

Range Avoidance, Remote Point, and Hard Partial Truth Tables via Satisfying-Pairs Algorithms

Yeyuan Chen
Xi'an Jiaotong University
yychen9961@gmail.com

Yizhi Huang
IIS, Tsinghua University
huangyizhi01@gmail.com

Jiatu Li
IIS, Tsinghua University
lijt19@mails.tsinghua.edu.cn

Hanlin Ren
University of Oxford
h4n1in.r3n@gmail.com

May 18, 2023

Abstract

The *range avoidance problem*, denoted as \mathcal{C} -AVOID, asks to find a non-output of a given \mathcal{C} -circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ with stretch $\ell > n$. This problem has recently received much attention in complexity theory for its connections with circuit lower bounds and other explicit construction problems. Inspired by the Algorithmic Method for circuit lower bounds, Ren, Santhanam, and Wang (FOCS'22) established a framework to design FP^{NP} algorithms for \mathcal{C} -AVOID via *slightly non-trivial* data structures related to \mathcal{C} . However, a major drawback of their approach is the lack of unconditional results even for $\mathcal{C} = \text{AC}^0$.

In this work, we present the first unconditional FP^{NP} algorithm for ACC^0 -AVOID. Indeed, we obtain FP^{NP} algorithms for the following stronger problems:

- (**ACC⁰-REMOTE-POINT**). Given $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ for some $\ell = \text{quasi-poly}(n)$ such that each output bit of C is computed by a quasi-poly(n)-size $\text{AC}^0[m]$ circuit, we can find some $y \in \{0, 1\}^\ell$ in FP^{NP} such that for every $x \in \{0, 1\}^n$, the relative Hamming distance between y and $C(x)$ is at least $1/2 - 1/\text{poly}(n)$. This problem is the “average-case” analogue of ACC^0 -AVOID.
- (**ACC⁰-PARTIAL-AVGHARD**). Given $x_1, \dots, x_\ell \in \{0, 1\}^n$ for some $\ell = \text{quasi-poly}(n)$, we can compute ℓ bits $y_1, \dots, y_\ell \in \{0, 1\}$ in FP^{NP} such that for every $2^{\log^c(n)}$ -size ACC^0 circuit C , $\Pr_i[C(x_i) \neq y_i] \geq 1/2 - 1/\text{poly}(n)$, where $c = O(1)$. This problem generalises the strong average-case circuit lower bounds against ACC^0 in a different way.

Our algorithms can be seen as natural generalisations of the best known almost-everywhere average-case lower bounds against ACC^0 circuits by Chen, Lyu, and Williams (FOCS'20). Note that both problems above have been studied prior to our work, and no FP^{NP} algorithm was known even for weak circuit classes such as $\text{GF}(2)$ -linear circuits and DNF formulas.

Our results follow from a strengthened algorithmic method: slightly non-trivial algorithms for the *Satisfying-Pairs* problem for \mathcal{C} implies FP^{NP} algorithms for \mathcal{C} -AVOID (as well as \mathcal{C} -REMOTE-POINT and \mathcal{C} -PARTIAL-AVGHARD). Here, given \mathcal{C} -circuits $\{C_i\}$ and inputs $\{x_j\}$, the \mathcal{C} -Satisfying-Pairs problem asks to (approximately) count the number of pairs (i, j) such that $C_i(x_j) = 1$.

A technical contribution of this work is a construction of a *short, smooth, and rectangular PCP of Proximity* that combines two previous PCP constructions, which may be of independent interest. It serves as a key tool that allows us to generalise the framework for AVOID to the average-case scenarios.

Contents

1	Introduction	1
1.1	From Circuit Lower Bounds to Range Avoidance	1
1.2	Our Results	4
1.3	Technical Overview	9
1.4	Further Related Work	11
1.5	Organisation	12
2	Preliminaries	12
3	Range Avoidance and Remote Point	23
3.1	Range Avoidance from SATISFYING-PAIRS	24
3.2	Remote Point from SATISFYING-PAIRS	31
3.3	Variants of Our Frameworks	42
4	Hard Partial Truth Tables	43
4.1	Hard Partial Truth Tables from SATISFYING-PAIRS	43
4.2	Average-Case Hard Partial Truth Tables	46
5	Unconditional Algorithms for Range Avoidance, Remote Point, and Hard Partial Truth Tables	52
5.1	Algorithms for Satisfying Pairs	52
5.2	An FP^{NP} Algorithm for XOR-REMOTE-POINT	54
5.3	Remote Point for ACC^0	55
5.4	Hard Partial Truth Tables for ACC^0	56
6	Construction of Smooth and Rectangular PCPP	57
6.1	Rectangular Neighbour Listing and Smoothness	58
6.2	A Rectangular PCPP with RNL Property	63
6.3	RNL-Preserving Composition Theorem	69
6.4	Soundness Amplification Preserving Smoothness and Rectangularity	73
6.5	Final Construction	76
7	Construction of Rectangular PCPPs with Low Query Complexity	80
7.1	A 3-Query PCPP for $\text{CIRCUIT-EVAL}^\perp$	80
7.2	A 3-Query Rectangular PCPP	81
7.3	A 2-Query Rectangular PCPP with Imperfect Completeness	83
	References	86
A	Missing Proofs in Section 3 and 4	90
A.1	Satisfying Pairs for $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ Circuits	91
A.2	Verifying PCPP with Satisfying Pairs	94
A.3	Proof of Claim 3.9	98
A.4	Proof of Lemma 3.10	100
A.5	An XOR Lemma in [CLW20]	100

1 Introduction

“You might wonder why should coming up with explicit construction be so difficult. After all, a proof of existence via the probabilistic method shows not only that an object with the desired property exists but in fact the vast majority of objects have the property.”

Sanjeev Arora and Boaz Barak [AB09]

1.1 From Circuit Lower Bounds to Range Avoidance

Proving unconditional lower bounds for non-uniform circuits is one of the grand challenges in theoretical computer science, with the holy grail of proving $\text{NP} \not\subseteq \text{P}_{/\text{poly}}$. Unfortunately, progress in unconditional circuit lower bounds has been slow, and the best lower bound for any explicit function against general circuits is only slightly above $3n$ [FGHK16, LY22]. A long-standing, and somewhat embarrassing, open problem is to find any language in EXP^{NP} (exponential time with an NP oracle) that cannot be computed by polynomial-size circuits. It seems unlikely that $\text{EXP}^{\text{NP}} \subseteq \text{P}_{/\text{poly}}$, but we appear to be very far from ruling out this possibility.

To add more embarrassment, it has been known since 1949 [Sha49] that *most* Boolean functions over n inputs require circuits of size $\Omega(2^n/n)$. 70 years later, we still struggle to spell out even a single such function from a plethora of them! It turns out that circuit lower bounds are not alone, and the difficulty of “finding hay in a haystack” ([AB09, Chapter 21]) is a general phenomenon in theoretical computer science. For example, most graphs are Ramsey graphs [Erd59] and most matrices are rigid matrices [Val77], but it remains major open problems to explicitly construct Ramsey graphs and rigid matrices with good parameters [CZ19, Li23, AC19, Ram20, BHPT20].

Our lack of progress in such explicit construction problems suggests the necessity of a systematic study of their difficulty. As a first step towards building a complexity theory for explicit construction problems, Korten [Kor21] studied the complexity class APEPP defined in [KKMP21], and argued that this is the complexity class that corresponds to explicit construction problems. APEPP is the class of total search problems that are polynomial-time reducible to the following problem:

Problem 1.1 (Range Avoidance Problem, denoted as AVOID). Given the description of a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, where $\ell > n$, output any string $y \in \{0, 1\}^\ell$ that is not in the range of C . That is, for every $x \in \{0, 1\}^n$, $C(x) \neq y$.

The existence of such y follows from the *dual weak pigeonhole principle*: if we throw 2^n pigeons into 2^ℓ holes, where $\ell \geq n + 1$, then there is an empty hole. Thus AVOID is a *total* search problem. Moreover, a random string $y \in \{0, 1\}^\ell$ is a valid solution w.p. $1 - 2^{n-\ell} \geq 1/2$, thus there is a trivial randomised algorithm for AVOID. Hence, the focus is to design *deterministic* algorithms for AVOID.

The following is a good example of how AVOID captures the complexity of explicit constructions:

Example 1.2 ([Kor21, Section 3.1]). Proving circuit lower bounds can be rephrased as solving the following total search problem, denoted as HARD: On input 1^N where $N = 2^n$, output the truth table of a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ that cannot be computed by circuits of size s (say $s = 2^{n/2}$).

Let $\text{TT} : \{0, 1\}^{O(s \log s)} \rightarrow \{0, 1\}^{2^n}$ be the circuit that takes as input the description of a size- s circuit and outputs the truth table of this circuit. (The circuit TT is sometimes called the *truth table generator*, hence the name TT.) If we could solve AVOID on the particular instance TT, we would find a truth table $tt \in \{0, 1\}^{2^n}$ without size- s circuits, thereby proving a circuit lower bound. It follows that HARD polynomial-time reduces to AVOID, and thus $\text{HARD} \in \text{APEPP}$.

More precisely, solving AVOID for TT in polynomial time is equivalent to proving a circuit lower bound for E, and solving AVOID for TT in FP^{NP} is equivalent to proving a circuit lower bound for E^{NP} .

1.1.1 Range Avoidance for Restricted Circuit Classes

In a recent paper, Ren, Santhanam, and Wang [RSW22] suggested studying the range avoidance problem for restricted circuit classes. Let \mathcal{C} be a circuit class and $\ell := \ell(n) > n$ be a stretch function. Consider the following problem:

Problem 1.3 (\mathcal{C} -AVOID). Given the description of a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell(n)}$, where each output bit of C is a \mathcal{C} circuit, output any string $y \in \{0, 1\}^{\ell(n)}$ that is not in the range of C . That is, for every $x \in \{0, 1\}^n$, $C(x) \neq y$.

There are lots of reasons for studying the problem \mathcal{C} -AVOID, but we only mention one of them here. Many interesting explicit construction problems reduce to \mathcal{C} -AVOID for restricted circuit classes \mathcal{C} and (sometimes) large stretch functions ℓ . For example:

- For any “nice” circuit class \mathcal{C} , the problem of proving circuit lower bounds against \mathcal{C} can be reduced to \mathcal{C} -AVOID via the truth table generator in Example 1.2, where the input of the truth table generator is replaced by a \mathcal{C} circuit (instead of a general circuit).
- Guruswami, Lyu, and Wang [GLW22] showed that the problem of finding rigid matrices and optimal binary linear codes can be reduced to NC^1 -AVOID. By a further result in [RSW22], these problems also reduce to NC^0 -AVOID. A recent work [GGNS23] showed that the problem of finding rigid matrices can even be reduced to NC_3^0 -AVOID.

In general, for any explicit construction problem Π , we can identify a circuit class \mathcal{C} that is as “simple” as possible, as well as a stretch function $\ell(n)$ that is as large as possible, such that Π reduces to \mathcal{C} -AVOID with stretch $\ell(n)$. The hope is that by making progress on the range avoidance problem for restricted circuits and by optimising the reduction (i.e., optimising \mathcal{C} and $\ell(n)$), we could solve many explicit construction problems systematically.

An “Algorithmic Method” for range avoidance. Inspired by the Algorithmic Method for proving circuit lower bounds (e.g. [Wil13, Wil14a, MW20, CW19b, Che19, CLW20]), [RSW22] proposed a framework that uses data structures to solve \mathcal{C} -AVOID in FP^{NP} . Consider the following data structure problem.

Problem 1.4 (Hamming Weight Estimation). Let \mathcal{C} be a circuit class and $\ell := \ell(n)$ be a stretch function. The data structure problem has two phases:

(Preprocessing) Given the description of a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell}$, where each output bit of C is a \mathcal{C} circuit, we need to preprocess the circuit in P^{NP} (i.e., in polynomial time with an NP oracle) and output a data structure $\text{DS} \in \{0, 1\}^{\text{poly}(\ell)}$.

(Query) Given an input x and oracle access (i.e., random access) to DS , we need to estimate the Hamming weight of $C(x)$ in “non-trivial” time, i.e., deterministic $\ell / \log^{\omega(1)} \ell$ time.

It was shown in [RSW22] that for “typical” circuit classes \mathcal{C} , a non-trivial data structure for the Hamming Weight Estimation problem for \mathcal{C} implies an FP^{NP} algorithm for \mathcal{C} -AVOID.

One drawback of [RSW22] is that their framework does not imply new unconditional algorithms for range avoidance.¹ For comparison, the original Algorithmic Method has made significant progress on proving *unconditional* circuit lower bounds that we do not know how to prove otherwise. One motivation for the current paper is to address this drawback by designing new and unconditional

¹Actually, [RSW22] provided an unconditional range avoidance algorithm for de Morgan formulas with non-trivial parameters. Subsequently, [GLW22] improved this result by using simpler techniques and achieving better parameters; in particular, the algorithm in [GLW22] does not require the Algorithmic Method.

range avoidance algorithms via the Algorithmic Method. In particular, can we solve $\text{ACC}^0\text{-AVOID}$ with parameters that match the circuit lower bounds in [CLW20]?

1.1.2 The Remote Point Problem

The Algorithmic Method is extremely good at proving average-case circuit lower bounds [CR22, CLW20, CL21]. Therefore, it is natural to wonder if there is an “average-case analogue” of [RSW22].

For two strings $x, y \in \{0, 1\}^n$, their *relative Hamming distance* is defined as the fraction of indices where x and y differ, formally $\delta(x, y) := \frac{1}{n} |\{i \in [n] : x_i \neq y_i\}|$. The “average-case analogue” of the range avoidance problem is the following problem:

Problem 1.5 (Remote Point Problem, denoted as \mathcal{C} -REMOTE-POINT). Given the description of a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ and a parameter $\delta > 0$, where each output bit of C is a \mathcal{C} circuit, output any string $y \in \{0, 1\}^\ell$ that is δ -far from the range of C . That is, for every $x \in \{0, 1\}^n$, $\delta(C(x), y) \geq \delta$.

By Chernoff bound, if $\delta < 1/2 - c\sqrt{n/\ell}$ for some absolute constant $c > 0$, then a random length- ℓ string is a valid solution for REMOTE-POINT w.h.p. Therefore, the challenge is to find deterministic algorithms for REMOTE-POINT.

It is not hard to see that \mathcal{C} -REMOTE-POINT for the truth table generator TT corresponds to average-case circuit lower bounds. In particular, the regime where δ is a small constant corresponds to proving “weak” average-case lower bounds (e.g. [COS18, Che19]), and the regime where δ is close to $1/2$ (say, $\delta = 1/2 - 1/n$) corresponds to proving “strong” average-case lower bounds (e.g. [CR22, CLW20]).²

The remote point problem was discussed in [KKMP21]. Indeed, an important special case of the problem has been studied by Alon, Panigrahy, and Yekhanin [APY09], namely the case that C is a linear transformation over $\text{GF}(2)$. In other words, we are given a linear code $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ and we want to find a string far from every codeword. They introduced this problem as an intermediate step towards constructing rigid matrices. In this paper, we call this special case XOR-REMOTE-POINT.

It is already quite hard to solve this special case deterministically. Alon, Panigrahy, and Yekhanin [APY09] designed a polynomial-time algorithm for XOR-REMOTE-POINT when $\ell > 2n$ and $\delta = O(\log n/n)$. For slightly larger δ , say $\delta = 0.1$, no deterministic algorithm is known even with an NP oracle. Arvind and Srinivasan [AS10] showed that for certain parameters, a polynomial-time algorithm for XOR-REMOTE-POINT implies a polynomial-time algorithm for $\text{AC}^0\text{-PARTIAL-HARD}$ (defined later in Section 1.1.3).

1.1.3 Hard Partial Truth Tables

Besides AVOID and REMOTE-POINT, we also consider the following problem that generalises the task of proving circuit lower bounds (in a different way from AVOID and REMOTE-POINT):

Problem 1.6 (Hard Partial Truth Tables against \mathcal{C} , denoted as \mathcal{C} -PARTIAL-HARD). Given a list of input strings $z_1, z_2, \dots, z_\ell \in \{0, 1\}^n$ and a parameter s , find a list of output bits $b_1, b_2, \dots, b_\ell \in \{0, 1\}$ such that the partial function defined by $\{(z_i, b_i)\}_{i \in [\ell]}$ cannot be computed by \mathcal{C} circuits of size s . In other words, for every size- s \mathcal{C} circuit C , there exists an index $i \in [\ell]$ such that $C(z_i) \neq b_i$.

²Typically, a strong average-case lower bound states that certain problems cannot be $(1/2 + 1/s)$ -approximated by size- s circuits. Suppose $\text{TT} : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ is the truth table generator, then n is roughly the size of the circuit (i.e., $n \approx s$). In this regard, strong average-case circuit lower bounds correspond to REMOTE-POINT where $\delta = 1/2 - 1/n$.

It is easy to see that \mathcal{C} -PARTIAL-HARD generalises the problem of proving circuit lower bounds against \mathcal{C} . Indeed, if we take $\ell := 2^n$ and z_1, z_2, \dots, z_ℓ be an enumeration of length- n strings, then \mathcal{C} -PARTIAL-HARD becomes exactly the problem of proving circuit lower bounds against \mathcal{C} . It is also easy to see that when $\ell > O(s \log s)$, this problem is in APEPP: given the input $(z_1, z_2, \dots, z_\ell)$, we can construct a circuit $\mathbb{T}\mathbb{T}' : \{0, 1\}^{O(s \log s)} \rightarrow \{0, 1\}^\ell$ which takes the description of a \mathcal{C} circuit C as input, and outputs the concatenation of $C(z_1), C(z_2), \dots, C(z_\ell)$. Finding a non-output of $\mathbb{T}\mathbb{T}'$ is equivalent to finding a solution of \mathcal{C} -PARTIAL-HARD, thus this problem reduces to AVOID.

This problem was introduced by Arvind and Srinivasan [AS10] under the name ‘‘circuit lower bounds with help functions.’’ Let $h_1, h_2, \dots, h_n : \{0, 1\}^m \rightarrow \{0, 1\}$ denote a sequence of *help functions*, \mathcal{C} be a circuit class, and $s \in \mathbb{N}$ be a size parameter. The goal is to construct the truth table of a function $f : \{0, 1\}^m \rightarrow \{0, 1\}$ that is hard to compute for size- s \mathcal{C} circuits, even when the circuit has access to these help functions. Formally, for any size- s circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$, there exists an input $x \in \{0, 1\}^m$ such that

$$C(h_1(x), h_2(x), \dots, h_n(x)) \neq f(x).$$

This problem is equivalent to PARTIAL-HARD with $\ell = 2^m$ inputs of length n , namely for every $x \in \{0, 1\}^m$, there is an input $h_1(x) \circ h_2(x) \circ \dots \circ h_n(x) \in \{0, 1\}^n$ in the PARTIAL-HARD instance.

This problem appears to be very hard. Neither [AS10] nor we are aware of an efficient deterministic solution for $\mathcal{C} = \text{AC}^0$ with (say) $\ell, s \in \text{quasi-poly}(n)$. That is, although exponential-size lower bounds against AC^0 are known [Ajt83, FSS84, Yao85, Hås89], we do not have any idea about how to prove such a lower bound for partial functions. Even when \mathcal{C} is the class of *polynomial-size DNF*, to the best of our knowledge, there is no known deterministic algorithm for \mathcal{C} -PARTIAL-HARD.

Besides being a natural problem itself, \mathcal{C} -PARTIAL-HARD also arises when we study the closure of non-uniform complexity classes (under reductions). Recall that AC^0 denotes the class of languages computable by a *non-uniform* family of polynomial-size constant-depth circuits; in particular, AC^0 contains undecidable languages such as unary versions of the halting problem. A language L Turing-reduces to some language in AC^0 if and only if $L \in \text{P}_{/\text{poly}}$ [Pip79], thus proving $\text{EXP} \not\leq_T^p \text{AC}^0$ is likely beyond current techniques. But what about *mapping reducibility*? Can we show that $\text{EXP} \not\leq_m^p \text{AC}^0$? It turns out that a deterministic algorithm for AC^0 -PARTIAL-HARD implies that $\text{EXP} \not\leq_m^p \text{AC}^0$ [AS10, Theorem 5]. Of course, there is nothing special with AC^0 , and it can be replaced by other non-uniform classes. Therefore, \mathcal{C} -PARTIAL-HARD sheds light on ruling out many-one reducibility of EXP (and other complexity classes) to non-uniform classes.

We can also define the average-case version of \mathcal{C} -PARTIAL-HARD, which is equivalent to proving average-case lower bounds with help functions.

Problem 1.7 (Average-Case Hard Partial Truth Tables against \mathcal{C} , denoted as \mathcal{C} -PARTIAL-AVGHARD). Given a list of input strings $z_1, z_2, \dots, z_\ell \in \{0, 1\}^n$ and parameters s, δ , find a list of output bits $b_1, b_2, \dots, b_\ell \in \{0, 1\}$ such that the partial function defined by $\{(z_i, b_i)\}_{i \in [\ell]}$ is δ -far from being computable by \mathcal{C} circuits of size s . In other words, for every size- s \mathcal{C} circuit C , there are at least $\delta\ell$ indices $i \in [\ell]$ such that $C(z_i) \neq b_i$.

1.2 Our Results

In this sub-section, we describe our results in detail.

1.2.1 Explicit Constructions from SATISFYING-PAIRS Algorithms

We start with the following observation: In the framework of solving AVOID via the Algorithmic Method [RSW22], the data structure for Problem 1.4 does not need to be *online*. Instead, it suffices

to design a data structure that preprocesses a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, receives a *batch* of inputs x_1, x_2, \dots, x_M , and estimates the Hamming weight of each $C(x_i)$ in *non-trivial total time*, i.e., $\ell M / \log^{\omega(1)}(\ell M)$ time. Moreover, we observe that it is not even necessary to estimate the individual Hamming weights $C(x_i)$; it suffices to estimate the *average* Hamming weight of $C(x_i)$ for $i \in [M]$. Indeed, we arrive at the following problem called *Satisfying Pairs*.³

Problem 1.8 (\mathcal{C} -SATISFYING-PAIRS). Let N, M, s, n be parameters. Given (single-output) \mathcal{C} circuits $C_1, C_2, \dots, C_N : \{0, 1\}^n \rightarrow \{0, 1\}$ of size s and input strings $x_1, x_2, \dots, x_M \in \{0, 1\}^n$, compute or estimate

$$\Pr_{i \leftarrow [M], j \leftarrow [N]} [C_j(x_i) = 1]. \quad (1)$$

We define the decisional and counting versions of the satisfying pairs problem as follows.

- $\text{Gap}_\delta\text{-}\mathcal{C}\text{-SATISFYING-PAIRS}$ is the problem of distinguishing between (1) = 1 and (1) < 1 - δ ;
- $\text{Approx}_\varepsilon\text{-}\mathcal{C}\text{-SATISFYING-PAIRS}$ is the problem of estimating (1) within additive error ε ;
- $\mathcal{C}\text{-SATISFYING-PAIRS}$ is the problem of deciding whether (1) > 0;
- $\#\mathcal{C}\text{-SATISFYING-PAIRS}$ is the problem of exactly computing (1).

We consider the regime where the input length n and the circuit size s are much smaller than N and M . In such case, a deterministic algorithm for \mathcal{C} -SATISFYING-PAIRS is said to be *non-trivial* if it runs in time $NM / \log^{\omega(1)}(NM)$.⁴

Remark 1.9. The circuit-analysis problems that arise in the Algorithmic Method⁵ are special cases of Satisfying Pairs problems. For instance, we can solve $\#\text{SAT}$ of the circuit C by solving $\#\text{SATISFYING-PAIRS}$ with $N = 2^{n/2}$ and $M = 2^{n/2}$, where the inputs (x_1, x_2, \dots, x_M) consists of all strings of length $n/2$, and the circuits are $\{C_y : y \in \{0, 1\}^{n/2}\}$, where $C_y(x) := C(x \circ y)$. Similarly, $\mathcal{C}\text{-SATISFYING-PAIRS}$ corresponds to $\mathcal{C}\text{-SAT}$, $\text{Gap-}\mathcal{C}\text{-SATISFYING-PAIRS}$ corresponds to $\mathcal{C}\text{-GapUNSAT}$, and $\text{Approx-}\mathcal{C}\text{-SATISFYING-PAIRS}$ corresponds to $\mathcal{C}\text{-CAPP}$.

Range Avoidance from SATISFYING-PAIRS. By plugging the observation above in [RSW22], we show that non-trivial algorithms for SATISFYING-PAIRS imply FP^{NP} algorithms for AVOID.

Theorem 1.10 (Theorem 3.2, Informal). *Let \mathcal{C} be a typical circuit class and $\mathcal{C}' := \text{OR}_2 \circ \mathcal{C}$.⁶ Suppose that there is a non-trivial algorithm for $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ for every constant $\varepsilon > 0$, then $\mathcal{C}\text{-AVOID}$ with certain parameters can be solved in FP^{NP} .*

This informal version of Theorem 3.2 hides the trade-off between the parameters of $\mathcal{C}\text{-AVOID}$ and $\mathcal{C}'\text{-SATISFYING-PAIRS}$. In general, to solve $\mathcal{C}\text{-AVOID}$ with smaller stretch ℓ (with respect to the input length n), we need to have non-trivial algorithms for $\mathcal{C}'\text{-SATISFYING-PAIRS}$ where the

³We remark that our definition of $\mathcal{C}\text{-SATISFYING-PAIRS}$ is different from the fine-grained complexity literature (e.g., [AHWW16, CW19a]). The input of the $\mathcal{C}\text{-SATISFYING-PAIRS}$ problem defined in [AHWW16, CW19a] consists of a circuit $C(-, -)$ and two sets of input strings $\{a_i\}$ and $\{b_j\}$, and one wants to compute or approximate the number of pairs (i, j) such that $C(a_i, b_j) = 1$; in our $\mathcal{C}\text{-SATISFYING-PAIRS}$ problem, we receive as input a list of circuits $\{C_i\}$ and a list of inputs $\{x_j\}$, and we want to compute or approximate the number of pairs (i, j) such that $C_i(x_j) = 1$. The new definition fits our purpose better. We also remark that for circuit classes that can “evaluate themselves” (such as AC^0 , ACC^0 , and TC^0), these two definitions are computationally equivalent.

⁴Analogous to the preprocessing phase in Problem 1.4, one could also add a P^{NP} -preprocessing phase that sees the circuits but not the inputs. Algorithms with such preprocessing phase would still imply our results, but the SATISFYING-PAIRS algorithms in this paper do not need this preprocessing phase.

⁵The definitions of circuit-analysis problems such as SAT or CAPP can be found in Lijie Chen’s PhD thesis [Che22].

⁶Here, $\text{OR}_d \circ \mathcal{C}$ refers to the composition of a single fan-in- d OR gate being the output gate of the circuit and (at most) d \mathcal{C} circuits feeding the top OR gate.

number of inputs N and the number of circuits M are smaller with respect to the circuit size s and the input length n . We highlight two typical choices of parameters of Theorem 3.2 as follows.

Corollary 1.11. *There is a constant $\varepsilon > 0$ such that the following holds. Let \mathcal{C} be a typical circuit class, $\mathcal{C}' := \text{OR}_2 \circ \mathcal{C}$, and $s = s(n)$ be a non-decreasing size parameter.*

- *Suppose that there is a non-trivial algorithm for $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ for $N = n^{1+\Omega(1)}$ \mathcal{C}' -circuits of size $2s(n)$ and $M = n^{1+\Omega(1)}$ inputs of length n . Then there is an FP^{NP} algorithm for \mathcal{C} -AVOID with stretch ℓ and circuit size s ,⁷ for some $\ell = n^{1+\Omega(1)}$.*
- *Suppose that there is a non-trivial algorithm for $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ for $N = \text{quasi-poly}(n)$ \mathcal{C}' -circuits of size $2s(n)$ and $M = \text{quasi-poly}(n)$ inputs of length n . Then there is an FP^{NP} algorithm for \mathcal{C} -AVOID with stretch ℓ and circuit size s , for some $\ell = \text{quasi-poly}(n)$.*

Remote Point from SATISFYING-PAIRS. With the help of a smooth and rectangular PCPP (see Section 6) and a linear-sum list-decodable code from [CLW20] (also see Appendix A.5), we show that non-trivial algorithms for SATISFYING-PAIRS imply FP^{NP} algorithms for REMOTE-POINT.

Theorem 1.12 (Theorem 3.5, Informal). *Let \mathcal{C} be a typical circuit class and $\mathcal{C}' := \text{AND}_{O(1)} \circ \mathcal{C}$. Suppose that there is a non-trivial algorithm for $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ for every constant $\varepsilon > 0$, then \mathcal{C} -REMOTE-POINT with certain parameters can be solved in FP^{NP} .*

In particular: suppose for every constant $\varepsilon > 0$, there is a non-trivial algorithm for $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ for $N = \text{quasi-poly}(n)$ \mathcal{C}' -circuits of size $O(s)$ and $M = \text{quasi-poly}(n)$ inputs of length n ; then for some stretch function $\ell = \text{quasi-poly}(n)$, there is an FP^{NP} algorithm for \mathcal{C} -REMOTE-POINT that takes as input a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ where each output bit of C is a \mathcal{C} -circuit of size s , and outputs a y that is 0.49-far from $\text{Range}(C)$.

Our framework provides REMOTE-POINT algorithms for the regime corresponding to “strong average-case lower bounds”, i.e., the distance between the output y and $\text{Range}(C)$ is close to $1/2$. In fact, the distance can be as large as $1/2 - 1/\text{poly}(n)$ given an $\text{Approx}\text{-}\mathcal{C}\text{-SATISFYING-PAIRS}$ algorithm with a small enough error. (see Theorem 3.5 for details).

Note that the stretch for \mathcal{C} -REMOTE-POINT that we can solve in FP^{NP} depends on both the parameters of the satisfying pairs algorithms and the rate of the linear-sum list-decodable code. Since the code from [CLW20] has a quasi-polynomial rate, our framework cannot solve REMOTE-POINT with stretch smaller than quasi-polynomial. It is an interesting open problem to improve the stretch of REMOTE-POINT that can be solved by our framework, possibly by designing new linear-sum decodable codes with a better rate; see, e.g., [CL21].

Hard Partial Truth Table from SATISFYING-PAIRS. Similar to the frameworks for AVOID and REMOTE-POINT, we can solve PARTIAL-HARD and PARTIAL-AVGHARD via non-trivial algorithms for SATISFYING-PAIRS.

Theorem 1.13 (Theorems 4.2 and 4.3, Informal). *Let \mathcal{C} be a typical circuit class.*

- *Suppose that there is a non-trivial algorithm for $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ for every $\varepsilon > 0$ and $\mathcal{C}' := \text{OR}_2 \circ \mathcal{C}$, then \mathcal{C} -PARTIAL-HARD with certain parameters can be solved in FP^{NP} .*
- *Suppose that there is a non-trivial algorithm for $\text{Approx}_\varepsilon\text{-}\mathcal{C}''\text{-SATISFYING-PAIRS}$ for every $\varepsilon > 0$ and $\mathcal{C}'' := \text{AND}_{O(1)} \circ \mathcal{C}$, then \mathcal{C} -PARTIAL-AVGHARD with certain parameters can be solved in FP^{NP} .*

⁷Note that the circuit size parameter of \mathcal{C} -AVOID refers to the maximum circuit size of each output bit of $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, instead of the total circuit size of C .

These results are proved using essentially the same approach as the framework for AVOID and REMOTE-POINT; subsequently, the trade-off between parameters for SATISFYING-PAIRS and PARTIAL-HARD (resp. PARTIAL-AVGHARD) is similar to that for AVOID (resp. REMOTE-POINT). We omit the details and refer the readers to Theorems 4.2 and 4.3.

Remark 1.14. It is not surprising to have a unified framework for AVOID and PARTIAL-HARD (and their average-case analogues REMOTE-POINT and PARTIAL-AVGHARD), since AVOID and PARTIAL-HARD can be considered as the *dual* problem of each other. Let $\text{Eval} : \{0, 1\}^{O(s \log s)} \times \{0, 1\}^n \rightarrow \{0, 1\}$ be the circuit-evaluation function that takes a circuit C of size s and an input of length n , and outputs $C(x)$. We can interpret AVOID and PARTIAL-HARD as follows:

- **(AVOID).** Given size- s circuits C_1, C_2, \dots, C_ℓ , find $y_1, y_2, \dots, y_\ell \in \{0, 1\}$ such that for every $x \in \{0, 1\}^n$, there is an $i \in [\ell]$ such that $\text{Eval}(C_i, x) \neq y_i$.
- **(PARTIAL-HARD).** Given inputs $x_1, x_2, \dots, x_\ell \in \{0, 1\}^n$, find $y_1, y_2, \dots, y_\ell \in \{0, 1\}$ such that for every size- s circuit C , there is an $i \in [\ell]$ such that $\text{Eval}(C, x_i) \neq y_i$.

Clearly, AVOID and PARTIAL-HARD are essentially the same problem on the table $\text{Eval}(\cdot, \cdot)$ with the rows and columns being exchanged.

1.2.2 Unconditional Results for Explicit Constructions

The seemingly marginal improvement of using SATISFYING-PAIRS instead of its online version Hamming Weight Estimation (see Problem 1.4) plays an important role in the design of unconditional FP^{NP} algorithms for ACC^0 -REMOTE-POINT and ACC^0 -PARTIAL-HARD, because we can indeed design non-trivial algorithms for ACC^0 -SATISFYING-PAIRS.

XOR-REMOTE-POINT from XOR-SATISFYING-PAIRS. We start from a simpler case where the circuit class $\mathcal{C} = \text{XOR}$, i.e., the circuit is an XOR of some of its input bits. Since an XOR circuit C can be represented by a vector $\vec{v} \in \{0, 1\}^n$ such that $C(x) = \langle v, x \rangle \bmod 2$, $\#\text{XOR-SATISFYING-PAIRS}$ is nothing but the counting version of the *Orthogonal Vector* problem over \mathbb{F}_2 , which admits a non-trivial algorithm [CW21, AC19]. By combining this with Theorem 1.10, we obtain an unconditional FP^{NP} algorithm for XOR-REMOTE-POINT.⁸

Theorem 1.15 ($\text{XOR-REMOTE-POINT} \in \text{FP}^{\text{NP}}$). *There is a constant $c_u \geq 1$ such that the following holds. Let $\varepsilon := \varepsilon(n) \geq 2n^{-c_u}$ be the error parameter and $\ell := \ell(n) \geq 2^{\log^{c_u+5} n}$ be the stretch function, then there is an FP^{NP} algorithm that takes as input a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, where each output bit of C is computed by an XOR gate, and outputs a string y that is $(1/2 - \varepsilon)$ -far from $\text{Range}(C)$.*

A non-trivial algorithm for ACC^0 -SATISFYING-PAIRS. By slightly adapting the technique introduced by Williams [Wil18c] to design non-trivial $\#\text{SAT}$ algorithms for ACC^0 circuits with an earlier quasi-polynomial size simulation of $\text{SYM} \circ \text{ACC}^0$ circuits by $\text{SYM} \circ \text{AND}$ circuits [BT94, AG91], we can obtain a non-trivial algorithm for $\#\text{ACC}^0$ -SATISFYING-PAIRS, formally stated as follows.

Theorem 1.16. *For every constants m, ℓ, c , there is a constant $\varepsilon \in (0, 1)$ such that the following holds. Let $n := 2^{\log^\varepsilon N}$ and $s := 2^{\log^c n}$. There is a deterministic algorithm running in $\tilde{O}((N/n)^2)$ time that given N strings $x_1, x_2, \dots, x_N \in \{0, 1\}^n$ and N $\text{AC}_\ell^0[m]$ circuits $C_1, C_2, \dots, C_N : \{0, 1\}^n \rightarrow \{0, 1\}$ of size s , outputs the number of pairs $(i, j) \in [N] \times [N]$ such that $C_i(x_j) = 1$.*

⁸The reduction from REMOTE-POINT to SATISFYING-PAIRS has a small overhead on the circuit class (i.e. the upper $\text{AND}_{O(1)}$ in Theorem 1.12). By a standard trick using Fourier analysis (see Theorem 2.17; also see [CW19b]), we can change the upper circuit class to be $\text{XOR}_{O(1)}$ so that we only need to design SATISFYING-PAIRS algorithms for $\text{XOR}_d \circ \text{XOR} = \text{XOR}$.

Explicit constructions for ACC^0 . The FP^{NP} algorithm for $\text{ACC}^0\text{-REMOTE-POINT}$ and $\text{ACC}^0\text{-PARTIAL-AVGHARD}$ follows from this algorithm together with Theorem 1.12 and Theorem 1.13.

Theorem 1.17 ($\text{ACC}^0\text{-REMOTE-POINT} \in \text{FP}^{\text{NP}}$). *There is a constant $c_u \geq 1$ such that for every constant $d, m \geq 1$, there is a constant $c_{\text{str}} := c_{\text{str}}(d, m) \geq 1$, such that the following holds.*

Let $n < s(n) \leq 2^{n^{o(1)}}$ be a size parameter, $\varepsilon := \varepsilon(n) \geq 2n^{-c_u}$ be an error parameter and $\ell := \ell(n) \geq 2^{\log^{c_{\text{str}}} s}$ be a stretch function, then there is an FP^{NP} algorithm that takes as input a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, where each output bit of C is computed by an $\text{AC}_d^0[m]$ circuit of size s , and outputs a string y that is $(1/2 - \varepsilon)$ -far from $\text{Range}(C)$.

Theorem 1.18 ($\text{ACC}^0\text{-PARTIAL-AVGHARD} \in \text{FP}^{\text{NP}}$). *There is a constant $c_u \geq 1$ such that for every constants $d, m \geq 1$, there is a constant $c_{\text{str}} := c_{\text{str}}(d, m) \geq 1$, such that the following holds.*

Let $n < s(n) \leq 2^{n^{o(1)}}$ be a size parameter, $\varepsilon := \varepsilon(n) \geq 2n^{-c_u}$ be an error parameter and $\ell := \ell(n) \geq 2^{\log^{c_{\text{str}}} s}$ be a stretch function, then there is an FP^{NP} algorithm that given inputs $x_1, \dots, x_\ell \in \{0, 1\}^n$, it outputs a string $y \in \{0, 1\}^\ell$ such that for any $s(n)$ -size $\text{AC}_d^0[m]$ circuit C , y is $(1/2 - \varepsilon)$ -far from $C(x_1) \circ \dots \circ C(x_\ell)$.

It is worth noting that the $\text{ACC}^0\text{-REMOTE-POINT}$ algorithm here recovers the best known almost-everywhere average-case circuit lower bounds against ACC^0 [CLW20]. This is done by considering the special case where the input circuit is the truth table generator $\text{TT} : \{0, 1\}^{O(s \log s)} \rightarrow \{0, 1\}^{2^n}$ that prints the truth table of a given ACC^0 circuit (see Section 5.3).

Corollary 1.19. *For every constants $d, m \geq 1$, there is an $\varepsilon > 0$ and a language $L \in \text{E}^{\text{NP}}$ such that L_n cannot be $(1/2 + 2^{-n^\varepsilon})$ -approximated by $\text{AC}_d^0[m]$ circuits of size 2^{n^ε} , for all sufficiently large n .*

Lower bounds on the many-one closure of ACC^0 . Following the observation of Arvind and Srinivasan [AS10], the FP^{NP} algorithm for $\text{ACC}^0\text{-PARTIAL-AVGHARD}$ can be used to prove *unconditionally* that E^{NP} cannot be mapping reduced to languages decidable by small-size *non-uniform* families of ACC^0 circuits.⁹ To the best of our knowledge, this is the first unconditional result ruling out the mapping reducibility from uniform classes to non-trivial non-uniform classes.

Corollary 1.20. *Let $d, m \in \mathbb{N}$ be constants, $\text{AC}_d^0[m]$ denote the class of languages computable by a non-uniform family of polynomial-size $\text{AC}_d^0[m]$ circuits. Then, there is a language $L^{\text{hard}} \in \text{E}^{\text{NP}}$ that does not have polynomial-time mapping reductions to any language in $\text{AC}_d^0[m]$.*

1.2.3 A Smooth and Rectangular PCPs of Proximity

One of the main technical ingredients in our framework for the average-case explicit construction problems (i.e. REMOTE-POINT and PARTIAL-AVGHARD) is a PCP of Proximity (PCPP) that is *short*, *smooth*, and (almost) *rectangular*.

A PCPP verifier V for a language L provides a super-efficient probabilistic proof system for checking whether $x \in L$ or x is far from being in L . Given an input x and a proof π , the verifier with access to some random bits only probes constantly many bits of x and π . If $x \in L$, then it accepts with an appropriate proof π ; if the relative Hamming distance between x and any $x' \in L$ is at least δ , then it rejects with constant probability regardless of the proof π . (The distance δ is called the *proximity parameter* of the PCPP.) In addition, our PCPP verifier is equipped with the following properties:

⁹In fact, it suffices to have an FP^{NP} algorithm for $\text{ACC}^0\text{-PARTIAL-HARD}$ (which is a trivial consequence of an FP^{NP} algorithm for $\text{ACC}^0\text{-PARTIAL-AVGHARD}$) for this application.

- **(Shortness).** For any language $L \in \text{NTIME}[T(n)]$, where $n \leq T(n) \leq 2^{\text{poly}(n)}$, the PCPP proof for L has length $T(n) \cdot \text{polylog}(T(n))$.
- **(Rectangularity).** The input and the proof are treated as matrices. Moreover, the queries of the verifier to the input and proof matrices can be done *rectangularly*, in the sense that there are a row verifier V_{row} and a column verifier V_{col} with (almost) independent random seeds that generate the row and column indices of the queries, respectively.
- **(Smoothness).** The queries of the verifier to the proof matrix are uniformly random. As a consequence, it means that the PCPP proof can tolerate errors in a correct proof.

We refer the readers to Section 2.5 for formal definitions of these properties.

Before our work, Bhangale, Harsha, Paradise, and Tal [BHPT20] constructed a short, smooth, and rectangular PCP (instead of PCPP) built upon [BGH⁺06] with an application of constructing *rigid matrices* (also see [Val77, AC19]). Ren, Santhanam, and Wang [RSW22] constructed a short and rectangular PCPP based on [BGH⁺06, BHPT20] for the Algorithmic Method for AVOID. It turns out that to generalise [RSW22] to the “average-case” explicit construction problems REMOTE-POINT and PARTIAL-AVGHARD, we need both *smoothness* (as in [BHPT20]) and PCPs of *proximity* (as in [RSW22]). A technical contribution of this work is to combine [BHPT20] and [RSW22] to obtain a *smooth PCPP*.

Theorem 1.21 (Theorem 2.14, Informal). *Let $T(n)$ be a good function. For every language $L \in \text{NTIME}[T(n)]$, there is a short, smooth, and (almost) rectangular PCP of proximity verifier V for L , with perfect completeness, constant soundness error, and constant query complexity.*

Following standard techniques in the algorithmic approach to lower bounds (see, e.g., [CW19b]), we also construct a short and rectangular (non-smooth) PCPP that makes at most two queries to the input and the proof matrices (see Theorem 2.13) to minimise the overhead on the circuit class when we reduce AVOID and PARTIAL-HARD to SATISFYING-PAIRS (i.e. the upper OR_2 in Theorem 1.10 and Theorem 1.13).

1.3 Technical Overview

As mentioned in Section 1.2.1, the range avoidance algorithm follows from slightly modifying the framework in [RSW22] and using an algorithm for SATISFYING-PAIRS. In what follows, we briefly illustrate our techniques for the remote point problem and for constructing hard partial truth tables. The high-level idea is to reduce these problems to AVOID and invoke our framework for AVOID to solve them.

Remote Point Problem. Our start point is the following reduction from REMOTE-POINT to AVOID. Suppose that $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ is the input circuit. Let $\text{Enc} : \{0, 1\}^{\ell'} \rightarrow \{0, 1\}^\ell$ be the encoding procedure of an error correcting code, and $\text{Dec} : \{0, 1\}^\ell \rightarrow \{0, 1\}^{\ell'}$ be the corresponding decoding procedure, where Dec can correct a δ fraction of errors. Define the circuit $C'(x) := \text{Dec}(C(x))$, and let z be any string not in the range of C' , then $\text{Enc}(z)$ is $(1 - \delta)$ -far from $\text{Range}(C)$. To see this, assume for contradiction that $\text{Enc}(z)$ is $(1 - \delta)$ -close to some $C(x)$, then $\text{Dec}(C(x))$ should return exactly z , contradicting that z is a non-output of C' .

Suppose that the function Dec can be implemented in the circuit class \mathcal{C}_{Dec} , then this is a reduction from \mathcal{C} -REMOTE-POINT to $(\mathcal{C}_{\text{Dec}} \circ \mathcal{C})$ -AVOID. Therefore, we would like the complexity of \mathcal{C}_{Dec} to be as small as possible. There are decoders that tolerate a small constant fraction of errors in AC^0 [GGH⁺07], so it might be possible to implement \mathcal{C}_{Dec} in AC^0 . However, when δ is very close to $1/2$ (say $\delta = 1/2 - \varepsilon$), we enter the list-decoding regime where \mathcal{C}_{Dec} seems to need the

power of majority [GR08]. Can we solve \mathcal{C} -REMOTE-POINT without invoking any circuit-analysis algorithms for $\text{MAJ} \circ \mathcal{C}$?

Fortunately, the required techniques already appeared in previous works on the Algorithmic Method for proving strong average-case circuit lower bounds. In [CLW20], they provided an error-correcting code that corrects a $1/2 - \varepsilon$ fraction of errors, where the decoder Dec_{CLW} can be implemented as a *linear sum*, i.e., each output is a linear combination of the input bits.¹⁰ Intuitively, this means that we can reduce \mathcal{C} -REMOTE-POINT to $(\text{Sum} \circ \mathcal{C})$ -AVOID, where Sum denotes the layer of Dec_{CLW} . Using the framework for range avoidance established above, it suffices to solve the SATISFYING-PAIRS problem for $\text{Sum} \circ \mathcal{C}$ circuits.¹¹ But it is easy to see that SATISFYING-PAIRS for $\text{Sum} \circ \mathcal{C}$ circuits directly reduces to SATISFYING-PAIRS for \mathcal{C} circuits! Therefore, the error-correcting code in [CLW20] allows us to use an algorithm for \mathcal{C} -SATISFYING-PAIRS to directly solve \mathcal{C} -REMOTE-POINT, with little or no circuit complexity overhead.

The above discussion omitted several important technical details:

- It turns out that Dec_{CLW} is only an *approximate* list-decoding algorithm: given a corrupted codeword that is $(1/2 - \varepsilon)$ -close to the correct codeword, we can only recover a message that is δ -close to the correct message (instead of perfectly recovering the correct message).

This drawback is handled by *smooth PCPPs* [Par21], which has the property that any slightly corrupted version of a correct proof is still accepted with good probability. As we need a rectangular PCPP in [RSW22], what we actually need is a *smooth and rectangular PCPP*, which we construct in Section 6. We remark that [CLW20] also encountered this difficulty; they got around it by combining a PCP and a PCPP for CIRCUIT-EVAL. It is not clear how to generalise this strategy to our case.

- Another technical complication is that Dec_{CLW} outputs *real values* instead of Boolean values. It is only guaranteed that the decoded message is close to the original message *in ℓ_1 -norm*. Consequently, after guessing the PCPP proof, we also need to verify that it is “close to Boolean”. This difficulty also appears in [CLW20]; however, we need to carefully define what it means by “close to Boolean” in our case.
- Since Dec_{CLW} works in the list-decoding regime, it also receives an advice string (specifying the index of the codeword in the list). In the above discussion, we omitted the advice string to highlight the main ideas. It turns out that the dependency of the decoder on the advice string *cannot* be captured by linear sums. Therefore, we need to define an ad hoc “linear sum” circuit class (in Section 2.4) that receives both an input and an advice string and computes a linear combination over the input, where the “linear combination” depends on the advice. It turns out that we need the dependency on the advice to be *local* (see Section 2.4 for details), which is fortunately satisfied by the code in [CLW20].

Another reduction via succinct dictionaries. We mention that there is another reduction from REMOTE-POINT to AVOID which appears in [Kor21, GLW22]. Let $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ be a circuit, $y \in \{0, 1\}^\ell$ be a string that is not δ -far from $\text{Range}(C)$. Then we can find a string $x \in \{0, 1\}^n$ and a “noise” string $e \in \{0, 1\}^m$ of relative Hamming weight at most δ such that $y = C(x) \oplus e$,

¹⁰[CLW20] stated this result as a non-standard XOR lemma in their Appendix A. We re-prove it in the form of error-correcting codes in Appendix A.5.

¹¹We made a simplification here. Actually, we need to solve SATISFYING-PAIRS for $\text{NC}^0 \circ \text{Sum} \circ \mathcal{C}$ circuits. Using the distributive property, we can push the NC^0 circuits below the Sum layer, thus it suffices to solve SATISFYING-PAIRS for $\text{Sum} \circ \text{NC}^0 \circ \mathcal{C}$ circuits. In this informal exposition, we may assume that \mathcal{C} is closed under top NC^0 gates, which means that a SATISFYING-PAIRS algorithm for $\text{Sum} \circ \mathcal{C}$ now suffices.

where \oplus refers to bit-wise XOR. Consider the circuit $C'(x, e) := C(x) \oplus e$. To solve the remote point problem for C , it suffices to solve the range avoidance problem for C' . Using a “succincter” dictionary to represent e [Pät08], [GLW22] managed to show that this reduction also preserves circuit complexity, and in particular reduces $\text{NC}^1\text{-REMOTE-POINT}$ to $\text{NC}^1\text{-AVOID}$.

A drawback of this approach is that it only reduces REMOTE-POINT to range avoidance instances with a small stretch. Indeed, suppose C' is a circuit from n' inputs to ℓ outputs, and $\delta = \Omega(1)$, then

$$n' \geq |\Pi(e)| \geq \log \binom{\ell}{\delta\ell} = \Omega(\ell).$$

In contrast, the algorithmic method in both [RSW22] and this paper could not solve range avoidance instances with such a small stretch ($\ell = c \cdot n$ for some constant c), even with the best possible algorithms for SATISFYING-PAIRS . Therefore we do not use this approach in this paper.

Hard Partial Truth Table. There is a simple reduction from PARTIAL-HARD to AVOID . Suppose we are given strings x_1, x_2, \dots, x_N . Let TT' be the circuit that receives a size- s circuit C as input, and outputs the concatenation of $C(x_1), C(x_2), \dots, C(x_N)$. If $N > O(s \log s)$ then the circuit TT' is stretching. It is also easy to see that solving the range avoidance of TT' is equivalent to solving the PARTIAL-HARD problem.

