aggregation. Let $w^{(0)} = w_U$, set $P^{(0),*} = P_U^*$, and define

$$P^{(i),*} = \Sigma_i P^{(i-1),*} \qquad (1 \leqslant i \leqslant k-1). \tag{45}$$

Let

$$\Omega^{(i)} = \prod_{j=i+1}^{k} S_j, \qquad M^{(i)} = |\Omega^{(i)}|,$$

with $\Omega^{(0)} = \Omega$.

We proceed by induction over $i = 1, \ldots, k-1$. The test on $w^{(i-1)}$ places it within relative distance $\epsilon_*/4$ of $\mathrm{ev}(P^{(i-1),*})$. The line-decoding calculation shows that only an $O(\epsilon_*)$-fraction of the aggregation outputs

can then be incorrect. Including the faulty gates, $w^{(i)}$ is within $Ck\epsilon_* M^{(i)}$ of $\mathrm{ev}(P^{(i),*})$. Its test identifies this prescribed reference and again places the word within relative distance $\epsilon_*/4$ of it.

In particular, the final univariate word is within relative distance $\epsilon_*/4 < \Delta_*/2$ of $\mathrm{ev}_{S_k}(P^{(k-1),*})$. The final gate therefore recovers this codeword and computes

$$\nu^* = \sum_{a \in H_k} P^{(k-1),*}(a) = \sum_{z \in D} U(z).$$

Since $B = B_A|_D$ and $U(z) = \mathbf{1}[B(z) \neq 0]$, this is the number of points on which B_A is nonzero on D, viewed as an element of $\mathbb{F}_q$. Because $q > |D|$, it vanishes exactly when $B_A|_D \equiv 0$. The final output is therefore

$$\mathbf{1}[\nu^* = 0] = f_I(A),$$

which proves assertion (2) and the lemma. □

8 Private and Zero-Knowledge PCPs for NP

Ben-Sasson and Sudan [BS08] reduced 3SAT to the Zero-on-Subcube problem for a Reed–Solomon code. Ben-Sasson et al. [BKK+16] adapted this reduction to tensor products of algebraic-geometry codes. We need the analogous reduction for tensor products of Reed–Solomon codes. Our construction follows the same outline but we include the proof in Appendix A for completeness.

8.1 Arithmetization: Reducing 3SAT to Zero-on-Subcube

Lemma 8.1 (Zero-on-Subcube arithmetization)**.** *Fix $k \geqslant 2$ and a prime q. Let Φ be a 3SAT[4] instance of size n, and let $D = \prod_{i=1}^{k} H_i$ be as in Definition 5.1. Suppose that $2^{\sum_i \lfloor \log_2 h_i \rfloor} \geqslant cn \log(n+1)$, for some absolute constant c. One can construct a k-dimensional zero-on-subcube instance $I_\Phi = (\mathbb{F}_q, D, \Psi, C_\Phi)$ such that:*

1. *every map in Ψ is the identity or a one-coordinate shift $\sigma_{i,t}$;*
2. *$|\Psi| = O(\log n)$;*
3. *Φ is satisfiable if and only if some table $A : D \to \mathbb{F}_q$ satisfies I_Φ;*
4. *C_Φ has an explicit polynomial-size description.*

Use the same formula for $\sigma_{i,t}$ on polynomial variables. For $P \in \mathbb{F}_q[X_1, \ldots, X_k]$ with $\deg_{X_i} P \leqslant a_i$, set

$$B_P(\mathbf{X}) = C_\Phi\big(\mathbf{X}, (P(\psi(\mathbf{X})))_{\psi \in \Psi}\big).$$

For every $\mathbf{x} \in \mathbb{F}_q^k$, this value depends on at most $O(\log n)$ distinct values of P. Moreover, writing $A = P|_D$,

$$B_P\big|_D = B_A\big|_D, \qquad \deg_{X_i} B_P \leqslant 3a_i + h_i - 1 \quad (i \in [k]).$$

Consequently,

$$\Phi \text{ is satisfiable} \iff \exists P \text{ with } \deg_{X_i} P < h_i \text{ for every } i \text{ and } B_P\big|_D \equiv 0.$$

For an absolute constant C_{alg}, the construction takes at most $(k + |D| + n + \log q)^{C_{\mathrm{alg}}}$ operations.

8.2 Private PCPs

We start by defining the formal model of a private PCP.