In Section 4, we essentially combine this reduction with the frameworks in Section 3. In other words, we could have reduced PARTIAL-HARD to AVOID in a black-box way and derived the main results in Section 4. However, this reduction only reduces $\mathcal{C}\text{-PARTIAL-HARD}$ to $\mathcal{C}'\text{-AVOID}$, where \mathcal{C}' is any circuit class that can solve $\mathcal{C}\text{-EVAL}$ in the following sense: for every fixed input x , there is a \mathcal{C}' circuit C' that takes as input the description of a \mathcal{C} circuit C , and outputs $C(x)$. For most circuit classes of interest (e.g., $\mathcal{C} \in \{\text{AC}^0, \text{ACC}^0, \text{NC}^1, \text{P}_{/\text{poly}}\}$), we could simply let $\mathcal{C}' = \mathcal{C}$; however, this is not necessarily true for more refined circuit classes (such as $\mathcal{C} = \text{ACC} \circ \text{THR}$). We choose to derive the main results in Section 4 from scratch instead of reducing it to Section 3, partly because we also want our framework to hold for these more refined circuit classes.

1.4 Further Related Work

SATISFYING-PAIRS and the Polynomial Method. The SATISFYING-PAIRS problems for restricted circuit classes nicely capture a wide range of algorithmic problems that have been extensively studied. For instance, the Orthogonal Vector Problem over \mathbb{F}_2 corresponds to $\text{XOR-SATISFYING-PAIRS}$, and the (decision version of) Nearest Neighbor Problem corresponds to the SATISFYING-PAIRS of polynomial threshold functions (see, e.g., [Wil14b, ACW16]).

There is a successful line of research on non-trivial algorithms for this kind of problems via the *polynomial method* [Raz87, Smo87] in circuit complexity. Williams [Wil18a] developed an $n^3/2^{(\log n)^{\Omega(1)}}$ -time algorithm for the All-Pairs Shortest Path problem using the Razborov-Smolensky polynomial representation of $\text{AC}^0[p]$ circuits [Raz87, Smo87, Smo93] and a fast batch evaluation of polynomials via fast rectangular matrix multiplication [Cop82] (also see Theorem 5.1). Similar techniques were used to design non-trivial algorithms for the Orthogonal Vector Problem over \mathbb{F}_2 [CW21, AWY15] and the (approximate) Nearest Neighbor Problems (with respect to Hamming distance, ℓ_1 -distance, and ℓ_2 -distance) [AW15, ACW16, ACW20]. Chen and Wang [CW19a] (following [AW17]) generalised the polynomial method in algorithm design by showing a connection between SATISFYING-PAIRS problems and quantum communication protocols, with an application in $\text{Approx}_\varepsilon\text{-XOR-SATISFYING-PAIRS}$ (which is called Approximate #OV in [CW19a]).

Explicit obstructions. Related to the PARTIAL-HARD problem is the notion of explicit obstructions [Mul11, CJW20]: on input 1^n , one wants to output a list of (x_i, y_i) deterministically, such that $x_i \neq x_j$ for $i \neq j$, and for all n -input circuit C from a certain circuit class \mathcal{C} , there is some i such that $C(x_i) \neq y_i$. This notion is weaker than deterministic algorithms for PARTIAL-HARD, as one has the freedom of choosing the inputs $\{x_i\}$. Chen, Jin, and Williams [CJW20] exhibited a “sharp threshold” phenomenon for explicit obstructions against de Morgan formulas: an explicit obstruction for $\text{Formula}[n^{1.99}]$ provably exists, while an explicit obstruction for $\text{Formula}[n^{2.01}]$ would imply very strong circuit lower bounds.

1.5 Organisation

ORGANISATION OF THE PAPER

- In Section 2, we introduce concepts and tools used in this paper and fix the notation.
- In Section 3, we demonstrate the framework of solving AVOID and REMOTE-POINT via non-trivial algorithms for SATISFYING-PAIRS.
- In Section 4, we demonstrate the framework of solving PARTIAL-HARD and PARTIAL-AVGHARD via non-trivial algorithms for SATISFYING-PAIRS.
- In Section 5, we present a non-trivial algorithm for ACC^0 -SATISFYING-PAIRS from [Wil18c], and combine it with the frameworks to obtain unconditional FP^{NP} algorithms for ACC^0 -REMOTE-POINT and ACC^0 -PARTIAL-AVGHARD. We also demonstrate the consequences of these algorithms.
- In Section 6, we construct short, smooth, and rectangular PCPs of Proximity, used in the frameworks of solving REMOTE-POINT and PARTIAL-AVGHARD.
- In Section 7, we construct short and rectangular PCPs of Proximity with query complexity only 2 or 3, used in the framework for AVOID and PARTIAL-HARD.
- In Appendix A, we prove some technical lemmas in Section 3 and 4.

2 Preliminaries

Notation. We use $\tilde{O}(f(n))$ to denote $f(n) \cdot (\log f(n))^{O(1)}$. The concatenation of the strings x and y is denoted by $x \circ y$. The relative Hamming distance of two strings x and y , denoted by $\delta(x, y)$, is the fraction of indices i such that $x_i \neq y_i$. A string x is said to be γ -far from (resp. γ -close to) a string y if $\delta(x, y) \geq \gamma$ (resp. $\delta(x, y) < \gamma$). We say $x \in \{0, 1\}^n$ is γ -far from $L \subseteq \{0, 1\}^n$ if x is γ -far from every $y \in L$; otherwise x is γ -close to L . For a vector $\vec{u} \in \mathbb{R}^n$ and an integer $d \geq 1$, the ℓ_d norm of \vec{u} is

$$\|\vec{u}\|_d := \left(\mathbb{E}_{i \leftarrow [n]} [|u_i|^d] \right)^{1/d}.$$

Given a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, the range of f , denoted by $\text{Range}(f)$, is defined as the set of outputs of f , i.e., $\text{Range}(f) := \{f(x) \mid x \in \{0, 1\}^n\}$.

A function $f : \mathbb{N} \rightarrow \mathbb{N}$ is said to be *good* if there is a Turing machine such that given n in binary, it runs in time $\text{poly}(\log n, \log f(n))$ and outputs $f(n)$ in binary.

A circuit class \mathcal{C} is said to be *typical* if it contains the identity circuit and is closed under negations and projections. More precisely, (1) every function that always outputs its input bits is computable by a constant size \mathcal{C} circuit; (2) for any \mathcal{C} circuit C of size s and projection proj , both $\neg C$ and $C \circ \text{proj}$ have \mathcal{C} circuits of size $\text{poly}(s)$, and the descriptions of these circuits can be computed in $\text{poly}(s)$ time.

Let S be a finite sample space and E be an event. We use $\Pr_{x \leftarrow S}[E]$ to denote the probability that E happens if x is sampled uniformly from S . Similarly, for a random variable Y , we use $\mathbb{E}_{x \leftarrow S}[Y]$ to denote the expectation of Y when x is sampled uniformly from S .

2.1 Circuit Classes

Throughout this paper, the *size* of a circuit is defined as the number of *wires* (instead of gates) in the circuit. We will use the following (single-output) circuit classes.

- AC_d^0 refers to depth- d circuits with AND and OR gates of unbounded fan-in, and NOT gates of fan-in 1. We define $\text{AC}^0 := \bigcup_d \text{AC}_d^0$.
- $\text{AC}_d^0[m]$ refers to depth- d circuits with AND, OR, and $\text{MOD}[m]$ gates of unbounded fan-in, and NOT gates of fan-in 1. A $\text{MOD}[m]$ gate outputs 1 if and only if m does not divide the number of 1 in its inputs. We define $\text{AC}^0[m] := \bigcup_d \text{AC}_d^0[m]$. Furthermore, we define $\text{ACC}^0 := \bigcup_m \text{AC}^0[m]$.
- $\text{CC}_d^0[m]$ refers to depth- d circuits with only $\text{MOD}[m]$ gates of unbounded fan-in. We define $\text{CC}^0[m] := \bigcup_d \text{CC}_d^0[m]$ and $\text{CC}^0 := \bigcup_m \text{CC}^0[m]$.
- NC_d^0 refers to constant-size circuits such that the output depends on at most d input bits. We define $\text{NC}^0 := \bigcup_d \text{NC}_d^0$.
- Assume that $F \in \{\text{AND}, \text{OR}, \text{XOR}, \text{MOD}[m], \dots\}$ is a gate, we define an F circuit as a circuit with only an F gate fed by some (or all) of the input bits. In particular, we define an F_d circuit as a circuit with an F gate of fan-in at most d fed by some (or all) of the input bits.

We define SYM as the class of any symmetric Boolean function, i.e., $f : \{0, 1\}^n \rightarrow \{0, 1\}$ such that $f(x) = g(x_1 + x_2 + \dots + x_n)$ for some function g .

Suppose that \mathcal{C}_1 and \mathcal{C}_2 are circuit classes, we denote $\mathcal{C}_1 \circ \mathcal{C}_2$ as the *composition* of these two classes: the input bits feed an n -input m -output \mathcal{C}_2 circuit C_2 , and the m output bits of C_2 feed an m -input single-output \mathcal{C}_1 circuit C_1 . For instance, a $\text{SYM} \circ \text{ACC}^0$ circuit contains a symmetric output gate whose inputs are ACC^0 circuits.

For a circuit class \mathcal{C} , we use $\mathcal{C}[s]$ to represent the sub-class of the \mathcal{C} circuits of size at most s .

2.2 Error-Correcting Codes

An *error-correcting code* with *message length* n , *rate* r , and *relative distance* δ is a function $\text{Enc} : \{0, 1\}^n \rightarrow \{0, 1\}^{rn}$ such that for every pair of distinct $x_1, x_2 \in \{0, 1\}^n$, $\delta(\text{Enc}(x_1), \text{Enc}(x_2)) \geq \delta$. It is said to *correct* γ *fraction of errors* if there is a $\text{Dec} : \{0, 1\}^{rn} \rightarrow \{0, 1\}^n$ such that for every y that is γ -close to $\text{Enc}(x)$ for some $x \in \{0, 1\}^n$, $\text{Dec}(y) = x$.

We need the following standard construction of error-correcting codes.

Theorem 2.1 ([Spi96]). *There is a GF(2)-linear error-correcting code (Enc, Dec) with a constant rate, constant relative distance, and can correct a constant fraction of errors. Moreover, both Enc and Dec are uniformly computable in linear time.*

2.3 An Almost-Everywhere NTIME Hierarchy with a Refuter

We need the almost-everywhere NTIME hierarchy against bounded nondeterminism [FS16], which has an FP^{NP} refuter as shown in [CLW20]. Let $T(n), G(n)$ be good functions, we define $\text{NTIME}[T(n)]$ to be the class of languages decidable by nondeterministic Turing machines in $T(n)$ time, and $\text{NTIMEGUESS}_{\text{RTM}}[T(n), G(n)]$ to be the class of languages decidable by nondeterministic Random-Access Turing Machines (RTMs) in $T(n)$ time with $G(n)$ nondeterministic bits.

Theorem 2.2 ([FS16, CLW20, RSW22]). *Let c be a large universal constant, $T : \mathbb{N} \rightarrow \mathbb{N}$ be a good function such that $n \log^{c+1} n \leq T(n) \leq 2^{\text{poly}(n)}$. There is a language*

$$L^{\text{hard}} \in \text{NTIME}[T(n)] \setminus \text{i.o.-NTIMEGUESS}_{\text{RTM}}[T(n)/\log^c T(n), n/10]_{/(n/10)}.$$

Moreover, there is an algorithm \mathcal{R} (the “refuter”) such that the following holds.

(Input) \mathcal{R} receives three inputs $(1^n, M, \alpha)$, where M is a nondeterministic RTM and $\alpha \in \{0, 1\}^{n/10}$ is an advice string. It is guaranteed that M runs in $T(n)/\log^c T(n)$ time and uses at most $n/10$ nondeterministic bits; moreover, the description length of M is $O(1)$.

(Output) For every fixed M , every sufficiently large n , and every advice $\alpha \in \{0, 1\}^{n/10}$, $\mathcal{R}(1^n, M, \alpha)$ outputs a string $x \in \{0, 1\}^n$ such that $M(x; \alpha) \neq L^{\text{hard}}(x)$.

(Complexity) \mathcal{R} runs in $\text{poly}(T(n))$ time with adaptive access to an NP oracle.

2.4 Linear Sum Circuits and Hardness Amplification with Them

We need an XOR lemma with “linear sum” decoders: given a corrupted codeword \tilde{f} that is $(1/2 - \varepsilon)$ -close to $\text{Amp}(f)$, there is an affine transformation A such that $A(\tilde{f})$ is δ -close to f .

The actual definition of *linear sum circuits* is more involved for the following reason. Our XOR lemma works in the *list-decoding* regime, therefore it also receives an advice string α (i.e., the index in the list) and outputs the α -th decoded message in the list. When α is fixed, $A(\tilde{f}; \alpha)$ is simply an affine function over \tilde{f} ; but the dependence on α can be more complicated. It turns out that we need an upper bound on the *locality* of the dependence on α , defined as follows.

Definition 2.3 (Linear Sum Circuits). Let $x \in \{0, 1\}^n$ and $\alpha \in \{0, 1\}^a$ be two inputs. A *linear sum* circuit on input x with advice α is a function $C : \{0, 1\}^n \times \{0, 1\}^a \rightarrow \mathbb{R}^m$ of the following form:

$$C(x, \alpha)_i = \sum_{k \in [A]} \text{coeff}_k(\alpha) \cdot x_{\text{idx}_k(\alpha, i)}.$$

Here, A is the *fan-in* of C . The circuit is described by two functions $\text{coeff}_k(\alpha)$ and $\text{idx}_k(\alpha, i)$; note that $\text{coeff}_k(\alpha)$ does not depend on i . For technical convenience, we will also allow $\text{idx}_k(\alpha, i)$ to take special values ZERO and ONE, where x_{ZERO} is always 0 and x_{ONE} is always 1.

Besides the fan-in A , the following complexity measures of C will also be important:

- We say the *coefficient sum* of C is at most U , if for every advice α , we have

$$\sum_{k \in [A]} |\text{coeff}_k(\alpha)| \leq U.$$

- We say that C has *locality* l , if for every fixed k , there is a subset S_k of l bits of α such that the functions $\text{coeff}_k(\alpha)$ and $\text{idx}_k(\alpha, i)$ only depends on $\alpha|_{S_k}$.

Example 2.4. Consider the following example (simplified from the proof of Theorem 2.5). Suppose the advice α consists of a list of sub-advice $(\alpha_1, \alpha_2, \dots, \alpha_{a'})$ where $a' \approx 1/\varepsilon^2$; given an index k , $\text{coeff}_k(\alpha)$ only depends in α_k , and $\text{idx}_k(\alpha, i)$ only depends on α_k and i . Suppose each α_k has length l , then regardless of the number a' , the linear sum has locality l .

We need the following XOR lemma with linear sum decoders. The XOR lemma was proved in [Lev87, GNW11] and it was shown in [CLW20, Section A] to admit linear sum decoders. For completeness, we provide a proof of Theorem 2.5 in Appendix A.5 and verify the locality of the linear sum. Note that the XOR lemma is stated below as an approximately locally list-decodable code.

Theorem 2.5. *Let $N \in \mathbb{N}$, $0 < \varepsilon, \delta < 1/10$, $k := O(\log(1/\varepsilon)/\delta)$, $\tilde{N} := N^k$, and $a := O(\log^2 N/(\varepsilon\delta)^2)$. There is an algorithm $\text{Amp} : \{0, 1\}^N \rightarrow \{0, 1\}^{\tilde{N}}$ computable in deterministic $\text{poly}(\tilde{N})$ time, and a linear sum circuit $C : \{0, 1\}^{\tilde{N}} \times \{0, 1\}^a \rightarrow \mathbb{R}^N$ such that the following hold.*

(List-decoding) *For every string $\tilde{f} \in \{0, 1\}^{\tilde{N}}$ that is $(1/2 - \varepsilon)$ -close to $\text{Amp}(f)$ for some hidden string f , there is an advice $\alpha \in \{0, 1\}^a$, such that (1) for every $i \in [N]$, $C(\tilde{f}, \alpha)_i \in [0, 1]$; and (2) $\|C(\tilde{f}, \alpha) - f\|_1 \leq \delta$.*

(Complexity) *The fan-in, coefficient sum, and locality of C are at most $O(\log N/(\varepsilon\delta)^2)$, $O(1/\varepsilon)$, and $\log \tilde{N}$ respectively.*

We will also use the notation $\text{dec}_\alpha(f)$ to denote $C(f, \alpha)$, emphasising that dec_α is an affine transformation that depends on α .

2.5 PCPs of Proximity

Now we will introduce Probabilistically Checkable Proofs of Proximity (PCPPs) [BGH⁺06] and two properties of PCPPs that will be useful in designing algorithms for explicit construction problems: *rectangularity* and *smoothness*.

In what follows, a *pair language* is simply a subset of $\{0, 1\}^* \times \{0, 1\}^*$. For an instance (z, x) of a pair language, we treat z as the *explicit input* (which the PCPP verifier can read entirely) and x as the *implicit input* (which the PCPP verifier could only read a few bits). For example, CIRCUIT-EVAL is a pair language with two inputs, i.e., a circuit C and an input x , and the task is to evaluate $C(x)$. A PCPP verifier for CIRCUIT-EVAL knows the input circuit C but can only access a few bits of x .

2.5.1 Basic Definitions

Definition 2.6 (PCP of Proximity Verifiers). Let $r = r(n)$, $q = q(n)$, $\ell = \ell(n)$, $d = d(n)$ be good functions and $L \subseteq \{0, 1\}^* \times \{0, 1\}^*$ be a pair language. A PCPP verifier VPCPP for L with *proof length* ℓ , *randomness complexity* r , *decision complexity* d , and *query complexity* q is a tuple of Turing machines $(V_{\text{type}}, V_{\text{index}}, V_{\text{dec}})$ that will verify a proof $\pi \in \{0, 1\}^\ell$ of the statement $(z, x) \in L \cap \{0, 1\}^* \times \{0, 1\}^n$ as follows.

- It randomly samples a *seed* $\in \{0, 1\}^r$ and generates

$$\begin{aligned} (\text{itype}[1], \text{itype}[2], \dots, \text{itype}[q]) &\leftarrow V_{\text{type}}(\text{seed}, z), \\ (i[1], i[2], \dots, i[q]) &\leftarrow V_{\text{index}}(\text{seed}, z). \end{aligned}$$

For every $j \in [q]$, $\text{itype}[j] \in \{\text{input}, \text{proof}\}$ determines the type of the j -th query: If $\text{itype}[j] = \text{input}$, the j -th query probes the $i[j]$ -th bit of the “implicit input” x ; otherwise (i.e., $\text{itype}[j] = \text{proof}$), the j -th query probes the $i[j]$ -th bit of the proof π .

- Let $\text{ans}_1, \dots, \text{ans}_q$ be the answers to the queries defined above, we say VPCPP accepts (z, x, π) , denoted by $\text{VPCPP}^{x \circ \Pi}(z, \text{seed}) = 1$, if and only if $V_{\text{dec}}(\text{seed}, z, \text{ans}_1, \dots, \text{ans}_q) = 1$. The machine V_{dec} is said to be the *decision predicate* of VPCPP, and has circuit complexity $d(n)$.

We may represent the “implicit input” x as $\Pi_{\text{input}} : [n] \rightarrow \{0, 1\}$ and the proof π as $\Pi_{\text{proof}} : [\ell] \rightarrow \{0, 1\}$ to emphasize that they are given as oracles to VPCPP. We sometimes denote the outputs of V_{type} and V_{index} as I and denote the answers $(\text{ans}_1, \dots, \text{ans}_q)$ as $(\Pi_{\text{input}} \circ \Pi_{\text{proof}})|_I$.

We will also consider the PCPP verifier of *pure* languages (i.e. the first part z of any input is always the empty string). In such case, we simply omit all the z in the definition above.

Definition 2.7 (PCP of Proximity). Let $s = s(n)$ and $\delta = \delta(n)$ be good functions, $L \subseteq \{0, 1\}^* \times \{0, 1\}^*$ be a pair language, and $\text{VPCPP} = (V_{\text{type}}, V_{\text{index}}, V_{\text{dec}})$ be a PCPP verifier for L . We say VPCPP is a PCPP verifier for L with *completeness error* $1 - c$, *soundness error* s , and *proximity parameter* δ if the following two conditions hold for every $(z, x) \in \{0, 1\}^* \times \{0, 1\}^n$.

- **(Completeness)**. If $(z, x) \in L$, then there is a proof $\pi \in \{0, 1\}^\ell$ such that VPCPP accepts (z, x, π) with probability at least c .
- **(Soundness)**. Denote $L(z)$ to be the set of $y \in \{0, 1\}^n$ such that $(z, y) \in L$. If x is δ -far from $L(z)$, then for every proof $\pi \in \{0, 1\}^\ell$, VPCPP accepts (z, x, π) with probability at most s .

For most of the constructions of PCPPs, the completeness error can be made 0, which means that for $(z, x) \in L$, there is a proof such that the verifier accepts with probability 1. Therefore we assume that the completeness error of a PCPP is 0 when it is not specified.

We need to define a stronger version of the soundness called *robust soundness* as follows, as an intermediate step to construct PCPPs with nice parameters.

Definition 2.8 (Robust PCP of Proximity [BGH⁺06]). Let $s = s(n)$, $\delta = \delta(n)$, and $\rho = \rho(n)$ be good functions, $L \subseteq \{0, 1\}^* \times \{0, 1\}^*$ be a pair language, and $\text{VPCPP} = (V_{\text{type}}, V_{\text{index}}, V_{\text{dec}})$ be a PCPP verifier for L . We say VPCPP is a robust PCPP verifier for L with *robust soundness error* s with robustness parameter ρ and *proximity parameter* δ if it satisfies the completeness property of PCPP and the following *robust soundness property*.

- **(Robust Soundness)**. The following holds for every $(z, x) \in \{0, 1\}^* \times \{0, 1\}^n$. Denote $L(z)$ to be the set of $y \in \{0, 1\}^n$ such that $(z, y) \in L$. If x is δ -far from $L(z)$, then for every proof $\pi \in \{0, 1\}^\ell$, with probability at least $1 - s$ over the random bits seed, the answer $(\text{ans}_1, \dots, \text{ans}_q)$ of the queries of VPCPP is ρ -far from being accepted (i.e. we need to flip at least a ρ fraction of the bits of the answers $(\text{ans}_1, \dots, \text{ans}_q)$ to make the verifier accept).

2.5.2 Rectangular PCPs of Proximity

Following [RSW22], one of the main technical ingredients in the algorithmic method for explicit construction problems is a variant of PCPPs, called *rectangular PCPPs*. Intuitively, a rectangular PCPP verifier treats the input as an $H_{\text{input}} \times W_{\text{input}}$ matrix and the proof as an $H_{\text{proof}} \times W_{\text{proof}}$ matrix, and can generate the query indices in a “rectangular” fashion. In particular, the random seed is split into two parts denoted as `seed.row` and `seed.col` respectively, and there are two algorithms V_{row} and V_{col} such that:

- V_{row} takes `seed.row` as input and generates $\text{irow}[1], \dots, \text{irow}[q]$;
- V_{col} takes `seed.col` as input and generates $\text{icol}[1], \dots, \text{icol}[q]$;
- The final indices of the queries $i[1], \dots, i[q]$ are defined as $i[j] := (\text{irow}[j] - 1) \cdot W + \text{icol}[j]$, where $W = W_{\text{input}}$ or $W = W_{\text{proof}}$ depending on the type of the j -th query.

In other words, the row verifier V_{row} (resp. the column verifier V_{col}) takes the row randomness `seed.row` (resp. the column randomness `seed.col`) and generates the row indices (resp. the column indices) of the queries. Ideally, a rectangular PCPP should satisfy the following properties:

- **(Perfect Rectangularity).** The row randomness `seed.row` and column randomness `seed.col` are independent random bits (i.e. the row and column query indices are independent).
- **(Randomness-Oblivious Type Predicate).** The type predicate V_{type} , which determines the types of the queries (i.e. whether a query is to the input or the proof oracle), does not depend on the row and column random seeds.
- **(Randomness-Oblivious Decision Predicate).** The decision predicate V_{dec} , which decides whether to accept the proof given the answers to the queries, does not depend on the row and column random seeds.

However, as in [BHPT20, RSW22], we do not know how to construct such rectangular PCPPs. Nevertheless, we could construct a weaker version where the row and column randomness are almost independent, and the dependency of the decision and type predicates on the random seeds are relatively simple. We formally define such rectangular PCPPs as follows (for simplicity, we only define rectangular PCPPs for pure languages).

Definition 2.9 (Rectangular PCPPs with Randomness-Oblivious Predicates). Suppose $H_{\text{input}} = H_{\text{input}}(n)$, $W_{\text{input}} = W_{\text{input}}(n)$, $H_{\text{proof}} = H_{\text{proof}}(n)$, $W_{\text{proof}} = W_{\text{proof}}(n)$, $\tau = \tau(n)$, $p = p(n)$ are good functions such that $H_{\text{input}} \cdot W_{\text{input}} = O(n)$ and $L \subseteq \{0, 1\}^*$ be a language. A PCPP verifier VPCPP is said to be a τ -almost rectangular PCPP that has a *randomness-oblivious predicate* (ROP) with parity-check complexity p if the following conditions hold.

- **(Randomness).** There are good functions $r_{\text{row}} = r_{\text{row}}(n)$, $r_{\text{col}} = r_{\text{col}}(n)$, and $r_{\text{shared}} = r_{\text{shared}}(n)$ such that the randomness complexity $r = r_{\text{row}} + r_{\text{col}} + r_{\text{shared}}$; i.e., the random seed can be partitioned into three independent parts: the *row randomness* `seed.row`, the *column randomness* `seed.col`, and the *shared randomness* `seed.shared`. We say r_{row} , r_{col} , and r_{shared} are the row, column, and shared randomness complexity of the PCPP verifier, respectively. Moreover, the shared randomness complexity $r_{\text{shared}}(n) \leq \tau \cdot r(n)$.
- **(Query Pattern).** There are algorithms V_{type} , V_{row} , and V_{col} to generate the queries in a rectangular fashion. Concretely speaking:
 - $(\text{itype}[1], \dots, \text{itype}[q]) \leftarrow V_{\text{type}}(\text{seed.shared})$, where $\text{itype}[j] \in \{\text{input}, \text{proof}\}$ for all $j \in [q]$.
 - $(\text{irow}[1], \dots, \text{irow}[q]) \leftarrow V_{\text{row}}(\text{seed.row}, \text{seed.shared})$
 - $(\text{icol}[1], \dots, \text{icol}[q]) \leftarrow V_{\text{col}}(\text{seed.col}, \text{seed.shared})$
 - For every $j \in [q]$, the index of the j -th query $i[j] := \text{irow}[j] \cdot W + \text{icol}[j]$, where $W = W_{\text{input}}$ if $\text{itype}[j] = \text{input}$ and $W = W_{\text{proof}}$ otherwise.

As normal PCPP verifiers, the j -th query is to the $i[j]$ -th bit of the input if $\text{itype}[j] = \text{input}$, and is to the $i[j]$ -th bit of the proof if $\text{itype}[j] = \text{proof}$. Note that since $H_{\text{input}} \cdot W_{\text{input}}$ may be larger than n , the query to the input is not well-defined when $i[j] > n$. In such case, we denote the answer to be \perp .

- **(Decision Predicate).** There are polynomial-time algorithms V_{dec} and V_{pc} such that the following holds.
 - The algorithm $V_{\text{dec}}(\text{seed.shared})$ generates a circuit $\text{VDec} : \{0, 1, \perp\}^{p+q} \rightarrow \{0, 1\}$.
 - The algorithm $V_{\text{pc}}(\text{seed.shared})$ generates p XOR circuits (i.e. circuits computing GF(2)-linear functions) $pc_1, \dots, pc_n : \{0, 1\}^{r_{\text{row}} + r_{\text{col}}} \rightarrow \{0, 1\}$.

Assume that $(\text{ans}_1, \dots, \text{ans}_q) \in \{0, 1, \perp\}^q$ are the answers to the queries. For every $i \in [p]$, we denote $pc_i := pc_i(\text{seed.row}, \text{seed.col})$. The PCPP verifier accepts the proof if

$$\text{VDec}(\text{ans}_1, \dots, \text{ans}_q, pc_1, \dots, pc_p) = 1.$$

The *decision complexity* of this PCPP verifier is said to be the circuit complexity of V_{dec} .

2.5.3 Smooth PCP of Proximity

Apart from rectangularity, another important property of the PCPPs is their *smoothness*. Roughly speaking, it means that the queries to the proof oracle are smooth, in the sense that each location is probed with equal probability.¹² The formal definition is as follows.

Definition 2.10 (Smooth PCPPs for Pure Languages). Let $r = r(n)$, $q = q(n)$ be good functions, $L \subseteq \{0, 1\}^*$ be a language, and $\text{VPCPP} = (V_{\text{type}}, V_{\text{index}}, V_{\text{dec}})$ be a PCPP verifier for L with randomness complexity r . It is said to be a *smooth PCPP verifier* if for all locations $\text{loc}_1, \text{loc}_2$ in the proof oracle, over a uniformly random $\text{seed} \in \{0, 1\}^r$ and a uniformly random index $j \in [q]$, loc_1 and loc_2 are probed by VPCPP with equal probability in the j -th query.

Smooth PCPPs can be viewed as PCPPs that can tolerate errors in the proof: since all the locations in the proof are queried with equal probability, a slightly corrupted version of a correct proof is still likely to be accepted, as shown in the following lemma.

Lemma 2.11. *Let $q = q(n)$, $\ell = \ell(n)$, $s = s(n)$ be good functions, $L \subseteq \{0, 1\}^*$ be a language, and VPCPP be a smooth PCPP verifier for L with soundness error s , proof length ℓ , and query complexity q . Assume that $x \in L \cap \{0, 1\}^n$ and $\pi \in \{0, 1\}^\ell$ is a correct proof for $x \in L$, i.e., $\text{VPCPP}^{x \circ \pi}(\text{seed})$ accepts with probability 1 over $\text{seed} \leftarrow \{0, 1\}^r$. Then for every π' such that the relative Hamming distance between π' and π is at most ε , $\text{VPCPP}^{x \circ \pi'}(\text{seed})$ accepts with probability at least $1 - q \cdot \varepsilon$ over $\text{seed} \sim \{0, 1\}^r$.*

Proof. We say a location $i \in [\ell]$ of the proof oracle is *bad* if $\pi[i] \neq \pi'[i]$. Let B_j be the event that the j -th query of VPCPP probes a bad location in the proof. By the smoothness, we know that

$$\Pr_{\text{seed} \in \{0, 1\}^r, j \in [q]} [B_j] \leq \varepsilon.$$

By a union bound, we can see that

$$\Pr_{\text{seed} \in \{0, 1\}^r} [\exists j \in [q], B_j] \leq \sum_{j \in [q]} \Pr_{\text{seed} \in \{0, 1\}^r} [B_j] \leq q \cdot \varepsilon. \quad (2)$$

Denote E be the event that there exists a $j \in [q]$ such that B_j happens. Then it follows that

$$\begin{aligned} & \Pr_{\text{seed} \in \{0, 1\}^r} [\text{VPCPP}^{x \circ \pi'}(\text{seed}) \text{ rejects}] \\ & \leq \Pr_{\text{seed} \in \{0, 1\}^r} [\text{VPCPP}^{x \circ \pi'}(\text{seed}) \text{ rejects} \mid \neg E] + \Pr_{\text{seed} \in \{0, 1\}^r} [E] \\ & \leq 0 + q \cdot \varepsilon \\ & = q \cdot \varepsilon, \end{aligned}$$

where the second inequality follows from Equation (2) and the perfect completeness of VPCPP . \square

¹²In some literature (e.g. [Par21]), the smoothness of PCPPs is defined differently: the queries to both the input oracle and the proof oracle need to be smooth, i.e., each location in the input (resp. the proof) is queried with equal probability. Here, we only require the queries to the proof oracle to be smooth and have no requirement on the query distribution over the input oracle.

Note that smoothness can be defined for rectangular PCPPs, in which case each location in the proof matrix is probed with equal probability. This further means that each row (resp. column) index is queried by the row (resp. column) verifier with equal probability.

Remark 2.12. A stronger definition of the smoothness is as follows: for every fixed $i \in [q]$, condition on the i -th query probing the proof oracle, the i -th query is uniformly random over the proof oracle. By randomly permuting the q queries, we can make a smooth PCPP satisfy this stronger definition of smoothness. In particular, if we have a smooth and rectangularity PCPP, we can make it satisfy this stronger definition of smoothness by adding $O(q \log q)$ bits in the shared randomness for a random permutation over the q queries.

2.5.4 Our Constructions

In this work, we will need two new constructions of rectangular PCPPs: for the range avoidance problem and worst-case hard partial truth tables, we need a rectangular PCPP with query complexity 3 or 2 (depending whether perfect completeness is required); for the remote point problem and average-case hard partial truth table, we need a smooth and rectangular PCPP with query complexity $O(1)$.¹³

Theorem 2.13 (3-Query and 2-Query Rectangular PCPPs). *For every constant $\delta \in (0, 1)$, there are constants $s_3 \in (0, 1)$ and $0 < s_2 < c_2 < 1$ such that the following holds. Let $m = m(n)$, $T(n)$, $w_{\text{proof}}(n)$, $w_{\text{input}}(n)$ be good functions such that $1 \leq m \leq (\log T(n))^{0.1}$, $n \leq T(n) \leq 2^{\text{poly}(n)}$, $w_{\text{proof}}(n) \leq \log T(n)$, and $w_{\text{input}}(n) \leq \log n$. Then there are good functions $h_{\text{proof}}^{3q}(n)$, $h_{\text{proof}}^{2q}(n)$, and $h_{\text{input}}(n)$ satisfying*

$$\begin{aligned} h_{\text{proof}}^{3q}(n), h_{\text{proof}}^{2q}(n) &= \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}(n), & \text{and} \\ h_{\text{input}}(n) &= \lceil \log n \rceil - w_{\text{input}}(n), \end{aligned}$$

such that the following holds.

Suppose that $w_{\text{proof}}, h_{\text{proof}}^{3q}, h_{\text{proof}}^{2q} \geq (5/m) \log T(n)$, and that for some absolute constant $C \geq 1$,

$$\frac{w_{\text{input}}(n)}{w_{\text{proof}}(n)}, \frac{h_{\text{input}}(n)}{h_{\text{proof}}^{3q}(n)}, \frac{h_{\text{input}}(n)}{h_{\text{proof}}^{2q}(n)} \leq 1 - \frac{Cm \log \log T(n)}{\log T(n)}.$$

Let $W_{\text{proof}}(n) := 2^{w_{\text{proof}}(n)}$, $H_{\text{proof}}^{3q}(n) := 2^{h_{\text{proof}}^{3q}(n)}$, $H_{\text{proof}}^{2q}(n) := 2^{h_{\text{proof}}^{2q}(n)}$, $W_{\text{input}}(n) := 2^{w_{\text{input}}(n)}$, and $H_{\text{input}}(n) := 2^{h_{\text{input}}(n)}$. Then $\text{NTIME}[T(n)]$ has:

- a rectangular PCP of proximity V_{3q} with perfect completeness, soundness error s_3 , an $H_{\text{proof}}^{3q}(n) \times W_{\text{proof}}(n)$ proof matrix and an $H_{\text{input}}(n) \times W_{\text{input}}(n)$ input matrix;
- a rectangular PCP of proximity V_{2q} with completeness error $1 - c_2$, soundness error s_2 , an $H_{\text{proof}}^{2q}(n) \times W_{\text{proof}}(n)$ proof matrix and an $H_{\text{input}}(n) \times W_{\text{input}}(n)$ input matrix.

Other parameters of V_{3q} and V_{2q} are specified in Table 1.

Furthermore, given the randomness $\text{seed} \in \{0, 1\}^r$, the total number of queries and parity-check bits is at most 3 for V^{3q} and 2 for V^{2q} , and the decision predicate $\text{VDec} \leftarrow \text{V}_{\text{dec}}(\text{seed}, \text{shared})$ of the rectangular PCPP verifier is an OR of the input bits (including queries and parity-check bits) or their negations for every $\text{seed}, \text{shared}$.

¹³Unfortunately, our smooth PCPP requires a large (although constant) number of queries, because of the arguments in Section 6.1.

PCPP Verifier	V^{3q}	V^{2q}
Completeness error	0	$1 - c_2$
Soundness error	s_3	s_2
Proximity parameter	δ	
Row randomness	$h_{\text{proof}}^{3q} - (5/m) \log T(n)$	$h_{\text{proof}}^{2q} - (5/m) \log T(n)$
Column randomness	$w_{\text{proof}} - (5/m) \log T(n)$	
Shared randomness	$(10/m) \log T(n) + O(\log \log T(n) + m \log m)$	
Query complexity	3	2
Parity check complexity		
Decision complexity	$\text{poly}(\log \log T)$	

Table 1: Parameters of the PCPPs constructed in Theorem 2.13.

Theorem 2.14 (Smooth and Rectangular PCPP). *For all constants $\delta \in (0, 1)$ and $s \in (0, 1)$, there is a constant $q \geq 1$ such that the following holds. Let $m = m(n)$, $T(n)$, $w_{\text{proof}}(n)$, $w_{\text{input}}(n)$ be good functions such that $1 \leq m(n) \leq (\log T(n))^{0.1}$, $n \leq T(n) \leq 2^{\text{poly}(n)}$, $w_{\text{proof}}(n) \leq \log T(n)$, and $w_{\text{input}}(n) \leq \log n$. Then there are good functions $h_{\text{proof}}(n)$ and $h_{\text{input}}(n)$ satisfying*

$$h_{\text{proof}}(n) := \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}(n), \quad \text{and}$$

$$h_{\text{input}}(n) := \lceil \log n \rceil - w_{\text{input}}(n).$$

such that the following holds.

Suppose that $w_{\text{proof}}, h_{\text{proof}} \geq (5/m) \log T(n)$, and that for some absolute constant $C \geq 1$,

$$\frac{w_{\text{input}}(n)}{w_{\text{proof}}(n)}, \frac{h_{\text{input}}(n)}{h_{\text{proof}}(n)} \leq 1 - \frac{Cm^2 \log \log T(n)}{\log T(n)}.$$

Let $W_{\text{proof}}(n) := 2^{w_{\text{proof}}(n)}$, $H_{\text{proof}}(n) := 2^{h_{\text{proof}}(n)}$, $W_{\text{input}}(n) := 2^{w_{\text{input}}(n)}$, and $H_{\text{input}}(n) := 2^{h_{\text{input}}(n)}$. Then $\text{NTIME}[T(n)]$ has a smooth and rectangular PCP of proximity with an $H_{\text{input}}(n) \times W_{\text{input}}(n)$ input matrix and an $H_{\text{proof}}(n) \times W_{\text{proof}}(n)$ proof matrix, whose other parameters are specified in Table 2.

Soundness error	s
Proximity parameter	δ
Row randomness	$r_{\text{row}} := h_{\text{proof}} - (5/m) \log T(n)$
Column randomness	$r_{\text{col}} := w_{\text{proof}} - (5/m) \log T(n)$
Shared randomness	$r_{\text{shared}} := (10/m) \log T(n) + O(\log \log T(n) + m \log m)$
Query complexity	$q = O_{s,\delta}(1)$
Parity check complexity	
Decision complexity	$\text{poly}(T(n)^{1/m})$

Table 2: Parameters of the PCPP constructed in Theorem 2.14.

Details of these two constructions are postponed to Section 6 and Section 7.

2.6 A Stretch Reduction for REMOTE-POINT and PARTIAL-AVGHARD

In our framework for solving REMOTE-POINT (Section 3.2), for technical convenience, we only consider circuits $C : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell(n)}$, where $\ell(n)$ is a certain stretch function. (For example, it

might be the case that $\ell(n)$ is rounded to a power of 2 for every n .) In this subsection, we show that such an algorithm can also solve REMOTE-POINT for circuits of larger stretches. This justifies that it is without loss of generality to only consider stretch functions that are *equal* to $\ell(n)$.

Lemma 2.15 (Stretch Reduction for REMOTE-POINT). *Let \mathcal{C} be a typical circuit class and s be a size parameter. Suppose that $\mathcal{C}[s]$ -REMOTE-POINT with stretch $\ell'(n)$ and distance parameter $1/2 - \varepsilon'(n)$ admits an FP^{NP} algorithm. Then for any stretch $\ell = \ell(n) \geq \ell'(n+1)/2$, $\mathcal{C}[s]$ -REMOTE-POINT with stretch $\ell(n)$ and distance parameter $1/2 - \varepsilon(n)$ also admits an FP^{NP} algorithm, where $\varepsilon(n) := 2 \cdot \varepsilon'(n+1)$.*

Proof. Denote $\ell' := \ell'(n+1)$, $\varepsilon' := \varepsilon'(n+1)$, $\ell := \ell(n)$, and $\varepsilon := \varepsilon(n)$, and let $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ be an input circuit. If ℓ is a multiple of ℓ' , we can split the ℓ -bit output of C into blocks of size ℓ' and add a dummy input bit to construct $m := \ell/\ell'$ circuits $C_1, C_2, \dots, C_m : \{0, 1\}^{n+1} \rightarrow \{0, 1\}^{\ell'(n+1)}$ such that for every $x \in \{0, 1\}^n$ and $b \in \{0, 1\}$,

$$C(x) = C_1(x, b) \circ C_2(x, b) \circ \dots \circ C_m(x, b).$$

Using the FP^{NP} algorithm for \mathcal{C} -REMOTE-POINT with stretch $\ell'(n+1)$ and error parameter $\varepsilon'(n+1)$, we can construct $y_1, y_2, \dots, y_m \in \{0, 1\}^{\ell'}$ such that each y_i is $(1/2 - \varepsilon')$ -far from $\text{Range}(C_i)$. It then follows that the concatenation $y_1 \circ y_2 \circ \dots \circ y_m$ is $(1/2 - \varepsilon')$ -far from $\text{Range}(C)$.

We now consider the case where ℓ is not a multiple of ℓ' . Let $I : \{0, 1\}^{n+1} \rightarrow \{0, 1\}$ be defined as the projection $I(x) = x_{n+1}$, that is, it always outputs the last bit. For any t , let $I^{\otimes t} : \{0, 1\}^{n+1} \rightarrow \{0, 1\}^t$ denote the concatenation of t copies of I . Therefore, $\text{Range}(I^{\otimes t}) = \{0^t, 1^t\}$. Since \mathcal{C} is typical, we have $I^{\otimes t} \in \mathcal{C}$.

Let $M = k \cdot \ell'$ be the nearest multiple of ℓ' larger than ℓ , and $\bar{\ell} := M - \ell$. For a multi-output \mathcal{C} circuit C , we define $\tilde{C} : \{0, 1\}^{n+1} \rightarrow \{0, 1\}^M$ as

$$\tilde{C}(x, b) = C(x) \circ I^{\otimes \bar{\ell}}(x, b),$$

where $x \in \{0, 1\}^n$ and $b \in \{0, 1\}$. Since \tilde{C} is of input length $n+1$ and output length being a multiple of ℓ' , we can get a remote point $s \in \{0, 1\}^M$ in FP^{NP} that is $(1/2 - \varepsilon')$ -far from $\text{Range}(\tilde{C})$.

Let $s = s_1 \circ s_2$, where s_1 and s_2 has length ℓ and $\bar{\ell}$, respectively. We then prove that s_1 is $(1/2 - \varepsilon)$ far from $\text{Range}(C)$. Towards a contradiction, we assume that s_1 is not $(1/2 - \varepsilon)$ -far from $\text{Range}(C)$. In other words, there is an $x \in \{0, 1\}^n$ such that $\delta(C(x), s_1) < 1/2 - \varepsilon$. By considering the Hamming weight of s_2 we know that there is a $b \in \{0, 1\}$ such that $\delta(I^{\otimes \bar{\ell}}(x, b), s_2) \leq 1/2$. It then follows that

$$\begin{aligned} \delta(s, \tilde{C}(x, b)) &= \delta(s_1 \circ s_2, C(x) \circ I^{\otimes \bar{\ell}}(x, b)) \\ &\leq \frac{\ell}{\ell + \bar{\ell}} \cdot \left(\frac{1}{2} - \varepsilon(n) \right) + \frac{\bar{\ell}}{\ell + \bar{\ell}} \cdot \frac{1}{2} \\ &\leq \frac{1}{2} - \frac{\ell}{\ell + \bar{\ell}} \cdot \varepsilon(n) \\ &< \frac{1}{2} - \varepsilon'(n+1). \end{aligned}$$

This leads to a contradiction as s is $(1/2 - \varepsilon'(n+1))$ -far from $\text{Range}(\tilde{C})$. □

Similar to REMOTE-POINT, another average-case problem PARTIAL-AVGHARD can also be reduced to the instances with smaller stretch in the same way.

Lemma 2.16 (Stretch Reduction for PARTIAL-AVGHARD). *Let \mathcal{C} be a typical circuit class and s be a size parameter. Suppose that $\text{NC}_2^0 \circ (\mathcal{C}[s])\text{-PARTIAL-AVGHARD}$ with stretch $\ell'(n)$ and the distance parameter $1/2 - \varepsilon'(n)$ admits an FP^{NP} algorithm, then for any stretch $\ell = \ell(n) \geq \ell'(n+1)/2$, $\mathcal{C}[s]\text{-PARTIAL-AVGHARD}$ with stretch $\ell(n)$ and distance parameter $1/2 - \varepsilon(n)$ admits an FP^{NP} algorithm, where $\varepsilon(n) := 2 \cdot \varepsilon'(n+1)$.*

Proof Sketch. The proof of this lemma is similar to that of Lemma 2.15. Here we use the same notation as the proof of Lemma 2.15.

Let $X = \{x_1, \dots, x_\ell\}$ denote input strings, and let $y_i := x_i \circ 0$. We create $\bar{\ell}$ copies of $0^n \circ 1$ and use $y_{\ell+1}, \dots, y_M$ to denote these copies.

we solve $\text{NC}_2^0 \circ \mathcal{C}\text{-PARTIAL-AVGHARD}$ on $\{y_1, \dots, y_M\}$ and get an average-case hard partial truth table $s = s_1 \circ s_2$ that is $(1/2 - \varepsilon'(n+1))$ -far from any truth table of $\text{NC}_2^0 \circ \mathcal{C}$ circuit, where s_1 and s_2 has length ℓ and $\bar{\ell}$. Then we prove s_1 is a solution for the original problem. For some \mathcal{C} circuit C , if s_1 is not $(1/2 - \varepsilon(n))$ far from partial truth table of C on X , we can define $\tilde{C}_1, \tilde{C}_2 : \{0, 1\}^n \times \{0, 1\} \rightarrow \{0, 1\}$ as $\tilde{C}_1(x; b) := C(x) \vee b$, $\tilde{C}_2(x; b) := C(x) \wedge (\neg b)$. Then one of \tilde{C}_1 and \tilde{C}_2 has partial truth table on $Y := \{y_1, \dots, y_M\}$ not $(1/2 - \varepsilon'(n+1))$ far from s , which leads to a contradiction. Therefore s_1 has to be a solution. The analysis is similar to Lemma 2.15. \square

2.7 Satisfying Pairs for $\text{NC}_d^0 \circ \mathcal{C}$ from Satisfying Pairs for $(\text{AND}_d^0/\text{XOR}_d^0/\text{OR}_d^0) \circ \mathcal{C}$

We show that satisfying pairs for $\text{NC}_d^0 \circ \mathcal{C}$ circuits can be reduced to the satisfying pairs of $\text{AND}_d^0 \circ \mathcal{C}$, $\text{XOR}_d^0 \circ \mathcal{C}$, or $\text{OR}_d^0 \circ \mathcal{C}$ via standard Fourier analysis (see, e.g., [CW19b, Section 4]). This will be beneficial for the unconditional results for weak circuit classes, such as the remote point algorithm for GF(2)-linear functions.

Theorem 2.17. *For every constants $\delta \in [0, 1]$ and $d \geq 1$, there is a constant δ' such that the following holds. Let $N = N(n), M = M(n), n, s = s(n)$ be parameters, and $C_d \in \{\text{AND}_d, \text{OR}_d, \text{XOR}_d\}$.*

Then $\#(\text{NC}_d^0 \circ \mathcal{C})\text{-SATISFYING-PAIRS}$ (resp. $\text{Approx}_\delta\text{-}(\text{NC}_d^0 \circ \mathcal{C})\text{-SATISFYING-PAIRS}$) with parameters (N, M, n, s) is $\tilde{O}(n)$ -time Turing-reducible to $\#(C_d \circ \mathcal{C})\text{-SATISFYING-PAIRS}$ (resp. $\text{Approx}_{\delta'}\text{-}(C_d \circ \mathcal{C})\text{-SATISFYING-PAIRS}$) with parameters $(\Theta(N), M, n, s)$, as long as each input circuit of the $\#(\text{NC}_d^0 \circ \mathcal{C})\text{-SATISFYING-PAIRS}$ (resp. $\text{Approx}_\delta\text{-}(\text{NC}_d^0 \circ \mathcal{C})\text{-SATISFYING-PAIRS}$) problem are given explicitly as a top NC_d^0 circuit C_{top} together with d circuits $C_1, C_2, \dots, C_d \in \mathcal{C}$ feeding C_{top} .

Moreover, the oracle algorithm for $\#(\text{NC}_d^0 \circ \mathcal{C})\text{-SATISFYING-PAIRS}$ (resp. $\text{Approx}_\delta\text{-}(\text{NC}_d^0 \circ \mathcal{C})\text{-SATISFYING-PAIRS}$) only makes $O(1)$ non-adaptive queries to the $\#(C_d \circ \mathcal{C})\text{-SATISFYING-PAIRS}$ (resp. $\text{Approx}_{\delta'}\text{-}(C_d \circ \mathcal{C})\text{-SATISFYING-PAIRS}$) oracle.

Proof. Let N, M, n, s be the parameters. Suppose that we are given $C_1, C_2, \dots, C_N \in \mathcal{C}[s]$ and $x_1, x_2, \dots, x_M \in \{0, 1\}^n$ as input. We assume that C_1, C_2, \dots, C_N share the same upper NC_d^0 function computing $f : \{0, 1\}^d \rightarrow \{0, 1\}$, that is for every $i \in [N]$, $C_i \equiv f \circ D_i$ for some d -output \mathcal{C} circuit D_i of size at most s . This is without loss of generality since there are at most $2^{2^d} = O(1)$ different NC_d^0 functions and we can (approximately) count the number of satisfying pairs for each of these cases separately.

We first consider the case for $C_d = \text{AND}_d$. We use the basis $\{0, 1\} \subseteq \mathbb{R}$ for Boolean values and write f as

$$f(x) = \sum_{S \subseteq [d]} \alpha_S \cdot \prod_{i \in S} x_i,$$

where each coefficient $\alpha_S \in [-2^d, 2^d] \cap \mathbb{Z}$. Note that we can compute the coefficients by writing the truth table of f in the canonical disjunctive normal form, represent x by x , $\neg x$ by $1 - x$, \wedge by

multiplication, and (disjoint) \vee by addition, and then expending the multi-linear polynomial using a brute-force algorithm in $O(1)$ time.