Definition 8.2 (Private PCP for Circuit-SAT)**.** *For a Boolean circuit $\Phi : \{0,1\}^m \to \{0,1\}$, define*

$$\mathrm{CircuitSAT} = \{\Phi : \exists w \in \{0,1\}^m,\ \Phi(w) = 1\}.$$

A private PCP for CircuitSAT *consists of a randomized prover Π, where for every circuit Φ and every (even possibly unsatisfying assignment) $w \in \{0,1\}^m$, $\Pi(\Phi, w)$ is a polynomial-time samplable distribution of honest encodings of w in Σ^M, and a nonadaptive probabilistic polynomial-time oracle verifier V making at most Q_{qry} queries. It has perfect completeness and soundness s if*

$$\Pr_{\pi \leftarrow \Pi(\Phi,w),V}[V^\pi(\Phi) = 1] = 1 \quad \textit{whenever } \Phi(w) = 1,$$

and

$$\Pr_V[V^\pi(\Phi) = 1] \leqslant s \quad \textit{for every } \Phi \notin \mathrm{CircuitSAT} \textit{ and every } \pi \in \Sigma^M.$$

To define privacy, let $\mu_{V,\Phi}$ be the natural measure induced by the verifier's queries: sample the verifier's randomness and then choose a uniformly random coordinate among those queried. The PCP is ν-private *if, for every Boolean circuit $\Phi : \{0,1\}^m \to \{0,1\}$, every $w, w' \in \{0,1\}^m$, and every $T \subseteq [M]$ satisfying $\mu_{V,\Phi}(T) < \nu$, the marginal distribution on T is the same for honest encodings of w, w',*

$$\Pi(\Phi, w)|_T \overset{\mathrm{d}}{=} \Pi(\Phi, w')|_T.$$

We now choose the parameters from Lemma 5.4, apply Lemma 8.1, use Theorem 5.5 to obtain a private and robust circuit for the resulting Zero-on-Subcube instance, and apply the usual Cook–Levin gate checks together with its tester $\mathcal{T}$.

Corollary 8.3 (Private PCP)**.** *Assume Conjecture 4.2, and put $\rho = \rho(k, 1/128, 4)$. For every $k \geqslant 2$, a* CircuitSAT *instance of size n has a private PCP of size $M \leqslant n \log(n+1) \exp(O(k \log k))$ over alphabet $\mathbb{F}_q$, where $q \leqslant M^6$. The verifier makes $O(kM^{1/k})$ queries, completeness is perfect, and the soundness error is at most $1 - \exp(-O(k))$. It is ν_k-private, where*

$$\tau_k = \frac{\rho^{k-1}}{128}, \qquad \nu_k = \frac{\tau_k}{3(2k+8)(|\Psi|+1)C(\rho,k)}, \qquad C(\rho,k) = \exp\bigl(O(k^3 + k^2 \log(1/\rho))\bigr),$$

where $|\Psi| = O(\log n)$. In particular,

$$\nu_k \geqslant \frac{1}{\log(n+1)} \exp\bigl(-O(k^3 + k^2 \log(1/\rho))\bigr).$$

Proof. Apply Lemma 5.4, and write $N = |\Omega|$ and $s_i = |S_i|$ for the resulting evaluation grid and its side lengths. Apply Lemma 8.1 and theorem 5.5. The proof contains the input, auxiliary, random-input, and computation-output wires of the resulting circuit, with each physical wire stored once. There are $4k + 5$ full-grid words, together with the shorter aggregation words and the final output bit. Thus its length is

$$M = (4k+5)N + \sum_{i=1}^{k-1} \frac{N}{s_1 \cdots s_i} + 1 = O(kN).$$

Since every $s_i \geqslant 2$, the sum is less than N. Consequently, Lemma 5.4 gives

$$M \leqslant n \log(n+1) \exp(O(k \log k)), \qquad N < q \leqslant N^6 \leqslant M^6.$$

With probability $1/3$ each, the verifier performs one of the following tests:

1. query the final output and accept if and only if it is 1;
2. choose a layer and then a computation gate uniformly, and check its claimed output against its inputs;
3. run the circuit tester $\mathcal{T}$ on the claimed transcript.

Let $\mu_{V,\Phi}$ be the measure induced by these queries as in Definition 8.2. By construction, it is exactly the transcript measure μ used in Lemma 6.3. Every test makes $O(kN^{1/k})$ queries by Theorem 5.5, and hence $O(kM^{1/k})$ queries since $N \leqslant M$. Honest correctness gives perfect completeness.