Let $\chi_S(x) := \prod_{i \in S} x_i$ for $S \subseteq [d]$. Then the number of $(i, j) \in [N] \times [M]$ such that $C_i(x_j) = 1$ is

$$\begin{aligned}
& \sum_{i \in [N]} \sum_{j \in [M]} f(D_i(x_j)) \\
&= \sum_{i \in [N]} \sum_{j \in [M]} \sum_{S \subseteq [d]} \alpha_S \cdot \chi_S(D_i(x_j)) \\
&= \sum_{S \subseteq [d]} \alpha_S \cdot \left[\sum_{i \in [N]} \sum_{j \in [M]} \chi_S(D_i(x_j)) \right] \\
&= \sum_{S \subseteq [d]} \alpha_S \cdot \left[\sum_{i \in [N]} \sum_{j \in [M]} \text{AND}_{|S|} \circ D_i|_S(x_j) \right],
\end{aligned}$$

where $D_i|_S : \{0, 1\}^{|S|} \rightarrow \{0, 1\}$ representing the circuit obtained from D_i by restricting to the output bits in S . Then our algorithm is as follows: We enumerate all $S \subseteq [d]$ and count (resp. approximately count) the number A_S of satisfying pairs for circuits $\text{AND}_{|S|} \circ D_1|_S, \dots, \text{AND}_{|S|} \circ D_N|_S$ and inputs x_1, \dots, x_M , then we output the answer $\sum_{S \subseteq [d]} \alpha_S \cdot A_S$.

For $C_d = \text{XOR}_d$ and $C_d = \text{OR}_d$, we only need to write f as

$$f(x) = \sum_{S \subseteq [d]} \alpha'_S \cdot \bigoplus_{i \in S} x_i, \quad (3)$$

$$f(x) = \sum_{S \subseteq [d]} \alpha''_S \cdot \bigvee_{i \in S} x_i, \quad (4)$$

where $\alpha'_S, \alpha''_S \leq 2^{O(d)}$. Note that Equation (3) can be obtained using the basis $\{\text{true} := -1, \text{false} := 1\}$ and Equation (4) can be obtained using the basis $\{\text{true} := 0, \text{false} := 1\}$. \square

3 Range Avoidance and Remote Point

In this section, we prove our main technical result: non-trivial algorithms for SATISFYING-PAIRS imply FP^{NP} algorithms for range avoidance and remote point. We first prove our results for range avoidance and remote point in Section 3.1 and Section 3.2, respectively, and then briefly discuss the variants of our algorithms in Section 3.3.

The naïve algorithm for SATISFYING-PAIRS is to evaluate every circuit on every input, which requires $O(NM \cdot \text{poly}(s, n))$ time. We will employ non-trivial algorithms (i.e. of time complexity $NM / \log^{\omega(1)}(NM)$) for SATISFYING-PAIRS to solve the range avoidance problem and the remote point problem. Indeed, we can even allow the Satisfying Pairs algorithm to have a *preprocessing phase*, in which a polynomial-time algorithm with access to an NP oracle is given the circuits C_1, \dots, C_N (but not the inputs) and produce a data structure of small (i.e. “fixed polynomial”) size.

Definition 3.1 (Algorithms for SATISFYING-PAIRS with P^{NP} Preprocessing on Circuits). Let P be one of the problems \mathcal{C} -SATISFYING-PAIRS, $\#\mathcal{C}$ -SATISFYING-PAIRS, Approx_δ - \mathcal{C} -SATISFYING-PAIRS, Gap_δ - \mathcal{C} -SATISFYING-PAIRS. A t -time algorithm for P with P^{NP} preprocessing of an ℓ -size data structure on circuits is a pair of algorithms (A_1, A_2) that solves P in two phases:

1. Given the circuits $C_1, C_2, \dots, C_N : \{0, 1\}^n \rightarrow \{0, 1\}$ of size s , the polynomial-time algorithm A_1 with oracle access to a SAT oracle computes a string $DS \in \{0, 1\}^\ell$.
2. Given the inputs $x_1, x_2, \dots, x_M \in \{0, 1\}^n$ and the string $DS \in \{0, 1\}^\ell$, the algorithm A_2 solves P on the instance $(C_1, \dots, C_N, x_1, \dots, x_M)$ in time t .

3.1 Range Avoidance from SATISFYING-PAIRS

In this sub-section, we establish the connection between the AVOID and SATISFYING-PAIRS. The main result is the following theorem.

Theorem 3.2. *There are constants $\varepsilon > 0$ and c_{imp} such that the following holds. Let $0 < \eta < 1/2$ be a constant, $\ell(n) > n^{1+4\eta}$ be a good function. Let $\mathcal{C}[s]$ be a typical circuit class where $s = s(n)$ is a size parameter, and $\mathcal{C}'[2s] := \text{OR}_2 \circ \mathcal{C}[s]$ (i.e. a \mathcal{C}' circuit of size $2s$ refers to the OR of at most two \mathcal{C} circuits of size s).*

Assumption: *Suppose that for some constant $c \geq 1$, there is an $(NM/\log^{c_{\text{imp}}}(NM))$ -time algorithm for $\text{Approx}_\varepsilon$ - \mathcal{C}' -SATISFYING-PAIRS with $N := \ell^{1-\eta} \cdot \text{polylog}(\ell)$ circuits of size $2s(n)$ and $M := \ell^{c+1-\eta} \cdot \text{polylog}(\ell)$ inputs of length $n \cdot \text{polylog}(\ell)$, allowing a P^{NP} preprocessing of an N^c -size data structure on circuits.*

Conclusion: *Then there is an FP^{NP} algorithm for $\mathcal{C}[s]$ -AVOID with stretch $\ell(n)$.*

3.1.1 Typical Choices of the Parameters

Before proving this general framework, we demonstrate two typical choices of the parameters as follows that deal with AVOID with polynomial stretch and quasi-polynomial stretch.

Corollary 1.11. *There is a constant $\varepsilon > 0$ such that the following holds. Let \mathcal{C} be a typical circuit class, $\mathcal{C}' := \text{OR}_2 \circ \mathcal{C}$, and $s = s(n)$ be a non-decreasing size parameter.*

- *Suppose that there is a non-trivial algorithm for $\text{Approx}_\varepsilon$ - \mathcal{C}' -SATISFYING-PAIRS for $N = n^{1+\Omega(1)}$ \mathcal{C}' -circuits of size $2s(n)$ and $M = n^{1+\Omega(1)}$ inputs of length n . Then there is an FP^{NP} algorithm for \mathcal{C} -AVOID with stretch ℓ and circuit size s ,¹⁴ for some $\ell = n^{1+\Omega(1)}$.*
- *Suppose that there is a non-trivial algorithm for $\text{Approx}_\varepsilon$ - \mathcal{C}' -SATISFYING-PAIRS for $N = \text{quasi-poly}(n)$ \mathcal{C}' -circuits of size $2s(n)$ and $M = \text{quasi-poly}(n)$ inputs of length n . Then there is an FP^{NP} algorithm for \mathcal{C} -AVOID with stretch ℓ and circuit size s , for some $\ell = \text{quasi-poly}(n)$.*

Proof. (Polynomial Case). Assume we have an algorithm for $\text{Approx}_\varepsilon$ - \mathcal{C}' -SATISFYING-PAIRS for $N_{\text{al}} = n_{\text{al}}^{1+c_n}$ \mathcal{C}' circuits of size $2s_{\text{al}}(n_{\text{al}})$ and $M_{\text{al}} = n_{\text{al}}^{1+c_m}$ inputs of length n_{al} that runs in non-trivial time (i.e. $N_{\text{al}} \cdot M_{\text{al}}/\log^{\omega(1)}(N_{\text{al}} \cdot M_{\text{al}})$), where c_n and c_m are some constants. Denote this algorithm by $(*)$. Note that $s(n) = \text{poly}(n)$ (otherwise there cannot be such an algorithm). We apply Theorem 3.2 with $\eta = 0.1$ and $s := s_{\text{al}}$. Let $c = 1$ be the constant in Theorem 3.2 and $\ell(n) = \text{poly}(n)$ to be determined later.

To obtain an FP^{NP} algorithm for $\mathcal{C}[s]$ -AVOID with stretch $\ell(n)$, we need to design an algorithm for \mathcal{C}' -SATISFYING-PAIRS with $N = \ell^{0.9} \cdot \text{polylog}(\ell)$ circuits of size $s(n_{\text{fw}})$ and $M = \ell^{1.9} \cdot \text{polylog}(\ell)$ inputs of length $n_{\text{fw}} = n \cdot \text{polylog}(\ell)$ that runs in time $NM/\log^{c_{\text{imp}}}(NM)$ for some absolute constant c_{imp} . Denote this problem by (\star) . We set $n_{\text{al}} = n_{\text{fw}}$ and $\ell(n) = \max\{(N_{\text{al}}s)^{1.1/0.9}, (M_{\text{al}}s)^{1.1/1.9}\} = n^{1+\Omega(1)}$, so that $N \geq N_{\text{al}}^{1.1} \cdot s^{1.1}$ and $M \geq M_{\text{al}}^{1.1} \cdot s^{1.1}$. The algorithm for (\star) works as follows.

¹⁴Note that the circuit size parameter of \mathcal{C} -AVOID refers to the maximum circuit size of each output bit of $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, instead of the total circuit size of C .

- We divide the N circuits into groups of size N_{al} and the M inputs into groups of size M_{al} .
- For every $i \in [N/N_{\text{al}}]$ and $j \in [M/M_{\text{al}}]$, we use the algorithm (*) to (approximately) count the number of satisfying pairs between the i -th circuit group and the j -th input group. (In particular, we use a brute-force algorithm with running time $(N_{\text{al}} \cdot M + M_{\text{al}} \cdot N) \cdot \tilde{O}(s) \leq NM / \log^{\omega(1)}(NM)$ to deal with the groups, if any, with less than N_{al} circuits or less than M_{al} inputs.) We add all the answers among groups together.

The correctness of the algorithm is obvious, and the running time is at most

$$\frac{N}{N_{\text{al}}} \cdot \frac{M}{M_{\text{al}}} \cdot \frac{N_{\text{al}} M_{\text{al}}}{\log^{\omega(1)}(N_{\text{al}} M_{\text{al}})} \leq \frac{NM}{\log^{c_{\text{imp}}}(NM)}.$$

(*Quasi-Polynomial Case.*) The proof is almost the same as the polynomial case. In particular, $s(n)$ and $\ell(n)$ as defined above are quasi-polynomial functions; the final running time is non-trivial since $\log(N_{\text{al}} M_{\text{al}}) = \log^{\Omega(1)}(n)$ and $\log(NM) = \log^{O(1)}(n)$. The details are omitted. \square

3.1.2 Proof of Theorem 3.2

Theorem 3.2. *There are constants $\varepsilon > 0$ and c_{imp} such that the following holds. Let $0 < \eta < 1/2$ be a constant, $\ell(n) > n^{1+4\eta}$ be a good function. Let $\mathcal{C}[s]$ be a typical circuit class where $s = s(n)$ is a size parameter, and $\mathcal{C}'[2s] := \text{OR}_2 \circ \mathcal{C}[s]$ (i.e. a \mathcal{C}' circuit of size $2s$ refers to the OR of at most two \mathcal{C} circuits of size s).*

Assumption: *Suppose that for some constant $c \geq 1$, there is an $(NM / \log^{c_{\text{imp}}}(NM))$ -time algorithm for $\text{Approx}_{\varepsilon}\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ with $N := \ell^{1-\eta} \cdot \text{polylog}(\ell)$ circuits of size $2s(n)$ and $M := \ell^{c+1-\eta} \cdot \text{polylog}(\ell)$ inputs of length $n \cdot \text{polylog}(\ell)$, allowing a P^{NP} preprocessing of an N^c -size data structure on circuits.*

Conclusion: *Then there is an FP^{NP} algorithm for $\mathcal{C}[s]\text{-AVOID}$ with stretch $\ell(n)$.*

Proof. Suppose that we are given a \mathcal{C} circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^{\ell}$. Without loss of generality, we may assume ℓ is a power of 2 and $c \geq 2$. We set the following parameters:

$$\begin{aligned} m &:= 5(c+2)/\eta = O(1), \\ w_{\text{proof}} &:= \log \ell, & W_{\text{proof}} &:= 2^{w_{\text{proof}}} = \ell, \\ h_{\text{proof}} &:= (c+1) \log \ell, & H_{\text{proof}} &:= 2^{h_{\text{proof}}} = \ell^{c+1}, \\ n_{\text{hard}} &:= 10H_{\text{proof}} \cdot n, \\ T &:= H_{\text{proof}} \cdot W_{\text{proof}} / \log^{c_{\text{tm}}}(\ell). \end{aligned}$$

The constants ε , c_{imp} , and c_{tm} will be determined later.

Let L^{hard} be the hard language constructed in Theorem 2.2. We use n_{hard} and T to denote the input length and the time complexity of L^{hard} , respectively, i.e.

$$L^{\text{hard}} \in \text{NTIME}_{\text{TM}}[T] \setminus \text{i.o.-NTIME}_{\text{GUESS}_{\text{RTM}}}[T / \log^{c_{\text{hard}}}(T), n_{\text{hard}}/10]_{(n_{\text{hard}}/10)},$$

where c_{hard} is some large universal constant. Note that since $T = \ell^{c+2} / \text{polylog}(\ell)$, $n_{\text{hard}} = O(\ell^{c+1} \cdot n)$, $\ell > n^{1+4\eta}$, we can see that $n_{\text{hard}}^{1+\Omega(1)} \leq T \leq n_{\text{hard}}^2$, which satisfies the technical condition of Theorem 2.2.

We describe a nondeterministic RAM M^{PCPP} that runs in $T / \log^{c_{\text{hard}}}(T)$ time, uses $n_{\text{hard}}/10$ advice bits, guesses $n_{\text{hard}}/10$ nondeterministic bits, and attempts to solve L^{hard} on n_{hard} -bit inputs.

By the definition of L^{hard} , M^{PCPP} has to fail on some input $x \in \{0, 1\}^{n_{\text{hard}}}$ when n_{hard} is sufficiently large. Our goal is to design such an algorithm M^{PCPP} that (1) rejects every $x \notin L^{\text{hard}}$, and (2) accepts every $x \in L^{\text{hard}}$ *with an easy witness*. Thus, if M^{PCPP} fails on some input x , then $x \in L$ and it has only “hard witnesses”, which will be exploited for finding a non-output of C .

Here, to define the inputs x “with an easy witness”, we will need the 2-query rectangular PCPP in Theorem 2.13 for the following language

$$L^{\text{enc}} := \{\text{Enc}(x) : x \in L^{\text{hard}}\},$$

where we fix an error-correcting code (Enc, Dec) as in Theorem 2.1. Let δ_{Enc} be the distance of the code. Suppose a string of length n_{hard} is encoded (via Enc) into a string of length $\tilde{n}_{\text{hard}} := O(n_{\text{hard}})$. We set the following parameters:

$$\begin{aligned} h_{\text{input}} &:= \left(1 - \frac{\Theta(\log \log T)}{\log T}\right) h_{\text{proof}}, & H_{\text{input}} &:= 2^{h_{\text{input}}} = H_{\text{proof}} / \text{polylog}(\ell), \\ w_{\text{input}} &:= \lceil \log \tilde{n}_{\text{hard}} \rceil - h_{\text{input}}, & W_{\text{input}} &:= 2^{w_{\text{input}}} = n \cdot \text{polylog}(\ell). \end{aligned}$$

We assume without loss of generality that $\tilde{n}_{\text{hard}} = H_{\text{input}} \cdot W_{\text{input}}$. (This can always be done by adding at most $W_{\text{input}} \leq \tilde{n}_{\text{hard}}$ dummy bits into the codeword of the error-correcting code, where the resulting code is still of constant rate and distance.)

We can check that the technical conditions of Theorem 2.13 for the 2-query rectangular PCPP construction holds:¹⁵

$$\begin{aligned} \log T(n) &= h_{\text{proof}} + w_{\text{proof}} - c_{\text{tm}} \log \log \ell = (c + 2) \log \ell - c_{\text{tm}} \log \log \ell \\ h_{\text{proof}} &= \log T(n) + c_{\text{tm}} \log \log \ell - w_{\text{proof}} = \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}} \\ \frac{h_{\text{input}}}{h_{\text{proof}}} &= 1 - \frac{\Theta(\log \log T)}{\log T} \leq 1 - \frac{Cm \log \log T}{\log T} \quad (\text{if the constant in } \Theta(\cdot) \text{ is large enough}) \\ \frac{w_{\text{input}}}{w_{\text{proof}}} &= \frac{\lceil \log \tilde{n}_{\text{hard}} \rceil - h_{\text{input}}}{\log \ell} \leq \frac{h_{\text{proof}} + \log n + O(1) - h_{\text{input}}}{\log \ell} \\ &= \frac{\log n + \Theta(h_{\text{proof}} \log \log T(n)) / \log T(n)}{\log \ell} \\ &= \frac{\log n}{\log \ell} + \frac{O(\log \log T(n))}{\log T(n)} \\ &\leq 1 - \Omega(1). \end{aligned}$$

By Theorem 2.13, there is a PCPP verifier VPCPP for L^{enc} with oracle access to $\Pi := \text{Enc}(x) \circ \pi$, where the input $\text{Enc}(x)$ is treated as a matrix of size $H_{\text{input}} \times W_{\text{input}}$, and the proof π is treated as a matrix of size $H_{\text{proof}} \times W_{\text{proof}}$. The PCPP verifier has the following parameters:

- completeness error = $1 - c_{\text{pcp}}$,
- soundness error = s_{pcp} ,
- proximity parameter = $\delta_{\text{Enc}}/3$,
- query complexity ≤ 2 ,

¹⁵Strictly speaking, we apply Theorem 2.13 with w_{proof} to obtain a PCPP with proof size $H'_{\text{proof}} \times W_{\text{proof}}$, where $W_{\text{proof}} = 2^{w_{\text{proof}}}$ and $H'_{\text{proof}} = 2^{h'_{\text{proof}}}$ for some $h'_{\text{proof}} = \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}$. We set c_{tm} to be sufficiently large so that $h_{\text{proof}} \geq h'_{\text{proof}}$, hence we can assume without loss of generality that the proof size is actually $H_{\text{proof}} \times W_{\text{proof}}$ (recall that $H_{\text{proof}} = 2^{h_{\text{proof}}}$) by adding dummy bits in the proof.

- parity-check bits ≤ 2 ,
- total randomness $= r := \log T + O(\log \log T + m \log m)$,
- row randomness $= r_{\text{row}} := h_{\text{proof}} - (5/m) \log T = (c + 1 - \eta) \log \ell - (5c_{\text{tm}}/m) \log \log \ell$,
- column randomness $= r_{\text{col}} := w_{\text{proof}} - (5/m) \log T = (1 - \eta) \log \ell - (5c_{\text{tm}}/m) \log \log \ell$,
- shared randomness $= r_{\text{shared}} := (10/m) \log T + c_{\text{tm}} \log \log T$
 $= 2\eta \log \ell + (c_{\text{tm}}/m) \log \log \ell + O(\log \log \ell + m \log m)$.

Moreover, the total number of parity-check bits and queries is at most 2, and the decision predicate $V_{\text{Dec}} \leftarrow V_{\text{dec}}(\text{seed.shared})$, which takes the parity-check bits and the answers to the queries as the input, is an OR of its input bits or their negations.

For an input $x \in L^{\text{hard}} \cap \{0, 1\}^{n_{\text{hard}}}$, we say that x has an easy witness if there is a proof matrix π for the statement “ $\text{Enc}(x) \in L^{\text{enc}}$ ” such that:

(completeness) $\Pr_{\text{seed} \leftarrow \{0, 1\}^r} [\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) \text{ accepts}] \geq c_{\text{pcp}}$; and

(easiness) for every row π_i of π , there exists a string w_i such that $\pi_i = C(w_i)$.

Description of M^{PCPP} . Now we define M^{PCPP} , which is a $(T/\log^{c_{\text{hard}}} T)$ -time non-deterministic algorithm that takes at most $\ell^{c+1} \leq n_{\text{hard}}/10$ bits of advice. M^{PCPP} aims to reject every $x \notin L$ and accept every $x \in L$ with easy witness when appropriate advice is given.

On input $x \in \{0, 1\}^{n_{\text{hard}}}$, we guess H_{proof} strings $w_1, w_2, \dots, w_{H_{\text{proof}}} \in \{0, 1\}^n$. Let π be the $H_{\text{proof}} \times W_{\text{proof}}$ proof matrix where for each $i \in [H_{\text{proof}}]$, the i -th row of π is equal to $C(w_i)$. Let p_{acc} be the acceptance probability of the PCPP verifier VPCPP for L^{enc} given the input $\text{Enc}(x)$ and the proof π , i.e.,

$$p_{\text{acc}} := \Pr_{\text{seed} \leftarrow \{0, 1\}^r} [\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) \text{ accepts}].$$

We need to distinguish between the case that $p_{\text{acc}} \geq c_{\text{pcp}}$ and the case that $p_{\text{acc}} \leq s_{\text{pcp}}$. We set $\varepsilon := (c_{\text{pcp}} - s_{\text{pcp}})/4$ so that this can be done by estimating p_{acc} with an additive error at most ε , which will be done by applying the $\text{Approx}_{\varepsilon}\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ algorithm in the assumption. (Recall that c_{pcp} and s_{pcp} are absolute constants that only depend on δ_{Enc} , which means that ε is also an absolute constant.)

In what follows, we reduce the problem of estimating p_{acc} to $2^{r_{\text{shared}}}$ instances of $\text{Approx}_{\varepsilon}\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$, where each instance consists of $2^{r_{\text{col}}} = \ell^{1-\eta} \cdot \text{polylog}(\ell)$ circuits and $2^{r_{\text{row}}} = \ell^{c+1-\eta} \cdot \text{polylog}(\ell)$ inputs. Then we will utilize the algorithm for $\text{Approx}_{\varepsilon}\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ to estimate p_{acc} , where the data structure in the preprocessing phase will be treated as an advice of M^{PCPP} .

For the simplicity of presentation, we define the notation:

$$\begin{aligned} (\text{itype}[1], \dots, \text{itype}[q]) &\leftarrow V_{\text{type}}(\text{seed.shared}), \\ (\text{row}[1], \dots, \text{row}[q]) &\leftarrow V_{\text{row}}(\text{seed.shared}, \text{seed.row}), \\ (\text{icol}[1], \dots, \text{icol}[q]) &\leftarrow V_{\text{col}}(\text{seed.shared}, \text{seed.col}), \quad \text{and} \\ (pc_1, \dots, pc_p) &\leftarrow V_{\text{pc}}(\text{seed.shared}), \end{aligned}$$

where $p + q \leq 2$ and $pc_i : \{0, 1\}^{r_{\text{row}} + r_{\text{col}}} \rightarrow \{0, 1\}$ is an XOR of (some of) its input bits (i.e. a $\text{GF}(2)$ -linear function) for every $i \in [p]$.

Reduction to SATISFYING-PAIRS. Our input strings in the $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ instance will be of the form $(a_1, \dots, a_q, pc_1^{\text{row}}, \dots, pc_p^{\text{row}})$. For each $j \in [q]$, the meaning of a_j is as follows:

- if $\text{itype}[j] = \text{input}$, then a_j is interpreted as a row of the input matrix, and we use $(a_j)_{\text{col}}$ to denote the col -th bit of a_j ;
- if $\text{itype}[j] = \text{proof}$, then a_j is interpreted as a “seed” such that $C(a_j)$ is a row of the proof matrix, and we use $(a_j)_{\text{col}}$ to denote the col -th bit of $C(a_j)$. (NOT the col -th bit of a_j !)

For each $j \in [p]$, pc_j^{row} is a bit representing the contribution of seed.row in the j -th parity-check bit, i.e. $pc_j^{\text{row}} := pc_j(\text{seed.row}, 0^{|\text{seed.col}|})$.

We first enumerate $\text{seed.shared} \in \{0, 1\}^{r_{\text{shared}}}$. For each seed.shared , we create an instance $\mathcal{I} := \mathcal{I}_{\text{seed.shared}}$ of $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ as follows. Let \tilde{x}_j be the j -th row of $\text{Enc}(x)$ (viewed as an $H_{\text{input}} \times W_{\text{input}}$ matrix). For each $\text{seed.row} \in \{0, 1\}^{r_{\text{row}}}$, we add the following input to \mathcal{I} :

$$\text{Input}_{\text{seed.shared,seed.row}} = (a_1, \dots, a_q, pc_1^{\text{row}}, \dots, pc_p^{\text{row}}),$$

where for every $j \in [q]$,

$$a_j := \begin{cases} \tilde{x}_{i_{\text{row}}[j]} & \text{if } \text{itype}[j] = \text{input}, \\ w_{i_{\text{row}}[j]} & \text{if } \text{itype}[j] = \text{proof}, \end{cases}$$

and pc_j^{row} is the contribution of seed.row to the j -th parity-check bit as defined above. Note that since $\tilde{n}_{\text{hard}} = H_{\text{input}} \cdot W_{\text{input}}$, $\tilde{x}_{i_{\text{row}}[j]} \in \{0, 1\}^{W_{\text{input}}}$ when $\text{itype}[j] = \text{input}$, i.e., $\tilde{x}_{i_{\text{row}}[j]}$ will not contain \perp (see Definition 2.9). The length of a_j is at most $\max\{W_{\text{input}}, n\} \leq n \cdot \text{polylog}(\ell)$, thus the total length of $\text{Input}_{\text{seed.shared,seed.row}}$ is also bounded by $n \cdot \text{polylog}(\ell)$.

Then, for every $\text{seed.col} \in \{0, 1\}^{r_{\text{col}}}$, we define a circuit $C_{\text{seed.shared,seed.col}}$ as follows. On input

$$(a_1, \dots, a_q, pc_1^{\text{row}}, \dots, pc_p^{\text{row}}),$$

it outputs

$$\text{VDec}\left((a_1)_{i_{\text{col}}[1]}, \dots, (a_q)_{i_{\text{col}}[q]}, pc_1^{\text{row}} \oplus pc_1^{\text{col}}, \dots, pc_p^{\text{row}} \oplus pc_p^{\text{col}}\right).$$

Here, $\text{VDec} \leftarrow V_{\text{dec}}(\text{seed.shared})$ is the decision predicate of VPCPP and pc_i^{col} represents the contribution of seed.col to the i -th parity-check bit, i.e., $pc_i^{\text{col}} := pc_i(0^{|\text{seed.row}|}, \text{seed.col})$. Note that by definition, $pc_i(\text{seed.row}, \text{seed.col}) = pc_i^{\text{row}} \oplus pc_i^{\text{col}}$. Also note that $C_{\text{seed.shared,seed.col}}$ is indeed an $\text{OR}_2 \circ \mathcal{C}$ circuit, since VDec is always the OR of its two input bits or their negation.



Figure 1: Examples of the circuit $C_{\text{seed.shared,seed.col}}$. In the left example, there are two queries and no parity-check bits, the first query has type **proof** and the second query has type **input**. In the right example, there are one query with type **proof** and one parity-check bit.

Now, our instance \mathcal{I} contains $M := 2^{r_{\text{row}}}$ inputs and $N := 2^{r_{\text{col}}}$ circuits. By definition, we have

$$\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) = C_{\text{seed.shared}, \text{seed.col}}(\text{Input}_{\text{seed.shared}, \text{seed.row}}).$$

Since $M = \ell^{c+1-\eta} \cdot \text{polylog}(\ell)$ and $N = \ell^{1-\eta} \cdot \text{polylog}(\ell)$, there is a non-trivial algorithm for $\text{Approx}_{\varepsilon}\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ with N circuits of size s and M inputs of length $n \cdot \text{polylog}(\ell)$. In particular, we can estimate $p_{\text{acc}}(\text{seed.shared})$ using this algorithm on $\mathcal{I}_{\text{seed.shared}}$ up to an additive error ε , where

$$p_{\text{acc}}(\text{seed.shared}) := \Pr_{\text{seed.row}, \text{seed.col}} \left[\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) \right].$$

In other words, we can obtain a $p'_{\text{acc}}(\text{seed.shared}) \in p_{\text{acc}}(\text{seed.shared}) \pm \varepsilon$. The overall acceptance probability of VPCPP on the input $\text{Enc}(x)$ and proof π is

$$\begin{aligned} p_{\text{acc}} &:= \Pr_{\text{seed} \leftarrow \{0,1\}^r} \left[\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) \right]. \\ &= \mathbb{E}_{\text{seed.shared}} \left[\Pr_{\text{seed.row}, \text{seed.col}} \left[\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) \right] \right] \\ &= \mathbb{E}_{\text{seed.shared}} \left[p_{\text{acc}}(\text{seed.shared}) \right] \\ &\in \mathbb{E}_{\text{seed.shared}} \left[p'_{\text{acc}}(\text{seed.shared}) \right] \pm \varepsilon. \end{aligned}$$

Hence we can estimate p_{acc} up to an additive error ε by taking average over all $p'_{\text{acc}}(\text{seed.shared})$ obtained by the $\text{Approx}_{\varepsilon}\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ algorithm over $\mathcal{I}_{\text{seed.shared}}$.

To summarise, our algorithm M^{PCPP} works as follows. It first computes $\text{Enc}(x)$ in $O(n)$ time. Then, it enumerates seed.shared , produces the instance $\mathcal{I}_{\text{seed.shared}}$, and feeds it to the algorithm for $\text{Approx}_{\varepsilon}\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ to obtain $p'_{\text{acc}}(\text{seed.shared})$. Let p'_{acc} be the average of $p'_{\text{acc}}(\text{seed.shared})$ over all $\text{seed.shared} \in \{0,1\}^{r_{\text{shared}}}$. It accepts if and only if $p'_{\text{acc}} \geq c_{\text{pcp}} - \varepsilon$.

Correctness of M^{PCPP} . For every $x \in \{0,1\}^{n_{\text{hard}}}$, we know by the discussion above that:

- If $x \notin L^{\text{hard}}$, we know that $\text{Enc}(x)$ is δ_{Enc} far from being in L^{enc} . By the soundness of VPCPP, $p_{\text{acc}} \leq s_{\text{pcp}}$, which further means that $p'_{\text{acc}} \leq p_{\text{acc}} + \varepsilon < c_{\text{pcp}} - \varepsilon$, hence M^{PCPP} will reject x .
- If $x \in L^{\text{hard}}$ has an *easy witness*, we can see by the definition of easiness that there is a proof π of $\text{Enc}(x) \in L^{\text{enc}}$ such that for every row $\pi_i \in \{0,1\}^{W_{\text{proof}}}$ of π , there is a string $w_i \in \{0,1\}^n$ such that $\pi_i = C(w_i)$. These w_i can be found by non-deterministic guessing at the beginning of M^{PCPP} . In such case, we know by the completeness of VPCPP that $p_{\text{acc}} \geq c_{\text{pcp}}$, which further means that $p'_{\text{acc}} \geq p_{\text{acc}} - \varepsilon \geq c_{\text{pcp}} - \varepsilon$. Therefore M^{PCPP} will accept x .

Complexity of M^{PCPP} . Each instance \mathcal{I} of $\text{Approx}_{\varepsilon}\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ contains $M := 2^{r_{\text{row}}}$ inputs and $N := 2^{r_{\text{col}}}$ circuits. Since each instance can be solved in $NM/\log^{c_{\text{imp}}}(NM)$ time, the total time are

$$\begin{aligned} &2^{r_{\text{shared}}} \cdot NM/\log^{c_{\text{imp}}}(NM) \\ &\leq 2^{r_{\text{shared}}} \cdot 2^{r_{\text{row}}} \cdot 2^{r_{\text{col}}}/r^{c_{\text{imp}}} \\ &\leq 2^r/r^{c_{\text{imp}}}. \end{aligned}$$

Recall that $r = \log T + O(\log \log T + m \log m)$, where $O(\cdot)$ hides some absolute constant, we can see that $2^r/r^{c_{\text{imp}}} = T \log^{O(1)} T / \log^{c_{\text{imp}}} T$. By setting c_{imp} to be an sufficiently large absolute constant depending on c_{hard} , we can make $2^r/r^{c_{\text{imp}}} \leq T/\log^{c_{\text{hard}}} T$. Also, we can compute $\text{Enc}(x)$ in

$O(n_{\text{hard}})$ time and this is not the bottleneck. Therefore, the total running time of M^{PCPP} is at most $T/\log^{c_{\text{hard}}} T$.

It then suffices to determine the advice and non-determinism complexity of M^{PCPP} . For every `seed.shared`, the machine M^{PCPP} needs the data structure $\text{DS}_{\text{seed.shared}}$ as advice to support the algorithm for Satisfying Pairs. Since $|\text{DS}_{\text{seed.shared}}| \leq N^c = 2^{cr_{\text{col}}}$ by the assumption, the advice complexity of M^{PCPP} is

$$2^{cr_{\text{col}}+r_{\text{shared}}} \leq \ell^{c-c\eta+2\eta} \leq \ell^{c+1} \leq n_{\text{hard}}/10.$$

Also, the number of nondeterministic bits that M^{PCPP} guesses is at most $H_{\text{proof}} \cdot n \leq n_{\text{hard}}/10$. Therefore we can see that

$$M^{\text{PCPP}} \in \text{NTIMEGUESS}_{\text{RTM}}[T/\log^{c_{\text{hard}}}(T), n_{\text{hard}}/10]_{/(n_{\text{hard}}/10)}.$$

The final algorithm. Given a multi-output circuit $C : \{0,1\}^n \rightarrow \{0,1\}^\ell$, our algorithm for finding a non-output of C works as follows. First, we construct the hard language L^{hard} and the algorithm M^{PCPP} . Since M^{PCPP} is a nondeterministic algorithm that runs in $T/\log^{c_{\text{hard}}}(T)$ time, uses at most $n_{\text{hard}}/10$ bits of nondeterminism and at most $n_{\text{hard}}/10$ bits of advice, it follows that there is an input $x_{\text{hard}} \in \{0,1\}^{n_{\text{hard}}}$ such that $M^{\text{PCPP}}(x_{\text{hard}}) \neq L^{\text{hard}}(x_{\text{hard}})$. Moreover, let α be the advice string fed to M^{PCPP} , i.e., the data structures $\text{DS}_{\text{seed.shared}}$ for each `seed.shared`. (Note that we can obtain α since the avoidance algorithm has an NP oracle.) We can find such an input x_{hard} by running $\mathcal{R}(1^{n_{\text{hard}}}, M^{\text{PCPP}}, \alpha)$, where \mathcal{R} is the refuter guaranteed by Theorem 2.2. Thus, we can find x_{hard} in $\text{poly}(T)$ time with an NP oracle.

If $x_{\text{hard}} \notin L^{\text{hard}}$, then M^{PCPP} also rejects x_{hard} , which means $M^{\text{PCPP}}(x_{\text{hard}}) = L^{\text{hard}}(x_{\text{hard}})$. Thus it has to be the case that $x_{\text{hard}} \in L^{\text{hard}}$ but M^{PCPP} rejects x_{hard} . Therefore, x_{hard} does not have an easy witness. We can then use the NP oracle to find the lexicographically first proof matrix π such that

$$\Pr_{\text{seed} \leftarrow \{0,1\}^r} [\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) \text{ accepts}] \geq c_{\text{pcp}}.$$

Treating π as a matrix of dimension $H_{\text{proof}} \times W_{\text{proof}}$, there has to be a row that is not in the range of C . We can pick such a row by using the NP oracle. \square

Remark 3.3. In Theorem 3.2, we assumed a non-trivial SATISFYING-PAIRS algorithm for the circuit class $\text{OR}_2 \circ \mathcal{C}$. By Theorem 2.17, a non-trivial SATISFYING-PAIRS algorithm for $\text{AND}_2 \circ \mathcal{C}$ or $\text{XOR}_2 \circ \mathcal{C}$ also suffices. This property might be useful for some circuit classes with a better closure property under top XOR_2 gates (or AND_2 gates).

By replacing the 2-query PCPP (with imperfect completeness) with the 3-query PCPP (with perfect completeness) in Theorem 2.13, we can show that non-trivial algorithms for $\text{Gap}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ where $\mathcal{C}' = \text{OR}_3 \circ \mathcal{C}$ also imply FP^{NP} algorithms for $\mathcal{C}\text{-AVOID}$. We state the result below but omit the proofs.

Corollary 3.4. *There are constants $\varepsilon > 0$ and c_{imp} such that the following holds. Let $0 < \eta < 1/2$ be a constant, $\ell(n) > n^{1+4\eta}$ be a good function. Let $s = s(n)$ be a size parameter, $\mathcal{C}[s]$ be a typical circuit class where s is a size parameter, and $\mathcal{C}'[3s] := \text{OR}_3 \circ \mathcal{C}[s]$ (i.e. a \mathcal{C}' circuit of size $3s$ refers to an OR_3 of at most two \mathcal{C} circuits of size s).*

Assumption: *Suppose that for some constant $c \geq 1$, there is an $(NM/\log^{c_{\text{imp}}}(NM))$ -time algorithm for $\text{Gap}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ with $N := \ell^{1-\eta} \cdot \text{polylog}(\ell)$ circuits of size $s(n)$ and $M := \ell^{c+1-\eta} \cdot \text{polylog}(\ell)$ inputs of length $n \cdot \text{polylog}(\ell)$, allowing a P^{NP} preprocessing of an N^c -size data structure.*

Conclusion: *Then there is an FP^{NP} algorithm for $\mathcal{C}[s]\text{-AVOID}$ with stretch $\ell(n)$.*

3.2 Remote Point from SATISFYING-PAIRS

Theorem 3.5. *There is a universal constant $c_u \geq 1$ such that the following holds. Let $N := N(n)$ be a parameter such that $2^{\log^{c_u} n} < N < 2^{n^{0.99}}$, $\varepsilon := \varepsilon(n) > n^{-c_u}$ be the error parameter, and $\ell := N^{c_u \log(1/\varepsilon)}$. Let $\mathcal{C}[s]$ be a typical circuit class, where $s := s(n) \leq N$ is a size parameter, and denote $\mathcal{C}'[c_u s] := \text{AND}_{c_u} \circ \mathcal{C}[s]$ (i.e. a \mathcal{C}' circuit of size $c_u s$ refers to the AND of at most c_u \mathcal{C} circuits of size s).*

Assumption: *Let $P := (\log N)^{\log(1/\varepsilon)}$. Suppose there is a deterministic algorithm running in time $T^{\text{alg}} := N^2/P^{c_u}$ that, given as input a list of N $\mathcal{C}'[c_u s]$ circuits $\{C_i\}$ and a list of N inputs $\{x_j\}$ with input length $n \cdot \text{polylog}(\ell)$, estimates $\Pr_{i,j \leftarrow [N]}[C_i(x_j)]$ with additive error $\eta := \varepsilon^{c_u}$.*

Conclusion: *Then there is an FP^{NP} algorithm that takes as input a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, where each output bit of C can be computed in $\mathcal{C}[s]$, and prints a string y that is $(1/2 - \varepsilon)$ -far from $\text{Range}(C)$.*

The rest of this section is devoted to proving Theorem 3.5.

OVERVIEW OF SECTION 3.2

- In Section 3.2.1, we define a circuit class called $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$, and show that a SATISFYING-PAIRS algorithm for $\text{AND}_d \circ \mathcal{C}$ implies a SATISFYING-PAIRS algorithm for this class. This will be a convenient tool for our subsequent arguments.
- To solve the remote point problem, we need to define a nondeterministic machine called M^{PCPP} trying to contradict the nondeterministic time hierarchy (Theorem 2.2). In Section 3.2.2, we set the framework for this machine: it uses the PCPP theorem in Theorem 2.14, guesses a “compressed” version of the PCPP proof, and verifies the validity of this PCPP proof without decompressing it.
- The first problem we encounter is the “non-Booleanness” of the PCPP proof. As we use Theorem 2.5, the decompressed proof consists of real numbers instead of Boolean values, and we need to check whether the decompressed proof is “close to Boolean” (in a carefully defined technical sense). This is done in Section 3.2.3 via the SATISFYING-PAIRS algorithm.
- In Section 3.2.4, we use the faster algorithm for SATISFYING-PAIRS to verify the PCPP proof. This step is straightforward but tedious.
- After we obtain a non-trivial algorithm for verifying the PCPP proof, we conclude the machine M^{PCPP} in Section 3.2.5. Then we use this machine to build an FP^{NP} algorithm for the remote point problem in Section 3.2.6.
- Finally, Appendix A contains postponed proofs.

3.2.1 SATISFYING-PAIRS for $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ Circuits

It turns out that as an intermediate step, we need a SATISFYING-PAIRS algorithm for the following class of multi-output circuits that output real numbers. Let $d \geq 1$ be a constant, $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ denotes the class of multi-output circuits which takes two inputs $x \in \{0, 1\}^n$ and α , and has the following components:

- Let $\ell_{\mathcal{C}}$ denote the number of bottom \mathcal{C} circuits. For each $i \in [\ell_{\mathcal{C}}]$, the i -th circuit is a \mathcal{C} circuit computing a function $C_i : \{0, 1\}^n \rightarrow \{0, 1\}$.
- Let ℓ_{Sum} denote the number of middle “linear sum” gates. For each $i \in [\ell_{\text{Sum}}]$, the i -th gate outputs

$$\text{Sum}_i(x, \alpha) := \sum_{k \in [A]} \text{coeff}_k(\alpha) \cdot C_{\text{id}_{\times k}(\alpha, i)}(x).$$

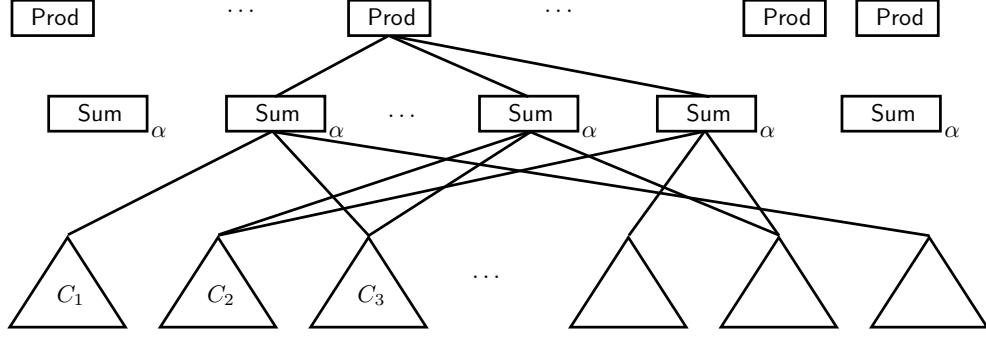


Figure 2: Example of a $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuit.

(See Definition 2.3 for the definition of linear sum circuits, in particular the coefficient sum and locality of a linear sum circuit. Note that this definition is different from the definitions in [Wil18b, CW19b].)

- Let ℓ_{Prod} denote the number of output gates. Each output gate is a product (i.e., multiplication) gate of fan-in d , and is connected to the $q_1(i), q_2(i), \dots, q_d(i)$ -th linear sum gate. Its output is

$$C_i^{\text{Prod}}(x, \alpha) := \prod_{t=1}^d \text{Sum}_{q_t(i)}(x, \alpha).$$

Remark 3.6. The important measures of a $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuit are:

- the number of gates in each level ($\ell_{\mathcal{C}}, \ell_{\text{Sum}}, \ell_{\text{Prod}}$);
- the fan-in of the top Prod gates (d);
- the fan-in (A), coefficient sum (U), and locality (l) of the linear sum layer.

We show that a SATISFYING-PAIRS algorithm for $\text{AND}_d \circ \mathcal{C}$ circuits implies a “SATISFYING-PAIRS algorithm” for $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuits that given a list of $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuits and a list of input strings, estimates the expected output value (as a real number) for a random circuit and a random input string in the lists. The proof is deferred to Appendix A.1.

Theorem 3.7. *Let \mathcal{C} be a typical circuit class, $M' \geq 1$ and $\eta \in (0, 1)$ be parameters. Suppose there is a deterministic algorithm running in time $T^{\text{alg}} = T^{\text{alg}}(N, M)$ that, given as input a list of $\hat{N} \leq N$ $\text{AND}_d \circ \mathcal{C}$ circuits $\{C_i\}$ and a list of $\hat{M} \leq M$ inputs $\{x_j\}$, estimates the following quantity with additive error η :*

$$\Pr_{i \leftarrow [\hat{N}], j \leftarrow [\hat{M}]} [C_i(x_j)].$$

Then, there is a deterministic algorithm running in time $A^d(2^{dl} + M'/M) \cdot (\ell_{\text{Prod}}/N) \cdot O(T^{\text{alg}})$ that, given as input a $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuit C^{Prod} with parameters specified in Remark 3.6, and a list of M' inputs $\{(x_j, \alpha_j)\}$, estimates the following quantity with additive error $\eta \cdot U^d$:

$$\mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow [M']} [C_i^{\text{Prod}}(x_j, \alpha_j)].$$

3.2.2 Set Up

Suppose that we are given a \mathcal{C} circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ as input. Let q, c_m, c_{tm} be constants that will be determined later. Define

$$\begin{aligned} \delta &:= (10^9 q)^{-q^2}, \\ m &:= c_m \log(1/\varepsilon)/\delta, \\ w_{\text{proof}} &:= (60q/m) \log \ell, & W_{\text{proof}} &:= 2^{w_{\text{proof}}} = \ell^{O(\delta/\log(1/\varepsilon))}, \\ h_{\text{proof}} &:= (5q + 1) \log \ell, & H_{\text{proof}} &:= 2^{h_{\text{proof}}} = \ell^{5q+1}, \\ n_{\text{hard}} &:= 20H_{\text{proof}} \cdot n, \\ T &:= H_{\text{proof}} \cdot W_{\text{proof}} / \log^{c_{\text{tm}}}(\ell). \end{aligned}$$

Let L^{hard} be the hard language constructed in Theorem 2.2, i.e.,

$$L^{\text{hard}} \in \text{NTIME}_{\text{TM}}[T] \setminus \text{i.o.-NTIME}_{\text{GUESS}_{\text{RTM}}}[T / \log^{c_{\text{hard}}}(T), n_{\text{hard}}/10]_{/(n_{\text{hard}}/10)},$$

where n_{hard} refers to the input length and c_{hard} is an absolute constant.

Since $T = \ell^{5q+1+\Theta(1/\log(1/\varepsilon))} / \log^{c_{\text{tm}}}(\ell) = \ell^{5q+1} \cdot n^{\omega(1)}$ and $n_{\text{hard}} = 20\ell^{5q+1}n$, we can see that $n_{\text{hard}} \cdot \text{polylog}(n_{\text{hard}}) \leq T \leq 2^{\text{poly}(n_{\text{hard}})}$, which satisfies the technical condition of Theorem 2.2.

Like in the proof of Theorem 3.2, we describe a nondeterministic RAM M^{PCPP} that runs in $T / \log^{c_{\text{hard}}}(T)$ time, guesses $n_{\text{hard}}/10$ nondeterministic bits, and attempts to solve L^{hard} . We will show that for every input $x \in \{0, 1\}^{n_{\text{hard}}}$, if $x \notin L^{\text{hard}}$ then $M^{\text{PCPP}}(x)$ rejects; while if $x \in L^{\text{hard}}$ and *has an easy witness*, then $M^{\text{PCPP}}(x)$ accepts. However, to solve \mathcal{C} -REMOTE-POINT, we need a slightly different definition for “easy witness”.

Let VPCPP be the verifier for the smooth and rectangular PCPP (Theorem 2.14) for the language

$$L^{\text{enc}} := \{\text{Enc}(x) : x \in L^{\text{hard}}\},$$

where we fix an error-correcting code (Enc, Dec) as in Theorem 2.1. Let δ_{Enc} be the (relative) distance of the error-correcting code. Suppose a string of length n_{hard} is encoded (via Enc) into a string of length $\tilde{n}_{\text{hard}} := \Theta(n_{\text{hard}})$. We set the following parameters:

$$\begin{aligned} h_{\text{input}} &:= \left(1 - \frac{\Theta(m^2 \log \log T)}{\log T}\right) h_{\text{proof}}, & H_{\text{input}} &:= 2^{h_{\text{input}}} = H_{\text{proof}} / \text{polylog}(\ell), \\ w_{\text{input}} &:= \lceil \log \tilde{n}_{\text{hard}} \rceil - h_{\text{input}}, & W_{\text{input}} &:= 2^{w_{\text{input}}} = n \cdot \text{polylog}(\ell). \end{aligned}$$

Again, we assume without loss of generality that $\tilde{n}_{\text{hard}} = H_{\text{input}} \cdot W_{\text{input}}$.

We invoke Theorem 2.14 for L^{enc} to obtain a verifier VPCPP with proof size $\hat{H}_{\text{proof}} \times W_{\text{proof}}$ and input size $H_{\text{input}} \times W_{\text{input}}$, where $\hat{H}_{\text{proof}} = 2^{\hat{h}_{\text{proof}}}$ for some¹⁶ $\hat{h}_{\text{proof}} = \log T + \Theta(m \log \log T) - w_{\text{proof}}$. We can check the technical requirements of Theorem 2.14 as follows:

$$\begin{aligned} T &\geq H_{\text{proof}} \cdot \ell^{\Theta(\delta/\log(1/\varepsilon))} \geq H_{\text{proof}} \cdot \Theta(n) = \tilde{n}_{\text{hard}}, \\ m &= \Theta(\log n / \delta) \leq (\log T)^{0.1}, \\ w_{\text{proof}} &= (60q/m) \log \ell \geq (5/m) \log T, \\ \hat{h}_{\text{proof}} &= h_{\text{proof}} + \Theta(\log n \log \log \ell) \geq (5q + 1) \log \ell \geq (5/m) \log T, \end{aligned}$$

¹⁶Note that the function \hat{h}_{proof} that Theorem 2.14 produces might not be exactly equal to h_{proof} . However, this difference is minor as these two quantities are close to each other.

$$\begin{aligned}
\frac{w_{\text{input}}}{w_{\text{proof}}} &= \frac{h_{\text{proof}} + \log n + O(1) - h_{\text{input}}}{w_{\text{proof}}} \\
&= \frac{\log n}{w_{\text{proof}}} + \frac{h_{\text{proof}}}{w_{\text{proof}}} \cdot \frac{\Theta(m^2 \log \log T)}{\log T} \\
&\leq \frac{m \log n}{\log \ell} + \frac{\Theta(m^3 \log \log T)}{\log T} \\
&\leq 1 - \Omega(1).
\end{aligned}$$

By Theorem 2.14, VPCPP has the following parameters:

- soundness error = $1/2$,
- proximity parameter = δ_{Enc} ,
- query complexity = $q := O(1)$,
- parity-check complexity = $q := O(1)$,
- total randomness = $r := \log T + O(\log \log T + m \log m)$,
- row randomness = $r_{\text{row}} := \hat{h}_{\text{proof}} - (5/m) \log T = \Theta(\log \ell)$,
- column randomness = $r_{\text{col}} := w_{\text{proof}} - (5/m) \log T = \Theta(\log \ell/m)$,
- shared randomness = $r_{\text{shared}} := (10/m) \log T + O(\log \log T + m \log m) = \Theta(\log \ell/m)$.

Here, all the $\Theta(\cdot)$ hides constants that may depend on q, c_m, c_{tm} . Moreover, as we choose the soundness error and the proximity parameters to be absolute constants, the query complexity q is also an absolute constant.

Note that if c_u is large enough, then we have $2^{\Omega(r_{\text{col}})} \leq N \leq 2^{r_{\text{col}}}$. Therefore we can solve the Approx_η -SATISFYING-PAIRS problem for $2^{r_{\text{col}}}$ inputs and $2^{r_{\text{col}}}$ $\text{AND}_{c_u} \circ \mathcal{C}$ circuits, by partitioning the inputs and circuits into groups of size N . The time complexity is still at most $2^{2r_{\text{col}}}/P^{c_u}$, where $P := (r_{\text{col}})^{\log(1/\varepsilon)}$. Without loss of generality, we may assume $N = 2^{r_{\text{col}}}$ in what follows.

We also fix the hardness amplification procedure $\text{Amp} : \{0, 1\}^{W_{\text{proof}}} \rightarrow \{0, 1\}^{\ell'}$ described by Theorem 2.5 that amplifies hardness δ to hardness $(1/2 - \varepsilon)$. Here, $\ell' := W_{\text{proof}}^{O(\log(1/\varepsilon)/\delta)} = \ell^{O(60q/c_m)}$. We set the parameter c_m such that $\ell' \leq \ell$. Without loss of generality, we may assume that $\ell' = \ell$.¹⁷ Let $(\text{idx}, \text{coeff})$ be the family of linear sum circuit described in Theorem 2.5, then $(\text{idx}, \text{coeff})$ has the following parameters:

$$\begin{aligned}
\text{advice complexity} &= a := O(\log^2 W_{\text{proof}}/(\varepsilon\delta)^2) = O(\log^2 \ell/\varepsilon^2), \\
\text{fan-in} &= A := O(\log W_{\text{proof}}/(\varepsilon\delta)^2) = O(\log \ell/\varepsilon^2), \\
\text{coefficient sum} &= U := O(1/\varepsilon), \\
\text{locality} &= l := \log \ell.
\end{aligned}$$

We say an input x has an easy witness if there is a proof matrix π such that:

(completeness) for every $\text{seed} \in \{0, 1\}^r$, $\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed})$ accepts;

(approximate easiness) for every row π_i of π , there exists an input $w_i \in \{0, 1\}^n$ and an advice $\alpha_i \in \{0, 1\}^a$ such that the decoding of $C(w_i)$ with advice α_i is δ -close to π_i with respect to ℓ_1 -norm. (Recall that $\text{dec}_\alpha(x)$ denotes the decoding of x under advice α .) In particular:

1. for every $j \in [W_{\text{proof}}]$, $(\text{dec}_{\alpha_i}(C(w_i)))_j \in [0, 1]$;

¹⁷If $\ell' \ll \ell$, we can partition the outputs of the circuit into blocks of size ℓ' , and solve the remote point problem for each block of output bits.

$$2. \|\text{dec}_{\alpha_i}(C(w_i)) - \pi_i\|_1 \leq \delta.$$

Recall that $P = (r_{\text{col}})^{\log(1/\varepsilon)}$. By our hypothesis, there is an algorithm that takes as input a list of N $\text{AND}_{4q} \circ \mathcal{C}$ circuits $\{C_i\}$ and a list of N inputs $\{x_j\}$, runs in deterministic $T^{\text{alg}} := 2^{2r_{\text{col}}}/P^{c_u}$ time, and estimates $\mathbb{E}_{i,j}[C_i(x_j)]$ within additive error $\eta := \varepsilon^{c_u} \leq U^{-10q}$.

3.2.3 Guessing and Verifying the PCPP

On input $x \in \{0,1\}^{n_{\text{hard}}}$, we guess \hat{H}_{proof} strings $w_1, w_2, \dots, w_{\hat{H}_{\text{proof}}} \in \{0,1\}^n$ as well as \hat{H}_{proof} advice strings $\alpha_1, \alpha_2, \dots, \alpha_{\hat{H}_{\text{proof}}} \in \{0,1\}^a$. Let $\pi_i^{\text{Real}} := \text{dec}_{\alpha_i}(C(w_i))$, and π_i^{Bool} be the Boolean string that is closest to π_i^{Real} . We will think of the matrix π^{Bool} as the PCPP proof, although our algorithm M^{PCPP} will operate on π^{Real} .

Therefore, before we proceed, we need to verify that π^{Real} and π^{Bool} are “close”, so that it does no harm to operate on π^{Real} even if the correct PCPP proof should be π^{Bool} . This verification phase also occurs in previous works proving lower bounds against linear combinations of circuits [Wil18b, CW19b, CR22, CLW20]. Like in previous work, we only provide an “approximate” verification algorithm: if the input has an easy witness, then the PCPP proof π^{Real} corresponding to this easy witness is accepted; on the other hand, we reject every π^{Real} that is “too far” from Boolean.

In what follows, denote

$$\begin{aligned} (\text{itype}[1], \text{itype}[2], \dots, \text{itype}[q]) &\leftarrow V_{\text{type}}(\text{seed.shared}), \\ (\text{irow}[1], \text{irow}[2], \dots, \text{irow}[q]) &\leftarrow V_{\text{row}}(\text{seed.shared}, \text{seed.row}), \quad \text{and} \\ (\text{icol}[1], \text{icol}[2], \dots, \text{icol}[q]) &\leftarrow V_{\text{col}}(\text{seed.shared}, \text{seed.col}). \end{aligned}$$

For each $\text{seed.shared} \in \{0,1\}^{r_{\text{shared}}}$ and each $\iota \in [q]$ such that $\text{itype}[\iota] = \text{proof}$, we define the following functions:

$$\begin{aligned} f_{\text{seed.shared},\iota}^{\text{Bool}}(\text{seed.row}, \text{seed.col}) &= \pi_{\text{irow}[\iota], \text{icol}[\iota]}^{\text{Bool}} \quad \text{and} \\ f_{\text{seed.shared},\iota}^{\text{Real}}(\text{seed.row}, \text{seed.col}) &= \pi_{\text{irow}[\iota], \text{icol}[\iota]}^{\text{Real}}. \end{aligned}$$

We will speak about the ℓ_d -norms of the above functions. For example, let $d \in \mathbb{N}$ be a constant, then

$$\|f_{\text{seed.shared},\iota}^{\text{Bool}} - f_{\text{seed.shared},\iota}^{\text{Real}}\|_d = \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\left| \pi_{\text{irow}[\iota], \text{icol}[\iota]}^{\text{Bool}} - \pi_{\text{irow}[\iota], \text{icol}[\iota]}^{\text{Real}} \right|^d \right]^{1/d}.$$

Lemma 3.8. *Let \mathcal{C} be a typical circuit class and $d \geq 2$ be an even number. Suppose there is an algorithm that takes as input a list of $N = 2^{r_{\text{col}}}$ $\text{AND}_{2d} \circ \mathcal{C}$ circuits $\{C_i\}$ and a list of N inputs $\{x_j\}$, runs in deterministic T^{alg} time, and estimates the following quantity with additive error η :*

$$\Pr_{i,j \leftarrow [2^{r_{\text{col}}}]}[C_i(x_j)].$$

Then there is an algorithm that takes the circuit C , $(w_1, w_2, \dots, w_{\hat{H}_{\text{proof}}})$, and $(\alpha_1, \alpha_2, \dots, \alpha_{\hat{H}_{\text{proof}}})$ as input, runs in deterministic $O((3A)^{2d} T^{\text{alg}}) \cdot (2^{2dl+r_{\text{shared}}} + T \log^{O(m)} T / 2^{2r_{\text{col}}})$ time, and satisfies the following:

(Completeness) *If for every $i \in [\hat{H}_{\text{proof}}]$, it holds that (1) for every $j \in [W_{\text{proof}}]$, $\pi_{i,j}^{\text{Real}} \in [0,1]$; (2) $\|\pi_i^{\text{Real}} - \pi_i^{\text{Bool}}\|_1 \leq \delta$, then the algorithm accepts.*

(Soundness) *If the algorithm accepts, then it holds that*

1. for every $\text{seed.shared} \in \{0, 1\}^{r_{\text{shared}}}$ and $\iota \in [q]$, $\|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d^d \leq 1 + 2\eta \cdot U^d$;
2. $\mathbb{E}_{i \leftarrow [\hat{H}_{\text{proof}}], j \leftarrow [W_{\text{proof}}]} \left[|\pi_{i,j}^{\text{Real}} - \pi_{i,j}^{\text{Bool}}|^d \right] \leq 4^d \cdot \delta + 2^{d+1} \eta (2U + 1)^{2d}$.

To prove this lemma, we need to reduce the task of checking whether the real proof is close to the Boolean proof to the satisfying pairs of $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuits, and then apply the non-trivial algorithm in Theorem 3.7. The details are given in Appendix A.2.

We substitute $d := 2q$ in the above lemma. If x has an easy witness, then there is some $(w_1, w_2, \dots, w_{\hat{H}_{\text{proof}}})$ and $(\alpha_1, \alpha_2, \dots, \alpha_{\hat{H}_{\text{proof}}})$ that passes the test; on the other hand, if the test is passed, then both soundness properties in Lemma 3.8 hold:

1. for every seed.shared and ι , $\|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_{2q}^{2q} \leq 1 + 2\eta \cdot U^{2q}$;
2. $\mathbb{E}_{i \leftarrow [H_{\text{proof}}], j \leftarrow [W_{\text{proof}}]} \left[|\pi_{i,j}^{\text{Real}} - \pi_{i,j}^{\text{Bool}}|^{2q} \right] \leq 16^q \cdot \delta + 128^q \eta U^{4q}$.

3.2.4 Estimating the Acceptance Probability

After checking that the PCPP proof is “close to Boolean”, the next step is to use it to speed up L^{hard} . We estimate

$$p_{\text{acc}} := \Pr_{\text{seed} \leftarrow \{0,1\}^r} \left[\text{VPCPP}^{\text{Enc}(x) \circ \pi^{\text{Bool}}}(\text{seed}) \text{ accepts} \right].$$

(Indeed, it suffices to distinguish between the case that $p_{\text{acc}} \geq 5/6$ and the case that $p_{\text{acc}} < 1/2$ as we will explain later.)

We enumerate seed.shared . After fixing seed.shared , each $\text{itype}[\iota]$ is completely fixed, each $\text{irow}[\iota]$ only depends on seed.row , and each $\text{icol}[\iota]$ only depends on seed.col . We now need to estimate

$$p_{\text{acc}}(\text{seed.shared}) := \Pr_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\text{VPCPP}^{\text{Enc}(x) \circ \pi^{\text{Bool}}}(\text{seed}) \text{ accepts} \right].$$

Suppose that we also fix seed.row . Then, we know the q rows of the input matrix $\text{Enc}(x)$ and the proof matrix π that could influence the PCPP verifier. We call them $\text{row}_1^{\text{Bool}}, \text{row}_2^{\text{Bool}}, \dots, \text{row}_q^{\text{Bool}}$. In particular, for each $\iota \in [q]$:

$$\text{row}_\iota^{\text{Bool}} = \begin{cases} \tilde{x}_{\text{irow}[\iota]} & \text{if } \text{itype}[\iota] = \text{input}, \\ \pi_{\text{irow}[\iota]}^{\text{Bool}} & \text{if } \text{itype}[\iota] = \text{proof}. \end{cases}$$

We also let $pc_1, pc_2, \dots, pc_q \leftarrow V_{\text{pc}}(\text{seed.shared})$ be the parity-check functions of the PCPP verifier, where each $pc_\iota : \{0, 1\}^{r_{\text{row}} + r_{\text{col}}} \rightarrow \{0, 1\}$. In particular, let pc_ι^{row} (resp. pc_ι^{col}) denote the contribution of seed.row (resp. seed.col) to pc_ι , i.e.,

$$\begin{aligned} pc_\iota^{\text{row}}(\text{seed.row}) &:= pc_\iota(\text{seed.row}, 0^{r_{\text{col}}}), \\ pc_\iota^{\text{col}}(\text{seed.col}) &:= pc_\iota(0^{r_{\text{row}}}, \text{seed.col}). \end{aligned}$$

Then $pc_\iota(\text{seed.row}, \text{seed.col}) = pc_\iota^{\text{row}}(\text{seed.row}) \oplus pc_\iota^{\text{col}}(\text{seed.col})$. For simplicity, we omit seed.row and seed.col when they are clear from the context.

Let VDec be the decision predicate of the PCPP verifier; note that as seed.shared is fixed, $\text{VDec} \leftarrow V_{\text{dec}}(\text{seed.shared})$ is also fixed. The input of VDec includes the answers to the q queries and the parity-check bits pc_1, \dots, pc_q . On seed.row and seed.col , the PCPP verifier outputs

$$\text{VDec} \left((\text{row}_1^{\text{Bool}})_{\text{icol}[1]}, (\text{row}_2^{\text{Bool}})_{\text{icol}[2]}, \dots, (\text{row}_q^{\text{Bool}})_{\text{icol}[q]}, pc_1, pc_2, \dots, pc_q \right).$$

As every Boolean function over $2q$ bits can be written as a degree- $2q$ polynomial over the reals, we write

$$\text{VDec}(a_1, a_2, \dots, a_q, pc_1, pc_2, \dots, pc_q) = \sum_{S \subseteq [q], S' \subseteq [q]} \theta_{S, S'} \left(\prod_{\iota \in S} a_\iota \right) \cdot \left(\prod_{\iota \in S'} pc_\iota \right),$$

where $\theta_{S, S'} \in [-2^{2q}, 2^{2q}]$. Now, define

$$p_{\text{acc}}(\text{seed.shared}, S, S') := \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\prod_{\iota \in S} (\text{row}_\iota^{\text{Bool}})_{\text{icol}[\iota]} \cdot \prod_{\iota \in S'} pc_\iota \right].$$

We have

$$p_{\text{acc}}(\text{seed.shared}) = \sum_{S \subseteq [q], S' \subseteq [q]} \theta_{S, S'} p_{\text{acc}}(\text{seed.shared}, S, S'),$$

thus it suffices to estimate each $p_{\text{acc}}(\text{seed.shared}, S, S')$.

Fix S and S' . Since we only have access to a real proof matrix π^{Real} instead of a Boolean proof matrix, we use the following number as an estimation of $p_{\text{acc}}(\text{seed.shared}, S, S')$, with the only difference being π_i^{Bool} being replaced by π^{Real} :

$$p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S') = \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\prod_{\iota \in S} (\text{row}_\iota^{\text{Real}})_{\text{icol}[\iota]} \cdot \prod_{\iota \in S'} pc_\iota \right],$$

where

$$\text{row}_\iota^{\text{Real}} = \begin{cases} \tilde{x}_{i_{\text{row}[\iota]}} & \text{if } \text{itype}[\iota] = \text{input}, \\ \pi_{i_{\text{row}[\iota]}}^{\text{Real}} & \text{if } \text{itype}[\iota] = \text{proof}. \end{cases}$$

The following claim bounds the accuracy of the estimation given the ℓ_d -distance between the functions $f_{\text{seed.shared}, \iota}^{\text{Bool}}$ and $f_{\text{seed.shared}, \iota}^{\text{Real}}$. The proof is deferred to Appendix A.3.

Claim 3.9. *For every $S, S' \subseteq [q]$,*

$$|p_{\text{acc}}(\text{seed.shared}, S, S') - p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S')| \leq (1 + \delta_{\text{seed.shared}})^{2q-1} \cdot \delta_{\text{seed.shared}}.$$

Here,

$$\begin{aligned} \delta_{\text{seed.shared}} &:= \sum_{\iota: \text{itype}[\iota] = \text{proof}} \|f_{\text{seed.shared}, \iota}^{\text{Bool}} - f_{\text{seed.shared}, \iota}^{\text{Real}}\|_{2q} \\ (\text{recall}) &= \sum_{\iota: \text{itype}[\iota] = \text{proof}} \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\left| \pi_{i_{\text{row}[\iota], \text{icol}[\iota]}}^{\text{Bool}} - \pi_{i_{\text{row}[\iota], \text{icol}[\iota]}}^{\text{Real}} \right|^{2q} \right]^{1/(2q)}. \end{aligned}$$

Now we fix S, S' and estimate $p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S')$. Let $d_S := |S|$, $d_{S'} := |S'|$, it is without loss of generality to assume that $S = \{1, 2, \dots, d_S\}$ and $S' = \{1, 2, \dots, d_{S'}\}$. We construct a $\text{Prod}_{d_S + d_{S'}} \circ \text{Sum} \circ \mathcal{C}$ circuit $C^{\text{Prod}} := C_{\text{seed.shared}, S, S'}^{\text{Prod}}$, as well as a list of inputs $(z_{\text{seed.row}}, \alpha_{\text{seed.row}})$, such that

$$C_{\text{seed.col}}^{\text{Prod}}(z_{\text{seed.row}}, \alpha_{\text{seed.row}}) = \prod_{\iota \in S} (\text{row}_\iota^{\text{Real}})_{\text{icol}[\iota]} \cdot \prod_{\iota \in S'} pc_\iota. \quad (5)$$

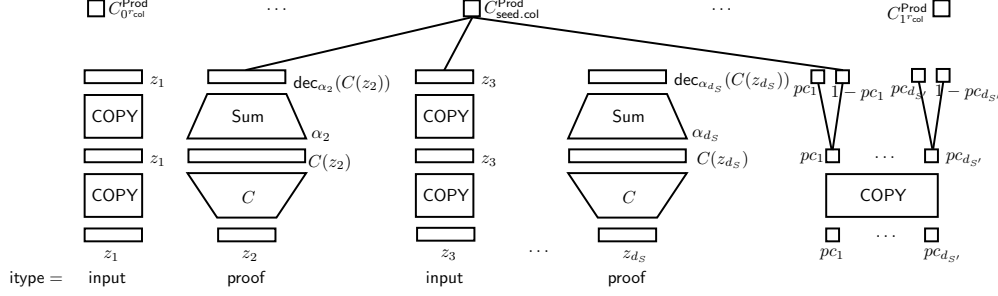


Figure 3: Construction of the circuit C^{Prod} . Note that for convenience, we only drew the “relevant” parts of this circuit, e.g., $C(z_i)$ when $\text{itype}[i] = \text{proof}$ and the copying circuit for z_i when $\text{itype}[i] = \text{input}$.

Construction of C^{Prod} and inputs. For each seed.row and each $\iota \in S$, define

$$(z_{\text{seed.row}})_\iota := \begin{cases} \tilde{x}_{\text{irow}[\iota]} & \text{if } \text{itype}[\iota] = \text{input}, \\ w_{\text{irow}[\iota]} & \text{if } \text{itype}[\iota] = \text{proof}. \end{cases}$$

Then we concatenate each $(z_{\text{seed.row}})_j$ and the (row-)parity-check bits to obtain

$$z_{\text{seed.row}} := \left((z_{\text{seed.row}})_1, (z_{\text{seed.row}})_2, \dots, (z_{\text{seed.row}})_{d_S}, pc_1^{\text{row}}, pc_2^{\text{row}}, \dots, pc_{d_{S'}}^{\text{row}} \right).$$

It is easy to check that given seed.row (and seed.shared), we can compute $z_{\text{seed.row}}$ easily. We also define

$$\alpha_{\text{seed.row}} := (\alpha_{\text{irow}[1]}, \alpha_{\text{irow}[2]}, \dots, \alpha_{\text{irow}[d_S]}).$$

Next, we define the circuit C^{Prod} that takes two inputs (z, α) and outputs $2^{|\text{seed.col}|}$ real numbers. Fix seed.col , we want that

$$\begin{aligned} & C_{\text{seed.col}}^{\text{Prod}}(z_{\text{seed.row}}, \alpha_{\text{seed.row}}) \\ &= \prod_{\iota \in S} (\text{row}_\iota^{\text{Real}})_{\text{icol}[\iota]} \cdot \prod_{\iota \in S'} pc_\iota \\ &= \prod_{\iota \in S^{\text{proof}}} (\text{dec}_{\alpha_{\text{irow}[\iota]}}(C(w_{\text{irow}[\iota]})))_{\text{icol}[\iota]} \cdot \prod_{\iota \in S^{\text{input}}} \tilde{x}_{\text{irow}[\iota], \text{icol}[\iota]} \cdot \prod_{\iota \in S'} (pc_\iota^{\text{row}} \oplus pc_\iota^{\text{col}}), \end{aligned}$$

where

$$(\text{dec}_{\alpha_{\text{irow}[\iota]}}(C(w_{\text{irow}[\iota]})))_{\text{icol}[\iota]} = \sum_{k \in [A]} \text{coeff}_k(\alpha_{\text{irow}[\iota]}) \cdot C_{\text{id} \times k}(\alpha_{\text{irow}[\iota], \text{icol}[\iota]})(w_{\text{irow}[\iota]})$$

denotes the $\text{icol}[\iota]$ -th bit of the string obtained by decoding $C(w_{\text{irow}[\iota]})$ with the advice $\alpha_{\text{irow}[\iota]}$ using the decoder dec , $S^{\text{proof}} := S \cap \{i \in [q] \mid \text{itype}[i] = \text{proof}\}$, $S^{\text{input}} := S \cap \{i \in [q] \mid \text{itype}[i] = \text{input}\}$.

This motivates the definition of the circuit C^{Prod} (see Figure 3 for graphic exposition and Figure 4 for detailed definition). The parameters of the circuit C^{Prod} are as follows.

- The number of gates: $\ell_\emptyset = d_S(\ell + W_{\text{proof}}) + d_{S'}$, $\ell_{\text{Sum}} = W_{\text{proof}} \cdot d_S + 2d_{S'}$, $\ell_{\text{Prod}} = 2^{r_{\text{col}}}$.
- The fan-in of the top Prod gates $d_S + d_{S'} \leq 2q$.
- The fan-in $A' := A \cdot d_S + 2$, coefficient sum $d_S \cdot U + 2$, and locality l of the linear sum layer.

Circuit C^{Prod}

(Inputs) The input z will have the form $z = (z_1, z_2, \dots, z_{d_S}, pc_1, pc_2, \dots, pc_{d_{S'}})$ and the input α will have the form $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_{d_S})$. The intended meanings are $z_i = (z_{\text{seed.row}})_i$, $pc_i = pc_i^{\text{ow}}$, and $\alpha_i = \alpha_{\text{irow}[i]}$.

(Bottom circuits) We make d_S copies to C , where the i -th copy is applied to the input z_i . (The i -th copy is useful only when $\text{itype}[i] = \text{proof}$, but we make all d_S copies for convenience.) We also add $W_{\text{proof}} \cdot d_S + d_{S'}$ gates to copy the input.

Thus, there are $\ell_{\mathcal{C}} := d_S \cdot \ell + d_S \cdot W_{\text{proof}} + d_{S'}$ output gates; we identify $[\ell_{\mathcal{C}}]$ with the disjoint union of $\{1\} \times [d_S] \times [\ell]$, $\{2\} \times [d_S] \times [W_{\text{proof}}]$, and $\{3\} \times [d_{S'}]$.

- For each $j \in [d_S]$ and $i \in [\ell]$, the $(1, j, i)$ -th gate is $C_{(1,j,i)}(z) := C(z_j)_i$.
- For each $j \in [d_S]$ and $i \in [W_{\text{proof}}]$, the $(2, j, i)$ -th gate is $C_{(2,j,i)}(z) := (z_j)_i$.
- For each $j \in [d_{S'}]$, the $(3, j)$ -th gate is $C_{3,j}(z) := pc_j$.

(Intermediate linear sum gates) There are $\ell_{\text{Sum}} := W_{\text{proof}} \cdot d_S + 2d_{S'}$ linear sum gates and we identify $[\ell_{\text{Sum}}]$ with the disjoint union of $[W_{\text{proof}}] \times [d_S]$ and $[d_{S'}] \times \{0, 1\}$.

Let $i \in [W_{\text{proof}}]$ and $j \in [d_S]$. If $\text{itype}[j] = \text{proof}$, then the (i, j) -th intermediate gate is

$$\text{Sum}_{(i,j)}(z, \alpha) = \sum_{k \in [A]} \text{coeff}_k(\alpha_j) \cdot C_{\text{idx}_k(\alpha_j, i) \cdot d_S + j}(z).$$

It is easy to verify that

$$\text{Sum}_{(i,j)}(z_{\text{seed.row}}, \alpha_{\text{seed.row}}) = (\text{dec}_{\alpha_{\text{irow}[j]}}(C(w_{\text{irow}[j]})))_i.$$

On the other hand, if $\text{itype}[j] = \text{input}$, then the (i, j) -th intermediate gate is $\text{Sum}_{(i,j)}(z, \alpha) := (z_j)_i$. (If $i > W_{\text{input}}$ then we simply set $\text{Sum}_{(i,j)}(z, \alpha) = 0$ and this intermediate gate would not be used.) Finally, for each $i \in [d_{S'}]$, we have two intermediate gates

$$\text{Sum}_{(i,0)}(z, \alpha) = pc_i, \quad \text{Sum}_{(i,1)}(z, \alpha) = 1 - pc_i.$$

Implementation of the linear sum layer: The linear sum has fan-in $A' := A \cdot d_S + 2$ and we identify $[A']$ with the disjoint union of $[A] \times [d_S]$ and $\{+, -\}$. Let idx' and coeff' be the idx and coeff functions of the linear sum layer of C^{Prod} , then

(Function $\text{idx}'_k(\alpha, i)$) Suppose $i = (i', j) \in [W_{\text{proof}}] \times [d_S]$. If $\text{itype}[j] = \text{proof}$ and $k = (k', j')$ where $j = j'$, then we return $\text{idx}'_k(\alpha, i) = (1, j, \text{idx}_{k'}(\alpha_j, i'))$; if $\text{itype}[j] = \text{input}$ and $k = +$, then $\text{idx}'_k(\alpha, i) = (2, j, i')$. Otherwise $\text{idx}'_k(\alpha, i) = \text{ZERO}$.

On the other hand, suppose $i = (j, b) \in [d_{S'}] \times \{0, 1\}$. If $(b = 0 \text{ and } k = +)$ or $(b = 1 \text{ and } k = -)$ then $\text{idx}'_k(\alpha, i) = (3, j)$. If $b = 1$ and $k = +$ then $\text{idx}'_k(\alpha, i) = \text{ONE}$. Otherwise $\text{idx}'_k(\alpha, i) = \text{ZERO}$.

(Function $\text{coeff}'_k(\alpha)$) If $k = +$ then $\text{coeff}'_k(\alpha) = 1$; if $k = -$ then $\text{coeff}'_k(\alpha) = -1$; otherwise, if $k = (k', j')$ then $\text{coeff}'_k(\alpha) = \text{coeff}_{k'}(\alpha_{j'})$.

The locality of $(\text{idx}', \text{coeff}')$ is still l . The coefficient sum becomes $d_S \cdot U + 2$.

(Output product gates) There are 2^{col} product gates. For each seed.col , the seed.col -th output gate is

$$C_{\text{seed.col}}^{\text{Prod}}(z, \alpha) = \prod_{i \in S} \text{Sum}_{\text{icol}[i] \cdot d_S + i}(z, \alpha) \cdot \prod_{i \in S'} \text{Sum}_{W_{\text{proof}} \cdot d_S + 2i + pc_i^{\text{col}}}(z, \alpha).$$

Figure 4: Detailed definition of C^{Prod} .

Given the above construction, it is easy to check that Eq. (5) holds for every `seed.row` and `seed.col`. We can see that

$$p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S') = \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[C_{\text{seed.col}}^{\text{Prod}}(z_{\text{seed.row}}, \alpha_{\text{seed.row}}) \right].$$

Since $d_S \leq q$ and $d_{S'} \leq q$, by Theorem 3.7, we can estimate $p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S')$ with additive error $\eta \cdot (qU + 2)^{2q}$ in deterministic $O((qA)^{2q}(2^{2ql} + 2^{r_{\text{row}}}/N) \cdot T^{\text{alg}})$ time.

Analysis. First, the verification step takes

$$\begin{aligned} & O((3A)^{4q} T^{\text{alg}}) \cdot \left(2^{4ql+r_{\text{shared}}} + T \log^{O(m)} T / 2^{2r_{\text{col}}} \right) \\ & \leq O((3A)^{4q}) \cdot (T \log^{O(m)} T) / (r_{\text{col}})^{c_u \log(1/\varepsilon)} \\ & \leq T (\log T)^{O(m) - c_u \log(1/\varepsilon)/2} \end{aligned}$$

time, which is at most $T / \log^{c_{\text{hard}}} T$ if c_u is a large enough constant.

Our algorithm estimates $p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S')$. By Claim 3.9, the same algorithm estimates $p_{\text{acc}}(\text{seed.shared}, S, S')$ within additive error $\eta(qU + 2)^{2q} + \delta'_{\text{seed.shared}}$, where

$$\delta'_{\text{seed.shared}} := (1 + \delta_{\text{seed.shared}})^{2q-1} \cdot \delta_{\text{seed.shared}}.$$

Running this algorithm for every possible (S, S') , we obtain an algorithm that runs in deterministic $(O(qA))^{2q}(2^{2ql} + 2^{r_{\text{row}}}/M) \cdot T^{\text{alg}}$ time and estimates $p_{\text{acc}}(\text{seed.shared})$ within additive error

$$\begin{aligned} & \leq \bar{\delta}_{\text{seed.shared}} := 4^q \cdot \sum_{S, S'} (\eta \cdot (qU + 2)^{2q} + \delta'_{\text{seed.shared}}) \\ & \leq \eta \cdot (4qU + 8)^{2q} + 16^q \cdot \delta'_{\text{seed.shared}}. \end{aligned}$$

Finally, running this algorithm for every $\text{seed.shared} \in \{0, 1\}^{r_{\text{shared}}}$, we obtain an algorithm that runs in deterministic

$$\begin{aligned} & O(qA)^{2q}(2^{2ql} + 2^{r_{\text{row}}}/N) \cdot 2^{r_{\text{shared}}} \cdot O(T^{\text{alg}}) \\ & \leq O(\log^{2q} \ell / \varepsilon^{4q}) 2^{r_{\text{row}}}/2^{r_{\text{col}}} \cdot 2^{r_{\text{shared}}} \cdot 2^{2r_{\text{col}}} / (r_{\text{col}})^{c_u \log(1/\varepsilon)} \\ & \leq 2^r / \log^{\Omega(c_u)} \ell < T / \log^{c_{\text{hard}}}(T) \end{aligned}$$

time that estimates p_{acc} within additive error of at most

$$\mathbb{E}_{\text{seed.shared}} [\bar{\delta}_{\text{seed.shared}}] \leq \eta \cdot (4qU + 8)^{2q} + 16^q \mathbb{E}_{\text{seed.shared}} [\delta'_{\text{seed.shared}}].$$

It remains to upper bound the quantity $\mathbb{E}_{\text{seed.shared}}[\delta'_{\text{seed.shared}}]$. We abstract this task in the following lemma and defer the proof to Appendix A.4.

Lemma 3.10. *Let $f : [N] \times [q] \rightarrow \mathbb{R}_{\geq 0}$ be a function and $d \geq 1$ be a constant. Suppose that*

1. *for every $s \in [N]$ and $i \in [q]$, $f(s, i) \leq \alpha$ (where $\alpha \geq 1$);*
2. $\mathbb{E}_{s,i}[f(s, i)^d] \leq \delta$.

Let $f(s) := \sum_{i \in [q]} f(s, i)$. Then

$$\mathbb{E}_s [(1 + f(s))^{d-1} \cdot f(s)] \leq q \delta^{1/d} (2q\alpha)^{d-1}.$$

To see how this lemma corresponds to our scenario: Let $d = 2q$ and $[N] = \{0, 1\}^{r_{\text{shared}}}$. For seed.shared and ι , if $\text{itype}[\iota] = \text{proof}$, then define $f(\text{seed.shared}, \iota) = \|f_{\text{seed.shared}, \iota}^{\text{Bool}} - f_{\text{seed.shared}, \iota}^{\text{Real}}\|_{2q}$; otherwise define $f(\text{seed.shared}, \iota) = 0$. Since the verification algorithm did not reject π^{Real} , we have

1. For every seed.shared and ι ,

$$f(\text{seed.shared}, \iota) \leq \|f_{\text{seed.shared}, \iota}^{\text{Bool}}\|_{2q} + \|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_{2q} \leq 2 + 2\eta \cdot U^{2q}.$$

2. Since the PCPP is smooth, the distribution of $(\text{irow}[\iota], \text{icol}[\iota])$ for random (seed, ι) (conditioned on $\text{itype}[\iota] = \text{proof}$) is the same as the uniform distribution over $[H_{\text{proof}}] \times [W_{\text{proof}}]$. Therefore

$$\begin{aligned} \mathbb{E}_{\text{seed.shared}, \iota} [f(\text{seed.shared}, \iota)^{2q}] &= \mathbb{E}_{\text{seed}, \iota: \text{itype}[\iota] = \text{proof}} [|\pi_{\text{irow}[\iota], \text{icol}[\iota]}^{\text{Bool}} - \pi_{\text{irow}[\iota], \text{icol}[\iota]}^{\text{Real}}|^{2q}] \\ &= \mathbb{E}_{i, j} [|\pi_{i, j}^{\text{Bool}} - \pi_{i, j}^{\text{Real}}|^{2q}] \\ &\leq 16^q \cdot \delta + 128^q \eta U^{4q}. \end{aligned}$$

It follows from Lemma 3.10 that

$$\begin{aligned} \mathbb{E}_{\text{seed.shared}} [\delta'_{\text{seed.shared}}] &\leq q(16^q \cdot \delta + 128^q \eta U^{4q})^{1/2q} (2q(2 + 2\eta \cdot U^{2q}))^{2q-1} \\ &\leq q(17^q \cdot \delta)^{1/2q} \cdot (100q)^{2q} < 100^{-q}. \end{aligned}$$

Therefore, the algorithm estimates p_{acc} within additive error at most

$$\eta(4qU + 8)^{2q} + 16^q \cdot 100^{-q} < 1/6,$$

thus successfully distinguishes between the case that $p_{\text{acc}} > 5/6$ and that $p_{\text{acc}} < 1/2$.

3.2.5 Wrap Up: Description of M^{PCPP}

On input x , we consider the smooth and rectangular PCPP for the language $L^{\text{enc}} = \{\text{Enc}(x) : x \in L^{\text{hard}}\}$. (Recall that M^{PCPP} aims to reject every $x \notin L$ and accepts every $x \in L$ with easy witness.) We guess $(w_1, \dots, w_{H_{\text{proof}}})$ and $(\alpha_1, \dots, \alpha_{H_{\text{proof}}})$, which implicitly defines the PCPP proof matrices π^{Bool} and π^{Real} . Then we verify π^{Real} using Lemma 3.8, reject immediately if π^{Real} did not pass the test. If π^{Real} passes the test (which means that it is “close” to a Boolean proof π^{Bool}), we use the algorithm described above to estimate p_{acc} . We accept x if and only if our estimation is above $2/3$.

The correctness of M^{PCPP} is easy to see:

Claim 3.11. *For every input x , if $x \notin L$ then M^{PCPP} rejects x ; while if $x \in L$ and x has an easy witness then M^{PCPP} accepts x .*

Proof. If $x \notin L$, then it always holds that $p_{\text{acc}} < 1/2$, so M^{PCPP} rejects. If $x \in L$ and x has an easy witness, then there exists a proof π , $(w_1, \dots, w_{H_{\text{proof}}})$ and $(\alpha_1, \dots, \alpha_{H_{\text{proof}}})$ such that

1. for every $j \in [W_{\text{proof}}]$, $\pi_{i, j}^{\text{Real}} \in [0, 1]$;
2. for every $i \in [W_{\text{proof}}]$, $j \in [H_{\text{proof}}]$, $\|\pi_i^{\text{Real}} - \pi_i\|_1 \leq \delta$.

Note that $\|\pi_i^{\text{Real}} - \pi_i^{\text{Bool}}\|_1 \leq \|\pi_i^{\text{Real}} - \pi_i\|_1 \leq \delta$ since π_i^{Bool} is the closest Boolean string to π_i^{Real} , and thus $\|\pi_i - \pi_i^{\text{Bool}}\|_1 \leq 2\delta$. Since the probability that VPCPP accepts π is 1, by Lemma 2.11, $p_{\text{acc}} \geq 1 - q \cdot 2\delta > 5/6$, so M^{PCPP} accepts. \square

The machine M^{PCPP} guesses $H_{\text{proof}}(n+a) < n_{\text{hard}}/10$ bits of nondeterminism, and uses $\tilde{O}(s\ell) < n_{\text{hard}}/10$ bits of advice. Thus $L(M^{\text{PCPP}}) \in \text{NTIMEGUESS}_{\text{RTM}}[T/\log^{c_{\text{hard}}}(T), n_{\text{hard}}/10]_{/(n_{\text{hard}}/10)}$.

3.2.6 The FP^{NP} Algorithm for REMOTE-POINT

Let $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ be the input circuit. We first construct the hard language L^{hard} and the algorithm M^{PCPP} . Since M^{PCPP} is a nondeterministic RAM algorithm that runs in $T/\log^{\text{c}_{\text{hard}}}(T)$ time, uses at most $n_{\text{hard}}/10$ nondeterministic bits and at most $n_{\text{hard}}/10$ advice bits, it follows that there is an input $x_{\text{hard}} \in \{0, 1\}^{n_{\text{hard}}}$ such that $M^{\text{PCPP}}(x_{\text{hard}}) \neq L^{\text{hard}}(x_{\text{hard}})$. Moreover, let α be the advice string fed to M^{PCPP} , i.e., the circuit C . We can find such an input x_{hard} by running $\mathcal{R}(1^{n_{\text{hard}}}, M^{\text{PCPP}}, \alpha)$, where \mathcal{R} is the refuter guaranteed by Theorem 2.2. Thus, we can find x_{hard} in deterministic $\text{poly}(T)$ time with an NP oracle.

It follows from Claim 3.11 that $x_{\text{hard}} \in L^{\text{hard}}$ but x_{hard} does not have an easy witness. Thus, we can use the NP oracle to find the lexicographically first PCPP proof matrix π such that

$$\Pr_{\text{seed} \leftarrow \{0,1\}^r} [\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) \text{ accepts}] = 1.$$

Then, there must exist a row π_i that is $(1/2 - \varepsilon)$ -far from $\text{Range}(C)$. To see this, suppose that for every i , the i -th row π_i is $(1/2 - \varepsilon)$ -close to $\text{Range}(C)$. Then there exists some $w_i \in \{0, 1\}^n$ such that $\delta(\text{Amp}(\pi_i), C(w_i)) \leq 1/2 - \varepsilon$. By Theorem 2.5, there is an advice α_i such that $\text{dec}_{\alpha_i}(C(w_i))$ satisfies (1) for every $j \in [W_{\text{proof}}]$, $(\text{dec}_{\alpha_i}(C(w_i)))_j \in [0, 1]$; (2) $\|\text{dec}_{\alpha_i}(C(w_i)) - \pi_i\|_1 \leq \delta$. It follows that π is an easy witness for x_{hard} , a contradiction.

Finally, we use the NP oracle to find the first row π_i , such that $\text{Amp}(\pi_i)$ is $(1/2 - \varepsilon)$ -far from $\text{Range}(C)$. The overall procedure takes deterministic $\text{poly}(T) \leq \text{poly}(\ell)$ time with an NP oracle.

3.3 Variants of Our Frameworks

In this sub-section, we discuss two variants of our frameworks in Theorems 3.2 and 3.5.

Allowing FP^{NP} preprocessing in Theorem 3.5. Notice that in Theorem 3.2, we allow the non-trivial algorithm for SATISFYING-PAIRS to have a polynomial-time preprocessing with an NP oracle on the circuits before seeing the inputs. In our framework for remote point problems (see Theorem 3.5), we did not consider the preprocessing phase, as we do not need any preprocessing in our unconditional results, and the trade-off between parameters will be too complicated. We now argue informally that by modifying our framework as follows, it suffices to have algorithms for SATISFYING-PAIRS with an FP^{NP} preprocessing phase that generates a “short” data structure.

- Fix any integer d . Recall that Theorem 3.7 states that non-trivial algorithms for (approximate counting version, omitted below) $\text{AND}_d \circ \mathcal{C}$ -SATISFYING-PAIRS imply non-trivial algorithms for $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ -SATISFYING-PAIRS. By verifying the proof of Theorem 3.7, we can show that if the $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuit has fan-in A , the algorithm will call $\text{AND}_d \circ \mathcal{C}$ -SATISFYING-PAIRS for A^d times. Moreover, if the $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuit is fixed, then the circuit parts of these A^d instances for $\text{AND}_d \circ \mathcal{C}$ -SATISFYING-PAIRS are also fixed.
- In the proof of Theorem 3.5, we need to apply the non-trivial algorithm for $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ -SATISFYING-PAIRS on three types of circuits: circuits C^{norm} and C^{diff} in Lemma 3.8 (see Appendix A.2), and circuit C^{Prod} in Section 3.2.4. There are $O(2^{r_{\text{shared}}}) \leq O(\ell^{1/m})$ circuits in total, all of which are constructible given the input circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$. Moreover, the fan-in of these $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ -SATISFYING-PAIRS circuits are $O(\log \ell)$, and $d = O(1)$. Therefore we will need to call the $\text{AND}_d \circ \mathcal{C}$ -SATISFYING-PAIRS for $O(\ell^{1/m}) \cdot (\log \ell)^d = \ell^{o(1)}$ times, and the circuits are fixed given the input circuit C . (That is, the circuits do not depend on the input of the machine M^{PCPP} .)

- Now we assume that our non-trivial algorithm for $\text{AND}_d \circ \mathcal{C}$ -SATISFYING-PAIRS has an FP^{NP} preprocessing phase. Similar to the proof of Theorem 3.2, the final FP^{NP} algorithm for \mathcal{C} -REMOTE-POINT works as follows. Given any input circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, we construct the circuits in all $\ell^{o(1)}$ instances of $\text{AND}_d \circ \mathcal{C}$ -SATISFYING-PAIRS that will be needed in the proof, run the FP^{NP} preprocessing on these instances, and put the generated data structure into the advice of M^{PCPP} . It remains to verify that there are enough space to store these data structures in the advice. The non-deterministic time hierarchy theorem allows us to have $n_{\text{hard}}/10$ bits of advice (see Theorem 2.2), where $n_{\text{hard}} = 20H_{\text{proof}} \cdot n \geq \ell^{5q+1}$. The circuit C of size $s \cdot \ell$, which can be properly encoded using $\tilde{O}(s \cdot \ell) \leq \ell^2$ bits, needs to be put in the advice; therefore the rest $n_{\text{hard}}/10 - \ell^2$ bits in advice can be used to store the data structure generated in the preprocessing phase.

Algorithm for “weak” remote point. We can see that Theorem 3.5 requires the input circuit C to have quasi-polynomial stretch. This is because the hardness amplification procedure **Amp** (i.e. the XOR lemma with linear-sum decoder, see Section 2.4 and Appendix A.5) requires quasi-polynomial stretch. An interesting open problem is then to improve the amplification procedure to achieve a better stretch in Theorem 3.5.

The application of the hardness amplification procedure is to reduce the “strong” remote point (i.e. the distance is $1/2 - n^{\Omega(1)}$) to “weak” remote point (i.e. the distance is $\Omega(1)$). If we only want to have an FP^{NP} algorithm in the “weak” remote point regime, it suffices to follow the proof of Theorem 3.2 and replace the rectangular PCPP to the smooth and rectangular PCPP in Theorem 2.14, which can deal with \mathcal{C} -REMOTE-POINT with polynomial stretch. (The key is that smooth PCPP can tolerate errors in the proof, see Lemma 2.11 and Section 3.2.5).

4 Hard Partial Truth Tables

In this section, we employ the same technique to construct algorithms for PARTIAL-HARD and PARTIAL-AVGHARD from non-trivial SATISFYING-PAIRS algorithms. We only consider circuit classes that are both *typical* and *complete*. Recall that a circuit class \mathcal{C} is typical if it is closed under negations and projections. We say \mathcal{C} is *complete* if given a truth table tt of length 2^n , we can compute a \mathcal{C} circuit of size $2^{O(n)}$ whose truth table is tt in deterministic $2^{O(n)}$ time.

4.1 Hard Partial Truth Tables from SATISFYING-PAIRS

Instead of allowing a P^{NP} preprocessing on the circuits, the algorithm for SATISFYING-PAIRS used to solve PARTIAL-HARD allows a P^{NP} preprocessing on *inputs*, formally defined as follows.

Definition 4.1 (Algorithms for SATISFYING-PAIRS with P^{NP} Preprocessing on Inputs). Let P be one of the problems \mathcal{C} -SATISFYING-PAIRS, $\#\mathcal{C}$ -SATISFYING-PAIRS, $\text{Approx}_\delta\text{-}\mathcal{C}$ -SATISFYING-PAIRS, $\text{Gap}_\delta\text{-}\mathcal{C}$ -SATISFYING-PAIRS. A t -time algorithm for P with P^{NP} preprocessing of an ℓ -size data structure on inputs is a pair of algorithms (A_1, A_2) that solves P in two phases:

1. Given the inputs $x_1, x_2, \dots, x_M \in \{0, 1\}^n$, the polynomial-time algorithm A_1 with oracle access to a SAT oracle computes a string $\text{DS} \in \{0, 1\}^\ell$.
2. Given the circuits $C_1, C_2, \dots, C_N : \{0, 1\}^n \rightarrow \{0, 1\}$ of size s and the string DS , the algorithm A_2 solves P on the instance $(C_1, \dots, C_N, x_1, \dots, x_M)$ in time t .

Theorem 4.2. *There are constants $\varepsilon > 0$ and c_{imp} such that the following holds. Let $0 < \eta < 1/2$ be a constant, $\mathcal{C}[s]$ be a typical and complete circuit class where $s = s(n) > n$ is a size parameter, and $\mathcal{C}'[2s] := \text{OR}_2 \circ \mathcal{C}[s]$. Let $\ell(n)$ be a good function such that $s(n)^{1+\Omega(1)} \leq \ell(n) \leq 2^n$.*

Assumption: *Suppose that for some constant $c \geq 1$, there is an $(NM/\log^{c_{\text{imp}}}(NM))$ -time algorithm for $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ with $N := \ell^{c+1-\eta} \cdot \text{polylog}(\ell)$ circuits of size $\text{poly}(s(n))$ and $M := \ell^{1-\eta} \cdot \text{polylog}(\ell)$ inputs of length $2n$, allowing a \mathbf{P}^{NP} preprocessing of an M^c -size data structure on inputs.*

Conclusion: *There is an FP^{NP} algorithm for $\mathcal{C}[s]\text{-PARTIAL-HARD}$ with $\ell(n)$ input strings. More precisely, given a list of inputs $z_1, z_2, \dots, z_\ell \in \{0, 1\}^n$, we can compute a list of bits b_1, b_2, \dots, b_ℓ such that for every \mathcal{C} circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$ of size s , there exists an $i \in [\ell]$ such that $C(z_i) \neq b_i$.*

Proof Sketch of Theorem 4.2. The proof is very similar to the proof of Theorem 3.2; in fact, it is (nearly) equivalent to first reducing PARTIAL-HARD to AVOID and then invoking Theorem 3.2. Therefore we only highlight the differences.

It is without loss of generality to assume ℓ is a power of 2 and $c \geq 2$. We set the following parameters:

$$\begin{aligned} m &:= 5(c+2)/\eta = O(1), \\ w_{\text{proof}} &:= \log \ell, & W_{\text{proof}} &:= 2^{w_{\text{proof}}} = \ell, \\ h_{\text{proof}} &:= (c+1) \log \ell, & H_{\text{proof}} &:= 2^{h_{\text{proof}}} = \ell^{c+1}, \\ n_{\text{hard}} &:= 100H_{\text{proof}} \cdot s \log s, \\ T &:= H_{\text{proof}} \cdot W_{\text{proof}} / \log^{c_{\text{tm}}}(\ell), \\ h_{\text{input}} &:= \left(1 - \frac{\Theta(\log \log T)}{\log T}\right) h_{\text{proof}}, & H_{\text{input}} &:= 2^{h_{\text{input}}} = H_{\text{proof}} / \text{polylog}(\ell), \\ w_{\text{input}} &:= \lceil \log \tilde{n}_{\text{hard}} \rceil - h_{\text{input}}, & W_{\text{input}} &:= 2^{w_{\text{input}}} = s \log s \cdot \text{polylog}(\ell). \end{aligned}$$

Here $\tilde{n}_{\text{hard}} = O(n_{\text{hard}})$ is the codeword length of a length- n_{hard} string encoded via Enc where (Enc, Dec) is a fixed error-correcting code in Theorem 2.1; and c_{tm} is a sufficiently large constant.

We can check the technical condition $n_{\text{hard}}^{1+\Omega(1)} \leq T \leq 2^{\text{poly}(n_{\text{hard}})}$, so it is valid to invoke Theorem 2.2. Also, $(5/m) \log T \leq w_{\text{proof}}$, so it is valid to invoke the 2-query rectangular PCPP in Theorem 2.13. There are other checks for technical conditions that we omit here.

The first difference is the definition of “easy witness”. We say x has an easy witness if there is a proof matrix π (of size $H_{\text{proof}} \times W_{\text{proof}}$) for the statement “ $\text{Enc}(x) \in L^{\text{enc}}$ ” such that:

(completeness) for every $\text{seed} \in \{0, 1\}^r$, $\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed})$ accepts with probability at least c_{pcp} ;

(easiness) for every row π_j of π , there exists a size- s \mathcal{C} circuit $C_j : \{0, 1\}^n \rightarrow \{0, 1\}$ such that for every i , $\pi_{j,i} = C_j(z_i)$.

Then, our machine M^{PCPP} guesses H_{proof} size- s \mathcal{C} circuits $C_1, C_2, \dots, C_{H_{\text{proof}}} : \{0, 1\}^n \rightarrow \{0, 1\}$. Let π be the $H_{\text{proof}} \times W_{\text{proof}}$ proof matrix where for each $j \in [H_{\text{proof}}]$, $i \in [W_{\text{proof}}]$, $\pi_{j,i} = C_j(z_i)$. We need to estimate

$$p_{\text{acc}} := \Pr_{\text{seed} \leftarrow \{0, 1\}^r} [\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) \text{ accepts}].$$

We still reduce the problem of estimating p_{acc} to $2^{r_{\text{shared}}}$ instances of $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$, where $\varepsilon := (c_{\text{pcp}} - s_{\text{pcp}})/4$. However, now, each instance consists of $N := 2^{r_{\text{row}}} = 2^{h_{\text{proof}} - (5/m) \log T}$ circuits and $M := 2^{r_{\text{col}}} = 2^{w_{\text{proof}} - (5/m) \log T}$ inputs.¹⁸

¹⁸That is, the role of inputs and circuits are swapped as compared to Theorem 3.2.

We enumerate seed.shared . For each seed.shared , we create an $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ instance $\mathcal{I}_{\text{seed.shared}}$ corresponding to seed.shared , which contains an input $\text{Input}_{\text{seed.shared,seed.col}}$ for every seed.col and a circuit $C_{\text{seed.shared,seed.row}}$ for every seed.row . We elaborate on how this instance is constructed as this is different from Theorem 3.2.

Each seed.col corresponds to an input $\text{Input}_{\text{seed.shared,seed.col}}$ of the following form:

$$(a_1, \dots, a_q, pc_1^{\text{col}}, \dots, pc_p^{\text{col}}),$$

where $p + q \leq 2$, for each $i \in [q]$,

$$a_i := \begin{cases} \text{icol}[i] & \text{if } \text{itype}[i] = \text{input}; \\ z_{\text{icol}[i]} & \text{if } \text{itype}[i] = \text{proof}, \end{cases}$$

and pc_i^{col} represents the contribution of seed.col in the i -th parity-check bit.