For soundness, suppose first that the claimed output is not 1. The first test rejects. Otherwise, if some layer has at least an $\epsilon_{\mathrm{lay},k}$-fraction of inconsistent computation gates, the second test rejects with probability at least $\epsilon_{\mathrm{lay},k}/(3L)$. If not, the proof describes an $\epsilon_{\mathrm{lay},k}$-faulty execution. If its input is far from C_{in}, Lemma 7.1 implies that $\mathcal{T}$ rejects with probability greater than γ_k. If it is close to $\mathrm{Enc}(A)$, then unsatisfiability gives $f_{I_\Phi}(A) = 0$, while the claimed output is 1, and the same lemma again implies that $\mathcal{T}$ rejects with probability greater than γ_k. Thus the rejection probability is at least

$$\frac{1}{3} \min\left\{1, \frac{\epsilon_{\mathrm{lay},k}}{L}, \gamma_k\right\} = \exp(-O(k)).$$

Finally, an honest proof is the circuit transcript, with each physical wire stored once. Its verifier-induced measure is therefore the measure μ from Lemma 6.3, which gives ν_k-privacy. The remaining parameter bounds follow from Lemmas 5.4 and 8.1 and theorem 5.5. □

Remark 8.4 (Lopsidedness)**.** *The verifier-induced measure $\mu_{V,\Phi}$ is not uniform over the PCP proof, which corresponds to the transcript of our circuit construction. Indeed, coordinates in the upper layer of the circuit are queried much more than coordinates in the lower layers (in fact, the output wire has measure at least $1/3$). Such lopsidedness seems essential for obtaining privacy while allowing the verifier to distinguish YES from NO instances. Thus Corollary 8.3 gives a PCP whose proof coordinates are non-uniformly weighted by the verifier.*

8.3 A Non-Lopsided Perfect-Zero-Knowledge PCP

We can remove this lopsidedness by relaxing to a zero-knowledge PCP.

Definition 8.5 (Perfect-zero-knowledge PCP)**.** *For an NP relation $\mathcal{R}$ and a possibly adaptive oracle algorithm V^*, let $\mathrm{View}_{V^*}(x;\pi)$ contain its random coins, ordered query–answer transcript, and output. A randomized PCP for $\mathcal{R}$ is τ-query* perfect zero knowledge *if, for every V^* making at most τ queries, there is a randomized simulator S_{V^*} such that, for every $(x,w) \in \mathcal{R}$,*

$$S_{V^*}(x) \stackrel{\mathrm{d}}{=} \mathrm{View}_{V^*}(x;\pi), \qquad \pi \leftarrow \Pi(x,w).$$

The equality is exact, and the simulator does not receive w. We call the PCP efficiently perfect zero knowledge if S_{V^*} is polynomial-time whenever V^* is. We will not require efficiency below.

There are many ways to modify our circuit construction to obtain such a PCP. The simplest way is to ask the honest prover in the YES case to just output a Nullstellensatz certificate. In particular, for a satisfying table $A : D \to \mathbb{F}_q$, sample P_A uniformly from $\mathcal{R}_{\mathbf{d}}(A)$, put

$$P_B = B_{P_A} = C_\Phi\big(\mathbf{X}, (P_A(\psi(\mathbf{X})))_{\psi\in\Psi}\big),$$

and sample uniformly a tuple $(Q_1, \dots, Q_k) \in \bigoplus_i \mathcal{Q}_i(\mathbf{e})$ satisfying $P_B = \sum_i Z_i Q_i$. Such tuples exist because $P_B|_D = B_A|_D \equiv 0$. The proof consists of the evaluations of $P_A, P_B, Q_1, \dots, Q_k$ on Ω.

With probability $1/3$ each, the verifier performs one of three tests:

1. choose one proof table uniformly and run the axis-line test for its tensor Reed–Solomon code;
2. choose $x \in \Omega$ uniformly, query and decode the k axis lines of P_A through x, and use these lines to check
$$P_B(x) = C_\Phi\big(x, (P_A(\psi(x)))_{\psi\in\Psi}\big);$$
3. choose $x \in \Omega$ uniformly and check $P_B(x) = \sum_i Z_i(x_i)Q_i(x)$.

The second test uses the same axis-line decoding as the constraint layer of the circuit. Since every shift preserves Ω, it recovers all the required shifted values. Every proof table has length N, and each test makes $O(kN^{1/k})$ queries.

Corollary 8.6 (Non-lopsided perfect zero knowledge)**.** *Under the hypotheses of Corollary 8.3, the preceding PCP has size N, query complexity $O(kN^{1/k})$, perfect completeness, and soundness error at most $1 - \exp(-O(k))$. It is τ-query perfect zero knowledge, with no efficiency guarantee for the simulator, for every*

$$\tau < q_{\mathrm{zk}} := \frac{\tau_k N}{(|\Psi|+1)C(\rho,k)} \geqslant \frac{N}{\log(n+1)} \exp\big(-O(k^3 + k^2\log(1/\rho))\big)\,.$$

The proof is essentially the same as that of Corollary 8.3 so we omit it.

Acknowledgements

We thank Umesh Vazirani for useful discussions and Anurag Anshu and Quynh Nguyen for discussions about [ABN24]. M.B. also thanks the Simons Institute for hosting her during the "Summer Cluster on Quantum Computing, 2026" program, where part of the work was carried out.

AI Disclosure. ChatGPT was used throughout the research and writing process. The proof of Theorem 4.11 and the main ideas in Sections 4.3 and 6.2 were generated by ChatGPT 5.6 Sol. The authors independently verified all AI-generated arguments and take full responsibility for all the results in this paper.

References

[AB97] Dorit Aharonov and Michael Ben-Or. Fault-tolerant quantum computation with constant error. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 176–188. ACM, 1997.

[ABG+26] Anurag Anshu, Nikolas P. Breuckmann, Louis Golowich, Quynh T. Nguyen, and Umesh Vazirani. Classical adversarial fault-tolerance and PCPs. *arXiv preprint arXiv:2608.16860*, 2026.

[ABN24] Anurag Anshu, Nikolas P. Breuckmann, and Quynh T. Nguyen. Circuit-to-hamiltonian from tensor networks and fault tolerance. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC '24, page 585–595. ACM, June 2024.

[ALM+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.

[Alo99] Noga Alon. Combinatorial nullstellensatz. *Combinatorics, Probability and Computing*, 8(1–2):7–29, 1999.

[AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.

[BGV26] Nikolas P. Breuckmann, Louis Golowich, and Umesh Vazirani. Fault-tolerant quantum computation with adversarial errors. *arXiv preprint arXiv:2608.16857*, 2026.

[BKK+16] Eli Ben-Sasson, Yohay Kaplan, Swastik Kopparty, Or Meir, and Henning Stichtenoth. Constant rate pcps for circuit-sat with sublinear query complexity. *J. ACM*, 63(4):32:1–32:57, 2016.

[BOGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, page 1–10, New York, NY, USA, 1988. Association for Computing Machinery.

[BS08] Eli Ben-Sasson and Madhu Sudan. Short pcps with polylog query complexity. *SIAM J. Comput.*, 38(2):551–607, 2008.

[CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC)*, pages 11–19. ACM, 1988.

[CZ13] Pietro Corvaja and Umberto Zannier. Greatest common divisors of $u-1, v-1$ in positive characteristic and rational points on curves over finite fields. *Journal of the European Mathematical Society*, 15(5):1927–1942, 2013.

[Din07] Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007.

[DLV24] Irit Dinur, Ting-Chun Lin, and Thomas Vidick. Expansion of high-dimensional cubical complexes: with application to quantum locally testable codes. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 379–385. IEEE, 2024.

[GG25] Louis Golowich and Venkatesan Guruswami. Near-asymptotically-good quantum codes with transversal CCZ gates and sublinear-weight parity-checks. In *66th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2025, Sydney, Australia, December 14-17, 2025*, pages 1561–1569. IEEE, 2025.

[GOS25] Tom Gur, Jack O'Connor, and Nicholas Spooner. A zero-knowledge PCP theorem. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC)*, pages 986–994. ACM, 2025.

[IMS12] Yuval Ishai, Mohammad Mahmoody, and Amit Sahai. On efficient zero-knowledge PCPs. In Ronald Cramer, editor, *Theory of Cryptography Conference (TCC)*, volume 7194 of *Lecture Notes in Computer Science*, pages 151–168. Springer, 2012.

[KP25] Gleb Kalachev and Pavel Panteleev. Maximally extendable product codes are good coboundary expanders. In *66th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2025, Sydney, Australia, December 14-17, 2025*, pages 1512–1524. IEEE, 2025.

[KPT97] Joe Kilian, Erez Petrank, and Gábor Tardos. Probabilistically checkable proofs with zero knowledge. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC)*, pages 496–505. ACM, 1997.

[PK22] Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical LDPC codes. In *STOC 2022*, pages 375–388. ACM, 2022.

[PS94] Alexander Polishchuk and Daniel A Spielman. Nearly-linear size holographic proofs. In *STOC 1994*, pages 194–203, 1994.

[RZW26] Noga Ron-Zewi and Mor Weiss. A note on the equivalence between zero-knowledge and quantum CSS codes, 2026.

[Smi01] Adam Smith. Multi-party quantum computation. *arXiv preprint quant-ph/0111030*, 2001.

[Vid15] Michael Viderman. A combination of testability and decodability by tensor products. *Random Structures & Algorithms*, 46(3):572–598, 2015.