The circuit $C_{\text{seed.shared,seed.row}}$ corresponding to seed.row is as follows:

- It receives input $(a_1, \dots, a_q, pc_1^{\text{col}}, \dots, pc_p^{\text{col}})$.
- For each $j \in [q]$, let

$$\text{ans}_j = \begin{cases} \text{Enc}(x)_{\text{irow}[j], a_j} & \text{if } \text{itype}[j] = \text{input} \\ C_{\text{irow}[j]}(a_j) & \text{if } \text{itype}[j] = \text{proof} \end{cases}.$$

Note that since \mathcal{C} is complete, we can compute a \mathcal{C} circuit of size $\text{poly}(W_{\text{input}}) = \text{poly}(s)$ whose truth table is the $\text{irow}[j]$ -th row of $\text{Enc}(x)$. That is, we can compute a \mathcal{C} circuit of size $\text{poly}(s)$ that on input a_j , outputs ans_j .

- For each $j \in [q]$, let pc_j^{row} be the contribution of seed.row in the j -th parity-check bit.
- It returns

$$\text{VDec}(\text{ans}_1, \dots, \text{ans}_q, pc_1^{\text{col}} \oplus pc_1^{\text{row}}, \dots, pc_p^{\text{col}} \oplus pc_p^{\text{row}}).$$

Here, VDec is the decision predicate of VPCPP , and is an OR_2 of its input bits or their negations. Since \mathcal{C} is typical, C is a $\text{OR}_2 \circ \mathcal{C}$ circuit. And one can easily verify that for each $\text{seed} = (\text{seed.shared}, \text{seed.row}, \text{seed.col})$, $C_{\text{seed.shared,seed.row}}(\text{Input}_{\text{seed.shared,seed.col}}) = 1$ if and only if $\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed})$ accepts. It follows that we can estimate p_{acc} by solving the instances $\mathcal{I}_{\text{seed.shared}}$ for every seed.shared .

To summarise, our algorithm M^{PCPP} works as follows. It first computes $\text{Enc}(x)$ and guesses $C_1, C_2, \dots, C_{H_{\text{proof}}}$. Then, it enumerates seed.shared , produces the instances $\mathcal{I}_{\text{seed.shared}}$, and feeds them to the algorithm for $\text{Approx}_\varepsilon\text{-}\mathcal{C}'\text{-SATISFYING-PAIRS}$ to obtain an estimation $p'_{\text{acc}}(\text{seed.shared})$. Let p'_{acc} be the average of $p'_{\text{acc}}(\text{seed.shared})$ over all $\text{seed.shared} \in \{0, 1\}^{r_{\text{shared}}}$.

We can still see that M^{PCPP} rejects every $x \notin L^{\text{hard}}$ and accepts every x with an easy witness. The machine M^{PCPP} runs in $T / \log^{\text{chard}} T$ time, guesses $H_{\text{proof}} \cdot 5s \log s < n_{\text{hard}}/10$ nondeterministic bits (since a size- s circuit can be encoded with at most $5s \log s$ bits), and uses at most $\ell^{c+1} < n_{\text{hard}}/10$ advice bits. By Theorem 2.2, M^{PCPP} cannot compute L^{hard} .

The hard partial truth table algorithm. Given a list of inputs $z_1, z_2, \dots, z_\ell \in \{0, 1\}^n$, our algorithm for finding a hard partial truth table $((z_1, b_1), (z_2, b_2), \dots, (z_\ell, b_\ell))$ works as follows. First, we construct the hard language L^{hard} and the algorithm M^{PCPP} . Let α be the advice string fed to M^{PCPP} and \mathcal{R} be the refuter in Theorem 2.2, we can use $\mathcal{R}(1^{n_{\text{hard}}}, M^{\text{PCPP}}, \alpha)$ to find an input

x_{hard} where M^{PCPP} fails on x_{hard} ; in particular, $M^{\text{PCPP}}(x_{\text{hard}}) = 0$ but $x_{\text{hard}} \in L^{\text{hard}}$. This takes deterministic $\text{poly}(H_{\text{proof}}) = \text{poly}(\ell)$ time with an NP oracle.

Then we find the lexicographically first proof matrix π such that $\text{VPCPP}^{\text{Enc}(x_{\text{hard}}) \circ \pi}$ accepts w.p. at least c_{pcp} , using the NP oracle. There has to be some $j \in [H_{\text{proof}}]$ such that for every size- s \mathcal{C} circuit C , there exists $i \in [W_{\text{proof}}]$ such that $C(z_i) \neq \pi_{j,i}$; moreover, the first such j can be found in $\text{poly}(H_{\text{proof}}) = \text{poly}(\ell)$ time with an NP oracle. We can pick

$$((z_1, \pi_{j,1}), (z_2, \pi_{j,2}), \dots, (z_{W_{\text{proof}}}, \pi_{j,W_{\text{proof}}}))$$

as the partial truth table hard for size- s \mathcal{C} circuits. \square

4.2 Average-Case Hard Partial Truth Tables

Theorem 4.3. *There is a universal constant $c_u \geq 1$ such that the following holds. Let $s = s(n) > n$ be a circuit size parameter, $N := N(n)$ be a parameter such that $2^{\log^{c_u} s} < N < 2^{s^{0.99}}$, $\varepsilon := \varepsilon(n) > s^{-c_u}$ be the error parameter, and $\ell := N^{c_u \log(1/\varepsilon)}$. Let $\mathcal{C}[s]$ be a typical and complete circuit class, and denote $\mathcal{C}^{\ell}[c_u s] := \text{AND}_{c_u} \circ \mathcal{C}[s]$ (i.e. a \mathcal{C}^{ℓ} circuit of size $c_u s$ refers to the AND of at most c_u \mathcal{C} circuits of size s).*

Assumption: *Let $P := (\log N)^{\log(1/\varepsilon)}$. Suppose there is a deterministic algorithm running in time $T^{\text{alg}} := N^2/P^{c_u}$ that, given as input a list of N $\mathcal{C}^{\ell}[c_u s]$ circuits $\{C_i\}$ and a list of N inputs $\{x_j\}$ with input length $n \cdot \text{polylog}(\ell)$, estimates $\Pr_{i,j \leftarrow [N]}[C_i(x_j)]$ with additive error $\eta := \varepsilon^{c_u}$.*

Conclusion: *There is an FP^{NP} algorithm for $\mathcal{C}[s]$ -PARTIAL-AVGHARD with $\ell(n)$ input strings. More precisely, given a list of inputs $w_1, w_2, \dots, w_{\ell} \in \{0, 1\}^n$, we can compute a list of bits $b_1, b_2, \dots, b_{\ell}$ such that for every \mathcal{C} circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$ of size s ,*

$$\Pr_{i \leftarrow [\ell]} [C(w_i) \neq b_i] \geq \frac{1}{2} - \varepsilon.$$

Proof Sketch of Theorem 4.3. The proof is similar to that of Theorem 3.5, so here we only highlight the difference. Roughly speaking, the main difference is that we swap the role of inputs and circuits.

For a circuit C and a list of inputs $w_1, w_2, \dots, w_{\ell}$, with slight abuse of notation, we define $C(w) := C(w_1) \circ C(w_2) \circ \dots \circ C(w_{\ell})$.

Analysing Prod \circ Sum circuits. Let $d \geq 1$ be a constant. We use $\text{Prod} \circ \text{Sum}$ to denote the class of multi-output circuits that take inputs $y \in \{0, 1\}^{\ell_y}$ and α , and has the following components:

- Let ℓ_{Sum} denote the number of middle “linear sum” gates. For each $i \in [\ell_{\text{Sum}}]$, the i -th gate outputs

$$\text{Sum}_i(y, \alpha) := \sum_{k \in [A]} \text{coeff}_k(\alpha) \cdot y_{\text{id}_{\times k}(\alpha, i)}.$$

- Let ℓ_{Prod} denote the number of output gates. The i -th output gate is a product gate of fan-in d , and is connected to the $q_1(i), q_2(i), \dots, q_d(i)$ -th linear sum circuits. Its output is

$$C_i^{\text{Prod}}(y, \alpha) := \prod_{t=1}^d \text{Sum}_{q_t(i)}(y, \alpha).$$

Remark 4.4. The important measures of a $\text{Prod}_d \circ \text{Sum}$ circuit are:

- the number of gates in each level ($\ell_{\text{Sum}}, \ell_{\text{Prod}}$);
- the fan-in of the top Prod gates (d);
- the fan-in (A), coefficient sum (U), and locality (l) of the linear sum layer.

It turns out that as an intermediate step, we need an algorithm provided by the following lemma.

Lemma 4.5. *Let \mathcal{C} be a typical circuit class, $M' \geq 1$ and $\eta \in (0, 1)$ be parameters. Suppose there is a deterministic algorithm running in time $T^{\text{alg}} = T^{\text{alg}}(N, M)$ that, given as input a list of $\hat{M} \leq M$ $\text{AND}_d \circ \mathcal{C}$ circuits $\{C_i\}$ and a list of $\hat{N} \leq N$ inputs $\{x_j\}$ of length $n \cdot \text{polylog}(\ell)$, estimates the following quantity with additive error η :*

$$\Pr_{i \leftarrow [\hat{M}], j \leftarrow [\hat{N}]} [C_i(x_j)].$$

Then, for any constant $\ell_{\mathcal{C}} > 0$, there is a deterministic algorithm running in time $A^d \cdot (\ell_{\mathcal{C}}^d + \ell_{\text{Prod}}/N) \cdot (2^{dl} + M'/M) \cdot O(T^{\text{alg}})$ that, given as input a $\text{Prod}_d \circ \text{Sum}$ circuit C^{Prod} with parameters specified in Remark 4.4, a list of ℓ_x strings $\{x_j\}$, a list of M' inputs $\{\alpha_j\}$, and a list of M' \mathcal{C} circuits $\{C_j\}$ from $\{0, 1\}^{|\mathcal{C}|}$ to $\{0, 1\}^{\ell_{\mathcal{C}}}$, estimates the following quantity with additive error $\eta \cdot U^d$:

$$\mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow [M']} [C_i^{\text{Prod}}(C_j(x), \alpha_j)].$$

Recall here that $C_j(x) = C_j(x_1) \circ C_j(x_2) \circ \dots \circ C_j(x_{\ell_x})$.

The proof is similar to that of Theorem 3.7 and is deferred to Appendix A.1.

Set up. We set the parameters as follows.

$$\begin{aligned} \delta &:= (10^9 q)^{-q^2}, \\ m &:= c_m \log(1/\varepsilon)/\delta, \\ w_{\text{proof}} &:= (60q/m) \log \ell, & W_{\text{proof}} &:= 2^{w_{\text{proof}}} = \ell^{O(\delta/\log(1/\varepsilon))}, \\ h_{\text{proof}} &:= (5q + 1) \log \ell, & H_{\text{proof}} &:= 2^{h_{\text{proof}}} = \ell^{5q+1}, \\ n_{\text{hard}} &:= 200H_{\text{proof}} \cdot s \log s, \\ T &:= H_{\text{proof}} \cdot W_{\text{proof}} / \log^{c_{\text{tm}}}(\ell). \\ h_{\text{input}} &:= \left(1 - \frac{\Theta(m^2 \log \log T)}{\log T}\right) h_{\text{proof}}, & H_{\text{input}} &:= 2^{h_{\text{input}}} = H_{\text{proof}} / \text{polylog}(\ell), \\ w_{\text{input}} &:= \lceil \log \tilde{n}_{\text{hard}} \rceil - h_{\text{input}}, & W_{\text{input}} &:= 2^{w_{\text{input}}} = s \log s \cdot \text{polylog}(\ell). \\ a &:= O(\log^2 W_{\text{proof}} / (\varepsilon \delta)^2) = O(\log^2 \ell / \varepsilon^2), \\ A &:= O(\log W_{\text{proof}} / (\varepsilon \delta)^2) = O(\log \ell / \varepsilon^2), \\ U &:= O(1/\varepsilon), \\ l &:= \log \ell. \end{aligned}$$

Here c_m and c_{tm} are sufficiently large constants, q is the query complexity of the smooth and rectangular PCPP in Theorem 2.14, and $\tilde{n}_{\text{hard}} = \Theta(n_{\text{hard}})$ is the length of $\text{Enc}(x)$ when the length of x is n_{hard} . Let \hat{H}_{proof} be the number of rows of the PCPP proof in Theorem 2.14, and let $\hat{h}_{\text{proof}} = \log \hat{H}_{\text{proof}}$. Also let $r, r_{\text{shared}}, r_{\text{col}}, r_{\text{row}}$ be the total, shared, column, row randomness in Theorem 2.14 respectively.

We use a different definition of “easy witness ” as follows. We say x has an easy witness if there is a proof matrix π such that:

(completeness) for every $\text{seed} \in \{0, 1\}^r$, $\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed})$ accepts;

(*approximate easiness*) for every row π_i of π , there exists a size- s \mathcal{C} circuit $C_i : \{0, 1\}^n \rightarrow \{0, 1\}$ and an advice $\alpha_i \in \{0, 1\}^a$ such that the decoding of the string $C_i(w)$ with advice α_i is δ -close to π_i with respect to ℓ_1 -norm. (Recall that $w = (w_1, w_2, \dots, w_\ell)$ is our input and $C(w)$ denotes the concatenation of $C(w_1), C(w_2), \dots, C(w_\ell)$.) In particular:

1. for every $j \in [W_{\text{proof}}]$, $(\text{dec}_{\alpha_i}(C_i(w)))_j \in [0, 1]$;
2. $\|\text{dec}_{\alpha_i}(C_i(w)) - \pi_i\|_1 \leq \delta$.

Our machine guesses \hat{H}_{proof} size- s \mathcal{C} circuits $C_1, C_2, \dots, C_{\hat{H}_{\text{proof}}} : \{0, 1\}^n \rightarrow \{0, 1\}$ as well as \hat{H}_{proof} advices $\alpha_1, \alpha_2, \dots, \alpha_{\hat{H}_{\text{proof}}}$. Let $\pi_i^{\text{Real}} := \text{dec}_{\alpha_i}(C_i(w))$, and π_i^{Bool} be the Boolean string that is closest to π_i^{Real} . For $\iota \in [q]$, we define

$$\begin{aligned} f_{\text{seed.shared}, \iota}^{\text{Bool}}(\text{seed.row}, \text{seed.col}) &= \pi_{\text{irow}[\iota], \text{icol}[\iota]}^{\text{Bool}} \quad \text{and} \\ f_{\text{seed.shared}, \iota}^{\text{Real}}(\text{seed.row}, \text{seed.col}) &= \pi_{\text{irow}[\iota], \text{icol}[\iota]}^{\text{Real}}. \end{aligned}$$

Verifying closeness of π^{Bool} and π^{Real} . The next lemma shows that we can verify whether a Boolean proof π^{Bool} and a real proof π^{Real} are close. The proof is deferred to Appendix A.3.

Lemma 4.6. *Let \mathcal{C} be a typical circuit class and $d \geq 2$ be an even number. Suppose there is an algorithm that takes as inputs a list of $2^{r_{\text{col}}}$ $\text{AND}_{2d} \circ \mathcal{C}$ circuits $\{C_i\}$ and a list of $2^{r_{\text{col}}}$ inputs $\{x_j\}$ of length $n \cdot \text{polylog}(\ell)$, runs in deterministic T^{alg} time, and estimates the following quantity with additive error η :*

$$\Pr_{i, j \leftarrow [2^{r_{\text{col}}}]}[C_i(x_j)].$$

Then there is an algorithm that takes the strings w_1, w_2, \dots, w_ℓ , circuits $(C_1, C_2, \dots, C_{\hat{H}_{\text{proof}}})$, and $(\alpha_1, \alpha_2, \dots, \alpha_{\hat{H}_{\text{proof}}})$ as inputs, runs in deterministic $O((3A)^{2d} T^{\text{alg}}) \cdot (2^{2dl+r_{\text{shared}}} + T \log^{O(m)} T / 2^{2r_{\text{col}}})$ time, and satisfies the following:

(Completeness) *If for every $i \in [\hat{H}_{\text{proof}}]$, it holds that (1) for every $j \in [W_{\text{proof}}]$, $\pi_{i,j}^{\text{Real}} \in [0, 1]$; (2) $\|\pi_i^{\text{Real}} - \pi_i^{\text{Bool}}\|_1 \leq \delta$, then the algorithm accepts.*

(Soundness) *If the algorithm accepts, then it holds that*

1. for every $\text{seed.shared} \in \{0, 1\}^{r_{\text{shared}}}$ and $\iota \in [q]$, $\|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d \leq 1 + 2\eta \cdot U^d$;
2. $\mathbb{E}_{i \leftarrow [\hat{H}_{\text{proof}}], j \leftarrow [W_{\text{proof}}]} \left[|\pi_{i,j}^{\text{Real}} - \pi_{i,j}^{\text{Bool}}|^d \right] \leq 4^d \cdot \delta + 2^{d+1} \eta (2U + 1)^{2d}$.

Estimating p_{acc} . Now we verified that π^{Real} is close to π^{Bool} using Lemma 4.6, with parameter $d = 2q$. After that, the next step is to use it to speed up L^{hard} . We estimate

$$p_{\text{acc}} := \Pr_{\text{seed} \leftarrow \{0, 1\}^r} \left[\text{VPCPP}^{\text{Enc}(x) \circ \pi^{\text{Bool}}}(\text{seed}) \text{ accepts} \right].$$

Actually, it suffices to distinguish between the case that $p_{\text{acc}} > 5/6$ and the case that $p_{\text{acc}} < 1/2$.

We still enumerate seed.shared , and we now need to estimate

$$p_{\text{acc}}(\text{seed.shared}) := \Pr_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\text{VPCPP}^{\text{Enc}(x) \circ \pi^{\text{Bool}}}(\text{seed}) \text{ accepts} \right].$$

Let $pc_1, pc_2, \dots, pc_q \leftarrow V_{\text{pc}}(\text{seed.shared})$ be the parity-check bits of the PCPP verifier, and let pc_ℓ^{row} (resp. pc_ℓ^{col}) denote the contribution of seed.row (resp. seed.col) to pc_ℓ , then $pc_\ell = pc_\ell^{\text{row}} \oplus pc_\ell^{\text{col}}$.

As in the proof of Theorem 3.5, here it suffices to estimate for every $S, S' \subseteq [q]$

$$p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S') = \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\prod_{\ell \in S} a_\ell^{\text{Real}} \cdot \prod_{\ell \in S'} pc_\ell \right],$$

where

$$a_\ell^{\text{Real}} := \begin{cases} \tilde{x}_{\text{irow}[\ell], \text{icol}[\ell]} & \text{itype}[\ell] = \text{input}, \\ \pi_{\text{irow}[\ell], \text{icol}[\ell]}^{\text{Real}} & \text{itype}[\ell] = \text{proof}. \end{cases}$$

We want to invoke Lemma 4.5 to estimate this, so we want to construct a $2^{r_{\text{col}}}$ -output $\text{Prod} \circ \text{Sum}$ circuit C^{Prod} , $2^{r_{\text{row}}}$ circuits $\{C_{\text{seed.row}}\}$ and $2^{r_{\text{row}}}$ strings $\{\alpha_{\text{seed.row}}\}$ such that

$$\begin{aligned} & C_{\text{seed.col}}^{\text{Prod}}(C_{\text{seed.row}}(w), \alpha_{\text{seed.row}}) \\ &= \prod_{\ell \in S} a_\ell^{\text{Real}} \cdot \prod_{\ell \in S'} pc_\ell \\ &= \prod_{\ell \in S^{\text{proof}}} \left(\sum_{k \in [A]} \text{coeff}_k(\alpha_{\text{irow}[\ell]}) \cdot C_{\text{irow}[\ell]}(w_{\text{id} \times_k(\alpha_{\text{irow}[\ell]}, \text{icol}[\ell])}) \right) \cdot \prod_{\ell \in S^{\text{input}}} \tilde{x}_{\text{irow}[\ell], \text{icol}[\ell]} \cdot \prod_{\ell \in S'} (pc_\ell^{\text{row}} \oplus pc_\ell^{\text{col}}) \end{aligned} \quad (6)$$

where $S^{\text{proof}} = \{\ell \in S : \text{itype}[\ell] = \text{proof}\}$ and $S^{\text{input}} = \{\ell \in S : \text{itype}[\ell] = \text{input}\}$. This motivates the following definitions.

For $i \in [\ell]$, let $z_i \in \{0,1\}^{n+w_{\text{input}}}$ be the string such that the first n bits of z_i is w_i , and the last w_{input} bit is

$$\begin{cases} i & i \in [W_{\text{input}}], \\ 1 & i \notin [W_{\text{input}}]. \end{cases}$$

Here we identify $[W_{\text{input}}]$ with $\{0,1\}^{w_{\text{input}}}$.

For any string v , define $\hat{C}_v : \{0,1\}^{\lceil \log |v| \rceil} \rightarrow \{0,1\}$ as the circuit that on input $i \leq |v|$, outputs v_i . Since \mathcal{C} is complete, \hat{C}_v is an efficiently computable \mathcal{C} circuit of size $\text{poly}(|v|)$. Let $\text{proj}_{\text{input}} : \{0,1\}^{n+w_{\text{input}}} \rightarrow \{0,1\}^{w_{\text{input}}}$ be the circuit that outputs the last w_{input} bits of its input, and let $\text{proj}_{\text{proof}} : \{0,1\}^{n+w_{\text{input}}} \rightarrow \{0,1\}^n$ be the circuit that outputs the first n bits of its input.

Now for fixed seed.row , define

$$C_\ell^a := \begin{cases} \hat{C}_{\tilde{x}_{\text{irow}[\ell]}} \circ \text{proj}_{\text{input}} & \text{itype}[\ell] = \text{input} \\ C_{\text{irow}[\ell]} \circ \text{proj}_{\text{proof}} & \text{itype}[\ell] = \text{proof} \end{cases}$$

for $\ell \in S$, and C_ℓ^{pc} be the circuit that outputs pc_ℓ^{row} for $\ell \in S'$, regardless of its input. Let $C_{\text{seed.row}} := (C_1^a, C_2^a, \dots, C_{d_S}^a, C_1^{\text{pc}}, C_2^{\text{pc}}, \dots, C_{d_{S'}}^{\text{pc}})$ where $d_S := |S|$ and $d_{S'} := |S'|$, that is, $C_{\text{seed.row}}$ is a circuit with $d_S + d_{S'}$ outputs and each of its outputs is a circuit C_ℓ^a or C_ℓ^{pc} .

Now we define the $\text{Prod} \circ \text{Sum}$ circuit C^{Prod} .

Circuit C^{Prod}

(Inputs) The input y has the form $y = y_{\text{seed.row}} = (y_1, y_2, \dots, y_\ell)$ and the input $\hat{\alpha}$ has the form $\hat{\alpha} = \hat{\alpha}_{\text{seed.row}} = (\hat{\alpha}_1, \hat{\alpha}_2, \dots, \hat{\alpha}_{d_S})$. The intended meanings are $y_i = C_{\text{seed.row}}(z_i)$, and $\hat{\alpha}_i = \alpha_{\text{row}[i]}$.

For convenience, we will use the following labels to refer to bits of y , assuming the intended meaning above:

- For $j \in S^{\text{proof}}$, $i \in [\ell]$, let $(y_i)_j := C_j^a(z_i) = C_{\text{row}[j]}(w_i)$;
- For $j \in S^{\text{input}}$, $i \in [W_{\text{input}}]$, let $(y_i)_j := C_j^a(z_i) = \hat{C}_{\tilde{x}_{\text{row}[j]}}(i) = \tilde{x}_{\text{row}[j],i}$;
- For $j \in S'$, let $(y_i)_{j+d_S} := C_j^{pc}(z_i) = pc_j^{\text{row}}$.

(Linear sum gates) There are $\ell_{\text{Sum}} := W_{\text{proof}} \cdot d_S + 2d_{S'}$ linear sum gates and we identify $[\ell_{\text{Sum}}]$ with the disjoint union of $[W_{\text{proof}}] \times S$ and $S' \times \{0, 1\}$.

Let $i \in [W_{\text{proof}}]$ and $j \in S$. If $\text{itype}[j] = \text{proof}$, then the (i, j) -th linear sum gate is

$$\text{Sum}_{(i,j)}(y, \alpha) = \sum_{k \in [A]} \text{coeff}_k(\alpha_j) \cdot (y_{\text{id}_{X_k}(\alpha_j, i)})_j.$$

It is easy to verify that

$$\text{Sum}_{(i,j)}(y, \alpha) = (\text{dec}_{\alpha_{\text{row}[j]}}(C_{\text{row}[j]}(w)))_i.$$

On the other hand, if $\text{itype}[j] = \text{input}$, then the (i, j) -th linear sum gate is $\text{Sum}_{(i,j)}(y, \alpha) := (y_i)_j$. (If $i > W_{\text{input}}$ then we simply set $\text{Sum}_{(i,j)}(y, \alpha) = 0$ and this gate would not be used.)

Finally, for each $j \in S'$, we have two intermediate gates

$$\text{Sum}_{(j,0)}(y, \alpha) = (y_1)_{j+d_S}, \quad \text{Sum}_{(j,1)}(y, \alpha) = 1 - (y_1)_{j+d_S}.$$

Implementation of the linear sum layer: The linear sum has fan-in $A' := A \cdot d_S + 2$ and we identify $[A']$ with the disjoint union of $[A] \times S$ and $\{+, -\}$. Also, the length of y is $\ell_y := \ell \cdot (d_S + d_{S'})$, and we identify $[\ell_y]$ with $[\ell] \times (S \cup S')$. Let $\text{id}_{X'}$ and coeff' be the $\text{id}_{X'}$ and coeff' functions of the linear sum layer of C^{Prod} , then

(Function $\text{id}_{X'_k}(\alpha, i)$) Suppose $i = (i', j) \in [W_{\text{proof}}] \times S$. If $\text{itype}[j] = \text{proof}$ and $k = (k', j')$ where $j = j'$, then we return $\text{id}_{X'_k}(\alpha, i) = (\text{id}_{X_{k'}}(\alpha_j, i'), j)$; if $\text{itype}[j] = \text{input}$ and $i' \in [W_{\text{input}}]$ and $k = +$, then $\text{id}_{X'_k}(\alpha, i) = (i', j)$. Otherwise $\text{id}_{X'_k}(\alpha, i) = \text{ZERO}$.

On the other hand, suppose $i = (j, b) \in S' \times \{0, 1\}$. If $(b = 0 \text{ and } k = +)$ or $(b = 1 \text{ and } k = -)$ then $\text{id}_{X'_k}(\alpha, i) = (1, j)$. If $b = 1$ and $k = +$ then $\text{id}_{X'_k}(\alpha, i) = \text{ONE}$. Otherwise $\text{id}_{X'_k}(\alpha, i) = \text{ZERO}$.

(Function $\text{coeff}'_k(\alpha)$) If $k = +$ then $\text{coeff}'_k(\alpha) = 1$; if $k = -$ then $\text{coeff}'_k(\alpha) = -1$; otherwise, if $k = (k', j')$ then $\text{coeff}'_k(\alpha) = \text{coeff}_{k'}(\alpha_{j'})$.

The locality of $(\text{id}_{X'}, \text{coeff}')$ is still l . The coefficient sum becomes $d_S \cdot U + 2$.

(Output product gates) There are $2^{r_{\text{col}}}$ product gates. For each seed.col , the seed.col -th output gate is

$$C_{\text{seed.col}}^{\text{Prod}}(y, \alpha) = \prod_{j \in S} \text{Sum}_{(\text{icol}[j], j)}(y, \alpha) \cdot \prod_{j \in S'} \text{Sum}_{(j, pc_j^{\text{col}})}(y, \alpha).$$

To summarise, the parameters of the circuit C^{Prod} are as follows.

- The number of gates in each layer: $\ell_{\text{Sum}} = W_{\text{proof}} \cdot d_S + 2d_{S'}$, $\ell_{\text{Prod}} = 2^{r_{\text{col}}}$.
- The length of input y : $\ell_y = \ell(d_S + d_{S'})$;
- The fan-in of the top Prod gates: $d_S + d_{S'} \leq 2q$.
- The fan-in $A' := A \cdot d_S + 2$, coefficient sum $d_S \cdot U + 2$, and locality l of the linear sum layer.

Given the above construction, it is easy to see that Eq. (6) holds for every `seed.row` and `seed.col`. We can thus see that

$$p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S') = \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[C_{\text{seed.col}}^{\text{Prod}}(C_{\text{seed.row}}(w), \alpha_{\text{seed.row}}) \right].$$

Since $d_S \leq q$, $d_{S'} \leq q$ and $\ell_{\mathcal{C}} \leq d_S + d_{S'} \leq 2q$, by Lemma 4.5, we can estimate $p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S')$ with additive error $\eta \cdot (qU + 2)^{2q}$ in deterministic time

$$(A \cdot d_S + 2)^{2q} \cdot ((2q)^{2q} + 2^{r_{\text{col}}}/N) \cdot (2^{2ql} + 2^{r_{\text{row}}}/N) \cdot O(T^{\text{alg}})$$

Analysis. First, the verification step takes

$$\begin{aligned} & O((3A)^{4q} T^{\text{alg}}) \cdot \left(2^{4ql+r_{\text{shared}}} + T \log^{O(m)} T / 2^{2r_{\text{col}}} \right) \\ & \leq O((3A)^{4q}) \cdot (T \log^{O(m)} T) / (r_{\text{col}})^{c_u \log(1/\varepsilon)} \\ & \leq T (\log T)^{O(m) - c_u \log(1/\varepsilon)/2} \end{aligned}$$

time, which is at most $T/(4 \log^{c_{\text{hard}}} T)$ if c_u is a large enough constant.

Then, the algorithm of Lemma 4.5 on C^{Prod} runs for every `seed.shared`, S, S' , and in total takes time

$$\begin{aligned} & (A \cdot d_S + 2)^{2q} \cdot ((2q)^{2q} + 2^{r_{\text{col}}}/N) \cdot (2^{2ql} + 2^{r_{\text{row}}}/N) \cdot O(T^{\text{alg}}) \cdot 2^{2q} \cdot 2^{r_{\text{shared}}} \\ & \leq O(\log^{2q} \ell / \varepsilon^{4q}) \cdot O(2^{r_{\text{col}}}/N) \cdot O(2^{r_{\text{row}}}/N) \cdot O(N^2 / \log^{c_u \log(1/\varepsilon)} N) \cdot 2^{r_{\text{shared}}} \\ & \leq 2^r / \log^{\Omega(c_u)} \ell < T / (4 \log^{c_{\text{hard}}} T), \end{aligned}$$

when c_u is sufficiently large. Therefore, the whole algorithm runs in $T / \log^{c_{\text{hard}}} T$ time.

Besides, by the same argument as in the proof of Theorem 3.5, which we omit here, the algorithm estimates p_{acc} within additive error at most

$$\eta(4qU + 8)^{2q} + 16^q \cdot 100^{-q} < 1/6,$$

thus successfully distinguishes between the case that $p_{\text{acc}} > 5/6$ and that $p_{\text{acc}} < 1/2$.

Description of M^{PCPP} We summarise the algorithm M^{PCPP} . On input x , we consider the smooth and rectangular PCPP for the language $L^{\text{enc}} = \{\text{Enc}(x) : x \in L^{\text{hard}}\}$. (Recall that M^{PCPP} aims to reject every $x \notin L$ and accepts every $x \in L$ with easy witness.) We guess $(C_1, \dots, C_{H_{\text{proof}}})$ and $(\alpha_1, \dots, \alpha_{H_{\text{proof}}})$, which implicitly defines the PCPP proof matrices π^{Bool} and π^{Real} . Then we verify π^{Real} using Lemma 3.8 and reject immediately if π^{Real} did not pass the test. If π^{Real} passes the test (which means that it is “close” to a Boolean proof π^{Bool}), we use the algorithm described above to estimate p_{acc} . We accept x if and only if our estimation is above $2/3$.

The correctness of M^{PCPP} is easy to see (and is exactly the same as Claim 3.11):

Claim 4.7. *For every input x , if $x \notin L$ then M^{PCPP} rejects x ; while if $x \in L$ and x has an easy witness then M^{PCPP} accepts x .*

The machine M^{PCPP} guesses $H_{\text{proof}}(5s \log s + a) < n_{\text{hard}}/10$ bits of nondeterminism (the number of size- s \mathcal{C} circuits is at most $2^{5s \log s}$), and uses $\ell^{4q} < n_{\text{hard}}/10$ bits of advice. Thus it computes a language in $\text{NTIMEGUESS}_{\text{RTM}}[T / \log^{c_{\text{hard}}}(T), n_{\text{hard}}/10]_{(n_{\text{hard}}/10)}$.

The FP^{NP} algorithm for average-case hard partial truth table. Let $w_1, w_2, \dots, w_\ell \in \{0, 1\}^n$ be the input. We first construct the hard language L^{hard} and the algorithm M^{PCPP} . Since M^{PCPP} is a nondeterministic RAM algorithm that runs in $T/\log^{c_{\text{hard}}}(T)$ time, uses at most $n_{\text{hard}}/10$ nondeterministic bits and at most $n_{\text{hard}}/10$ advice bits, it follows that there is an input $x_{\text{hard}} \in \{0, 1\}^{n_{\text{hard}}}$ such that $M^{\text{PCPP}}(x_{\text{hard}}) \neq L^{\text{hard}}(x_{\text{hard}})$. Moreover, let α be the advice string fed to M^{PCPP} , i.e., the circuit C . We can find such an input x_{hard} by running $\mathcal{R}(1^{n_{\text{hard}}}, M^{\text{PCPP}}, \alpha)$, where \mathcal{R} is the refuter guaranteed by Theorem 2.2. Thus, we can find x_{hard} in deterministic $\text{poly}(T)$ time with an NP oracle.

It follows from Claim 4.7 that $x_{\text{hard}} \in L^{\text{hard}}$ but x_{hard} does not have an easy witness. Thus, we can use the NP oracle to find the lexicographically first PCPP proof matrix π such that

$$\Pr_{\text{seed} \leftarrow \{0,1\}^r} [\text{VPCPP}^{\text{Enc}(x) \circ \pi}(\text{seed}) \text{ accepts}] = 1.$$

Then, there must exist a row π_i such that $\text{Amp}(\pi_i)$ is $(1/2 - \varepsilon)$ -far from $C(w) = C(w_1) \circ \dots \circ C(w_\ell)$ for any size- s \mathcal{C} circuit C . To see this, suppose that for every i , there exists a size- s \mathcal{C} circuit C_i such that $\text{Amp}(\pi_i)$ is $(1/2 - \varepsilon)$ -close to $C_i(w)$. By Theorem 2.5, there is an advice α_i such that $\text{dec}_{\alpha_i}(C_i(w))$ satisfies (1) for every $j \in [W_{\text{proof}}]$, $(\text{dec}_{\alpha_i}(C_i(w)))_j \in [0, 1]$; (2) $\|\text{dec}_{\alpha_i}(C_i(w)) - \pi_i\|_1 \leq \delta$. It follows that π is an easy witness for x_{hard} , a contradiction.

Finally, we use the NP oracle to find the first row π_i , such that $\text{Amp}(\pi_i)$ is $(1/2 - \varepsilon)$ -far from $C(w)$ for any size- s \mathcal{C} circuit C , and output $\text{Amp}(\pi_i)$. The overall procedure takes deterministic $\text{poly}(T) \leq \text{poly}(\ell)$ time with an NP oracle. \square

5 Unconditional Algorithms for Range Avoidance, Remote Point, and Hard Partial Truth Tables

In this section, we apply the frameworks in Section 3 and Section 4 to obtain *unconditional* results for XOR-REMOTE-POINT, ACC⁰-REMOTE-POINT, and ACC⁰-PARTIAL-AVGHARD.

5.1 Algorithms for Satisfying Pairs

In this section, we present some algorithms for \mathcal{C} -SATISFYING-PAIRS for various circuit classes \mathcal{C} that run in non-trivial time.

The algorithms in this section require the following algorithm for the batch evaluation of low-degree polynomials via fast rectangular matrix multiplication. This algorithm has been extensively used in previous works on the polynomial method and circuit complexity (see, e.g., [Wil14a, Wil18a]). We provide a proof sketch for completeness.

Theorem 5.1. *Let $x_1, x_2, \dots, x_N \in \{0, 1\}^n$ be N input strings, and $p_1, p_2, \dots, p_N : \{0, 1\}^n \rightarrow \mathbb{N}$ be N integer polynomials of degree at most d . Suppose that $n^{20d} \leq N$. Then there is a deterministic algorithm running in $\tilde{O}(N^2)$ time that outputs the table of $p_j(x_i)$ for every $i, j \in [N]$.*

Theorem 5.2 ([Cop82]; see also [Wil18a, Appendix C]). *There is a (deterministic) algorithm for multiplying an $N \times N^{0.1}$ matrix and an $N^{0.1} \times N$ matrix using $\tilde{O}(N^2)$ arithmetic operations.*

Proof of Theorem 5.1. There are $m := \sum_{i=0}^d \binom{n}{i} \leq (en/d)^d \leq N^{0.1}$ monomials of degree at most d . We number these monomials from 1 to m . Let S_j denote the set of indices in the j -th monomial. That is, the j -th monomial is $\prod_{k \in S_j} x_k$.

We construct two matrices $M_1 \in \mathbb{Z}^{N \times m}$ and $M_2 \in \mathbb{Z}^{m \times N}$. For each $i \in [N]$ and $j \in [m]$, $M_1[i, j]$ is the evaluation of the j -th monomial on input x_i . (That is, $M_1[i, j] = \prod_{k \in S_j} (x_i)_k$.) For each $j \in [m]$ and $k \in [N]$, $M_2[j, k]$ is the coefficient of the j -th monomial in p_k .

Let $M := M_1 \cdot M_2$. It follows that for every $i, j \in [N]$, $M[i, j] = p_j(x_i)$. Since $m \leq N^{0.1}$, we can compute M in $\tilde{O}(N^2)$ time using Theorem 5.2. \square

5.1.1 An Algorithm for #XOR-SATISFYING-PAIRS

We start by demonstrating a non-trivial algorithm for #XOR-SATISFYING-PAIRS, i.e., the exact counting version of SATISFYING-PAIRS problem where each circuit consists of a single XOR gate (with unbounded fan-in). We observe that #XOR-SATISFYING-PAIRS is essentially \mathbb{F}_2 -#OV (i.e., the counting version of Orthogonal Vectors problem over \mathbb{F}_2), which has been studied in the literature before (see, e.g., [CW21, AC19]).

In more details, an unbounded fan-in XOR gate over n input bits can be represented by a vector $y \in \{0, 1\}^n$ and a bit $b \in \{0, 1\}$ such that on input $x \in \{0, 1\}^n$, the gate outputs 1 if $\langle x, y \rangle = b$ and outputs 0 otherwise. Suppose that we have N inputs $x_1, x_2, \dots, x_N \in \{0, 1\}^n$ and N XOR gates $(y_1, b_1), \dots, (y_N, b_N)$. Let $x'_i \in \{0, 1\}^{n+1}$ be the vector that is the concatenation of x_i and the bit 1, and let $y'_j \in \{0, 1\}^{n+1}$ be the vector that is the concatenation of y_j and b_j . Then the gate (y_j, b_j) accepts the input x_i if and only if $\langle x'_i, y'_j \rangle = 0$. Hence we can reduce #XOR-SATISFYING-PAIRS to the counting version of \mathbb{F}_2 -OV, which has a non-trivial algorithm [AC19, CW21].

Theorem 5.3 ([AC19, Theorem 2.9]). *Let $n \leq N^{o(1)}$. There is a deterministic algorithm running in $N^{2-\Omega(1/\log(n/\log N))}$ time that given $2N$ length- n vectors $x_1, x_2, \dots, x_N, y_1, y_2, \dots, y_N \in \mathbb{F}_2^n$, outputs the number of pairs (i, j) such that $\langle x_i, y_i \rangle = 0$.*

5.1.2 An Algorithm for #ACC⁰-SATISFYING-PAIRS

We present a non-trivial algorithm for #ACC⁰-SATISFYING-PAIRS, generalising the algorithm for XOR-circuits (since $\text{XOR} \subseteq \text{AC}^0[2] \subseteq \text{ACC}^0$). This algorithm utilises a quasi-polynomial simulation of $\text{SYM} \circ \text{ACC}^0$ circuits with $\text{SYM} \circ \text{AND}$ circuits.

Theorem 5.4 (From $\text{SYM} \circ \text{ACC}^0$ to $\text{SYM} \circ \text{AND}$ [BT94, AG91, Wil18c]). *Let m, ℓ be any constants, there exists an integer c' such that every $\text{SYM} \circ \text{AC}^0_\ell[m]$ circuit of size s can be simulated by a $\text{SYM} \circ \text{AND}$ circuit of $2^{(\log s)^{c'}}$ size. Moreover, the AND gates of the final circuit have only $(\log s)^{c'}$ fan-in, the final circuit can be constructed from the original one in $2^{O((\log s)^{c'})}$ time, and the final symmetric function at the output can be computed in $2^{O((\log s)^{c'})}$ time.*

Combining Theorem 5.4 with Theorem 5.1, we can derive the #ACC⁰-SATISFYING-PAIRS algorithm in non-trivial time as follows.

Theorem 1.16. *For every constants m, ℓ, c , there is a constant $\varepsilon \in (0, 1)$ such that the following holds. Let $n := 2^{\log^\varepsilon N}$ and $s := 2^{\log^c n}$. There is a deterministic algorithm running in $\tilde{O}((N/n)^2)$ time that given N strings $x_1, x_2, \dots, x_N \in \{0, 1\}^n$ and N $\text{AC}^0_\ell[m]$ circuits $C_1, C_2, \dots, C_N : \{0, 1\}^n \rightarrow \{0, 1\}$ of size s , outputs the number of pairs $(i, j) \in [N] \times [N]$ such that $C_i(x_j) = 1$.*

Proof. Let ε be a constant to be determined. We divide C_1, C_2, \dots, C_N into N/n groups where each group has size n . Let C_{ij} denote the j -th circuit in the i -th group. We also partition the inputs x_1, x_2, \dots, x_N into N/n groups of size n and define x_{ij} similarly. Let $X_i := x_{i1} \circ x_{i2} \circ \dots \circ x_{in}$.

For each group i , we can construct $g := \lceil 2 \log n \rceil$ $\text{SYM} \circ \text{AC}_\ell^0[m]$ circuits $D_{i1}, D_{i2}, \dots, D_{ig} : \{0, 1\}^{n^2} \rightarrow \{0, 1\}$, each of size $s' := O(n^2 \cdot s)$, such that for any group j , we have:

$$\sum_{i'=1}^n \sum_{j'=1}^n C_{ii'}(x_{jj'}) = \sum_{k=0}^g 2^k D_{ik}(X_j).$$

That is, $D_{ik}(X_j)$ computes the k -th bit of the number of satisfying pairs between the i -th group of circuits and the j -th group of inputs.

Let c' be the constant in Theorem 5.4 depending on ℓ and m . We can transform each $\text{SYM} \circ \text{AC}_\ell^0[m]$ circuit D_{ij} into a $\text{SYM} \circ \text{AND}$ circuit D'_{ij} of size $2^{(\log s')^{c'}}$ such that each AND gate has fan-in at most $d := (\log s')^{c'}$. We can write each $D'_{ij}(x)$ as $f_{ij}(p_{ij}(x))$, where $p_{ij}(x) : \{0, 1\}^{n^2} \rightarrow \{0, 1, \dots, 2^{(\log s')^{c'}}\}$ is a polynomial of degree at most d that only outputs integers upper bounded by $2^{(\log s')^{c'}}$ on Boolean inputs, and f_{ij} is some function that can be evaluated in $2^{O((\log s')^{c'})}$ time. We can construct the polynomials p_{ij} and (the truth tables of) the functions f_{ij} in $(N/n)g2^{O((\log s')^{c'})}$ time. Let $\varepsilon := 1/(10cc')$ (and recall $n = 2^{\log^\varepsilon N}$ and $s = 2^{\log^c n}$), this time bound is at most $(N/n)^2$.

Then, for each $k = 1, 2, \dots, g$, since $(n^2)^{20d} \leq N/n$, we can compute the table of $p_{ik}(X_j)$ for every $i, j \in [N/n]$ in $\tilde{O}((N/n)^2)$ time by invoking Theorem 5.1. In fact, by checking the truth-tables of f_{ij} , we actually get the table for $D'_{ik}(X_j) = D_{ik}(X_j)$. Finally it follows that:

$$\sum_{i=1}^N \sum_{j=1}^N C_i(x_j) = \sum_{i=1}^{N/n} \sum_{j=1}^{N/n} \sum_{i'=1}^n \sum_{j'=1}^n C_{ii'}(x_{jj'}) = \sum_{i=1}^{N/n} \sum_{j=1}^{N/n} \sum_{k=0}^g 2^k D_{ik}(X_j).$$

The total run-time is bounded by $2(N/n)^2 + g\tilde{O}((N/n)^2) = \tilde{O}((N/n)^2)$. \square

5.2 An FP^{NP} Algorithm for XOR-REMOTE-POINT

As a warm-up, we show an FP^{NP} algorithm for XOR-REMOTE-POINT. This result is subsumed by the algorithm for ACC-REMOTE-POINT (Theorem 1.17) as well as previous results on rigid matrices [AC19, BHPT20, CLW20, CL21, HV21], but we still mention it for several reasons:

1. This is an important special case for REMOTE-POINT.
2. This result only requires the algorithm for $\#\mathbb{F}_2\text{-OV}$, and is a nice illustration that many special cases of SATISFYING-PAIRS problems are widely studied by algorithm designers.

Theorem 1.15 (XOR-REMOTE-POINT $\in \text{FP}^{\text{NP}}$). *There is a constant $c_u \geq 1$ such that the following holds. Let $\varepsilon := \varepsilon(n) \geq 2n^{-c_u}$ be the error parameter and $\ell := \ell(n) \geq 2^{\log^{c_u+5} n}$ be the stretch function, then there is an FP^{NP} algorithm that takes as input a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, where each output bit of C is computed by an XOR gate, and outputs a string y that is $(1/2 - \varepsilon)$ -far from $\text{Range}(C)$.*

Proof. Let c_u be the constant in Theorem 3.5. In order to make XOR a typical circuit class (closed under negation and projection), we here allow XOR using bias, denoted by $\widetilde{\text{XOR}}$. That is, an $\widetilde{\text{XOR}}$ gate g can be represented by parameters $w \in \{0, 1\}^n$ and $b \in \{0, 1\}$, such that for any $x \in \{0, 1\}^n$, we have $g(x) := \langle w, x \rangle \oplus b$. To obtain an FP^{NP} algorithm of XOR-REMOTE-POINT by Theorem 3.5, we need to design an efficient algorithm for $\text{AND}_{c_u} \circ \widetilde{\text{XOR}}$ -SATISFYING-PAIRS.

Since an $\widetilde{\text{XOR}}$ gate can be regarded as an XOR gate with (or without) a NOT gate on its output wire, it follows that $\text{AND}_{c_u} \circ \widetilde{\text{XOR}} \subseteq \text{NC}^0 \circ \text{XOR}$. Furthermore, by Theorem 2.17, It suffices to solve

$\#\text{XOR}_{c_u} \circ \text{XOR-SATISFYING-PAIRS}$ with roughly the same parameters. As $\text{XOR}_{c_u} \circ \text{XOR} \subseteq \text{XOR}$, we only need to consider $\#\text{XOR-SATISFYING-PAIRS}$.

Let $d := \max\{c_u, 2.02\}$ be a constant, we apply Theorem 5.3 with $N := N(n) = 2^{\log^d n}$, $n_{\text{sat}} := 2^{\log^{0.499} N} \geq 2^{\log^{1.001} n} \geq \text{poly}(n)$ as parameters. (That is, we pad $n_{\text{sat}} - n$ bits to the inputs of the SATISFYING-PAIRS instance.) It gives an algorithm for $\#\text{XOR-SATISFYING-PAIRS}$ (and therefore also for $\text{AND}_{c_u} \circ \text{XOR}$) with running time $T = N^{2-\Omega(1/\log(n_{\text{sat}}/\log N))}$. We can check the technical requirement for T :

$$\begin{aligned} T &= N^{2-\Omega(1/\log(n_{\text{sat}}/\log N))} \leq N^2 2^{-\Omega(\log N/\log n_{\text{sat}})} \leq N^2 2^{-\log^{0.5001} N} \\ &\leq N^2 2^{-\log^{1.01} n} \leq N^2 2^{-c_u^2 \log n \log \log N} \leq N^2 / P^{c_u}, \end{aligned}$$

where $P := (\log N)^{\log(1/\varepsilon)}$. Other technical requirements for Theorem 3.5 can be easily verified, so we can invoke it and get an FP^{NP} Remote Point algorithm for XOR gates with error $\varepsilon'(n) > n^{-c_u}$ and stretch $\ell'(n) = N^{c_u \log(1/\varepsilon)}$. Finally, we can check $\ell(n) > \ell'(n+1)$ and $\varepsilon(n) > 2\varepsilon'(n+1)$, then get a desired FP^{NP} algorithm for original parameters by Lemma 2.15. \square

5.3 Remote Point for ACC^0

Theorem 1.17 ($\text{ACC}^0\text{-REMOTE-POINT} \in \text{FP}^{\text{NP}}$). *There is a constant $c_u \geq 1$ such that for every constant $d, m \geq 1$, there is a constant $c_{\text{str}} := c_{\text{str}}(d, m) \geq 1$, such that the following holds.*

Let $n < s(n) \leq 2^{n^{o(1)}}$ be a size parameter, $\varepsilon := \varepsilon(n) \geq 2n^{-c_u}$ be an error parameter and $\ell := \ell(n) \geq 2^{\log^{c_{\text{str}}} s}$ be a stretch function, then there is an FP^{NP} algorithm that takes as input a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, where each output bit of C is computed by an $\text{ACC}_d^0[m]$ circuit of size s , and outputs a string y that is $(1/2 - \varepsilon)$ -far from $\text{Range}(C)$.

Proof. Let c_u be the constant from Theorem 3.5, and c_{str} be a constant to be determined later. We then set parameters for invoking Theorem 3.5.

We set $n_{\text{sat}} := \max\{2^{\log^{c_u+2} n}, 2^{\log^{c_u+1} s}\}$. Then we can invoke Theorem 1.16 with input length n_{sat} and size parameter also n_{sat} to get a $\#\text{ACC}_{d+1}^0[m]\text{-SATISFYING-PAIRS}$ algorithm for N circuits and N inputs, where $N := N(n) = 2^{\log^{1/\varepsilon_{\text{sat}}} n_{\text{sat}}}$ for some constant $\varepsilon_{\text{sat}} \in (0, 1)$. This algorithm runs in time $T = \tilde{O}((N/n_{\text{sat}})^2)$.