[Xyl09] Triantafyllos Xylouris. On linnik's constant, 2009.

A NP-Hardness of Zero-on-Subcube

This section proves Lemma 8.1. We use the hypercube arithmetization of Ben-Sasson and Sudan (BSS) [BS08, Section 4.2 in ECCC version] and embed the cube into the product domain D. Like them we start from a $3\mathrm{SAT}[4]$ instance Φ which consists of Boolean variables $x_1, \dots, x_v$ and clauses with at most three signed literals, where every variable appears in at most 4 clauses. We write $n = v + m$, where m is the number of clauses.

A.1 The BSS hypercube instance

We begin with the hypercube instance extracted from the proof of Theorem 2 of [BS08].

Lemma A.1 (BSS hypercube instance)**.** *From a $3\mathrm{SAT}[4]$ instance Φ of size n, one can construct in deterministic polynomial time an integer $r_{\mathrm{cube}} = O(\log n)$, with $2^{r_{\mathrm{cube}}} = O(n \log n)$, a set $\mathcal{T}$ of clause types, and the following pairwise-disjoint subsets of $\{0,1\}^{r_{\mathrm{cube}}}$:*

$$\mathrm{Var}, \qquad E_{\ell,b} \quad (\ell \in [r_{\mathrm{cube}}],\ b \in \{0,1\}), \qquad \mathrm{Clause}_\tau \quad (\tau \in \mathcal{T}).$$

Each clause type has the form

$$\tau = (d; \ell_1, \dots, \ell_d; b_1, \dots, b_d; c_1, \dots, c_d),$$

where $1 \leqslant d \leqslant 3$, $\ell_t \in [r_{\mathrm{cube}}]$, *and* $b_t, c_t \in \{0,1\}$. *Moreover,* $|\mathcal{T}| \leqslant m$. *Every* $u \in E_{\ell,b}$ *satisfies* $u_\ell = b$, *and every* $u \in \mathrm{Clause}_\tau$ *satisfies* $u_{\ell_t} = b_t$ *for every* $t \in [d]$. *The bit* c_t *is the value that satisfies the* t*-th literal of the clause.*

The formula Φ *is satisfiable if and only if there is a table* $A : \{0,1\}^{r_{\mathrm{cube}}} \to \mathbb{F}_q$ *satisfying*

$$\begin{aligned} A(u)(A(u)-1) &= 0 && (u \in \mathrm{Var}), \\ A(u) - A(u+e_\ell) &= 0 && (u \in E_{\ell,b}), \\ \prod_{t=1}^{d} \big(A(u+e_{\ell_t}) - c_t\big) &= 0 && (u \in \mathrm{Clause}_\tau), \end{aligned} \tag{46}$$

where addition is in $\mathbb{F}_2^{r_{\mathrm{cube}}}$.

A.2 Embedding the cube into the product domain

Recall that $H_i = \zeta_i \Gamma_i$, where $\Gamma_i = \langle \omega_i \rangle$ has order h_i, and $D = \prod_{i=1}^k H_i$. Lemma A.1 gives $2^{r_{\mathrm{cube}}} = O(n \log(n+1))$, so the hypothesis of Lemma 8.1 gives

$$\sum_{i=1}^{k} \lfloor \log_2 h_i \rfloor \geqslant r_{\mathrm{cube}}. \tag{47}$$

Choose integers $r_i \leqslant \lfloor \log_2 h_i \rfloor$ with $\sum_i r_i = r_{\mathrm{cube}}$, and partition the r_{cube} cube coordinates into blocks of sizes $r_1, \dots, r_k$. We index the bits in block i by $j \in \{0, \dots, r_i - 1\}$. For $u \in \{0,1\}^{r_{\mathrm{cube}}}$, write $u^{(i)}$ for its i-th block and put

$$[u^{(i)}] = \sum_{j=0}^{r_i - 1} u_j^{(i)} 2^j.$$

Define $\phi(u) = \big(\zeta_1 \omega_1^{[u^{(1)}]}, \dots, \zeta_k \omega_k^{[u^{(k)}]}\big) \in D$. This is the blockwise analogue of the map $u \mapsto \omega^{[u]}$ in [BS08, Section 4.2]. If cube coordinate ℓ is bit j of block i, define, for $b \in \{0,1\}$,

$$\beta_{\ell,b}(x_1, \dots, x_k) = \big(x_1, \dots, \omega_i^{(-1)^b 2^j} x_i, \dots, x_k\big).$$

We regard the $2r_{\mathrm{cube}}$ maps $\beta_{\ell,b}$ as labeled shifts, allowing two labels to define the same map. If ℓ is bit j of block i, then $\beta_{\ell,b} = \sigma_{i,t}$ for $t \equiv (-1)^b 2^j \pmod{h_i}$.