Set $c_{\text{str}} := (c_u + 2)/\varepsilon_{\text{sat}} + 3$. Let $\varepsilon' := \varepsilon'(n) \geq n^{-c_u}$ be the error parameter and $\ell' := \ell'(n) = N^{c_u \log(1/\varepsilon')}$ be the stretch, then we can check these parameters satisfy the requirements of Theorem 3.5 as follows.

$$\begin{aligned} 2^{\log^{c_u} n} &\leq N \leq 2^{n^{0.99}} \\ \ell'(n) &= N^{c_u \log(1/\varepsilon')} \geq c_u s \\ T &\leq N^2 / 2^{\log^{c_u+2} n} \leq N^2 / n^{c_u^2 \log^{c_u} n} \leq N^2 / (\log N)^{c_u^2 \log^{c_u} n} \leq N^2 / P^{c_u} \end{aligned}$$

(That is, we use the aforementioned algorithm to solve SATISFYING-PAIRS with N circuits of size s and N inputs of length n by padding $n_{\text{sat}} - n$ dummy bits to each input, and then apply Theorem 3.5). By Theorem 3.5, we get an FP^{NP} algorithm for $\text{ACC}^0\text{-REMOTE-POINT}$ with error $\varepsilon'(n) > n^{-c_u}$, and stretch $\ell'(n) = N^{c_u \log(1/\varepsilon')}$. We can check that both $\ell(n) > \ell'(n+1)$ and $\varepsilon(n) > 2\varepsilon'(n+1)$ hold, so we can invoke Lemma 2.15 and get a desired FP^{NP} algorithm for remote point with the original parameters. \square

We can easily recover the state-of-the-art almost-everywhere average-case circuit lower bounds against ACC^0 [CLW20] by giving the truth table generator as the input.

Corollary 5.5. *For every constants $d, m \geq 1$, there is an $\varepsilon > 0$ and a language $L \in \mathbf{E}^{\text{NP}}$ such that L_n cannot be $(1/2 + 2^{-n^\varepsilon})$ -approximated by $\text{AC}_d^0[m]$ circuits of size 2^{n^ε} , for all sufficiently large n .*

Proof Sketch. Let $\text{TT}_s : \{0, 1\}^{O(s \log s)} \rightarrow \{0, 1\}^{2^n}$ be the truth table generator of $\text{AC}_d^0[m]$ circuits, where $s = 2^{n^\varepsilon}$ for some constant ε to be determined later. Each output bit of TT_s is computable by an $\text{AC}_{d'}^0[m]$ circuit of size $s' = \text{poly}(s)$ for some $d' = O(1)$.

For clarity we define $n_{\text{tt}} = O(s \log s)$ to be the input length of TT_s , $s_{\text{tt}}(n_{\text{tt}}) := s'$ and $d_{\text{tt}} := d'$ to be the size and depth of TT_s , respectively. Let c_u and $c_{\text{str}} := c_{\text{str}}(d_{\text{tt}}, m)$ be the constants in Theorem 1.17. Then there is an FP^{NP} algorithm A_{hard} that takes as input a circuit $C : \{0, 1\}^{n_{\text{tt}}} \rightarrow \{0, 1\}^{\ell_{\text{tt}}}$ and outputs a string y that is $(1/2 - \varepsilon_{\text{tt}})$ -far from $\text{Range}(C)$, where $\ell_{\text{tt}} \geq 2^{\log^{c_{\text{str}}} s_{\text{tt}}}$ and $\varepsilon_{\text{tt}} := 2n_{\text{tt}}^{-c_u}$. By choosing ε to be a sufficiently small constant, we can make

$$2^n \geq 2^{\log^{c_{\text{str}}} s_{\text{tt}}} \quad \text{and} \quad 2^{-n^\varepsilon} > \varepsilon_{\text{tt}}.$$

We then fix the input of the FP^{NP} algorithm A_{hard} above to be TT_s to obtain an FP^{NP} algorithm A that takes 1^{2^n} as input and produces a truth table of length 2^n that cannot be $(1/2 + 2^{-n^\varepsilon})$ -approximated by any size- s circuits. The required hard language is then defined as

$$L := \left\{ x \in \{0, 1\}^n : n \in \mathbb{N}, \text{tt} \leftarrow A(1^{2^n}) \in \{0, 1\}^{2^n}, \text{tt}_x = 1 \right\}. \quad \square$$

5.4 Hard Partial Truth Tables for ACC^0

Theorem 1.18 ($\text{ACC}^0\text{-PARTIAL-AVGHARD} \in \text{FP}^{\text{NP}}$). *There is a constant $c_u \geq 1$ such that for every constants $d, m \geq 1$, there is a constant $c_{\text{str}} := c_{\text{str}}(d, m) \geq 1$, such that the following holds.*

Let $n < s(n) \leq 2^{n^{o(1)}}$ be a size parameter, $\varepsilon := \varepsilon(n) \geq 2n^{-c_u}$ be an error parameter and $\ell := \ell(n) \geq 2^{\log^{c_{\text{str}}} s}$ be a stretch function, then there is an FP^{NP} algorithm that given inputs $x_1, \dots, x_\ell \in \{0, 1\}^n$, it outputs a string $y \in \{0, 1\}^\ell$ such that for any $s(n)$ -size $\text{AC}_d^0[m]$ circuit C , y is $(1/2 - \varepsilon)$ -far from $C(x_1) \circ \dots \circ C(x_\ell)$.

Proof Sketch. The proof is similar to Theorem 1.17, so we only sketch the proof.

Let c_u be the constant from Theorem 4.3, and then we set values¹⁹ of $c_{\text{str}}, n_{\text{sat}}, N, \varepsilon_{\text{sat}}, T, \varepsilon'(n)$ and $\ell'(n)$ in the same way as proof Theorem 1.17.

Since parameter constraints of Theorem 4.3 are similar to those of Theorem 3.5, These parameter settings can be used to invoke Theorem 4.3 and get an FP^{NP} algorithm for $O(s(n))$ -size $\text{AC}_{d+O(1)}^0[m]$ -PARTIAL-AVGHARD with stretch $\ell'(n)$ and error $\varepsilon'(n)$. We can check that both $\ell(n) > \ell'(n+1)/2$ and $\varepsilon(n) > 2\varepsilon'(n+1)$ hold, so it is valid to invoke Lemma 2.16 and get a desired FP^{NP} algorithm for average-case hard partial truth tables with the original parameters.

Alternatively, we can reduce $\text{ACC}^0\text{-PARTIAL-AVGHARD}$ to $\text{ACC}^0\text{-REMOTE-POINT}$ (see Section 1.3) and simply apply Theorem 1.17, since the evaluation of ACC^0 circuits can be implemented in ACC^0 . \square

As a consequence, we show (following the observation in [AS10]) that there is no efficient mapping reduction from \mathbf{E}^{NP} to any language decidable by small-size non-uniform ACC^0 circuits.

Corollary 1.20. *Let $d, m \in \mathbb{N}$ be constants, $\text{AC}_d^0[m]$ denote the class of languages computable by a non-uniform family of polynomial-size $\text{AC}_d^0[m]$ circuits. Then, there is a language $L^{\text{hard}} \in \mathbf{E}^{\text{NP}}$ that does not have polynomial-time mapping reductions to any language in $\text{AC}_d^0[m]$.*

¹⁹In order to invoke Lemma 2.16, we actually use $s'(n) = O(s(n))$ as size function and $d' := d + O(1)$ as depth. These are rather minor changes, so we can still use the same parameter settings strategy.

Proof. Our E^{NP} language L^{hard} receives two inputs: a Turing machine R and a string y . Here, the lengths of $\langle R \rangle$ (the encoding of R) and y are $\lceil n/2 \rceil$ and $n' := \lfloor n/2 \rfloor$ respectively, thus L^{hard} receives n -bit strings as inputs. The machine R is interpreted as a reduction that runs in $T(n) := n^{\log n}$ time (which we diagonalise against).

We run R on all inputs of the form $(\langle R \rangle, x')$, where $|x'| = n'$. Let $x_1, x_2, \dots, x_{2^{n'}}$ be an enumeration of length- n' strings, and $z_i := R(\langle R \rangle, x_i)$ be a string of length at most $T(n)$. Note that the strings z_i may not be of the same length, but the length of each z_i is at most $T(n)$. By an averaging argument, there is an $\ell \leq T(n)$ such that there are at least $2^{n'}/T(n) \geq 2^{n^{0.99}}$ strings z_i with length exactly ℓ . Let N be the number of strings z_i with length exactly ℓ and denote these strings to be $z_{i_1}, z_{i_2}, \dots, z_{i_N}$. We can check the technical constraints and invoke Theorem 1.18 to get an FP^{NP} algorithm for solving the $\text{AC}_d^0[m]$ -PARTIAL-HARD problem on inputs $z_{i_1}, z_{i_2}, \dots, z_{i_N}$. We obtain a sequence of bits $y_{i_1}, y_{i_2}, \dots, y_{i_N} \in \{0, 1\}$ such that for every size- $\ell^{\log \ell}$ $\text{AC}_d^0[m]$ circuit C , there is some $j \in [N]$ such that $C(z_{i_j}) \neq y_{i_j}$. This can be done in deterministic $2^{O(n)}$ time with an NP oracle. Finally, we define L^{hard} as follows: suppose x is the i -th string of length n' (i.e., $x = x_i$), then $x \in L^{\text{hard}}$ if and only if $|z_i| = \ell$ and $y_i = 1$.

Clearly, L^{hard} runs in deterministic $2^{O(n)}$ time with an NP oracle. We still need to show that for every language $L \in \text{AC}_d^0[m]$, there is no polynomial time reduction from L^{hard} to L . Suppose, for the sake of contradiction, that there is a polynomial-time reduction R from L^{hard} to L . Let n be a sufficiently large number such that $n/2 > \langle R \rangle$ and $T(n) = n^{\log n}$ is larger than the running time of R . Consider running R on inputs of the form $(\langle R \rangle, x)$ where $|x| = \lfloor n/2 \rfloor$. Let x_i, y_i, z_i, ℓ , and N be defined as above, and C be an $\text{AC}_d^0[m]$ circuit that decides L on input length ℓ . Since the size of C is at most $\text{poly}(\ell) \leq \ell^{\log \ell}$, there is some $j \leq N$ such that $C(z_{i_j}) \neq y_{i_j}$. In other words,

$$\exists i \in N, C(R(\langle R \rangle, x_i)) \neq L^{\text{hard}}(\langle R \rangle, x_i).$$

It follows that R is not a correct reduction from L^{hard} to L . □

6 Construction of Smooth and Rectangular PCPP

Recall that a PCPP verifier is *smooth* if every bit of the proof is equally likely to be queried, i.e., the distribution of a random query position over a random seed is uniformly random. We do not impose any smoothness requirement on the input oracle.

In this section, we construct a smooth and rectangular PCP of proximity. There is essentially nothing new in our construction: With a careful inspection of their techniques, we can combine [RSW22] (a rectangular PCP of proximity) and [BHPT20] (a smooth and rectangular PCP) to obtain our PCPPs. This is not a coincidence as [BHPT20, RSW22] are both variants of [BGH+06]. Nevertheless, we present an (almost) self-contained proof of the construction for the convenience of the reader, as there are several components and many parameters in [BHPT20, RSW22]. As this section is quite technical, we give a brief overview of the construction.

OVERVIEW OF SECTION 6

- Instead of constructing a smooth PCPP verifier directly, we will construct a rectangular PCPP verifier with the *rectangular neighbour listing* (RNL) property, following [BHPT20]. In Section 6.1, we will give the definition of RNL property and verify that the transformation from rectangular PCP verifiers with RNL property to smooth and rectangular PCP verifiers in [BHPT20] also works for PCPP verifiers.
- In Section 6.2, we show that the PCPP verifier in [BGH+06] is a robust and rectangular PCPP

verifier with RNL property. This combines the observations in [BHPT20] (for constructing robust and rectangular PCP verifiers with RNL property) and [RSW22] (for constructing robust and rectangular PCPP verifiers).

- In Section 6.3, we prove a *composition theorem* that generalizes the counterparts in [BHPT20, RSW22]: we can compose a robust and rectangular PCPP verifier for L (called the outer PCPP verifier) and a PCPP verifier for a variant of the circuit-evaluation problem (called the inner PCPP verifier) to obtain a rectangular PCPP verifier for L whose query complexity is at most the query complexity of the inner PCPP verifier. Moreover, the composed PCPP verifier has RNL property if the outer PCPP verifier has RNL property. Due to technical limitations, this rectangular PCPP verifier will also take some ROP parity-check bits (see Definition 2.9).
- By composing the robust PCPP verifier in Section 6.2 with a PCPP with query complexity $O(1)$, we can construct a rectangular PCPP verifier with query complexity $O(1)$ and RNL property. Moreover, we can transform it into a smooth and rectangular PCPP verifier with query complexity $O(1)$. However, the soundness error of this PCPP verifier could be a constant that is very close to 1. In Section 6.4, we show that the soundness error can be reduced to an arbitrarily small constant with an $O(1)$ blow-up to the query complexity, using an expander-walk-based technique developed in [RSW22, Section 7.1.3].
- We wrap all these components up and set the parameters in Section 6.5.

6.1 Rectangular Neighbour Listing and Smoothness

Definition 6.1. Let V be a rectangular PCPP verifier for some language L with randomness complexity r and query complexity q . A *configuration* is defined as a pair $(\text{seed}, k) \in \{0, 1\}^r \times [q]$. It is said to be a proof (resp. input) configuration if the verifier with randomness seed will query the proof (resp. input) oracle on the k -th query.

Two configurations (seed_1, k_1) and (seed_2, k_2) are said to be neighbours if the verifier will access the same bit of the same oracle with randomness seed_1 on the k_1 -th query, or with randomness seed_2 on the k_2 -th query.

Definition 6.2 (Rectangular Neighbour Listing). Let L be a language and V be a rectangular PCPP verifier for L with row randomness complexity r_{row} , column randomness complexity r_{col} , and shared randomness complexity r_{shared} . We say V has $t_{\text{RNL}}(n)$ -time rectangular neighbour listing property if there are two $t_{\text{RNL}}(n)$ -time algorithms A_{row} and A_{col} such that the following conditions hold:

1. The shared randomness $\text{seed.shared} \in \{0, 1\}^{r_{\text{shared}}}$ consists of $\text{seed.shared.row} \in \{0, 1\}^{r_{\text{shared}}/2}$ and $\text{seed.shared.col} \in \{0, 1\}^{r_{\text{shared}}/2}$, i.e., $\text{seed.shared} = (\text{seed.shared.row}, \text{seed.shared.col})$.
2. Let $(\text{seed}, k) = (\text{seed.row}, \text{seed.col}, \text{seed.shared}, k)$ be a configuration, where $\text{seed.shared} = (\text{seed.shared.row}, \text{seed.shared.col})$. The algorithms A_{row} and A_{col} can list all the neighbours of (seed, k) in a “rectangular and synchronized” fashion:
 - Given the row-part randomness $(\text{seed.row}, \text{seed.shared}, k)$ as input, A_{row} will output an ordered list $\text{NList}_{\text{row}}(\text{seed}, k) := \{(\text{seed}_i.\text{row}, \text{seed}_i.\text{shared.row}, k_i.\text{row})\}_{i \in [\ell_{\text{row}}]}$ and an index $\text{self}_{\text{row}}(\text{seed}, k) \in [\ell_{\text{row}}]$
 - Given the column part randomness $(\text{seed.col}, \text{seed.shared}, k)$ as input, A_{col} will output an ordered list $\text{NList}_{\text{col}}(\text{seed}, k) := \{(\text{seed}_i.\text{col}, \text{seed}_i.\text{shared.col}, k_i.\text{col})\}_{i \in [\ell_{\text{col}}]}$ and an index $\text{self}_{\text{col}}(\text{seed}, k) \in [\ell_{\text{col}}]$.

Verifier	V^{old}	V^{new}
Soundness error	s	$s + \mu$
Proximity parameter	δ	δ
Row randomness	r_{row}	r_{row}
Column randomness	r_{col}	r_{col}
Shared randomness	r_{shared}	r_{shared}
Proof matrix height	H_{proof}	$q \cdot 2^{r_{\text{row}} + r_{\text{shared}}/2}$
Proof matrix width	W_{proof}	$2^{r_{\text{col}} + r_{\text{shared}}/2}$
Query complexity	q	$\text{poly}(q/\mu)$
Parity check complexity	p	p
Decision complexity	d	$\text{poly}(d, q/\mu, t_{\text{RNL}}(n))$

Table 3: The parameters of the “smoothened” PCPP V^{new} .

- It holds that $\ell_{\text{row}} = \ell_{\text{col}}$ and $k_i.\text{row} = k_i.\text{col}$ for every $i \in [\ell_{\text{row}}]$. Let $\ell := \ell_{\text{row}}$ and $k_i := k_i.\text{row}$ for every $i \in [\ell]$. Then the “zipped” list of $\text{NList}_{\text{row}}$ and $\text{NList}_{\text{col}}$

$$\text{NList}(\text{seed}, k) := \left\{ (\text{seed}_i.\text{row}, \text{seed}_i.\text{col}, \text{seed}_i.\text{shared}.\text{row}, \text{seed}_i.\text{shared}.\text{col}, k_i) \right\}_{i \in [\ell]}$$

is a list of all the neighbours of (seed, k) . Moreover, $\text{self}_{\text{row}}(\text{seed}, k) = \text{self}_{\text{col}}(\text{seed}, k)$ and the $\text{self}_{\text{row}}(\text{seed}, k)$ -th entry of $\text{NList}(\text{seed}, k)$ is the configuration (seed, k) itself.

- For every pair of neighbours (seed_1, k_1) and (seed_2, k_2) , $\text{NList}(\text{seed}_1, k_1) = \text{NList}(\text{seed}_2, k_2)$ (i.e., these two ordered lists are the same).

By slightly generalizing the technique of [BHPT20], we can smoothen a rectangular PCPP with the rectangular neighbour listing property.

Theorem 6.3. *Suppose that L has a rectangular PCPP verifier V^{old} with ROP and $t_{\text{RNL}}(n)$ -time RNL property, where the parameters are specified in Table 3. Then for every $\mu \in (0, 1)$, L has a smooth and rectangular PCPP verifier V^{new} with ROP and the parameters specified in Table 3.*

Proof. Let $\Pi^{\text{old}} : \{0, 1\}^\ell \rightarrow \{0, 1\}$ be the proof oracle of V^{old} . Assume that $V^{\text{old}}(\text{seed}, i)$ outputs the index of the i -th query of V^{old} given the randomness $\text{seed} \in \{0, 1\}^r$. The “smoothened” verifier V^{new} expects the proof $\Pi^{\text{new}} : \{0, 1\}^{2^r} \times [q] \rightarrow \{0, 1\}$ to be

$$\Pi^{\text{new}}(\text{seed}, i) := \Pi^{\text{old}}[V^{\text{old}}(\text{seed}, i)].$$

Concretely, V^{new} works as follows: It firstly checks that Π^{new} is (close to being) defined as above, i.e., there is a proof matrix Π^{old} such that $\Delta(\Pi^{\text{new}}(\text{seed}, i), \Pi^{\text{old}}[V^{\text{old}}(\text{seed}, i)])$ is sufficiently small; then it runs V^{old} using Π^{new} as the proof oracle, i.e., the verifier randomly chooses a $\text{seed} \in \{0, 1\}^r$, queries $\Pi^{\text{new}}(\text{seed}, 1), \Pi^{\text{new}}(\text{seed}, 2), \dots, \Pi^{\text{new}}(\text{seed}, q)$, and decides whether to accept using the decision predicate of V^{old} . In fact, as in [BHPT20, Section 4.1], the first step can be combined into the second step: we only need to check the consistency of Π^{new} on the fly during the simulation of V^{old} .

The verifier V^{new} . For $\alpha \in (0, 1)$, we say a graph $G = (V, E)$ is an α -sampler if for every $S \subseteq V$,

$$\Pr_{v \leftarrow V} \left[\left| \frac{|S|}{|V|} - \frac{|\Gamma(v) \cap S|}{|\Gamma(v)|} \right| > \alpha \right] < \alpha,$$

where $\Gamma(v)$ is the set of neighbours of v in G . By [Gol11], there is a poly(n)-time algorithm that constructs a $(4/\alpha^4)$ -regular graph on n vertices that is an α -sampler, given n and $\alpha \in (0, 1)$. The new PCPP verifier works as follows:

- Let $\text{seed} \in \{0, 1\}^r$ be the random bits and $i \in [q]$ be the index of a query. If $V^{\text{old}}(\text{seed}, i)$ makes a query to the input oracle, it firstly makes the same query to the input oracle, and then probes $\Pi^{\text{new}}(\text{seed}, i)$ for Δ times. The last Δ queries to $\Pi^{\text{new}}(\text{seed}, i)$ are for the smoothness property.
- Now we assume that $V^{\text{old}}(\text{seed}, i)$ makes a query to the proof oracle. Let $\text{NList} := \text{NList}(\text{seed}, k)$ be the ordered list of neighbours of (seed, k) from the rectangular neighbour listing property, $m := |\text{NList}|$, and $\text{self} \in [m]$ be the index of (seed, i) in the list. Let $\alpha := \mu/(10q)$, and G_α^{NList} be an explicit $(\Delta - 1)$ -regular α -sampler with m nodes from [Gol11]. Let $j_2, j_3, \dots, j_\Delta \in [m]$ be the neighbours of self in G_α^{NList} , and $j_1 := \text{self}$. The verifier probes

$$\Pi^{\text{new}}(\text{NList}[j_1]), \Pi^{\text{new}}(\text{NList}[j_2]), \dots, \Pi^{\text{new}}(\text{NList}[j_\Delta]),$$

and rejects if the answers are not the same. Otherwise, the verifier treats the consistent answer as the answer to the i -th query of V^{old} and simulates V^{old} .

Rectangularity. Recall that the proof oracle is $\Pi^{\text{new}} : \{0, 1\}^{2^r} \times [q] \rightarrow \{0, 1\}$, where $r = r_{\text{row}} + r_{\text{col}} + r_{\text{shared}}$ is the randomness complexity. By the rectangular neighbour listing property of V^{old} , we know that the shared randomness can be partitioned into $(\text{seed.shared.row}, \text{seed.shared.col}) \in \{0, 1\}^{r_{\text{shared}}/2} \times \{0, 1\}^{r_{\text{shared}}/2}$. We define $W^{\text{new}} := 2^{r_{\text{col}} + r_{\text{shared}}/2}$, $H^{\text{new}} := q \cdot 2^{r_{\text{row}} + r_{\text{shared}}/2}$, and the $H^{\text{new}} \times W^{\text{new}}$ proof matrix

$$\begin{aligned} \Pi^{\text{new}}[u, v] &:= \Pi^{\text{new}}(\text{seed}, i); \\ \text{where } u &:= (\text{seed.row}, \text{seed.shared.row}, i) \in \{0, 1\}^{r_{\text{row}} + r_{\text{shared}}/2 + \log q}, \\ v &:= (\text{seed.col}, \text{seed.shared.col}) \in \{0, 1\}^{r_{\text{col}} + r_{\text{shared}}/2}, \\ \text{seed} &:= (\text{seed.row}, \text{seed.col}, \text{seed.shared} := (\text{seed.shared.row}, \text{seed.shared.col})). \end{aligned}$$

Now it suffices to construct the type predicate $V_{\text{type}}^{\text{new}}$ and the row and column verifiers $V_{\text{row}}^{\text{new}}$ and $V_{\text{col}}^{\text{new}}$. Recall that the new verifier V^{new} simulates V^{old} as follows: If V^{old} makes a query to the proof oracle, it makes Δ queries to the proof oracle using the RNL property; otherwise, it makes the same query to the input oracle and Δ queries to the same bit of the proof oracle.

- The type predicate $V_{\text{type}}^{\text{new}}$ (given the shared randomness) calls the type predicate $V_{\text{type}}^{\text{old}}$ of the old PCPP verifier, obtains the list of types of the queries, replaces each “proof” by Δ continuous “proof” and replaces each “input” by an “input” and Δ continuous “proof”.
- For a query of V^{old} to the proof oracle, the row verifier $V_{\text{row}}^{\text{new}}$ (resp. the column verifier $V_{\text{col}}^{\text{new}}$) calls the row verifier $V_{\text{row}}^{\text{old}}$ (resp. the column verifier $V_{\text{col}}^{\text{old}}$) of the old PCPP. By the rectangular neighbour listing property, it can list the “row-part” (resp. the “column-part”) of the neighbour list NList and know the index self in the list. It then constructs the sampler G_α , finds the Δ selected neighbours of self (including itself), and outputs the “row-parts” (resp. the “column-parts”) of them.
- For a query of V^{old} to the input oracle, the row verifier $V_{\text{row}}^{\text{new}}$ (resp. the column verifier $V_{\text{col}}^{\text{new}}$) calls the row verifier $V_{\text{row}}^{\text{old}}$ (resp. the column verifier $V_{\text{col}}^{\text{old}}$) of the old PCPP to obtain the query to the input oracle rectangulary. It is easy to see that the remaining Δ queries to the proof oracle can be done rectangulary.

Smoothness. We need to show that for uniformly random $\text{seed} \in \{0, 1\}^r$ and $i \in [q \cdot \Delta]$, each bit of the proof oracle Π^{new} is equally likely to be probed given randomness seed on the i -th query. Let $i_1 := \lfloor (i-1)/\Delta \rfloor + 1$ and $i_2 := (i-1) \bmod \Delta + 1$. Let $\hat{G} = (\hat{V}, \hat{E})$ be the “union” of all G_α^{NList} , that is:

- $\hat{V} := \{0, 1\}^{2r} \times [q]$.
- Let $(\text{seed}, i_1), (\text{seed}', i_1')$ be two configurations given which V^{old} will probe the proof oracle. Then $((\text{seed}, i_1), (\text{seed}', i_1')) \in \hat{E}$ if and only if the configurations (seed, i_1) and (seed', i_1') are neighbours, and there is an edge between them in G_α^{NList} , where NList is the neighbourhood containing these two configurations.
- For each (seed, i_1) such that $V^{\text{old}}(\text{seed}, i_1)$ probes the input oracle, we add Δ self-loops on the node $(\text{seed}, i_1) \in \hat{V}$.

Assume that $\text{seed} \in \{0, 1\}^r$ and $i \in [q \cdot \Delta]$ are uniformly chosen. The query pattern of V^{new} to the proof oracle is as follows: It firstly selects a node $(\text{seed}, i_1) \in \hat{V}$ uniformly, and then chooses a uniform neighbour of it. It is easy to see that each bit of the proof oracle is probed with probability

$$\frac{\Delta}{2^r \cdot q \cdot \Delta} = \frac{1}{2^r \cdot q}.$$

Soundness. The soundness of V^{old} follows from [BHPT20, Appendix A.1] (which is for PCP instead of PCPP); for completeness, we present a self-contained proof here. Assume that x is δ -far from being in L and $\Pi^{\text{new}} : \{0, 1\}^{2r} \times [q]$ is a proof, we need to show that the verifier accepts with probability at most $s + \mu$. Let $\Pi^{\text{old}} : \{0, 1\}^\ell \rightarrow \{0, 1\}$ be defined as follows:

$$\Pi^{\text{old}}[j] := \text{Majority}_{(\text{seed}, i) \in \{0, 1\}^{2r} \times [q]} \left\{ \Pi^{\text{new}}(\text{seed}, i) : V^{\text{old}}(\text{seed}, i) = j \right\},$$

By the soundness of V^{old} , we know that V^{old} will accept (x, Π^{old}) with probability at most s .

Let $\text{idx}^i(\text{seed}) \in [\ell]$ be the i -th query of V^{old} given randomness seed . An index $j \in [\ell]$ is said to be β -consistent if for at least β fraction of (seed, i) such that $\text{idx}^i(\text{seed}) = j$, $\Pi^{\text{new}}(\text{seed}, i) = \Pi^{\text{old}}[j]$. We define the following events over the random variable seed :

- H is the event that for every $i \in [q]$, $\text{idx}^i(\text{seed})$ is $(1 - 2\alpha)$ -consistent (recall that $\alpha := \mu/(10q)$ is the parameter of the sampler).
- M is the event that for every $i \in [q]$, $\Pi^{\text{new}}(\text{seed}, i) = \Pi^{\text{old}}(\text{idx}^i(\text{seed}))$.
- A is the event that V^{new} accepts (x, Π^{new}) on the randomness seed .
- C_i is the event that the Δ queries made by V^{new} corresponding to the i -th query of V^{old} returns the same answer (i.e. the “consistency check” passes on the simulation of the i -th query of V^{old}).

Claim 6.4. $\Pr[A \wedge \overline{H}] \leq q\alpha$.

Claim 6.5. $\Pr[\overline{M} \wedge H] \leq 2q\alpha$.

Claim 6.6. $\Pr[A \wedge H] \leq 2q\alpha + s$.

From the claims above we can see that

$$\begin{aligned} \Pr_{\text{seed} \leftarrow \{0, 1\}^r} [A] &= \Pr_{\text{seed} \leftarrow \{0, 1\}^r} [A \wedge \overline{H}] + \Pr_{\text{seed} \leftarrow \{0, 1\}^r} [A \wedge H] \\ &\leq s + q\alpha + 2q\alpha \\ &\leq s + \mu. \end{aligned}$$

Proof of Claim 6.4. Let (seed, i) be a configuration, $\text{NList} := \text{NList}(\text{seed}, i)$ be the list of all its neighbours, $G_\alpha^{\text{NList}} = (V, E)$ be the explicit sampler graph corresponding to the neighbourhood NList (i.e. V contains the configurations in NList), and self be the node corresponding to (seed, i) . We say that a configuration (seed, i) is an *error configuration* if

$$\left| \frac{|S|}{|V|} - \frac{|\Gamma(\text{self}) \cap S|}{|\Gamma(\text{self})|} \right| > \alpha,$$

where $S := \{(\text{seed}', i') \in V \mid \Pi^{\text{new}}(\text{seed}', i') \neq \Pi^{\text{new}}(\text{seed}, i)\}$. Since G_α^{NList} is an α -sampler, there are at most α fraction of error configurations in each neighbourhood. Suppose \overline{H} happens, then there is some i such that $\text{idx}^i(\text{seed})$ is not $(1 - 2\alpha)$ -consistent, which means that $|S|/|V| \geq 2\alpha$. Suppose in addition that C_i happens (i.e., the ‘‘consistency check’’ on the i -th query passes), then $\Gamma(\text{self}) \cap S = \emptyset$, which means that (seed, i) is an error configuration.

Let E be the set of error configurations. We can then calculate that

$$\begin{aligned} \Pr_{\text{seed} \leftarrow \{0,1\}^r} [A \wedge \overline{H}] &\leq \Pr_{\text{seed} \leftarrow \{0,1\}^r} [C_1 \wedge \cdots \wedge C_q \wedge \overline{H}] \\ &\leq \Pr_{\text{seed} \leftarrow \{0,1\}^r} [\exists i \in [q] (\text{seed}, i) \in E \wedge \overline{H}] \\ &\leq q \cdot \Pr_{\substack{\text{seed} \leftarrow \{0,1\}^r \\ i \leftarrow [q]}} [(\text{seed}, i) \in E \wedge \overline{H}] \\ &\leq q \cdot \alpha. \end{aligned} \quad \diamond$$

Proof of Claim 6.5. For every $j \in [q]$, we denote H_j to be the event that $\text{idx}^j(\text{seed})$ is $(1 - 2\alpha)$ -consistent (and thus $H = \bigwedge_{j \in [q]} H_j$). Let H^* be the event that $\text{idx}^i(\text{seed})$ is $(1 - 2\alpha)$ -consistent over the random variable $(\text{seed}, i) \in \{0, 1\}^r \times [q]$. We can see that:

$$\begin{aligned} \Pr_{\text{seed} \leftarrow \{0,1\}^r} [\overline{M} \wedge H] &\leq q \cdot \Pr_{\text{seed} \leftarrow \{0,1\}^r, i \in [q]} \left[\Pi^{\text{new}}(\text{seed}, i) \neq \Pi^{\text{old}}(\text{idx}^i(\text{seed})) \wedge \bigwedge_{j \in [q]} H_j \right] \\ &\leq q \cdot \Pr_{\text{seed} \leftarrow \{0,1\}^r, i \in [q]} \left[\Pi^{\text{new}}(\text{seed}, i) \neq \Pi^{\text{old}}(\text{idx}^i(\text{seed})) \wedge H_i \right] \\ &= q \cdot \Pr_{\text{seed} \leftarrow \{0,1\}^r, i \in [q]} \left[\Pi^{\text{new}}(\text{seed}, i) \neq \Pi^{\text{old}}(\text{idx}^i(\text{seed})) \wedge H^* \right]. \end{aligned} \quad (7)$$

Let N be the set of all neighbourhoods that contain a configuration $(\text{seed}, i) \in H^*$ (i.e. $\text{idx}^i(\text{seed})$ is $(1 - 2\alpha)$ -consistent). By the definition of H^* and the neighbours of configurations, we can see that for each $h \in N$, all the configurations in h are also in H^* . Thus the uniform distribution over H^* is identical to the following distribution: we first sample a neighbourhood $h \in N$ (with probability proportional to the size of h), then uniformly sample a configuration $(\text{seed}, i) \in h$. Thus we have:

$$\begin{aligned} (7) &\leq q \cdot \Pr_{(\text{seed}, i) \leftarrow \{0,1\}^r \times [q]} \left[\Pi^{\text{new}}(\text{seed}, i) \neq \Pi^{\text{old}}(\text{idx}^i(\text{seed})) \mid H^* \right] \\ &= q \cdot \mathbb{E}_{h \leftarrow \mathcal{N}} \left[\Pr_{(\text{seed}, i) \leftarrow h} \left[\Pi^{\text{new}}(\text{seed}, i) \neq \Pi^{\text{old}}(\text{idx}^i(\text{seed})) \right] \right] \\ &\leq q \cdot \mathbb{E}_{h \in \mathcal{N}} [2\alpha], \\ &\leq 2q\alpha, \end{aligned} \quad (8)$$

where \mathcal{N} is some distribution over N , and Eq. (8) holds by the definition of $(1 - 2\alpha)$ -consistency. \diamond

Proof of Claim 6.6. We can see that:

$$\begin{aligned}
\Pr_{\text{seed} \leftarrow \{0,1\}^r} [A \wedge H] &\leq \Pr_{\text{seed} \leftarrow \{0,1\}^r} [\overline{M} \wedge H] + \Pr_{\text{seed} \leftarrow \{0,1\}^r} [A \wedge H \wedge M] \\
&\leq 2q\alpha + \Pr_{\text{seed} \leftarrow \{0,1\}^r} \left[V^{\text{old}} \text{ accepts } (x, \Pi^{\text{old}}) \wedge H \wedge M \right] \\
&\leq 2q\alpha + \Pr_{\text{seed} \in \{0,1\}^r} \left[V^{\text{old}} \text{ accepts } (x, \Pi^{\text{old}}) \right] \\
&\leq 2q\alpha + s.
\end{aligned}$$

Note that the first inequality follows from the definition of V^{new} and Π^{old} . \diamond

Other Properties. The query complexity, parity-check complexity, and decision complexity can be easily checked by the definition. \square

6.2 A Rectangular PCPP with RNL Property

In this section, we construct a rectangular PCPP with rectangular neighbour listing property, by combining [BHPT20] and [RSW22, Section 7.1].

Lemma 6.7 (Lemma 7.3 of [RSW22]). *Let $\lambda < 0.1$, q, m be integers such that $q \geq \log \frac{4}{\lambda}$, and let $\mathbb{F} = \text{GF}(2^q)$. There is a deterministic polynomial-time algorithm that on input $(1^m, 1^q, 1^{\lceil 1/\lambda \rceil})$, outputs a λ -biased set $S_\lambda \subseteq (\mathbb{F} \setminus \{0\}) \times \mathbb{F}^{m-1}$ of size $O((qm/\lambda)^2)$.*

Theorem 6.8. *For all constants $\delta > 0$, there is a constant $\rho \in (0, 1)$ such that the following holds. Let $m = m(n)$, $T(n)$, $w_{\text{proof}}(n)$, $w_{\text{input}}(n)$ be good functions such that $1 \leq m \leq (\log T(n))^{0.1}$, $n \leq T(n) \leq 2^{\text{poly}(n)}$, $w_{\text{proof}}(n) \leq \log T(n)$, and $w_{\text{input}}(n) \leq \log n$. Then there are good functions $h_{\text{proof}}(n)$ and $h_{\text{input}}(n)$ satisfying*

$$\begin{aligned}
h_{\text{proof}}(n) &= \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}(n), & \text{and} \\
h_{\text{input}}(n) &= \lceil \log n \rceil - w_{\text{input}}(n).
\end{aligned}$$

such that the following holds.

Suppose that $h_{\text{proof}}, w_{\text{proof}} \geq (4/m) \log T(n)$, and that for some absolute constant $C \geq 1$,²⁰

$$\frac{w_{\text{input}}(n)}{w_{\text{proof}}(n)}, \frac{h_{\text{input}}(n)}{h_{\text{proof}}(n)} \leq 1 - \frac{Cm \log \log T(n)}{\log T(n)}.$$

Let $W_{\text{proof}}(n) := 2^{w_{\text{proof}}(n)}$, $H_{\text{proof}}(n) := 2^{h_{\text{proof}}(n)}$, $W_{\text{input}}(n) := 2^{w_{\text{input}}(n)}$, and $H_{\text{input}}(n) := 2^{h_{\text{input}}(n)}$. Then $\text{NTIME}[T(n)]$ has a rectangular neighbour listable, robust, and rectangular PCP of proximity with an $H_{\text{proof}}(n) \times W_{\text{proof}}(n)$ proof matrix and an $H_{\text{input}}(n) \times W_{\text{input}}(n)$ input matrix, whose other parameters are specified in Table 4.

²⁰Note that in this theorem, we allow m to be a super-constant. Plugging in the inequalities $w_{\text{input}} \leq k$ and $k + h_{\text{input}} \leq hm$ in [RSW22, Section 7.1.2, Eq. (7)], it turns out that both $\frac{w_{\text{input}}}{w_{\text{proof}}}$ and $\frac{h_{\text{input}}}{h_{\text{proof}}}$ needs to be at most $1 - \frac{Cm \log \log T}{\log T}$ instead of simply $1 - \frac{C \log \log T}{\log T}$.

Soundness error	$1 - \rho$
Proximity parameter	δ
Robustness parameter	ρ
Row randomness	$h_{\text{proof}} - (4/m) \log T(n)$
Column randomness	$w_{\text{proof}} - (4/m) \log T(n)$
Shared randomness	$(7/m) \log T(n) + O(\log \log T(n) + m \log m)$
Query complexity	$T(n)^{1/m} \cdot \text{polylog}(T(n))$
Decision complexity	
RNL time complexity	$\text{poly}(\log T(n), m^m)$

Table 4: Parameters of the PCPP constructed in Theorem 6.8.

Basic Definitions. Let α, t, h, f be defined as in [RSW22, Section 7.1.1], where: α is a universal constant, $t := \log T(n)$, $h := \lceil (t+3)/m \rceil$, $f := h + \alpha \log_2 t$. Let $\lambda := \min\{1/(ct), 1/m^{2cm}\}$ for some universal constant c .²¹ We work with the field $\mathbb{F} := \text{GF}(2^f)$.

By field theory, we know that \mathbb{F} is an f -dimensional vector space over \mathbb{F}_2 . Let e_1, e_2, \dots, e_f be a basis, then each element in \mathbb{F} can be uniquely represented by $\sum_{i=1}^f e_i b_i$ for $(b_1, b_2, \dots, b_f) \in \{0, 1\}^f$. Let $H := \text{span}\{e_1, e_2, \dots, e_h\}$. Let $\text{bin}_{H^m} : H^m \rightarrow \{0, 1\}^{hm}$ and $\text{bin}_{\mathbb{F}^m} : \mathbb{F}^m \rightarrow \{0, 1\}^{fm}$ be the two bijections in [RSW22, Section 7.1.1], where

$$\text{bin}_{\mathbb{F}^m} : \left(\sum_{i=1}^f e_i b_i, \sum_{i=1}^f e_i b_{f+i}, \dots, \sum_{i=1}^f e_i b_{(m-1)f+i} \right) \mapsto (b_1, b_2, \dots, b_{mf}).$$

The bijection bin_{H^m} will be used to define the queries to the input oracle (see [RSW22, Section 7.1.2]), which is not needed for our purpose of verifying the rectangular neighbour listing property.

We identify the binary strings as numbers where the leftmost bit is the least significant bit. Let $I_t : [n] \rightarrow H^m$ be the map to project the input to H^m defined as $I_t(i) := \text{bin}_{H^m}^{-1}(2^{t+1} + i)$, and $I := \{I_t(k) : k \leq |\Pi_{\text{input}}|\}$.

The Proof. The PCPP proof is a Boolean string of length $\ell \cdot |\mathbb{F}|^m$ for some $\ell = \text{polylog}(T)$. We consider the PCPP proof as an oracle $\Pi_{\text{proof}} : \mathbb{F}^m \rightarrow \{0, 1\}^\ell$, where the i -th bit of the proof is the j' -th bit of $\Pi_{\text{proof}}[\text{bin}_{\mathbb{F}^m}^{-1}(j)]$ for $j := \lfloor i/\ell \rfloor$ and $j' := i \bmod \ell$.²² Without loss of generality we assume that ℓ is a power of 2. Note that to ensure rectangularity, the actual proof is treated as an $H_{\text{proof}} \times W_{\text{proof}}$ Boolean matrix, where $H_{\text{proof}} = \ell \cdot |\mathbb{F}|^m / W_{\text{proof}}$ and thus

$$\begin{aligned} h_{\text{proof}} &= \log \ell + m \log |\mathbb{F}| - w_{\text{proof}} \\ &= O(\log \log T(n)) + mh + m \cdot \alpha \log \log T(n) - w_{\text{proof}} \\ &= \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}. \end{aligned}$$

The layout of the proof matrix which will be discussed later.

Query Pattern. The *line* over \mathbb{F}^m with *intercept* $\vec{x} \in \mathbb{F}^m$ and *direction* $\vec{y} \in \mathbb{F}^m$ is the set $\{\vec{x} + t\vec{y} : t \in \mathbb{F}\}$, denoted by $\vec{x} + \mathbb{F}\vec{y}$. We now describe the query pattern of the PCPP verifier.

²¹In [RSW22, Section 7.1.2] they set $\lambda = 1/(ct)$, since when m is a constant $1/(ct) \leq 1/m^{2cm}$ always holds in sufficiently large n . We set λ as this to satisfy the technical requirement in the soundness proof, which will be discussed later.

²²The ℓ -bit output of Π_{proof} is actually a word of a non-Boolean PCPP encoded by some linear-time computable error-correcting code, see [BHPT20, Section 5.2]. We omit the details as it is not important to us.

Let $S_\lambda \subseteq \mathbb{F}^m$ be the explicit λ -biased set in Lemma 6.7 and seed be the randomness of the verifier, defined as $\text{seed} := (R_2, R_3, \dots, R_m, R_y) \in \mathbb{F}^{m-1} \times [|S_\lambda|]$. Note that $|S_\lambda| = O((fm/\lambda)^2)$. We define $\vec{x} := (0, R_2, R_3, \dots, R_m)$, $\vec{y}_e := (1, 0, \dots, 0)$, $\vec{y} := S_\lambda[R_y]$, $\mathcal{L}_0 := \vec{x} + \mathbb{F}\vec{y}_e$, and $\mathcal{L}_1 := \vec{x} + \mathbb{F}\vec{y}$. Denote

$$\text{shift}(x_1, x_2, \dots, x_m) := (x_2, x_3, \dots, x_m, x_1),$$

and $\text{shift}(\mathcal{L}) := \{\text{shift}(\vec{x}) : x \in \mathcal{L}\}$ for a line \mathcal{L} . The PCPP verifier will query the following locations (see [BHPT20, Section C]):

- For every $\vec{x} \in \mathcal{L}_0 \cup \text{shift}(\mathcal{L}_0) \cup \mathcal{L}_1$, it makes a query to $\Pi_{\text{proof}}[\vec{x}]$.
- For every $\vec{x} \in \mathcal{L}_1 \cap I$, it makes a query to $\Pi_{\text{input}}[I_t^{-1}(\vec{x})]$.

It is easy to see that the type predicate V_{type} is pretty simple: the first $3\ell \cdot |\mathbb{F}|$ queries are to the proof (i.e., the first $3|\mathbb{F}|$ queries to Π_{proof} for \mathcal{L}_0 , $\text{shift}(\mathcal{L}_0)$, and \mathcal{L}_1), and the remaining $|\mathbb{F}|$ queries are to the input (i.e., the $|\mathbb{F}|$ queries to $\mathcal{L}_1 \cap I$). To verify the rectangular neighbour listing property, we only need to consider the $3|\mathbb{F}|$ queries to the proof oracle. Denote these queries to be $(\vec{a}_1, \dots, \vec{a}_{|\mathbb{F}|})$, $(\vec{a}_{|\mathbb{F}|+1}, \dots, \vec{a}_{2|\mathbb{F}|})$, and $(\vec{a}_{2|\mathbb{F}|+1}, \dots, \vec{a}_{3|\mathbb{F}|})$, for \mathcal{L}_0 , $\text{shift}(\mathcal{L})$, and \mathcal{L}_1 , respectively.

Rectangularity of the Proof. Let $k := \lceil (w_{\text{proof}} - \log \ell) \cdot (h/f) \rceil$ and $c_2 := \lceil k/h \rceil$. We partition the random bits $\text{seed} = (R_2, R_3, \dots, R_m, R_y) \in \mathbb{F}^{m-1} \times [|S_\lambda|]$ as follows:

$$\begin{aligned} \text{seed.col} &:= (R_3, R_4, \dots, R_{c_2-1}), \\ \text{seed.row} &:= (R_{c_2+2}, R_{c_2+3}, \dots, R_{m-1}), \\ \text{seed.shared} &:= (R_2, R_{c_2}, R_{c_2+1}, R_m, R_y). \end{aligned}$$

Lemma 6.9 ([RSW22], Lemma 7.5). *Let $\text{seed} = (R_2, R_3, \dots, R_m, R_y)$ as defined above, $R_1 = R_{m+1} = 0^f$, $\vec{\alpha}_i = (a_{i,1}, \dots, a_{i,m}) \in \mathbb{F}^m$ be the queries. For every $j \in [m]$, $(a_{1,j}, a_{2,j}, \dots, a_{3|\mathbb{F}|,j})$ can be efficiently computed given R_j, R_{j+1} , and R_y .*

Proof. Let $\mathbb{F} = \{h_1, h_2, \dots, h_{|\mathbb{F}|}\}$ and $j \in [m]$. We can see that

$$\begin{aligned} (a_{1,j}, a_{2,j}, \dots, a_{3|\mathbb{F}|,j}) &= (R_j + h_1 \cdot y_{1,j}, R_j + h_2 \cdot y_{1,j}, \dots, R_j + h_{|\mathbb{F}|} \cdot y_{1,j}, \\ &\quad R_{j+1} + h_1 \cdot y_{2,j}, R_{j+1} + h_2 \cdot y_{2,j}, \dots, h_{|\mathbb{F}|} \cdot y_{2,j}, \\ &\quad R_j + h_1 \cdot y_{3,j}, R_j + h_2 \cdot y_{3,j}, \dots, R_j + h_{|\mathbb{F}|} \cdot y_{3,j}). \end{aligned}$$

where $\vec{y}_1 := (1, 0, \dots, 0) \in \mathbb{F}^m$, $\vec{y}_2 := (0, \dots, 0, 1) \in \mathbb{F}^m$, and $\vec{y}_3 = S_\lambda[R_y]$. \square

By Lemma 6.9, we can see that: Given seed.shared and seed.col , we can recover $a_{1,j}, \dots, a_{3|\mathbb{F}|,j}$ for every $j \in [1, c_2]$; given seed.shared and seed.row , we can recover $a_{1,j}, a_{2,j}, \dots, a_{3|\mathbb{F}|,j}$ for every $j \in [c_2, m]$. Note that

$$\begin{aligned} |\text{seed.col}| &= (c_2 - 3)f \geq w_{\text{proof}} - 4t/m \geq 0, \\ |\text{seed.row}| &= (m - c_2 - 2)f \geq h_{\text{proof}} - 4t/m \geq 0. \end{aligned}$$

We assume without loss of generality that t/m is an integer. Similar to [RSW22, Section 7.1.2], we move some bits to the shared randomness so that $|\text{seed.col}| = w_{\text{proof}} - 4t/m$ and $|\text{seed.row}| = h_{\text{proof}} - 4t/m$, then

$$\begin{aligned} |\text{seed.shared}| &= (m - 1)f + |R_y| - (w_{\text{proof}} + h_{\text{proof}}) + 8t/m. \\ &\leq (m - 1)f + O(\log(fm/\lambda)) - (\log \ell + mf) + 8t/m \\ &\leq 8t/m + O(\log t + m \log m) - f \\ &\leq 7t/m + O(\log t + m \log m). \end{aligned}$$

Let $i \in [3\ell \cdot |\mathbb{F}|]$ be an index, $j := \lfloor (i-1)/\ell \rfloor + 1$, and $j' := (i-1) \bmod \ell$. The i -th query asks for the j' -th bit of $\Pi_{\text{proof}}[\vec{a}_j]$, or equivalently, the $(\text{bin}_{\mathbb{F}^m}(\vec{a}_j) \cdot \ell + j')$ -th bit of the proof. We will construct V_{row} and V_{col} such that

$$\begin{aligned} V_{\text{row}}(\text{seed.row}, \text{seed.shared}, i) &\rightarrow \text{irow}[i], \\ V_{\text{col}}(\text{seed.col}, \text{seed.shared}, i) &\rightarrow \text{icol}[i], \\ \text{irow}[i] \cdot W_{\text{proof}} + \text{icol}[i] &= \text{bin}_{\mathbb{F}^m}(\vec{a}_j) \cdot \ell + j'. \end{aligned}$$

That is, if we place the proof into a matrix such that the (p, q) -th entry represents the $p \cdot W_{\text{proof}} + q$ bit of the proof, the first $3\ell \cdot |\mathbb{F}|$ queries to the proof can be made rectangular.

- $V_{\text{col}}(\text{seed.col}, \text{seed.shared}, i)$ computes $a_{j,1}, a_{j,2}, \dots, a_{j,c_2}$ by Lemma 6.9; it outputs $\text{icol}[i]$ as the concatenation of j' and $a_{j,1}, a_{j,2}, \dots, a_{j,c_2-1}$ and the lowest $w_{\text{proof}} - (c_2 - 1)f - \log \ell$ bits of a_{j,c_2} .
- $V_{\text{row}}(\text{seed.row}, \text{seed.shared}, i)$ computes $a_{j,c_2}, a_{j,c_2+1}, \dots, a_{j,m}$ by Lemma 6.9; it outputs $\text{irow}[i]$ as the concatenation of the highest $c_2 f + \log \ell - w_{\text{proof}}$ bits of a_{j,c_2} and $a_{j,c_2+1}, a_{j,c_2+1}, \dots, a_{j,m}$.




Figure 5: The binary representation of $\text{bin}_{\mathbb{F}^m}(\vec{a}_j) \cdot \ell + j'$, where the leftmost bits are the least significant ones.

Recall that we identify binary strings and numbers with the leftmost bit as the least significant bit. It is then easy to see that $\text{irow}[i] \cdot W_{\text{proof}} + \text{icol}[i]$ is the concatenation of $\text{irow}[i]$ and $\text{icol}[i]$, and $\text{bin}_{\mathbb{F}^m}(\vec{a}_j) \cdot \ell + j'$ is the concatenation of j' and $\text{bin}_{\mathbb{F}^m}(\vec{a}_j)$. Both of them are the concatenation of $j', a_{j,1}, a_{j,2}, \dots, a_{j,m}$, hence $\text{irow}[i] \cdot W_{\text{proof}} + \text{icol}[i] = \text{bin}_{\mathbb{F}^m}(\vec{a}_j) \cdot \ell + j'$. Similarly, we arrange the input matrix as in [RSW22] so that the queries to the input oracle are also rectangular.

neighbours of (seed, i) . Let $\mathbb{F} = \{h_1, h_2, \dots, h_{|\mathbb{F}|}\}$. Let $\text{seed} = (R_2, R_3, \dots, R_y)$, $i \in [3\ell \cdot |\mathbb{F}|]$, $j := \lfloor (i-1)/\ell \rfloor + 1$, and $j' := (i-1) \bmod \ell$. Assume that this query probes the j' -th bit of $\Pi_{\text{proof}}[\vec{a}_j]$, where $\vec{a}_j = (a_{j,1}, a_{j,2}, \dots, a_{j,m})$. We define the *canonical neighbour* (seed_1, i_1) of (seed, i) as follows:

$$\text{seed}_1 := (R_2^1 := a_{j,2}, R_3^1 := a_{j,3}, \dots, R_m^1 := a_{j,m}, R_y^1 := 0) \quad (9)$$

$$i_1 := (j_1 - 1) \cdot \ell + j'_1 + 1 \quad (10)$$

where $j_1 \in [1, |\mathbb{F}|]$ such that $h_{j_1} = a_{j,1}$, and $j'_1 = j'$. It is easy to see that the canonical neighbours are the representative elements of the equivalence class induced by the neighbourhood relation. Denote \mathcal{S} as the set of canonical neighbours, then $(\text{seed}, i) \in \mathcal{S}$ if and only if $i \in [1, |\mathbb{F}|]$ and $R_y = 0$.

To list all the neighbours of (seed, i) , it suffices to find its canonical neighbour (seed_1, i_1) and list all the neighbours of (seed_1, i_1) . Let $\text{seed}_2 = (R_2^2, R_3^2, \dots, R_m^2, R_y^2)$, $i_2 \in [3\ell \cdot |\mathbb{F}|]$, $j_2 := \lfloor (i_2-1)/\ell \rfloor + 1$, $j'_2 := (i_2-1) \bmod \ell$. Suppose that (seed_2, i_2) is a neighbour of (seed_1, i_1) , then they represent the

queries to the same bit of the same entry of Π_{proof} . This means that $j'_1 = j'_2$, and one of the following conditions holds:

$$j_2 \in [1, |\mathbb{F}|] \quad \text{and} \quad (h_{j_2}, R_2^2, R_3^2, \dots, R_m^2) = (h_{j_1}, R_2^1, \dots, R_{m-1}^1, R_m^1) \quad (11)$$

$$j_2 \in [|\mathbb{F}| + 1, 2|\mathbb{F}|] \quad \text{and} \quad (R_2^2, R_3^2, \dots, R_m^2, h_{j_2-|\mathbb{F}|}) = (h_{j_1}, R_2^1, \dots, R_{m-1}^1, R_m^1); \quad (12)$$

$$j_2 \in [2|\mathbb{F}| + 1, 3|\mathbb{F}|] \quad \text{and} \quad (h \cdot y_1, R_2^2 + h \cdot y_2, \dots, R_m^2 + h \cdot y_m) = (h_{j_1}, R_2^1, \dots, R_m^1) \\ \text{where } h := h_{j_2-2|\mathbb{F}|}, (y_1, y_2, \dots, y_m) := S_\lambda[R_y^2]. \quad (13)$$

We will list the neighbours of (seed_1, i_1) in the following order: the $|S_\lambda|$ configurations satisfying (11) in the lexicographic order of R_y^2 (note that this includes (seed_1, i_1)), the $|S_\lambda|$ configurations satisfying (12) in the lexicographic order of R_y^2 , and then the $|S_\lambda|$ configurations satisfying (13) in the lexicographic order of R_y^2 .²³

Rectangular neighbour Listing of \mathcal{S} . Now we need to verify that the aforementioned listing of the neighbours satisfies the rectangular neighbour listing property (see Definition 6.2). To start with, we consider the case when $(\text{seed}, i) = (\text{seed}_1, i_1)$, i.e., $(\text{seed}, i) \in \mathcal{S}$. Recall that

$$\begin{aligned} \text{seed}_1.\text{col} &:= (R_3^1, R_4^1, \dots, R_{c_2-1}^1), \\ \text{seed}_1.\text{row} &:= (R_{c_2+2}^1, R_{c_2+3}^1, \dots, R_{m-1}^1), \\ \text{seed}_1.\text{shared} &:= (R_2^1, R_{c_2}^1, R_{c_2+1}^1, R_m^1, R_y^1). \end{aligned}$$

Let $\text{low}(\cdot)$ and $\text{high}(\cdot)$ denote the lower and higher halves of a Boolean string, respectively. We partition $\text{seed}.\text{shared} = (R_2^1, R_{c_2}^1, R_{c_2+1}^1, R_m^1, R_y^1)$ into two parts $(\text{seed}.\text{shared}.\text{row}, \text{seed}.\text{shared}.\text{col})$, where $\text{seed}.\text{shared}.\text{col} := (R_2^1, R_{c_2}^1, \text{low}(R_y^1))$ and $\text{seed}.\text{shared}.\text{row} := (R_{c_2+1}^1, R_m^1, \text{high}(R_y^1))$.

- The column algorithm $A_{\text{col}}(\text{seed}_1.\text{col}, \text{seed}_1.\text{shared}, i_1)$ outputs $\text{NList}_{\text{col}}(\text{seed}_1, i_1)$ as follows:
 1. It firstly outputs the column-part of the $|S_\lambda|$ neighbours (seed_2, i_2) satisfying (11), including (seed_1, i_1) itself. We can see that $i_2 = i_1$, $\text{seed}_2.\text{col} = \text{seed}_1.\text{col}$, and $\text{seed}_2.\text{shared} = (R_2^1, R_{c_2}^1, R_{c_2+1}^1, R_m^1, R_y^2)$, where R_y^2 enumerates over all $|S_\lambda|$ possibilities in lexicographic order.
 2. It then aims to output the column-part of the $|S_\lambda|$ neighbours (seed_2, i_2) satisfying (12). We will enumerate all R_y^2 in lexicographic order and output:

$$\begin{aligned} \text{seed}_2.\text{col} &= (R_3^2 = R_2^1, \dots, R_{c_2-1}^2 = R_{c_2-2}^1), \\ \text{seed}_2.\text{shared}.\text{col} &= (R_2^2 = h_{j_1}, R_{c_2}^2 = R_{c_2-1}^1, \text{low}(R_y^2)). \end{aligned}$$

Let $j_2 := \lfloor (i_2 - 1) / \ell \rfloor + 1$ and $j'_2 := (i_2 - 1) \bmod \ell$, we can see that $j'_2 = j'_1$ and $h_{j_2} = R_m^1$, hence i_2 can also be computed efficiently.

3. Finally, it aims to output the column-part of the neighbours (seed_2, i_2) satisfying (13). We enumerate R_y^2 in lexicographic order. Let $(y_1, y_2, \dots, y_m) = S_\lambda[R_y^2]$. Denote $h := y_1^{-1} \cdot h_{j_1}$, and let j_2 be the unique number in $[2|\mathbb{F}| + 1, 3|\mathbb{F}|]$ such that $h = h_{j_2-2|\mathbb{F}|}$, we output:

$$\begin{aligned} i_2 &= (j_2 - 1) \cdot \ell + j'_1 + 1 \\ \text{seed}_2.\text{col} &= (R_3^2 = R_3^1 - h \cdot y_3, \dots, R_{c_2-1}^2 = R_{c_2-1}^1 - h \cdot y_{c_2-1}) \\ \text{seed}_2.\text{shared}.\text{col} &= (R_2^2 = R_2^1 - h \cdot y_2, R_{c_2}^2 = R_{c_2}^1 - h \cdot y_{c_2}, \text{low}(R_y^2)). \end{aligned}$$

²³Recall that for every $(y_1, y_2, \dots, y_m) = S_\lambda[R_y^2]$, $y_1 \neq 0$. Thus for every R_y^2 , there is exactly one $j_2 \in [2|\mathbb{F}| + 1, 3|\mathbb{F}|]$ that satisfies (13), namely the j_2 such that $h_{j_2-2|\mathbb{F}|} = y_1^{-1} \cdot h_{j_1}$.

- Similarly, the row algorithm $A_{\text{row}}(\text{seed}_1.\text{row}, \text{seed}_1.\text{shared}, i_1)$ can output $\text{NList}_{\text{row}}(\text{seed}_1, i_1)$. We omit the details since it can be adapted from A_{col} directly.

It is clear that the “zipped” list of $\text{NList}_{\text{row}}$ and $\text{NList}_{\text{col}}$ is the list of neighbours of (seed_1, i_1) . Since (seed_1, i_1) appears at the head of the list, A_{row} and A_{col} can simply output $\text{self}_{\text{row}} = \text{self}_{\text{col}} = 1$.

Rectangular neighbour Listing for $\bar{\mathcal{S}}$. For the general case when $(\text{seed}, i) \notin \mathcal{S}$, the algorithms A_{row} and A_{col} need to find its canonical neighbour (seed_1, i_1) “rectangularly”: Let (seed_1, i_1) be the canonical neighbour of (seed, i) , then A_{row} (resp. A_{col}) can output the row-part (resp. the column-part) of (seed_1, i_1) given $(\text{seed}.\text{row}, \text{seed}.\text{shared}, i)$ (resp. $(\text{seed}.\text{col}, \text{seed}.\text{shared}, i)$). This can be done by checking Equation (9), Equation (10), and Lemma 6.9.

Finally, we need to verify that both A_{row} and A_{col} know the index of (seed, i) itself in this list. Let $j := \lfloor (i-1)/\ell \rfloor + 1$ and $j' := (i-1) \bmod \ell$. The list of (seed_1, i_1) contains three parts: the neighbours specified by (11), (12), and (13). By checking j , which is known by both A_{row} and A_{col} , we can find the part that contains (seed, i) . The index of (seed, i) within the part is then determined by the lexicographic order of R_y in $\text{seed}.\text{shared}$.

Complexity of A_{row} and A_{col} . Recall that $\log |\mathbb{F}| = f = (\log T(n) + 3)/m + O(\log \log T(n))$, $|S_\lambda| = O((mf/\lambda)^2) = \text{poly}(m^m, \log T(n))$. It is then easy to check that the running time of both A_{row} and A_{col} is $\text{poly}(m^{O(m)}, f, \log T(n)) = \text{poly}(m^m, \log T(n))$.

Robust Soundness. Finally, we prove a *weak* version of robust soundness here. This PCPP only guarantees an *expected* version of robust soundness: the expected fraction of bits that we need to flip in order to make the verifier accept is at least ρ , where the expectation is over the choice of seed . (See [BGH⁺06, Lemma 8.11].) We use a Markov bound to turn this into a standard robust soundness property, but only with soundness parameter very close to 1. Since the robust soundness amplification (Section 6.4) preserves smoothness but does not seem to preserve RNL, we do not apply it here.

Let $\delta_{\text{proof}}(\text{seed})$ (resp. $\delta_{\text{input}}(\text{seed})$) be the fraction of bits of Π_{proof} (resp. Π_{input}) read by the verifier that we need to flip to make the verifier accept given the randomness seed . Let $\hat{\delta}(\text{seed})$ be the fraction of bits (of both Π_{proof} and Π_{input}) read by the verifier that we need to flip to make the verifier accept given the randomness seed . By [BGH⁺06, Lemma 8.11] (also see [RSW22, Proof of Theorem 7.1]),²⁴ there is a constant $\rho_0 \in (0, 1)$ such that for every constant $\delta \in (0, 1)$, if Π_{input} is δ -far from being in L , then for any proof oracle Π_{proof} , either $\mathbb{E}_{\text{seed}}[\delta_{\text{proof}}(\text{seed})] \geq \rho_0$ or $\mathbb{E}_{\text{seed}}[\delta_{\text{input}}(\text{seed})] \geq \delta/3$. Fix this constant ρ_0 .

Recall that the verifier V^{new} makes $|\mathbb{F}|$ queries to Π_{input} and $3\ell \cdot |\mathbb{F}|$ queries to Π_{proof} . We repeat each query to the input oracle for $9(\rho_0/\delta)\ell$ times. Then if Π_{input} is δ -far from being in L , the fraction of bits read by the verifier that we need to flip on expectation to make the verifier accept is

$$\mathbb{E}_{\text{seed}} \left[\hat{\delta}(\text{seed}) \right] = \frac{\min\{\rho_0 \cdot 3\ell \cdot |\mathbb{F}|, (\delta/3) \cdot 9(\rho_0/\delta)\ell \cdot |\mathbb{F}|\}}{3\ell \cdot |\mathbb{F}| + 9(\rho_0/\delta)\ell \cdot |\mathbb{F}|} \geq \frac{3\ell\rho_0}{3\ell + 9(\rho_0/\delta)\ell} \geq \frac{\rho_0}{1 + 3\rho_0/\delta}.$$

Let $\rho := \frac{\rho_0}{1 + 3\rho_0/\delta}$. By a Markov bound,

$$\Pr_{\text{seed}} \left[\hat{\delta}(\text{seed}) \leq \rho \right] \leq 1 - \rho.$$

²⁴Recall that we assume $m \leq (\log T(n))^{0.1}$ and set $\lambda = \min\{1/(ct), 1/m^{2cm}\}$, so that $m^m \leq T(n)^{1/m^2}$ and $\lambda \leq \min\{1/(ct), 1/m^{cm}\}$, which satisfies the technical requirement of [BGH⁺06, Lemma 8.11].

Verifier	V^{out}	V^{in}	V^{comp}
Soundness error	$1 - \varepsilon^{\text{out}}$	$1 - \varepsilon^{\text{in}}$	$1 - \varepsilon^{\text{out}} \cdot \varepsilon^{\text{in}}$
Proximity parameter	δ^{out}	δ^{in}	δ^{out}
Robustness parameter	ρ^{out}	-	-
Row randomness	$r_{\text{row}}^{\text{out}}$	-	$r_{\text{row}}^{\text{out}}$
Column randomness	$r_{\text{col}}^{\text{out}}$	-	$r_{\text{col}}^{\text{out}}$
Shared randomness	$r_{\text{shared}}^{\text{out}}$	r^{in}	$r_{\text{shared}}^{\text{out}} + r^{\text{in}}$
Proof matrix height	$H_{\text{proof}}^{\text{out}}$	-	$H_{\text{proof}}^{\text{out}} + 2^{r_{\text{col}}^{\text{out}} + r^{\text{in}}} / W_{\text{proof}}^{\text{out}}$
Proof matrix width	$W_{\text{proof}}^{\text{out}}$	-	$W_{\text{proof}}^{\text{out}}$
Query complexity	q^{out}	q^{in}	q^{in}
Parity check complexity	-	-	q^{in}
Decision complexity	d^{out}	d^{in}	d^{in}
Proof length	ℓ^{out}	ℓ^{in}	$H_{\text{proof}}^{\text{out}} \cdot W_{\text{proof}}^{\text{out}} + 2^{r_{\text{col}}^{\text{out}} + r^{\text{in}}}$

Table 5: The parameters of the PCPPs in the composition theorem. Note that the input length of the inner PCPP is $d^{\text{out}} = d^{\text{out}}(n)$, e.g., r^{in} in the table actually refers to $r^{\text{in}}(d^{\text{out}}(n))$.

Thus, the PCPP verifier has robust soundness error $1 - \rho$ with robustness parameter ρ and proximity parameter δ .

6.3 RNL-Preserving Composition Theorem

Now we verify that the composition theorem in [RSW22, Section 7.2] preserves the rectangular neighbour listing property, using essentially the same approach of [BHPT20, Section 7.2]. This will be used to reduce the number of queries of our rectangular PCPP with RNL property.

Theorem 6.10. *Let $n \leq T(n) \leq 2^{\text{poly}(n)}$. Suppose that $\text{NTIME}[T(n)]$ has a robust and rectangular PCPP verifier V^{out} and $\text{CIRCUIT-EVAL}^\perp$ has a (not necessarily rectangular) PCPP verifier V^{in} with parameters specified in Table 5. Moreover, assume that $q^{\text{in}} = O(1)$, $\rho^{\text{out}} \geq \delta^{\text{in}}$, $\ell^{\text{in}} = 2^{r^{\text{in}}}$,²⁵ $W_{\text{proof}}^{\text{out}}$ is a power of 2, and $r_{\text{col}}^{\text{out}} \leq \log W_{\text{proof}}^{\text{out}} \leq r_{\text{col}}^{\text{out}} + r_{\text{shared}}^{\text{out}}$. Then $\text{NTIME}[T(n)]$ has a rectangular PCPP verifier V^{comp} with parameters specified in Table 5.*

Furthermore, if V^{out} has $t_{\text{RNL}}^{\text{out}}(n)$ -time rectangular neighbour listing property, then V^{comp} has $t_{\text{RNL}}(n)$ -time rectangular neighbour listing property, where $t_{\text{RNL}}(n) := \text{poly}(t_{\text{RNL}}^{\text{out}}(n), \ell^{\text{in}}, q^{\text{in}}, d^{\text{in}})$.

The composed PCPP verifier. Assume that we have a robust and rectangular PCPP verifier V^{out} for $L \in \text{NTIME}[T(n)]$ and a PCPP verifier V^{in} for $\text{CIRCUIT-EVAL}^\perp$. We now describe the composed PCPP verifier V^{comp} for L (also see [BGH⁺06, Section 2.4], [BHPT20, Section 7.2], [RSW22, Section 7.2]). In a nutshell, we will reduce the verification of the outer PCPP V^{out} to $\text{CIRCUIT-EVAL}^\perp$, where the circuit represents the decision predicate of V^{out} and the input consists of the input of L and the proof for the outer PCPP. As in [BHPT20, RSW22], we need to carefully arrange the proof matrix to maintain the rectangularity.

Assume that $\Pi_{\text{input}}^{\text{out}}$, $\Pi_{\text{proof}}^{\text{out}}$, and seed^{out} are the proof matrix, the input matrix, and the random seed of V^{out} , respectively. The input matrix of V^{comp} is simply the input matrix of V^{out} , denoted

²⁵This is without loss of generality, because $\ell^{\text{in}} \leq 2^{r^{\text{in}}} \cdot q^{\text{in}}$, and in our case q^{in} will be a constant. We could always add $O(\log q^{\text{in}}) = O(1)$ dummy bits to the inner verifier's randomness and pad the inner verifier's proof oracle to length $2^{r^{\text{in}}}$.

by Π^{in} . The proof of V^{comp} is the concatenation of $\Pi_{\text{proof}}^{\text{out}}$ and $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$ for every $\text{seed}^{\text{out}} \in \{0, 1\}^{r^{\text{out}}}$, where each $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$ is a PCPP proof for “ V^{in} accepts seed^{out} .” The random seed is $\text{seed} := \text{seed}^{\text{out}} \circ \text{seed}^{\text{in}}$. The verifier V^{comp} works as follows:

1. Obtain the decision circuit Dec^{out} and the list of query indices $I^{\text{out}} \leftarrow V^{\text{out}}(\text{seed}^{\text{out}})$ of V^{out} .
2. Use the inner PCPP to verify the following $\text{CIRCUIT-EVAL}^\perp$ instance $(C, \Pi_{\text{input}}^{\text{in}}(\text{seed}^{\text{out}}))$: the (explicit) circuit $C : \{0, 1, \perp\}^{q^{\text{out}}(n)} \times \{0, 1\}^{\hat{r}^{\text{out}}} \rightarrow \{0, 1, \perp\}$ and the (implicit) input $\Pi_{\text{input}}^{\text{in}}(\text{seed}^{\text{out}}) \in \{0, 1, \perp\}^{q^{\text{out}}(n)} \times \{0, 1\}^{\hat{r}^{\text{out}}}$ are defined as follows:

$$C(u, v) := \text{Dec}^{\text{out}}(u, \text{Dec}(v)),$$

$$\Pi_{\text{input}}^{\text{in}}(\text{seed}^{\text{out}}) := ((\Pi_{\text{input}} \circ \Pi_{\text{proof}}^{\text{out}})|_{I^{\text{out}}}, \text{Enc}(\text{seed}^{\text{out}})),$$

where (Enc, Dec) is a linear-time encodable and decodable error-correcting code such that $\text{Enc} : \{0, 1\}^{r^{\text{out}}} \rightarrow \{0, 1\}^{\hat{r}^{\text{out}}}$ is linear over $\text{GF}(2)$.²⁶ Specifically:

- (a) The decision circuit Dec^{comp} of V^{comp} is defined as the decision circuit Dec^{in} of V^{in} .
- (b) The queries are sampled using $V^{\text{in}}(\text{seed}^{\text{in}})$ for the $\text{CIRCUIT-EVAL}^\perp$ instance defined above with the proof $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$, i.e., we sample the queries $I^{\text{in}} \leftarrow V^{\text{in}}(\text{seed}^{\text{in}})$ and “redirect” them to the input oracle and the proof of the composed PCPP to obtain I^{comp} .²⁷

Rectangularity of V^{comp} . We now verify the rectangularity of the composed PCPP verifier. Recall that: the proof $\Pi_{\text{proof}}^{\text{out}} \in \{0, 1\}^{\ell^{\text{out}}}$ of V^{out} is arranged as an $H_{\text{proof}}^{\text{out}} \times W_{\text{proof}}^{\text{out}}$ matrix, where the i -th row and the j -th column $\Pi_{\text{proof}}^{\text{out}}[i, j] := \Pi_{\text{proof}}^{\text{out}}[(i-1) \cdot W_{\text{proof}}^{\text{out}} + j]$; the inner proofs $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}}) \in \{0, 1\}^{\ell^{\text{in}}}$ for every $\text{seed}^{\text{out}} \in \{0, 1\}^{r^{\text{out}}}$.

Let $W_{\text{proof}}^{\text{in}} := \min\{W_{\text{proof}}^{\text{out}}/2^{r^{\text{col}}}, \ell^{\text{in}}\}$ and $H_{\text{proof}}^{\text{in}} := \ell^{\text{in}}/W_{\text{proof}}^{\text{in}}$. Note that both $W_{\text{proof}}^{\text{in}}$ and $H_{\text{proof}}^{\text{in}}$ are powers of 2, and $W_{\text{proof}}^{\text{out}} \geq 2^{r^{\text{col}}}$. We arrange the proof matrix as follows: The first $H_{\text{proof}}^{\text{out}}$ rows contain the $H_{\text{proof}}^{\text{out}} \times W_{\text{proof}}^{\text{out}}$ proof matrix of the outer PCPP; the rest part of the matrix is divided into blocks of size $W_{\text{proof}}^{\text{in}} \times H_{\text{proof}}^{\text{in}}$, each of which contains a proof of the inner PCPP $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$ for some $\text{seed}^{\text{out}} \in \{0, 1\}^{\ell^{\text{in}}}$, sorted in the lexicographic order of seed^{out} . Clearly, the proof matrix height of the composed PCPP verifier is

$$H_{\text{proof}}^{\text{out}} + \ell^{\text{in}} \cdot 2^{r^{\text{out}}}/W_{\text{proof}}^{\text{out}} = H_{\text{proof}}^{\text{out}} + 2^{r^{\text{out}}+r^{\text{in}}}/W_{\text{proof}}^{\text{out}}.$$

Recall that the seed of the composed PCPP verifier is $\text{seed} := (\text{seed}^{\text{out}}, \text{seed}^{\text{in}})$. Assume that the partition of random bits of V^{out} is $\text{seed}^{\text{out}} = (\text{seed}^{\text{out}}.\text{row}, \text{seed}^{\text{out}}.\text{col}, \text{seed}^{\text{out}}.\text{shared})$. We partition the random bits as follows:

$$\begin{aligned} \text{seed}.\text{shared} &:= (\text{seed}^{\text{out}}.\text{shared}, \text{seed}^{\text{in}}). \\ \text{seed}.\text{row} &:= \text{seed}^{\text{out}}.\text{row}. \\ \text{seed}.\text{col} &:= \text{seed}^{\text{out}}.\text{col}. \end{aligned}$$

²⁶The reason to apply an error-correcting code on the randomness is that we want Dec^{out} to have *robust* soundness. Let $(\Pi', \text{Enc}(\text{seed}^{\text{out}}))$ be the input of Dec^{out} , if given seed^{out} , Π' is far from being accepted by Dec^{out} , then $(\Pi', \text{Enc}(\text{seed}^{\text{out}}))$ is also far from being accepted by Dec^{out} . This is not true if we do not encode seed^{out} .

²⁷In fact, there are three kinds of queries: the queries to $(\Pi_{\text{input}} \circ \Pi_{\text{proof}}^{\text{out}})|_{I^{\text{out}}}$, $\text{Enc}(\text{seed}^{\text{out}})$, and $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$. The queries of the first and the third kinds will be redirected as queries, and the second kind will be treated as a parity-check bit, since Enc is a linear function over $\text{GF}(2)$. Details are contained in the verification of the rectangularity, also see [RSW22, Algorithm 2].

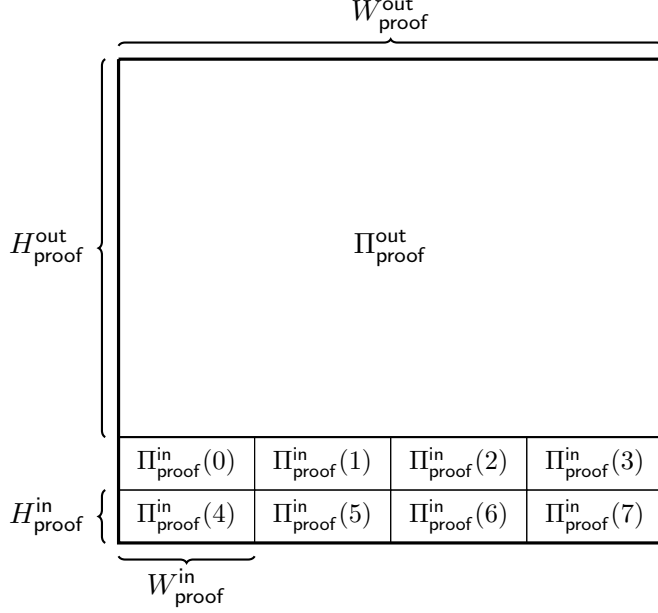


Figure 6: The layout of the proof matrix of the composed PCPP.

Now we describe the type predicate V_{type} and rectangular verifiers $V_{\text{row}}^{\text{comp}}$ and $V_{\text{col}}^{\text{comp}}$ of the composed PCPP.

The type predicate and the row/column verifier firstly obtain seed^{in} in $\text{seed}^{\text{shared}}$ and compute the queries $I^{\text{in}} \leftarrow V^{\text{in}}(\text{seed}^{\text{in}})$ of the inner PCPP for $\text{CIRCUIT-EVAL}^{\perp}$. There can be three cases for each query in I^{in} :

1. It probes the i -th cell of $\Pi_{\text{input}}^{\text{in}}(\text{seed}^{\text{out}})$ and $i \leq q^{\text{out}}(n)$, i.e., it queries $(\Pi_{\text{input}} \circ \Pi_{\text{proof}}^{\text{out}})|_{I^{\text{out}}}$.
 - The type predicate can call the type predicate of V^{out} to compute the type of the query, since it has $\text{seed}^{\text{out}}.\text{shared}$ and the index of the query in hand.
 - The row/column verifier of the composed PCPP runs the row/column verifier of the outer PCPP to obtain the row/column index of the query in $(\Pi_{\text{input}}, \Pi_{\text{proof}}^{\text{out}})$. By the definition of $\Pi_{\text{proof}}^{\text{comp}}$, the row/column index is also the row/column index of the query in $(\Pi_{\text{input}}, \Pi_{\text{proof}}^{\text{comp}})$.
2. It probes the i -th cell of $\Pi_{\text{input}}^{\text{in}}(\text{seed}^{\text{out}})$ and $i > q^{\text{out}}(n)$, i.e., it queries $\text{Enc}(\text{seed}^{\text{out}})$. Instead of making a query, we fix this input of Dec^{comp} to be $\text{Enc}(\text{seed}^{\text{out}})[i - q^{\text{out}}(n)]$. This bit will be considered as a parity-check bit.
3. It probes the i -th cell of $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$. This is a query to the proof, so the type predicate always outputs **proof**. Recall that $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$ is placed in some $W_{\text{proof}}^{\text{in}} \times H_{\text{proof}}^{\text{in}}$ size block in the proof matrix. Let $N_{\text{proof}} := W_{\text{proof}}^{\text{out}}/W_{\text{proof}}^{\text{in}}$, i.e., there are N_{proof} blocks of inner PCPP proofs in a row of the proof matrix. It is then easy to see that the column (resp. row) index of the query depends on seed^{in} and the lowest $\log N_{\text{proof}}$ bits (resp. the highest $r^{\text{out}} - \log N_{\text{proof}}$ bits) of seed^{out} . Note that since

$$\begin{aligned} \log N_{\text{proof}} &\geq \log \left(W_{\text{proof}}^{\text{out}} / (W_{\text{proof}}^{\text{out}} / 2^{r^{\text{col}}}) \right) \geq |\text{seed}^{\text{out}}.\text{col}|, \\ \log N_{\text{proof}} &\leq \log W_{\text{proof}}^{\text{out}} \leq |\text{seed}^{\text{out}}.\text{col}| + |\text{seed}^{\text{out}}.\text{shared}|, \end{aligned}$$

we can arrange

$$\text{seed}^{\text{out}} := \text{seed}^{\text{out}}.\text{col} \circ \text{seed}^{\text{out}}.\text{shared} \circ \text{seed}^{\text{out}}.\text{row}$$

such that the lowest $\log N_{\text{proof}}$ bits (resp. the highest $r^{\text{out}} - \log N_{\text{proof}}$ bits) can be obtained from $\text{seed}^{\text{out}}.\text{col}$ (resp. $\text{seed}^{\text{out}}.\text{row}$) and $\text{seed}^{\text{out}}.\text{shared}$. Concretely, the row/column verifier will firstly identify the row/column index of the block that contains $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$ and then obtain the row/column index of the query within the block by running the inner PCPP verifier $V^{\text{in}}(\text{seed}^{\text{in}})$.

Rectangular neighbour Listing of V^{comp} . Now we verify that the composed PCPP verifier V^{comp} has the rectangular neighbour listing property with $t_{\text{RNL}}(n) := \text{poly}(t_{\text{RNL}}^{\text{out}}(n), \ell^{\text{in}}, q^{\text{in}}, d^{\text{in}})$. Let (seed, k) be a configuration of V^{comp} , where

$$\text{seed} = (\text{seed}.\text{row} := \text{seed}^{\text{out}}.\text{row}, \text{seed}.\text{col} := \text{seed}^{\text{out}}.\text{col}, \text{seed}.\text{shared} := (\text{seed}^{\text{out}}.\text{shared}, \text{seed}^{\text{in}})).$$

and $k \in [q^{\text{comp}}]$. Assume that the verifier probes the proof matrix $\Pi_{\text{proof}}^{\text{comp}}$ on the k -th query given the randomness seed . By the discussion above, we know that the k -th query of the composed PCPP verifier can be one of the following two cases: a query to $\Pi_{\text{proof}}^{\text{out}}|_{I_{\text{out}}}$ for $I_{\text{out}} \leftarrow V^{\text{out}}(\text{seed}^{\text{out}})$, or a query to $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$.

Assume that the rectangular neighbour listing algorithm for V^{out} partitions $\text{seed}^{\text{out}}.\text{shared}$ into $(\text{seed}^{\text{out}}.\text{shared}.\text{row}, \text{seed}^{\text{out}}.\text{shared}.\text{col})$. We now partition $\text{seed}.\text{shared}$ as follows:

$$\text{seed}.\text{shared}.\text{row} := (\text{seed}^{\text{out}}.\text{shared}.\text{row}, \text{low}(\text{seed}^{\text{in}})),$$

$$\text{seed}.\text{shared}.\text{col} := (\text{seed}^{\text{out}}.\text{shared}.\text{col}, \text{high}(\text{seed}^{\text{in}})).$$

The row and column algorithms A_{row} and A_{col} for the rectangular neighbour listing of V^{comp} work as follows.

Case 1. Given the configuration (seed, k) , the verifier V^{comp} probes the i -th bit of $\Pi_{\text{proof}}^{\text{out}}|_{I_{\text{out}}}$, where the index i depends on the seed^{in} . In other words, the composed PCPP verifier probes the answer of the i -th query made by the outer PCPP verifier when it is ‘‘simulating’’ the outer verifier using the inner PCPP verifier. A neighbour $(\text{seed}' = (\text{seed}'^{\text{out}}, \text{seed}'^{\text{in}}), k')$ of (seed, k) must be a query of the same type, i.e., it is a query to the i' -th bit of $\Pi_{\text{proof}}^{\text{out}}|_{I'_{\text{out}}}$ where the index i' depends on seed'^{in} . Furthermore, the i -th query index in I_{out} must be the same as the i' -th query index in I'_{out} . In such case, the row/column algorithms for rectangular neighbour listing will generate the following list:

1. The row algorithm $A_{\text{row}}^{\text{comp}}$ (resp. the column algorithm $A_{\text{col}}^{\text{comp}}$) firstly runs the outer PCPP verifier and generates the row part (resp. the column part) of I_{out} , and runs the inner PCPP verifier using seed^{in} to obtain the index i defined above.
2. Then it runs the row algorithm $A_{\text{row}}^{\text{out}}$ (resp. the column algorithm $A_{\text{col}}^{\text{out}}$) for the outer PCPP to generate the row part (resp. the column part) of the list of neighbours of $(\text{seed}^{\text{out}}, i)$, denoted by L_{row} (resp. L_{col}).
3. For every $(\text{seed}_j^{\text{out}}.\text{row}, \text{seed}_j^{\text{out}}.\text{shared}.\text{row}, i_j)$ (resp. $(\text{seed}_j^{\text{out}}.\text{col}, \text{seed}_j^{\text{out}}.\text{shared}.\text{col}, i_j)$) in the list L_{row} (resp. L_{col}), we enumerate $(\text{seed}'^{\text{in}}, k') \in \{0, 1\}^{r^{\text{in}}} \times [q^{\text{in}}]$ in lexicographic order. If the k' -th query of V^{in} given the seed'^{in} as the seed is a query to the i' -th bit of $(\Pi_{\text{input}} \circ \Pi_{\text{proof}}^{\text{out}})|_{I'_{\text{out}}}$ (where i' depends on seed'^{in} , and $i' = i_j$, then append

$$(\text{seed}_j^{\text{out}}.\text{row}, \text{seed}_j^{\text{out}}.\text{shared}.\text{row}, \text{low}(\text{seed}'^{\text{in}}), k') \quad (\text{for } A_{\text{row}}^{\text{comp}})$$

$$(\text{seed}_j^{\text{out}}.\text{col}, \text{seed}_j^{\text{out}}.\text{shared}.\text{col}, \text{high}(\text{seed}'^{\text{in}}), k') \quad (\text{for } A_{\text{col}}^{\text{comp}})$$

to L_{row} (resp. L_{col}).

It is easy to check that the requirements of RNL property are satisfied.

Case 2. Given the configuration (seed, k) , the verifier V^{comp} probes the i -th cell of $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$. Recall that for every $\text{seed}^{\text{out}} \in \{0, 1\}^{r^{\text{out}}}$, $\Pi_{\text{proof}}^{\text{in}}(\text{seed}^{\text{out}})$ is arranged in a block of size $H_{\text{proof}}^{\text{in}} \times W_{\text{proof}}^{\text{in}}$ in the proof matrix. The neighbours of (seed, k) need to query the same block, therefore the neighbours must have the same random seed for the outer PCPP verifier. Hence the row/column algorithms will work as follows:

1. The row algorithm $A_{\text{row}}^{\text{comp}}$ (resp. the column algorithm $A_{\text{col}}^{\text{comp}}$) firstly finds the list $\mathcal{L}^{\text{in}} := \{(\text{seed}_j^{\text{in}}, k_j) \in \{0, 1\}^{r^{\text{in}}(d^{\text{out}}(n))} \times [q^{\text{comp}}]\}$ sorted in lexicographic order such that the inner PCPP will query the i_j -th bit of the inner proof on the k_j -th query given $\text{seed}_j^{\text{in}}$ as randomness. This can be done in $\text{poly}(\ell^{\text{in}}, q^{\text{in}}, d^{\text{in}})$ -time by enumerating all possible $(\text{seed}_j^{\text{in}}, k_j)$ and running the inner PCPP verifier.
2. We define the final list of neighbours as

$$\mathcal{L} := \left\{ \left(\text{seed}_j := (\text{seed}^{\text{out}}, \text{seed}_j^{\text{in}}, k_j) \right) : (\text{seed}_j^{\text{in}}, k_j) \in \mathcal{L}^{\text{in}} \right\}.$$

It is easy to check that the list satisfies the promises of the rectangular neighbour listing property.

Other properties. The soundness error and proximity parameter can be found in [BGH⁺06, Section 2.4]. The query complexity, ROP parity-check complexity, and decision complexity can be found in [RSW22, Section 7.2]. We can see that the proof matrix of the composed PCPP verifier has width $W_{\text{proof}}^{\text{comp}} = W_{\text{proof}}^{\text{out}}$ and height $H_{\text{proof}}^{\text{comp}} = W_{\text{proof}}^{\text{out}} + 2^{r^{\text{out}}(n) + r^{\text{in}}(d^{\text{out}}(n))} / W_{\text{proof}}^{\text{out}}$ (recall that $\ell^{\text{in}} = 2^{r^{\text{in}}}$). By the definitions of the random seeds, we can see that: The row and column randomness complexity of V^{comp} is the same as the row and column randomness complexity of V^{out} , respectively; the shared randomness complexity of V^{comp} is the sum of the shared randomness complexity of V^{out} and the randomness complexity of V^{in} .

Remark 6.11. The composed PCPP verifier V^{comp} will use the inner PCPP verifier V^{in} to simulate the outer PCPP verifier V^{out} . This means that the total number of queries and parity-check functions is at most the query complexity of the inner PCPP verifier. Moreover, the decision predicate of V^{comp} (after fixing the random seed) is the decision predicate of V^{in} , where the input bits of the decision circuit of V^{comp} are the parity-check bits and the answers to the queries. For instance, if the decision predicate of V^{in} given seed^{in} is an OR of the answers or their negations, then the decision predicate of V^{comp} given $\text{seed} = (\text{seed}^{\text{in}}, \text{seed}^{\text{out}})$ is also the same OR of its input bits (i.e. the answers to the queries and the parity-check bits).

6.4 Soundness Amplification Preserving Smoothness and Rectangularity

Now we use the technique from [RSW22, Section 7.1.3] to boost the soundness error of a rectangular PCPP. In addition to their original analysis, we need to verify that their expander-walk construction preserves smoothness.

Lemma 6.12 ([VW18, RSW22]). *For every $\lambda \in (0, 1)$, there is some $d = \text{poly}(\lambda^{-1})$ such that the following holds. For every n , there is an expander graph $G_n = (V_n, E_n)$ with second largest eigenvalue at most λ , where $V_n := \{0, 1\}^n$. Moreover, there are d explicit projections (i.e., NC_1^0 circuits) $C_1, C_2, \dots, C_d : \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that for every $x \in V_n$, the d neighbours of x are $C_1(x), C_2(x), \dots, C_d(x)$.*

Verifier	V^{old}	V^{new}
Soundness error	s	μ
Proximity parameter	δ	δ
Row randomness	r_{row}	r_{row}
Column randomness	r_{col}	r_{col}
Shared randomness	r_{shared}	$r_{\text{shared}} + O((1-s)^{-2} \log(\mu^{-1}) \log((1-s)^{-1}))$
Proof matrix height	H_{proof}	H_{proof}
Proof matrix width	W_{proof}	W_{proof}
Query complexity	q	$O(q \cdot (1-s)^{-2} \cdot \log(\mu^{-1}))$
Parity check complexity	p	$O(p \cdot (1-s)^{-2} \cdot \log(\mu^{-1}))$
Decision complexity	d	$O(d \cdot (1-s)^{-2} \log(\mu^{-1}) + \text{poly}(r_{\text{shared}}, r_{\text{row}}, r_{\text{col}}))$

Table 6: The parameters of the soundness amplification ($O(\cdot)$ hides absolute constants).

Lemma 6.13 (Expander Walk, [AB09, Theorem 21.12]). *Let $G = (V, E)$ be a d -regular graph with second largest eigenvalue λ . For every $S \subseteq V$ such that $|S| \leq \beta \cdot |V|$ for some $\beta \in (0, 1)$, let $(X_1, X_2, \dots, X_\ell)$ be a random walk in G with random starting point, then*

$$\Pr [\forall i \in [\ell], X_i \in S] \leq \left((1-\lambda)\sqrt{\beta} + \lambda \right)^{\ell-1}.$$

Lemma 6.14 (Expander Chernoff Bound, [Vad12, Theorem 4.22]). *Let $G = (V, E)$ be a d -regular graph with second largest eigenvalue λ , $B \subseteq V$ be a set of size $|B| = \beta|V|$. Let X_1, X_2, \dots, X_ℓ be random variables denoting a length- ℓ random walk from a random starting point. For every $i \in [\ell]$, we define $B_i = 1$ if $X_i \in B$ and $B_i = 0$ otherwise. Then:*

$$\Pr \left[\left| \frac{1}{\ell} \sum_{i=1}^{\ell} B_i - \beta \right| \geq 2\lambda \right] < 2 \exp(-\Omega(\lambda^2 \ell)).$$

Theorem 6.15. *Suppose that L has a rectangular PCPP verifier V^{old} (resp. a rectangular PCPP verifier V^{old} with ROP), where the parameters are specified in Table 6. Then for every $\mu \in (0, 1)$, L has a rectangular PCPP verifier V^{new} (resp. a rectangular PCPP verifier with ROP), whose parameters are specified in Table 6.*

Moreover, if V^{old} is smooth, then V^{new} is also smooth; if V^{old} has robust soundness (instead of soundness) s with robustness parameter ρ , then V^{old} has robustness soundness μ with robustness parameter $(1-s)\rho/3$.

Proof. Let $\varepsilon := 1-s$, $\lambda := \varepsilon/3$, and $r := r_{\text{row}} + r_{\text{col}} + r_{\text{shared}}$. We construct the following d -regular expander graphs with second largest eigenvalue λ by Lemma 6.12: $G_{\text{row}} = (V_{\text{row}}, E_{\text{row}})$ with $V_{\text{row}} := \{0, 1\}^{r_{\text{row}}}$, $G_{\text{col}} = (V_{\text{col}}, E_{\text{col}})$ with $V_{\text{col}} := \{0, 1\}^{r_{\text{col}}}$, and $G_{\text{shared}} = (V_{\text{shared}}, E_{\text{shared}})$ with $V_{\text{shared}} := \{0, 1\}^{r_{\text{shared}}}$. Let $G = (V, E)$ be the tensor product of these expanders:

$$\begin{aligned} V &:= V_{\text{row}} \times V_{\text{col}} \times V_{\text{shared}} = \{0, 1\}^{r_{\text{row}}} \times \{0, 1\}^{r_{\text{col}}} \times \{0, 1\}^{r_{\text{shared}}}; \\ E &:= \{((u, v, w), (u', v', w')) : (u, u') \in E_{\text{row}}, (v, v') \in E_{\text{col}}, (w, w') \in E_{\text{shared}}\}. \end{aligned}$$

Note that G is a d^3 -regular graph with second largest eigenvalue λ (see [AB09, Lemma 21.17]).

The Construction of V^{new} . The new verifier has the same proof matrix, row randomness, and column randomness as the old verifier V^{old} . The shared random seed of the new verifier V^{new} consists of the shared random seed seed.shared of V^{old} and seed.walk , which is used to sample a random walk in G of length $\ell := O(\lambda^{-2} \log(\mu^{-1}))$. Concretely:

- The random seed seed.walk will be used to sample $\sigma_1, \sigma_2, \dots, \sigma_{3(\ell-1)} \in [d]$. We can see that

$$|\text{seed.walk}| = O(\ell \cdot \log d) = O(\lambda^{-2} \log(\mu^{-1}) \log(\lambda^{-1})).$$

- Let $u_1 := \text{seed.row}$, $v_1 := \text{seed.col}$, and $w_1 := \text{seed.shared}$. We use $\sigma_1, \sigma_2, \dots, \sigma_{\ell-1}$ to specify a length- ℓ random walk $(u_1, u_2, \dots, u_\ell)$ in G_{row} . In particular, let C_1, C_2, \dots, C_d be the projections in Lemma 6.12 for G_{row} . For every $j \in \{1, 2, \dots, \ell-1\}$, we define $u_{j+1} := C_{\sigma_j}(u_j)$. Similarly, we can use the remaining $2(\ell-1)$ bits to specify a random walk $(v_1, v_2, \dots, v_\ell)$ in G_{col} and a random walk $(w_1, w_2, \dots, w_\ell)$ in G_{shared} .

The verifier V^{old} will run the verifier V^{new} for ℓ times with the seeds:

$$(u_1, v_1, w_1), (u_2, v_2, w_2), \dots, (u_\ell, v_\ell, w_\ell),$$

and will accept the proof if V^{old} accepts given all these ℓ seeds. Since seed.walk is treated as the shared randomness of V^{new} and G is obtained from the tensor product of G_{row} , G_{col} , and G_{shared} , it is easy to see that V^{new} is still a rectangular PCPP verifier. The query complexity (and parity-check complexity when V^{old} has ROP) increases by an $\ell = O(\varepsilon^{-2} \log(\mu^{-1}))$ multiplicative factor.

Smoothness. Let $\text{idx} \in [H_{\text{proof}} \cdot W_{\text{proof}}]$ be an index in the proof. Assume that $V^{\text{new}}(\text{seed}, i)$ (resp. $V^{\text{old}}(\text{seed}, i)$) denotes the index in the proof probed by V^{new} (resp. V^{old}) for the i -th query. We can see that

$$\begin{aligned} \gamma &:= \Pr_{\text{seed}, \text{seed.walk}, i \in [q\ell]} [V^{\text{new}}(\text{seed} \circ \text{seed.walk}, i) = \text{idx}] \\ &= \mathbb{E}_{j \in [\ell]} \left[\Pr_{\text{seed}, \text{seed.walk}, i \in [q]} [V^{\text{new}}(\text{seed} \circ \text{seed.walk}, (j-1)\ell + i) = \text{idx}] \right]. \end{aligned}$$

Fix a $j \in [\ell]$. By the definition of V^{new} , we know that $V^{\text{new}}(\text{seed} \circ \text{seed.walk}, (j-1)\ell + i)$ will work as follows: Let $(u, v, w) := (\text{seed.row}, \text{seed.col}, \text{seed.shared})$, and $\sigma_1, \sigma_2, \dots, \sigma_{3\ell}$ be defined as above; V^{new} will choose the j -th node in the random walk on G seeded by seed.walk starting from $(u_1 := u, v_1 := v, w_1 := w)$ as the seed for V^{old} , and probe the proof according to the i -th query of V^{old} . Since the expander graph is regular, each $\text{seed} \in \{0, 1\}^r$ is equally likely to be selected from a random walk with a random starting point. Hence

$$\begin{aligned} &\Pr_{\text{seed}, \text{seed.walk}, i \in [q]} [V^{\text{new}}(\text{seed} \circ \text{seed.walk}, (j-1)\ell + i) = \text{idx}] \\ &= \Pr_{\text{seed}, i \in [q]} [V^{\text{old}}(\text{seed}, i) = \text{idx}] \\ &= \frac{1}{H_{\text{proof}} \cdot W_{\text{proof}}}. \end{aligned}$$

This means that $\gamma = 1/(H_{\text{proof}} \cdot W_{\text{proof}})$, i.e., every bit in the proof is equally likely to be probed.

Soundness. Assume that $x \in \{0, 1\}^n$ is δ -far from being in L and Π be an arbitrary proof. We say a node (u, v, w) in the expander graph $G = (V, E)$ to be *bad* if V^{old} accepts (x, Π) with the random seed $\text{seed.row} := u$, $\text{seed.col} := v$, and $\text{seed.shared} := w$. Let B be the set of all bad nodes, then $|B| \leq s \cdot |V|$. Note that the new verifier accepts (x, Π) if and only if a length- ℓ random walk on G from a random starting point only accesses bad nodes. By Lemma 6.13, we can see that

$$\Pr[V^{\text{new}} \text{ accepts } (x, \Pi)] \leq ((1 - \lambda)\sqrt{s} + \lambda)^{\ell-1} \leq \left(1 - \frac{\varepsilon}{5}\right)^{\ell-1} \leq \exp\left(-\frac{\varepsilon(\ell-1)}{5}\right) \leq \mu,$$

when $\ell \geq 10 \cdot \varepsilon^{-1} \ln(\mu^{-1})$.

Robust Soundness. Assume that the original PCPP verifier V^{old} has robust soundness s with robustness parameter ρ (instead of only soundness s), we need to show that V^{new} has robustness soundness μ . Let $x \in \{0, 1\}^n$ be δ -far from L and Π be an arbitrary proof. We say a node (u, v, w) in the expander graph $G = (V, E)$ is *bad* if given the randomness $\text{seed} = (\text{seed.row}, \text{seed.col}, \text{seed.shared}) := (u, v, w)$, the fraction of bits read by the old PCPP verifier that we need to change to make V^{old} accepts (x, Π) , denoted by $\hat{\delta}(\text{seed})$, is at most ρ .

Let B be the set of bad nodes and X_1, X_2, \dots, X_ℓ be the random variables denoting a random walk from a random starting point (equivalently, denoting the randomness V^{new} used to simulate V^{old}). By the robustness soundness of V^{old} , we know that $|B| \leq s \cdot |V|$. Let $B_i = 1$ when $X_i \in B$ and 0 otherwise. By Lemma 6.14, we can see that

$$\Pr\left[\frac{1}{\ell} \sum_{i=1}^{\ell} B_i \geq s + 2\lambda\right] \leq 2 \exp(-\Omega(\lambda^2 \ell)) \leq \mu$$

when $\ell = O(\lambda^{-2} \log(\mu^{-1}))$. As a result, with probability at least $1 - \mu$, the fraction of bits read by V^{new} that we need to change to make V^{new} accepts (x, Π) is at least

$$(1 - (s + 2\lambda))\rho \geq \varepsilon\rho/3.$$

This satisfies the requirement of robust soundness μ with robustness parameter $\varepsilon\rho/3$. \square

6.5 Final Construction

Theorem 6.16 ([Mie09, RSW22]). *Let L be a pair language in $\text{NTIME}[T(n)]$ for some non-decreasing function $T : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$. For all constants $s, \delta > 0$, L has a PCPP verifier with randomness complexity $\log T(n) + O(\log \log T(n))$, soundness error s , proximity parameter δ , query complexity $O(1)$, and decision complexity $\text{polylog}(T(n))$.*

Theorem 2.14 (Smooth and Rectangular PCPP). *For all constants $\delta \in (0, 1)$ and $s \in (0, 1)$, there is a constant $q \geq 1$ such that the following holds. Let $m = m(n)$, $T(n)$, $w_{\text{proof}}(n)$, $w_{\text{input}}(n)$ be good functions such that $1 \leq m(n) \leq (\log T(n))^{0.1}$, $n \leq T(n) \leq 2^{\text{poly}(n)}$, $w_{\text{proof}}(n) \leq \log T(n)$, and $w_{\text{input}}(n) \leq \log n$. Then there are good functions $h_{\text{proof}}(n)$ and $h_{\text{input}}(n)$ satisfying*

$$\begin{aligned} h_{\text{proof}}(n) &:= \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}(n), & \text{and} \\ h_{\text{input}}(n) &:= \lceil \log n \rceil - w_{\text{input}}(n). \end{aligned}$$

such that the following holds.