Lemma A.2 (Cube embedding into D)**.** *The map* ϕ *is injective. If* $u_\ell = b$, *then*

$$\phi(u + e_\ell) = \beta_{\ell,b}(\phi(u)).$$

Every $\beta_{\ell,b}$ *preserves* D. *Moreover, if* $\Gamma_i \leqslant S_i$ *and* $\Omega = \prod_i S_i$, *then every* $\beta_{\ell,b}$ *also preserves* Ω.

Proof. For every i, the integer $[u^{(i)}]$ lies between 0 and $2^{r_i} - 1 < h_i$. Thus equality of the i-th coordinates of $\phi(u)$ and $\phi(v)$ implies $[u^{(i)}] = [v^{(i)}]$, and hence $u^{(i)} = v^{(i)}$. This proves injectivity. If $u_\ell = b$ and ℓ is bit j of block i, flipping that bit changes $[u^{(i)}]$ by $(-1)^b 2^j$, which proves the displayed identity. Finally, the multiplier $\omega_i^{(-1)^b 2^j}$ belongs to Γ_i. It therefore preserves the coset $H_i = \zeta_i \Gamma_i$, and it also preserves S_i when $\Gamma_i \leqslant S_i$. □

A.3 Proof of Lemma 8.1

The proof of the lemma now follows along the lines of Theorem 2 in BSS, but we reproduce it here for completeness. Map the sets of the hypercube in Lemma A.1 to D using ϕ, and use the same symbols for their indicator functions on D. Thus Var, $E_{\ell,b}$, and Clause_τ agree with the corresponding cube indicators on $\phi(\{0,1\}^{r_{\mathrm{cube}}})$ and vanish outside this image. For a table $T : D \to \mathbb{F}_q$, let $\widehat{T}(\mathbf{X})$ denote its unique interpolation with $\deg_{X_i} \widehat{T} < h_i$ for every i.

The list Ψ contains the identity map and, for every $\ell \in [r_{\mathrm{cube}}]$ and $b \in \{0,1\}$, one labeled copy of $\beta_{\ell,b}$. Denote the corresponding formal variables by Y_0 and $Y_{\ell,b}$. Thus $|\Psi| = 1 + 2r_{\mathrm{cube}} = O(\log n)$. Two labels remain separate inputs even when they define the same map. Define

$$\begin{aligned} C_\Phi\big(\mathbf{X}, Y_0, (Y_{\ell,b})_{\ell,b}\big) = {} & \widehat{\mathrm{Var}}(\mathbf{X}) Y_0(Y_0 - 1) \\ & + \sum_{\ell \in [r_{\mathrm{cube}}],\ b \in \{0,1\}} \widehat{E}_{\ell,b}(\mathbf{X})(Y_0 - Y_{\ell,b}) \\ & + \sum_{\tau \in \mathcal{T}} \widehat{\mathrm{Clause}}_\tau(\mathbf{X}) \prod_{t=1}^{d} \big(Y_{\ell_t,b_t} - c_t\big). \end{aligned} \tag{48}$$

Each selector interpolation has degree $< h_i$ in X_i. The Booleanity term has degree two, every routing term is linear, and every clause term has degree at most three. Thus C_Φ has total degree at most three in the query variables and has the syntax required in Definition 5.1.

We next verify its semantics. At $\phi(u)$ for $u \in \mathrm{Var}$, the local equation is $A(\phi(u))(A(\phi(u)) - 1) = 0$. At $u \in E_{\ell,b}$, Lemma A.2 makes the local equation

$$A(\phi(u)) - A(\phi(u + e_\ell)) = 0.$$

At $u \in \mathrm{Clause}_\tau$, it makes the local equation

$$\prod_{t=1}^{d} \big(A(\phi(u + e_{\ell_t})) - c_t\big) = 0.$$

All selectors vanish outside $\phi(\{0,1\}^{r_{\mathrm{cube}}})$. Consequently, a table $A : D \to \mathbb{F}_q$ satisfies the resulting zero-on-subcube instance if and only if $A \circ \phi$ satisfies (46). Lemma A.1 therefore gives the satisfiability equivalence with Φ.

It remains to prove the polynomial, locality, and degree claims. The polynomial extension of $\beta_{\ell,b}$ multiplies one variable by a nonzero field element and fixes the other variables, so it preserves every coordinate degree. Let P satisfy $\deg_{X_i} P \leqslant a_i$, and define B_P as in Lemma 8.1. Put $A = P|_D$. For every $\mathbf{x} \in D$, the polynomial extensions agree with their maps, and hence

$$B_P(\mathbf{x}) = B_A(\mathbf{x}).$$

Moreover, for every $\mathbf{x} \in \mathbb{F}_q^k$, the value $B_P(\mathbf{x})$ depends only on

$$P(\mathbf{x}) \quad \text{and} \quad \{P(\beta_{\ell,b}(\mathbf{x})) : \ell \in [r_{\mathrm{cube}}],\ b \in \{0,1\}\}.$$

Thus it depends on at most $1 + 2r_{\mathrm{cube}} = O(\log n)$ distinct values of P.

In coordinate X_i, a selector contributes degree at most $h_i - 1$, and a query monomial contributes degree at most $3a_i$. Therefore

$$\deg_{X_i} B_P \leqslant 3a_i + h_i - 1.$$

Taking $P = \widetilde{A}$, where $\widetilde{A}$ is the unique interpolation of $A : D \to \mathbb{F}_q$ with coordinate degrees $< h_i$, gives

$$\deg_{X_i} B_A < 4h_i.$$

Finally, each selector interpolation has at most $|D|$ monomials, there are $2r_{\mathrm{cube}}$ routing selectors, and $|\mathcal{T}| \leqslant m$. Each clause product has at most eight monomials. Hence C_Φ has a polynomial-size monomial description. The selectors are computable in polynomial time by Lemma A.1, and so are their transports, interpolations, and the construction of C_Φ. This completes the proof of Lemma 8.1. □

B Axis-line Soundness

We use the following form of Viderman's tensor test [Vid15, Theorem A.5],

Theorem B.1 (Viderman's hyperplane theorem)**.** *Let $C \subseteq \mathbb{F}_q^n$ be a linear code of relative distance Δ, and let $m \geqslant 3$. If H is a uniformly random $(m-1)$-dimensional axis-parallel hyperplane of $[n]^m$, then every $w \in \mathbb{F}_q^{[n]^m}$ satisfies*

$$\mathbb{E}_H \left[\frac{\mathit{dist}(w|_H, C^{\otimes(m-1)})}{n^{m-1}} \right] \geqslant \frac{\Delta^m}{2m^2} \frac{\mathit{dist}(w, C^{\otimes m})}{n^m}.$$

It is well-known that using Viderman's theorem one can show the local testability of tensor codes with respect to the axis-line test. We prove Lemma 7.2 by adapting Theorem B.1 to unequal component codes and lengths, and then composing the resulting tests down to an axis line.

Lemma B.2 (Rectangular three-block test)**.** *Let $D_j \subseteq \mathbb{F}_q^{U_j}$, $j \in [3]$, be nonzero linear codes of relative distance at least δ. Given a word w on $U_1 \times U_2 \times U_3$, choose $j \in [3]$ uniformly, fix a uniformly random coordinate of U_j, and let H be the resulting two-dimensional slice. Then*

$$\mathbb{E}_H \left[\frac{\mathit{dist}(w|_H, \bigotimes_{\ell \neq j} D_\ell)}{\prod_{\ell \neq j} |U_\ell|} \right] \geqslant \frac{\delta^3}{486} \frac{\mathit{dist}(w, D_1 \otimes D_2 \otimes D_3)}{|U_1||U_2||U_3|}.$$

Proof. Let L be a common multiple of $|U_1|, |U_2|, |U_3|$, and repeat every coordinate of D_j exactly $L/|U_j|$ times. This preserves every relative distance in the statement. Let $D = D_1' \oplus D_2' \oplus D_3'$, where the primes denote the repeated codes. Then D has length $3L$ and relative distance at least $\delta/3$.

Since tensor products distribute over direct sums, $D^{\otimes 3}$ is the direct sum of the 27 codes $D_a' \otimes D_b' \otimes D_c'$ for $a, b, c \in [3]$. Embed the repeated word w' in the $D_1' \otimes D_2' \otimes D_3'$ block of $D^{\otimes 3}$ and set it to zero on the other 26 blocks. Its global relative distance from $D^{\otimes 3}$ is 3^{-3} times the corresponding distance for w'. The same factor appears locally: the fixed coordinate lies in its designated summand with probability $1/3$, and the designated block occupies a $1/9$ fraction of the remaining two-dimensional view. Theorem B.1 with $m = 3$ gives robustness $(\delta/3)^3/18 = \delta^3/486$ which proves the lemma. □

The recursion ends with the following elementary two-dimensional bound.

Lemma B.3 (Two-dimensional line consistency)**.** *Let $C_i \subseteq \mathbb{F}_q^{T_i}$ have relative distance Δ_i. For a word w on $T_1 \times T_2$, let b_1 be the fraction of columns outside C_1 and b_2 the fraction of rows outside C_2. Then*

$$\frac{\mathit{dist}(w, C_1 \otimes C_2)}{|T_1||T_2|} \leqslant \frac{b_1 + b_2}{\min\{\Delta_1, \Delta_2\}}.$$

Proof. If $b_2 < \Delta_1$ and $b_1 < \Delta_2$, the good rows contain an information set I_1 for C_1 and the good columns contain an information set I_2 for C_2. Let $c \in C_1 \otimes C_2$ be the unique codeword agreeing with w on $I_1 \times I_2$. Each row indexed by I_1 agrees with the corresponding row of c, since both are in C_2 and agree on I_2. Thus $c = w$ on $I_1 \times T_2$. Every good column then agrees with the corresponding column of c, since both are in C_1 and agree on I_1. Finally, every good row agrees with the corresponding row of c, since both are in C_2 and agree on I_2. Hence c differs from w only at intersections of bad rows and columns, on at most a $b_1 b_2$ fraction of the coordinates. This is at most $(b_1 + b_2)/\min\{\Delta_1, \Delta_2\}$. If $b_2 \geqslant \Delta_1$ or $b_1 \geqslant \Delta_2$, that ratio is at least one. □

We now compose this three-block test. For each current view v, let $\delta(v)$ denote its relative distance from the corresponding tensor code. Suppose v varies in $q \geqslant 3$ coordinate directions. Randomly permute these directions and partition them into three consecutive blocks whose sizes differ by at most one. Regard the current tensor code as $D_1 \otimes D_2 \otimes D_3$, where D_j is the tensor product over the directions in the jth block. Each D_j has relative distance at least $\Delta^{\lceil q/3 \rceil}$. If v' is the random two-block slice selected at this step, Lemma B.2 gives

$$\mathbb{E}[\delta(v') \mid v] \geqslant \frac{\Delta^{3\lceil q/3 \rceil}}{486}\, \delta(v).$$

Fix the block selected by that test and recurse on the directions in the other two blocks. Stop when two directions i, j remain. Choose a random row or column. Lemma B.3 gives

$$\Pr[\text{the selected line is not in its code} \mid v] \geqslant \frac{\Delta}{2}\, \delta(v).$$

Applying conditional expectation inductively at each three-block step and then using this final rejection bound gives

$$\Pr[\text{reject}] \geqslant \min_{\substack{q_0 = r,\ q_s = 2 \\ q_{t+1} \leqslant \lceil 2q_t/3 \rceil}} \frac{\Delta}{2} \prod_{t<s} \frac{\Delta^{3\lceil q_t/3 \rceil}}{486} \cdot \delta(w). \tag{*}$$

The choices determining which directions remain do not depend on the coordinates fixed along the way. Moreover, the procedure is invariant under relabeling the r coordinate directions, so each direction is selected with probability $1/r$. Conditional on selecting direction i, every coordinate outside i is fixed exactly once by an independent uniform choice. Thus the fixed tuple is uniform in $\prod_{j \neq i} T_j$, and the resulting line has exactly the distribution of a uniformly random axis line. It remains only to bound the coefficient in $(*)$. For any sequence in the minimum, q_t is the number of directions before the tth three-block test, and $q_{t+1} \leqslant \lceil 2q_t/3 \rceil$. There are at most $\min\{r - 2, 1 + \lceil \log_{3/2} r \rceil\}$ tests. Moreover, $q_{t+1} - 2 \leqslant \frac{2}{3}(q_t - 2)$ and $3\lceil q_t/3 \rceil \leqslant q_t + 2$. Hence $1 + \sum_t 3\lceil q_t/3 \rceil \leqslant 7r$, where the initial 1 accounts for the final factor $\Delta/2$. Since $\log_{3/2} 486 < 16$, substituting the preceding bounds into $(*)$ gives

$$\Pr_{\ell}\big[w|_\ell \notin C_i\big] \geqslant \frac{\Delta^{7r}}{2 \cdot 486^2 r^{16}}\, \frac{\operatorname{dist}(w, C_1 \otimes \cdots \otimes C_r)}{|T|},$$

where ℓ is a uniformly random axis line and i is its direction. The right-hand side is at least the same expression with coefficient $\tau_r(\Delta) = \Delta^{8r}/(10^6 r^{16})$. Since the rejection probability is $r^{-1} \sum_i b_i(w)$, this proves Lemma 7.2.

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il