Soundness error	s
Proximity parameter	δ
Row randomness	$r_{\text{row}} := h_{\text{proof}} - (5/m) \log T(n)$
Column randomness	$r_{\text{col}} := w_{\text{proof}} - (5/m) \log T(n)$
Shared randomness	$r_{\text{shared}} := (10/m) \log T(n) + O(\log \log T(n) + m \log m)$
Query complexity	q
Parity check complexity	q
Decision complexity	$\text{poly}(T(n)^{1/m})$

Table 7: Parameters of the PCPP constructed in Theorem 2.14.

Suppose that $w_{\text{proof}}, h_{\text{proof}} \geq (5/m) \log T(n)$, and that for some absolute constant $C \geq 1$,

$$\frac{w_{\text{input}}(n)}{w_{\text{proof}}(n)}, \frac{h_{\text{input}}(n)}{h_{\text{proof}}(n)} \leq 1 - \frac{Cm^2 \log \log T(n)}{\log T(n)}.$$

Let $W_{\text{proof}}(n) := 2^{w_{\text{proof}}(n)}$, $H_{\text{proof}}(n) := 2^{h_{\text{proof}}(n)}$, $W_{\text{input}}(n) := 2^{w_{\text{input}}(n)}$, and $H_{\text{input}}(n) := 2^{h_{\text{input}}(n)}$. Then $\text{NTIME}[T(n)]$ has a smooth and rectangular PCP of proximity with an $H_{\text{input}}(n) \times W_{\text{input}}(n)$ input matrix and an $H_{\text{proof}}(n) \times W_{\text{proof}}(n)$ proof matrix, whose other parameters are specified in Table 2.

Proof. The high-level roadmap of the proof is as follows.

1. By Theorem 6.8, we can obtain a robust and rectangular PCPP verifier V^{out} with RNL property for $t_{\text{RNL}} = \text{poly}(\log T(n), m^m)$ and query complexity $T(n)^{1/m} \cdot \text{polylog}(T(n))$.
2. Let V^{in} be a PCPP verifier for $\text{CIRCUIT-EVAL}^\perp$ with constant query complexity. We compose V^{out} and V^{in} by Theorem 6.10 to obtain a rectangular PCPP verifier V^{comp} with RNL property.
3. We smoothen V^{comp} by Theorem 6.3 to obtain a smooth and rectangular PCPP V^{smth} with constant query complexity, whose soundness error is some constant $s^{\text{smth}} \in (0, 1)$.
4. By Theorem 6.15, we reduce the soundness error to s while still maintaining the query complexity to be a (larger) constant.

Robust and Rectangular PCPP. Let $\delta \in (0, 1)$ and $s \in (0, 1)$ be some constants; q be a large constant to be determined that only depends on δ and s ; C be a large constant; $m = m(n)$, $T(n)$, $w_{\text{input}}(n)$, $h_{\text{input}}(n)$, $w_{\text{proof}}(n)$, and $h_{\text{proof}}(n)$ be defined as above. Let $w_{\text{proof}}^{\text{out}}(n) := w_{\text{proof}}(n) - O(\log \log T(n) + m \log m)$ where the concrete value will be determined later. We will set $h_{\text{proof}}(n) = h_{\text{proof}}^{\text{out}}(n) + O(\log \log T(n) + m \log m)$ for some good function $h_{\text{proof}}^{\text{out}}(n)$ (which is actually the proof height parameter of the outer PCPP). We check the technical conditions of Theorem 6.8 holds; in particular we need to ensure that

Claim 6.17. For some constant C' that could be made large enough (depending on C),

$$\frac{w_{\text{input}}}{w_{\text{proof}}^{\text{out}}}, \frac{h_{\text{input}}}{h_{\text{proof}}^{\text{out}}} \leq 1 - \frac{C' m \log \log T(n)}{\log T(n)}.$$

Proof. Since $w_{\text{proof}} \geq (5/m) \log T(n)$, and $w_{\text{proof}} - w_{\text{proof}}^{\text{out}} \leq \alpha_1(\log \log T(n) + m \log m)$ for some constant α_1 , it follows that

$$\begin{aligned} \frac{w_{\text{input}}}{w_{\text{proof}}^{\text{out}}} &\leq \frac{w_{\text{input}}}{w_{\text{proof}}} \left(1 + \frac{\alpha_1(\log \log T(n) + m \log m)}{w_{\text{proof}} - \alpha_1(\log \log T(n) + m \log m)} \right) \\ &\leq 1 - \frac{Cm^2 \log \log T(n)}{\log T(n)} + \frac{\Theta(m) \cdot (\log \log T(n) + m \log m)}{\log T(n)} \\ &\leq 1 - \frac{C'm \log \log T(n)}{\log T(n)}. \end{aligned}$$

The same argument works for $\frac{h_{\text{input}}}{h_{\text{proof}}^{\text{out}}}$. ◇

By Theorem 6.8, we can construct a robust and rectangular PCPP verifier V^{out} for L with RNL property and other parameters as follows:

- Proximity parameter $\delta^{\text{out}} := \delta$.
- Robust soundness error $s^{\text{out}} := 1 - \rho^{\text{out}}$ with robustness parameter ρ^{out} , where $\rho^{\text{out}} \in (0, 1)$ is some constant depending on δ .
- Proof matrix size $H_{\text{proof}}^{\text{out}}(n) \times W_{\text{proof}}^{\text{out}}(n)$, where $H_{\text{proof}}^{\text{out}} = 2^{h_{\text{proof}}^{\text{out}}}$, $W_{\text{proof}}^{\text{out}} = 2^{w_{\text{proof}}^{\text{out}}}$. The proof height parameter $h_{\text{proof}}^{\text{out}}$, which is given by Theorem 6.8, satisfies

$$h_{\text{proof}}^{\text{out}} = \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}^{\text{out}}(n).$$

- Row randomness complexity $r_{\text{row}}^{\text{out}} = h_{\text{proof}}^{\text{out}} - (4/m) \log T(n)$.
- Column randomness complexity $r_{\text{col}}^{\text{out}} = w_{\text{proof}}^{\text{out}} - (4/m) \log T(n)$.
- Shared randomness complexity $r_{\text{shared}}^{\text{out}} = (7/m) \log T(n) + O(\log \log T(n) + m \log m)$.
- Query complexity $q^{\text{out}}(n) = T(n)^{1/m} \cdot \text{polylog}(T(n))$.
- Decision complexity $d^{\text{out}}(n) = T(n)^{1/m} \cdot \text{polylog}(T(n))$.
- RNL time complexity $t_{\text{RNL}}^{\text{out}}(n) = \text{poly}(\log T(n), m^m)$.

Reducing the Query Complexity. By Theorem 6.16, we can construct a PCPP verifier V^{in} for $\text{CIRCUIT-EVAL}^\perp$ with input length $d^{\text{out}}(n)$ and other parameters specified as follows.

- Randomness complexity $r^{\text{in}}(n) = \log d^{\text{out}}(n) + O(\log \log d^{\text{out}}(n)) = \frac{1}{m} \log T(n) + O(\log \log T(n))$.
- Soundness error $s^{\text{in}} := \rho^{\text{out}}/2$.
- Proximity parameter $\delta^{\text{in}} := \rho^{\text{out}}/2$.
- Query complexity $q^{\text{in}} = O(1)$ is a constant depends on s^{in} and δ^{in} , which further means that it only depends on δ .
- Decision complexity $d^{\text{in}}(d^{\text{out}}(n)) = \text{polylog}(T(n))$.

Without loss of generality, we assume that the proof length $\ell^{\text{in}}(n) = 2^{r^{\text{in}}(n)}$.

We now construct V^{comp} by composing V^{out} and V^{in} by Theorem 6.10. We firstly check the requirements of the composition theorem.

- $q^{\text{in}} = O(1)$, $\rho^{\text{out}} \geq \delta^{\text{in}} = \rho^{\text{out}}/2$, $\ell^{\text{in}} = 2^{r^{\text{in}}}$.
- $\log W_{\text{proof}}^{\text{out}} = r_{\text{col}}^{\text{out}} + (4/m) \log T(n) \geq r_{\text{col}}^{\text{out}}$.
- $\log W_{\text{proof}}^{\text{out}} = r_{\text{col}}^{\text{out}} + (4/m) \log T(n) \leq r_{\text{col}}^{\text{out}} + r_{\text{shared}}^{\text{out}}$.

Hence we can obtain a rectangular PCPP V^{comp} with ROP that has RNL property. The parameters of the composed PCPP are as follows.

- Soundness error $s^{\text{comp}} := 1 - (1 - s^{\text{out}}) \cdot (1 - s^{\text{in}}) < 1$ that only depends on δ .
- Proximity parameter $\delta^{\text{comp}} := \delta^{\text{out}} = \delta$.
- Row randomness complexity $r_{\text{row}}^{\text{comp}} = r_{\text{row}}^{\text{out}} = h_{\text{proof}}^{\text{out}} - (4/m) \log T(n)$.
- Column randomness complexity $r_{\text{col}}^{\text{comp}} = r_{\text{col}}^{\text{out}} = w_{\text{proof}}^{\text{out}} - (4/m) \log T(n)$.
- Shared randomness complexity $r_{\text{shared}}^{\text{comp}} = r_{\text{shared}}^{\text{out}} + r^{\text{in}} = (7/m) \log T(n) + O(\log \log T(n) + m \log m) + (1/m) \log T(n) + O(\log \log T(n)) = (8/m) \log T(n) + O(\log \log T(n) + m \log m)$.
- Proof matrix height $H_{\text{proof}}^{\text{comp}} = H_{\text{proof}}^{\text{out}} + 2^{r^{\text{out}} + r^{\text{in}}} / W_{\text{proof}}^{\text{out}}$.
- Proof matrix width $W_{\text{proof}}^{\text{comp}} = W_{\text{proof}}^{\text{out}}$.
- Query complexity $q^{\text{comp}} = q^{\text{in}} = O(1)$ that only depends on δ .
- ROP parity check complexity $p^{\text{comp}} = q^{\text{in}} = O(1)$.
- Decision complexity $d^{\text{comp}}(n) = d^{\text{in}}(d^{\text{out}}(n)) = \text{polylog}(T(n))$.
- RNL time complexity $t_{\text{RNL}}^{\text{comp}}(n) = \text{poly}(t_{\text{RNL}}^{\text{out}}(n), \ell^{\text{in}}, q^{\text{in}}, d^{\text{in}}) = \text{poly}(T(n)^{1/m})$, where $\text{poly}(\cdot)$ hides some absolute constant on the exponent. Note that $t_{\text{RNL}}^{\text{out}}(n) = \text{poly}(\log T(n), m^m) \leq \text{poly}(T(n)^{1/m})$, since $m \leq (\log T(n))^{0.1}$.

Smoothing via RNL. Now we apply Theorem 6.3 to obtain a smooth and rectangular PCPP V^{smth} with $\mu := (1 - s^{\text{comp}})/2$ and other parameters as follows.

- Soundness error $s^{\text{smth}} := s^{\text{comp}} + \mu < 1$ that only depends on δ .
- Proximity parameter $\delta^{\text{smth}} := \delta^{\text{comp}} = \delta$.
- Row randomness complexity $r_{\text{row}}^{\text{smth}} := r_{\text{row}}^{\text{comp}} = h_{\text{proof}}^{\text{out}} - (4/m) \log T(n)$.
- Column randomness complexity $r_{\text{col}}^{\text{smth}} := r_{\text{col}}^{\text{comp}} = w_{\text{proof}}^{\text{out}} - (4/m) \log T(n)$.
- Shared randomness complexity $r_{\text{shared}}^{\text{smth}} = r_{\text{shared}}^{\text{comp}} = (8/m) \log T(n) + O(\log \log T(n) + m \log m)$.
- Proof matrix width $W_{\text{proof}}^{\text{smth}} = 2^{r_{\text{col}}^{\text{comp}} + r_{\text{shared}}^{\text{comp}}/2} = 2^{w_{\text{proof}}^{\text{out}}} \cdot \text{poly}(\log T(n), m^m)$. Note that here we set $w_{\text{proof}}^{\text{out}}$ carefully so that $r_{\text{col}}^{\text{comp}} + r_{\text{shared}}^{\text{comp}}/2 = w_{\text{proof}}^{\text{out}} + O(\log \log T(n) + m \log m) = w_{\text{proof}}^{\text{out}}$. This means that the proof matrix width is exactly $2^{w_{\text{proof}}^{\text{out}}}$.
- Proof matrix height $H_{\text{proof}}^{\text{smth}} = q^{\text{comp}} \cdot 2^{r_{\text{row}}^{\text{comp}} + r_{\text{shared}}^{\text{comp}}/2} = 2^{h_{\text{proof}}^{\text{out}}} \cdot \text{poly}(\log T(n), m^m)$. Since $q^{\text{comp}} = O(1)$, we can add $O(1)$ dummy queries to the composed PCPP V^{comp} so that q^{comp} becomes a power of two. We then set

$$\begin{aligned}
h_{\text{proof}} &= \log H_{\text{proof}}^{\text{smth}} \\
&= h_{\text{proof}}^{\text{out}} + O(\log \log T + m \log m) \\
&= \log T(n) + \Theta(m \log \log T(n)) + O(\log \log T + m \log m) \\
&\quad - (w_{\text{proof}}^{\text{out}}(n) - O(\log \log T + m \log m)) \\
&= \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}^{\text{out}}(n).
\end{aligned}$$

- Query complexity $q^{\text{smth}} = \text{poly}(q^{\text{comp}}/\mu) = O(1)$ that only depends on δ .
- ROP parity check complexity $p^{\text{smth}} = p^{\text{comp}} = O(1)$.
- Decision complexity $d^{\text{smth}}(n) = \text{poly}(d^{\text{comp}}(n), q^{\text{comp}}/\mu, t_{\text{RNL}}^{\text{comp}}(n)) = \text{poly}(T(n)^{1/m})$, where $\text{poly}(\cdot)$ hides some absolute constant on the exponent.

Amplifying the Soundness Error. Finally, we boost the soundness error of V^{smth} to be s by Theorem 6.15, to obtain a smooth and rectangular PCPP with parameters specified as follows.

- Soundness error s .
- Proximity parameter δ .
- Row randomness complexity $h_{\text{proof}}^{\text{out}} - (4/m) \log T(n) \geq h_{\text{proof}} - (5/m) \log T(n)$.
- Column randomness complexity $w_{\text{proof}}^{\text{out}} - (4/m) \log T(n) \geq w_{\text{proof}} - (5/m) \log T(n)$.
- Shared randomness complexity $(8/m) \log T(n) + O(\log \log T(n) + m \log m)$.
- Proof matrix height $2^{h_{\text{proof}}}$ and proof matrix width $2^{w_{\text{proof}}}$.
- Query complexity $q = \text{poly}(q^{\text{smth}}) = O(1)$ that depends on δ and s .
- ROP parity check complexity $\text{poly}(p^{\text{smth}}) = O(1)$ that depends on δ and s .
- Decision complexity $\text{poly}(d^{\text{smth}}(n)) = \text{poly}(T(n)^{1/m})$.

We can move some bits from the row and column randomness to the shared randomness, so that the row and column randomness complexity become exactly $h_{\text{proof}} - (5/m) \log T(n)$ and $w_{\text{proof}} - (5/m) \log T(n)$, respectively, and the shared randomness complexity becomes $(10/m) \log T(n) + O(\log \log T(n) + m \log m)$. This completes the construction. \square

7 Construction of Rectangular PCPPs with Low Query Complexity

Recall that in our framework of solving range avoidance and hard partial truth table, the query complexity of the PCPPs will affect the circuit class for which we need to construct satisfying-pair algorithms. In this section, we construct a rectangular (but not necessarily smooth) PCPP with query complexity 3. We further construct a 2-query PCPP with a constant gap between the completeness and soundness parameters (instead of having perfect completeness).

7.1 A 3-Query PCPP for CIRCUIT-EVAL $^\perp$

Theorem 7.1 ([CW19b], Lemma 24). *For every constant $\delta > 0$, there is a constant $s \in (0, 1)$ and a PCP of proximity for CIRCUIT-EVAL with proximity δ , soundness error s , randomness complexity $O(\log n)$, query complexity $q = 3$, and decision complexity $\text{polylog}(n)$. Moreover, the decision predicate is an OR of the 3 answers to the queries or their negations.*

We need the following standard composition theorem for PCP of Proximity from [BGH⁺06] to construct 3-query PCPPs for any pair language in $\text{NTIME}[T(n)]$.

Theorem 7.2 ([BGH⁺06]). *Let $r^{\text{out}}, r^{\text{in}}, d^{\text{out}}, d^{\text{in}}, q^{\text{in}} : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon^{\text{out}}, \varepsilon^{\text{in}}, \rho^{\text{out}}, \delta^{\text{in}}, \delta^{\text{out}} : \mathbb{N} \rightarrow [0, 1]$. Suppose that:*

- Language L has a robust PCPP verifier V^{out} with randomness complexity $r^{\text{out}}(n)$, decision complexity $d^{\text{out}}(n)$, robust soundness error $1 - \varepsilon^{\text{out}}(n)$, robustness parameter $\rho^{\text{out}}(n)$, and proximity parameter $\delta^{\text{out}}(n)$.
- CIRCUIT-EVAL has a PCPP verifier V^{in} with randomness complexity $r^{\text{in}}(n)$, query complexity $q^{\text{in}}(n)$, decision complexity $d^{\text{in}}(n)$, soundness error $1 - \varepsilon^{\text{in}}(n)$, and proximity parameter $\delta^{\text{in}}(n)$.
- $\delta^{\text{in}}(d^{\text{out}}(n)) \leq \rho^{\text{out}}(n)$ for every n .

Then L has a PCPP Verifier V^{comp} with randomness complexity $r^{\text{out}}(n) + r^{\text{in}}(d^{\text{out}}(n))$, query complexity $q^{\text{in}}(d^{\text{out}}(n))$, decision complexity $d^{\text{in}}(d^{\text{out}}(n))$, soundness error $1 - \varepsilon^{\text{out}}(n) \cdot \varepsilon^{\text{in}}(d^{\text{out}}(n))$, and proximity parameter $\delta^{\text{out}}(n)$.

Theorem 7.3. *Let L be a pair language in $\text{NTIME}[T(n)]$ for some non-decreasing function $T : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$. For every constant δ , there is a constant $s \in (0, 1)$ and a PCPP of proximity for L with randomness complexity $\log T(n) + O(\log \log T(n))$, decision complexity $\text{poly}(\log \log T(n))$, soundness error s , proximity parameter δ , and query complexity $q = 3$.*

Proof. Let L be a pair language in $\text{NTIME}[T(n)]$ and $\delta > 0$. We will compose the following two PCPP verifiers with Theorem 7.2:

- By Theorem 7.1, for every $\delta^{\text{in}} > 0$, there is a constant $s^{\text{in}} \in (0, 1)$ and a PCPP verifier $V_{\delta^{\text{in}}}^{\text{in}}$ for CIRCUIT-EVAL with randomness complexity $r^{\text{in}} = O(\log n)$, soundness error s^{in} , proximity parameter δ^{in} , query complexity $q^{\text{in}} = 3$, and decision complexity $d^{\text{in}} = \text{polylog}(n)$.
- By Theorem 6.16, for all constants $\delta^{\text{out}}, s^{\text{out}} > 0$, there is a constant q^{out} and a PCPP V^{out} with randomness complexity $r^{\text{out}} = \log T(n) + O(\log \log T(n))$, soundness error s^{out} , proximity parameter δ^{out} , query complexity q^{out} , and decision complexity $d^{\text{out}} = \text{polylog}(T(n))$. Since $q^{\text{out}} = O(1)$, V^{out} is trivially a robust PCPP with robustness parameter $\rho^{\text{out}} = 1/q^{\text{out}}$.

Fix $\delta^{\text{out}} = \delta$, $s^{\text{out}} = 0.5$, $\delta^{\text{in}} = 1/(2q^{\text{out}})$, and $s = 1 - 0.5 \cdot (1 - s^{\text{in}})$. It is clear that $\delta^{\text{in}}(d^{\text{out}}(n)) \leq \rho^{\text{out}}(n)$. By Theorem 7.2, we can obtain a PCPP verifier V^{comp} for L with the following parameters:

- Randomness complexity $r^{\text{out}} + r^{\text{in}}(d^{\text{out}}(n)) = \log T(n) + O(\log \log T(n))$.
- Decision complexity $d^{\text{in}}(d^{\text{out}}(n)) = \text{poly}(\log \log T(n))$.
- Soundness error $1 - (1 - s^{\text{out}}(n)) \cdot (1 - s^{\text{in}}(d^{\text{out}}(n))) = s$.
- Proximity parameter $\delta^{\text{out}}(n) = \delta$.
- Query complexity $q^{\text{in}}(d^{\text{out}}(n)) = 3$.

This satisfies our requirements. □

7.2 A 3-Query Rectangular PCPP

Now we construct a 3-query rectangular PCPP by composing the PCPP constructions in Theorem 6.8 and Theorem 7.3 using the composition theorem (see Theorem 6.10).

Theorem 7.4 (3-Query Rectangular PCPP). *For every constant $\delta \in (0, 1)$, there is a constant $s \in (0, 1)$ such that the following holds. Let $m = m(n)$, $T(n)$, $w_{\text{proof}}(n)$, $w_{\text{input}}(n)$ be good functions such that $1 \leq m \leq (\log T(n))^{0.1}$, $n \leq T(n) \leq 2^{\text{poly}(n)}$, $w_{\text{proof}}(n) \leq \log T(n)$, and $w_{\text{input}}(n) \leq \log n$. Then there are good functions $h_{\text{proof}}(n)$ and $h_{\text{input}}(n)$ satisfying*

$$\begin{aligned} h_{\text{proof}}(n) &= \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}(n), & \text{and} \\ h_{\text{input}}(n) &= \lceil \log n \rceil - w_{\text{input}}(n). \end{aligned}$$

such that the following holds.

Suppose that $w_{\text{proof}}, h_{\text{proof}} \geq (5/m) \log T(n)$, and that for some absolute constant $C \geq 1$,

$$\frac{w_{\text{input}}(n)}{w_{\text{proof}}(n)}, \frac{h_{\text{input}}(n)}{h_{\text{proof}}(n)} \leq 1 - \frac{Cm \log \log T(n)}{\log T(n)}.$$

Let $W_{\text{proof}}(n) := 2^{w_{\text{proof}}(n)}$, $H_{\text{proof}}(n) := 2^{h_{\text{proof}}(n)}$, $W_{\text{input}}(n) := 2^{w_{\text{input}}(n)}$, and $H_{\text{input}}(n) := 2^{h_{\text{input}}(n)}$. Then $\text{NTIME}[T(n)]$ has a rectangular PCPP of proximity with an $H_{\text{proof}}(n) \times W_{\text{proof}}(n)$ proof matrix and an $H_{\text{input}}(n) \times W_{\text{input}}(n)$ input matrix, whose other parameters are specified in Table 8.

Moreover, the total number of queries and parity-check bits is at most 3; and for every seed.shared, the decision predicate $\text{VDec} \leftarrow V_{\text{dec}}(\text{seed.shared})$ of the rectangular PCPP verifier is an OR of its 3 input bits or their negations, where each input is either a query answer or a parity-check bit.

Soundness error	s
Proximity parameter	δ
Row randomness	$h_{\text{proof}} - (5/m) \log T(n)$
Column randomness	$w_{\text{proof}} - (5/m) \log T(n)$
Shared randomness	$(10/m) \log T(n) + O(\log \log T(n) + m \log m)$
Query complexity	3
Parity check complexity	
Decision complexity	$\text{poly}(\log \log T(n))$

Table 8: Parameters of the PCPP constructed in Theorem 7.4.

Proof. Let $L \in \text{NTIME}[T(n)]$ and $m \geq 1, \delta > 0$ be constants; $T(n), w_{\text{proof}}(n), h_{\text{proof}}(n), w_{\text{input}}(n), h_{\text{input}}(n)$, and C be defined as above. In one sentence, we compose the robust and rectangular PCPP verifier (Theorem 6.8) with the 3-query PCPP verifier (Theorem 7.3) using the composition theorem (Theorem 6.10).

Outer PCPP. Let $w_{\text{proof}}^{\text{out}}(n) := w_{\text{proof}}(n)$. By Theorem 6.8, we can construct a robust and rectangular PCPP verifier V^{out} for L with parameters as follows:

- Robust soundness error $s^{\text{out}} \in (0, 1)$ with robustness parameter $\rho^{\text{out}} := 1 - s^{\text{out}}$.
- Proximity parameter $\delta^{\text{out}} := \delta$.
- Proof matrix size $H_{\text{proof}}^{\text{out}}(n) \times W_{\text{proof}}^{\text{out}}(n)$, where $H_{\text{proof}}^{\text{out}} = 2^{h_{\text{proof}}^{\text{out}}}$, $W_{\text{proof}}^{\text{out}} = 2^{w_{\text{proof}}^{\text{out}}}$, and $h_{\text{proof}}^{\text{out}} = \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}(n)$.
- Row randomness complexity $r_{\text{row}}^{\text{out}} = h_{\text{proof}}^{\text{out}} - (4/m) \log T(n)$.
- Column randomness complexity $r_{\text{col}}^{\text{out}} = w_{\text{proof}}^{\text{out}} - (4/m) \log T(n)$.
- Shared randomness complexity $r_{\text{shared}}^{\text{out}} = (7/m) \log T(n) + O(\log \log T(n) + m \log m)$.
- Query complexity $q^{\text{out}}(n) = T(n)^{1/m} \cdot \text{polylog}(T(n))$.
- Decision complexity $d^{\text{out}}(n) = T(n)^{1/m} \cdot \text{polylog}(T(n))$.

Inner PCPP. Let $\delta^{\text{in}} := \rho^{\text{out}}/2$. By Theorem 7.3, there is a constant $s^{\text{in}} \in (0, 1)$ and a PCPP verifier V^{in} for $\text{CIRCUIT-EVAL}^\perp$ with randomness complexity $\log T(n) + O(\log \log T(n))$, soundness error s^{in} , proximity parameter δ^{in} , query complexity $q = 3$, and decision complexity $d^{\text{in}} = \text{poly}(\log \log T(n))$. Without loss of generality, we assume that the proof length is $\ell^{\text{in}} = 2^{r^{\text{in}}}$.

Composition. We now compose V^{out} with the inner PCPP V^{in} by Theorem 6.10. We first check that the technical conditions are satisfied.

- $q^{\text{in}} = 3 = O(1)$, $\rho^{\text{out}} \geq \delta^{\text{in}}$, $\ell^{\text{in}} = 2^{r^{\text{in}}}$.
- Since $r_{\text{col}}^{\text{out}} = w_{\text{proof}}^{\text{out}} - (4/m) \log T(n)$ and $r_{\text{shared}}^{\text{out}} \geq (7/m) \log T(n)$, we know that $r_{\text{col}}^{\text{out}} \leq w_{\text{proof}}^{\text{out}} \leq r_{\text{col}}^{\text{out}} + r_{\text{shared}}^{\text{out}}$.

By Theorem 6.10, we can obtain a rectangular PCPP V^{comp} with ROP, whose parameters are as follows:

- Soundness error $s^{\text{comp}} = 1 - (1 - s^{\text{in}}) \cdot (1 - s^{\text{out}}) < 1$.
- Proximity parameter $\delta^{\text{comp}} = \delta^{\text{out}} = \delta$.
- Query complexity and ROP parity checking complexity $q^{\text{comp}} = q = 3$.
- Proof matrix width $W_{\text{proof}} = W_{\text{proof}}^{\text{out}} = 2^{w_{\text{proof}}^{\text{out}}}$.

- Proof matrix height $H_{\text{proof}} = H_{\text{proof}}^{\text{out}} + 2^{r^{\text{out}}(n) + r^{\text{in}}(d^{\text{out}}(n))} / W_{\text{proof}}^{\text{out}}$. Note that

$$\begin{aligned} H_{\text{proof}}^{\text{out}} &= T(n) \log^{\Theta(m)} T(n) / W_{\text{proof}}^{\text{out}}(n), \\ r^{\text{out}} &= (1 - 1/m) \log T(n) + \Theta(m \log \log T(n)), \\ r^{\text{in}}(d^{\text{out}}) &= (1/m) \log T(n) + O(\log \log(T(n))), \end{aligned}$$

hence $H_{\text{proof}} = T(n) \cdot \log^{\Theta(m)}(T(n)) / W_{\text{proof}}$. Without loss of generality, we assume that H_{proof} is a power of two. We then define

$$h_{\text{proof}} := \log H_{\text{proof}} = \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}.$$

- Decision complexity $d^{\text{in}}(d^{\text{out}}(n)) = \text{poly}(\log \log T(n))$.

Now we determine the randomness complexity of the composed PCPP verifier. Note that

$$\begin{aligned} r_{\text{shared}} &= r_{\text{shared}}^{\text{out}} + r^{\text{in}} \\ &= (7/m) \log T(n) + O(\log \log T(n) + m \log m) + (1/m) \log T(n) + O(\log \log T(n)) \\ &= (8/m) \log T(n) + O(\log \log T(n) + m \log m), \\ r_{\text{row}} &= r_{\text{row}}^{\text{out}} = h_{\text{proof}} - (4/m) \log T(n) - \Theta(m \log \log T(n)) \geq h_{\text{proof}} - (5/m) \log T(n), \\ r_{\text{col}} &= r_{\text{col}}^{\text{out}} = w_{\text{proof}} - (4/m) \log T(n). \end{aligned}$$

Since we can always move some portion of `seed.row` or `seed.col` into `seed.shared`, we can simply assume that $r_{\text{row}} = h_{\text{proof}} - (5/m) \log T(n)$, $r_{\text{col}} = w_{\text{proof}} - (5/m) \log T(n)$, and $r_{\text{shared}} = (10/m) \log T(n) + O(\log \log T(n) + m \log m)$.

Moreover, by Remark 6.11 and the fact that the decision predicate of V^{in} is an OR of the answers or their negations (see Theorem 7.3), we know that the total number of queries and parity-check bits of V^{comp} is at most 3, and that for every `seed.shared`, the decision predicate $\text{VDec} \leftarrow V_{\text{dec}}^{\text{comp}}(\text{seed.shared})$ of V^{comp} is an OR of its input bits (i.e., query answers and parity-check bits) or their negations. \square

7.3 A 2-Query Rectangular PCPP with Imperfect Completeness

Following the construction in [CW19b, Appendix A], we can also construct a 2-query rectangular PCPP with a constant gap between the completeness and soundness parameters, using the following classical gadget due to [GJS76].

Lemma 7.5. *Let $x_1, x_2, x_3 \in \{0, 1\}$ be Boolean variables. If $x_1 \vee x_2 \vee x_3$, then there is an $y \in \{0, 1\}$ such that at least 7 of the following 10 constraints are satisfied:*

$$x_1, x_2, x_3, \overline{x_1} \vee \overline{x_2}, \overline{x_1} \vee \overline{x_3}, \overline{x_2} \vee \overline{x_3}, y, x_1 \vee \overline{y}, x_2 \vee \overline{y}, x_3 \vee \overline{y}. \quad (14)$$

Otherwise, at most 6 of the constraints in Eq. (14) are satisfied for any $y \in \{0, 1\}$. Moreover, every $x_1, x_2, x_3, y \in \{0, 1\}$ satisfies at most 7 of the above 10 constraints.

Theorem 7.6 (2-Query Rectangular PCPP). *For every constant $\delta \in (0, 1)$, there are constants $0 < s < c < 1$ such that the following holds. Let $m = m(n)$, $T(n)$, $w_{\text{proof}}(n)$, $w_{\text{input}}(n)$ be good functions such that $1 \leq m \leq (\log T(n))^{0.1}$, $n \leq T(n) \leq 2^{\text{poly}(n)}$, $w_{\text{proof}}(n) \leq \log T(n)$, and $w_{\text{input}}(n) \leq \log n$. Then there are good functions $h_{\text{proof}}(n)$ and $h_{\text{input}}(n)$ satisfying*

$$\begin{aligned} h_{\text{proof}}(n) &= \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}(n), & \text{and} \\ h_{\text{input}}(n) &= \lceil \log n \rceil - w_{\text{input}}(n), \end{aligned}$$

such that the following holds.

Suppose that $w_{\text{proof}}, h_{\text{proof}} \geq (5/m) \log T(n)$, and that for some absolute constant $C \geq 1$,

$$\frac{w_{\text{input}}(n)}{w_{\text{proof}}(n)}, \frac{h_{\text{input}}(n)}{h_{\text{proof}}(n)} \leq 1 - \frac{Cm \log \log T(n)}{\log T(n)}.$$

Let $W_{\text{proof}}(n) := 2^{w_{\text{proof}}(n)}$, $H_{\text{proof}}(n) := 2^{h_{\text{proof}}(n)}$, $W_{\text{input}}(n) := 2^{w_{\text{input}}(n)}$, and $H_{\text{input}}(n) := 2^{h_{\text{input}}(n)}$. Then $\text{NTIME}[T(n)]$ has a rectangular PCP of proximity with an $H_{\text{proof}}(n) \times W_{\text{proof}}(n)$ proof matrix and an $H_{\text{input}}(n) \times W_{\text{input}}(n)$ input matrix, whose other parameters are specified in Table 9.

Furthermore, given the randomness $\text{seed} \in \{0, 1\}^r$, the total number of queries and parity-check bits is at most 2, and the decision predicate $\text{VDec} \leftarrow V_{\text{dec}}(\text{seed}, \text{shared})$ of the rectangular PCPP verifier is an OR of the 2 input bits (including queries and parity-check bits) or their negations for every $\text{seed}, \text{shared}$.

Completeness error	$1 - c$
Soundness error	s
Proximity parameter	δ
Row randomness	$h_{\text{proof}} - (5/m) \log T(n)$
Column randomness	$w_{\text{proof}} - (5/m) \log T(n)$
Shared randomness	$(10/m) \log T(n) + O(\log \log T(n) + m \log m)$
Query complexity	2
Parity check complexity	
Decision complexity	$\text{poly}(\log \log T(n))$

Table 9: Parameters of the PCPP constructed in Theorem 7.6.

Proof. Let δ be an arbitrary constant. By Theorem 7.4, that there is a rectangular PCPP verifier V^{3q} with perfect completeness and parameters:

- Soundness error $s^{3q} \in (0, 1)$.
- Proximity parameter δ .
- Query complexity and parity-check complexity 3.
- Proof matrix size $H_{\text{proof}}^{3q} \times W_{\text{proof}}^{3q}$, with $w_{\text{proof}}^{3q} = w_{\text{proof}} = \log W_{\text{proof}}$ and $h_{\text{proof}}^{3q} = \log H_{\text{proof}}^{3q} = \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}(n)$.²⁸
- Shared randomness complexity $r_{\text{shared}} = (10/m) \log T(n) + O(\log \log T(n) + m \log m)$.
- Row randomness complexity $r_{\text{row}} = h_{\text{proof}}^{3q} - (5/m) \log T(n)$.
- Column randomness complexity $r_{\text{col}} = w_{\text{proof}}^{3q} - (5/m) \log T(n)$.
- Decision complexity $\text{poly}(\log \log T(n))$.

Let $r = r_{\text{row}} + r_{\text{col}} + r_{\text{shared}}$ be the length of total randomness. Moreover, we know that the total number of queries and parity-check bits is at most 3, and that the decision circuit of V is an OR of its input bits (i.e. the answers to the queries and parity-check bits) or their negations after fixing the random seed. We will now combine V^{3q} and the gadget in Lemma 7.5 to construct a 2-query PCPP.

²⁸Note that the final matrix height is $h_{\text{proof}} \leq h_{\text{proof}}^{3q} + O(\log \log T(n))$, hence the technical requirement $h_{\text{input}}(n)/h_{\text{proof}}^{3q} \leq 1 - C' \log \log T(n)/\log T(n)$ for large C' holds, given the assumption that $h_{\text{input}}(n)/h_{\text{proof}} \leq 1 - C \log \log T(n)/\log T(n)$ for large C , $h_{\text{proof}} \geq (5/m) \log T(n)$, and $m \leq (\log T(n))^{0.1}$, as shown in Claim 6.17.

Suppose, for the simplicity of presentation, that the PCPP verifier V always probes 2 bits of the input and proof oracles, and has 1 parity-check bit. (The other cases can be considered similarly and we omit the details.) Then the decision predicate $\text{VDec} \leftarrow V_{\text{dec}}(\text{seed}.\text{shared})$ for every fixed $\text{seed}.\text{shared} \in \{0, 1\}^{r_{\text{shared}}}$ is a function

$$\text{VDec}(\text{ans}_1, \text{ans}_2, pc_1(\text{seed})) := (\text{ans}_1 \oplus b_1) \vee (\text{ans}_2 \oplus b_2) \vee (pc_1(\text{seed}) \oplus b_3).$$

where $b_1, b_2, b_3 \in \{0, 1\}$. The new PCPP verifier V is defined as follows.

- The proof of the new PCPP verifier V is the concatenation of the proof for V^{3q} and an $y : \{0, 1\}^r \rightarrow \{0, 1\}$ of length 2^r used as the additional variable y in Lemma 7.5.
- The randomness of V is the concatenation of the randomness seed for V^{3q} and an $j \in [10]$.

Queries and parity-check bits. Assume that $(\text{seed}, j) \in \{0, 1\}^r \times [10]$ is given as the randomness. The verifier V first generates the indices i_1, i_2 of the queries to the input and proof oracles (denoted by a single oracle Π for simplicity) and the parity-check function pc_1 . Instead of making all these queries and doing the parity-check, we identify $\text{ans}_1 \oplus b_1, \text{ans}_2 \oplus b_2, pc_1(\text{seed}) \oplus b_3, y(\text{seed})$ with x_1, x_2, x_3, y in the gadget given by Lemma 7.5, respectively, and queries the j -th gadget. (For instance, if $j = 5$, the corresponding constraint is $\overline{x_1} \vee \overline{x_3}$, so that we will query the i_1 -th of Π and do the parity-check pc_1 ; if $j = 8$, the constraint is $x_1 \vee \overline{y}$, so that we will query the i_1 -th bit of Π and the $y(\text{seed})$.) The decision predicate will accept if and only if either the j -th constraint is satisfied when identifying $\text{ans}_1 \oplus b_1, \text{ans}_2 \oplus b_2, pc_1(\text{seed}) \oplus b_3, y(\text{seed})$ with x_1, x_2, x_3, y , respectively.

Completeness. For every input $x \in L$, by the completeness of V^{3q} , there is a proof oracle Π_{proof}^{3q} such that V^{3q} accepts given the oracle $x \circ \Pi_{\text{proof}}^{3q}$ with probability 1, which means that for every $\text{seed} \in \{0, 1\}^r$, the answers $\text{ans}_1, \text{ans}_2$ to the queries and the parity-check bits $pc_1(\text{seed})$ satisfies

$$\text{VDec}(\text{ans}_1, \text{ans}_2, pc_1(\text{seed})) = (\text{ans}_1 \oplus b_1) \vee (\text{ans}_2 \oplus b_2) \vee (pc_1(\text{seed}) \oplus b_3) = 1.$$

By Lemma 7.5, there is an y_{seed} such that at least 7 of the 10 constraints in the gadgets are satisfied. This means that given the proof oracle $\Pi \circ y$ for $y(\text{seed}) := y_{\text{seed}}$, the verifier will accept with probability at least $c := 7/10$.

Soundness. Assume that $x \in \{0, 1\}^n$ that is δ -far from being in L , and $\Pi_{\text{proof}} = \Pi_{\text{proof}}^{3q} \circ y$ is any proof, where Π_{proof}^{3q} is a proof for V^{3q} and $y : \{0, 1\}^r \rightarrow \{0, 1\}$. By the soundness of V^{3q} , we know that for each least $1 - s_1$ fraction of $\text{seed} \in \{0, 1\}^n$,

$$\text{VDec}(\text{ans}_1, \text{ans}_2, pc_1(\text{seed})) = (\text{ans}_1 \oplus b_1) \vee (\text{ans}_2 \oplus b_2) \vee (pc_1(\text{seed}) \oplus b_3) = 0.$$

By Lemma 7.5, we can see that for these seed , the accept probability of V is at most $6/10$, whereas in other cases the accept probabilistic of V is at most $7/10$. Thus the accept probability of V is at most $s := (7/10) \cdot s_1 + (6/10) \cdot (1 - s_1) < c$.

Rectangularity. Since we only need to introduce $O(1)$ bits of randomness representing $j \sim [10]$, we can put it into the shared randomness. We only need to show that the new proof $\Pi_{\text{proof}}^{3q} \circ y$ can be arranged as a matrix so that the queries can be done rectangulary. Let $W_{\text{proof}} := W_{\text{proof}}^{3q}$ and

$H_{\text{proof}} := H_{\text{proof}}^{3q} + 2^r/W_{\text{proof}}^{3q}$. Without loss of generality, we assume that H_{proof} is a power of two, therefore we define

$$\begin{aligned} h_{\text{proof}} &:= \log H_{\text{proof}} = \log \left(H_{\text{proof}}^{3q} + 2^r/W_{\text{proof}}^{3q} \right) \\ &\leq h_{\text{proof}}^{3q} + O(\log \log T(n)) \\ &= \log T(n) + \Theta(m \log \log T(n)) - w_{\text{proof}}. \end{aligned}$$

The final proof matrix will be of size $H_{\text{proof}} \times W_{\text{proof}}$, arranged as follows: The first H_{proof}^{3q} rows will contain the proof Π_{proof}^{3q} of V^{3q} ; The remaining $2^r/W_{\text{proof}}$ contains the proof $y : \{0, 1\}^r \rightarrow \{0, 1\}$, represented as the string $y(0) \circ y(1) \circ \dots \circ y(2^r - 1)$ of length 2^r . Recall that there will be two kinds of queries to the proof oracle.

1. If the query is to the proof oracle Π_{proof}^{3q} of V^{3q} or to the input oracle, we can use the row and column verifier of V^{3q} to generate the queries rectangularly.
2. Otherwise, the query is to the proof $y(\text{seed})$ for the randomness $\text{seed} \in \{0, 1\}^r$ of V^{3q} . Then the column (resp. row) index of this query only depends on the lowest w_{proof} bits (resp. the highest $r - w_{\text{proof}}$ bits) of the random seed of V . Recall that the random seed of V is the concatenation of seed and a $j \in [10]$. If we arrange the randomness as

$$\text{seed.col} \circ \text{seed.shared} \circ j \circ \text{seed.row},$$

then the lowest w_{proof} bits (resp. the highest $r - w_{\text{proof}}$ bits) of the random seed only depends on the $(\text{seed.col}, \text{seed.shared})$ (resp. $(\text{seed.shared}, j, \text{seed.row})$), since

$$\begin{aligned} r_{\text{col}} &= w_{\text{proof}} - (5/m) \log T(n) \leq w_{\text{proof}} \\ r_{\text{col}} + r_{\text{shared}} &= w_{\text{proof}} + (5/m) \log T(n) + O(\log \log n + m \log m) \geq w_{\text{proof}}. \end{aligned}$$

As a result, the queries can be done rectangularly. □

Acknowledgements

We got the initial idea of this work when Jiayu was a research intern at Igor Carboni Oliveira's group and Hanlin was visiting Igor. Hanlin wants to thank his supervisor Rahul Santhanam for introducing the problems PARTIAL-HARD [AS10] and XOR-REMOTE-POINT [APY09] to him and for helpful discussions. We thank Lijie Chen and an anonymous STOC reviewer for pointing out that the technique in [Wil18c] can be slightly adapted to obtain a non-trivial algorithm for #ACC-SATISFYING-PAIRS. We also thank Zhikun Wang and Tianqi Yang for discussion on this work.

References

- [AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity - A Modern Approach*. Cambridge University Press, 2009. URL: <http://www.cambridge.org/catalogue/catalogue.asp?isbn=9780521424264>. (cit. on p. 1, 74)
- [AC19] Josh Alman and Lijie Chen. Efficient construction of rigid matrices using an NP oracle. In *FOCS*, pages 1034–1055. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00067. (cit. on p. 1, 7, 9, 53, 54)

- [ACW16] Josh Alman, Timothy M. Chan, and R. Ryan Williams. Polynomial representations of threshold functions and algorithmic applications. In *FOCS*, pages 467–476. IEEE Computer Society, 2016. doi:10.1109/FOCS.2016.57. (cit. on p. 11)
- [ACW20] Josh Alman, Timothy M. Chan, and R. Ryan Williams. Faster deterministic and Las Vegas algorithms for offline approximate nearest neighbors in high dimensions. In *SODA*, pages 637–649. SIAM, 2020. doi:10.1137/1.9781611975994.39. (cit. on p. 11)
- [AG91] Eric Allender and Vivek Gore. On strong separations from AC^0 (extended abstract). In *Fundamentals of Computation Theory, 8th International Symposium, FCT '91*, volume 529 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 1991. doi:10.1007/3-540-54458-5_44. (cit. on p. 7, 53)
- [AHWW16] Amir Abboud, Thomas Dueholm Hansen, Virginia Vassilevska Williams, and Ryan Williams. Simulating branching programs with edit distance and friends: or: a polylog shaved is a lower bound made. In *STOC*, pages 375–388. ACM, 2016. doi:10.1145/2897518.2897653. (cit. on p. 5)
- [Ajt83] Miklós Ajtai. Σ_1^1 -formulae on finite structures. *Ann. Pure Appl. Log.*, 24(1):1–48, 1983. doi:10.1016/0168-0072(83)90038-6. (cit. on p. 4)
- [APY09] Noga Alon, Rina Panigrahy, and Sergey Yekhanin. Deterministic approximation algorithms for the nearest codeword problem. In *APPROX-RANDOM*, volume 5687 of *Lecture Notes in Computer Science*, pages 339–351. Springer, 2009. doi:10.1007/978-3-642-03685-9_26. (cit. on p. 3, 86)
- [AS10] Vikraman Arvind and Srikanth Srinivasan. Circuit lower bounds, help functions, and the remote point problem. In *ICS*, pages 383–396. Tsinghua University Press, 2010. URL: <http://conference.iis.tsinghua.edu.cn/ICS2010/content/papers/30.html>. (cit. on p. 3, 4, 8, 56, 86)
- [AW15] Josh Alman and R. Ryan Williams. Probabilistic polynomials and Hamming nearest neighbors. In *FOCS*, pages 136–150. IEEE Computer Society, 2015. doi:10.1109/FOCS.2015.18. (cit. on p. 11)
- [AW17] Josh Alman and R. Ryan Williams. Probabilistic rank and matrix rigidity. In *STOC*, pages 641–652. ACM, 2017. doi:10.1145/3055399.3055484. (cit. on p. 11)
- [AWY15] Amir Abboud, R. Ryan Williams, and Huacheng Yu. More applications of the polynomial method to algorithm design. In *SODA*, pages 218–230. SIAM, 2015. doi:10.1137/1.9781611973730.17. (cit. on p. 11)
- [BGH⁺06] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil P. Vadhan. Robust PCPs of proximity, shorter PCPs, and applications to coding. *SIAM J. Comput.*, 36(4):889–974, 2006. doi:10.1137/S0097539705446810. (cit. on p. 9, 15, 16, 57, 68, 69, 73, 80)
- [BHPT20] Amey Bhangale, Prahladh Harsha, Orr Paradise, and Avishay Tal. Rigid matrices from rectangular PCPs or: Hard claims have complex proofs. In *FOCS*, pages 858–869. IEEE, 2020. doi:10.1109/FOCS46700.2020.00084. (cit. on p. 1, 9, 17, 54, 57, 58, 59, 61, 63, 64, 65, 69)
- [BT94] Richard Beigel and Jun Tarui. On ACC. *Comput. Complex.*, 4:350–366, 1994. doi:10.1007/BF01263423. (cit. on p. 7, 53)
- [Che19] Lijie Chen. Non-deterministic quasi-polynomial time is average-case hard for ACC circuits. In *FOCS*, pages 1281–1304. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00079. (cit. on p. 2, 3)
- [Che22] Lijie Chen. *Better Hardness via Algorithms, and New Forms of Hardness versus Randomness*. PhD thesis, Massachusetts Institute of Technology, 2022. (cit. on p. 5)
- [CJW20] Lijie Chen, Ce Jin, and R. Ryan Williams. Sharp threshold results for computational complexity. In *STOC*, pages 1335–1348. ACM, 2020. doi:10.1145/3357713.3384283. (cit. on p. 12)

- [CL21] Lijie Chen and Xin Lyu. Inverse-exponential correlation bounds and extremely rigid matrices from a new derandomized XOR lemma. In *STOC*, pages 761–771. ACM, 2021. doi:10.1145/3406325.3451132. (cit. on p. 3, 6, 54)
- [CLW20] Lijie Chen, Xin Lyu, and R. Ryan Williams. Almost-everywhere circuit lower bounds from non-trivial derandomization. In *FOCS*, pages 1–12. IEEE, 2020. doi:10.1109/FOCS46700.2020.00009. (cit. on p. ii, 2, 3, 6, 8, 10, 13, 14, 15, 35, 54, 55, 100)
- [Cop82] Don Coppersmith. Rapid multiplication of rectangular matrices. *SIAM J. Comput.*, 11(3):467–471, 1982. doi:10.1137/0211037. (cit. on p. 11, 52)
- [COS18] Ruiwen Chen, Igor Carboni Oliveira, and Rahul Santhanam. An average-case lower bound against ACC^0 . In *LATIN*, volume 10807 of *Lecture Notes in Computer Science*, pages 317–330. Springer, 2018. doi:10.1007/978-3-319-77404-6_24. (cit. on p. 3)
- [CR22] Lijie Chen and Hanlin Ren. Strong average-case circuit lower bounds from nontrivial derandomization. *SIAM J. Comput.*, 51(3):STOC20–115–STOC20–173, 2022. doi:10.1137/20M1364886. (cit. on p. 3, 35)
- [CW19a] Lijie Chen and Ruosong Wang. Classical algorithms from quantum and Arthur-Merlin communication protocols. In *ITCS*, volume 124 of *LIPICs*, pages 23:1–23:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ITCS.2019.23. (cit. on p. 5, 11)
- [CW19b] Lijie Chen and R. Ryan Williams. Stronger connections between circuit analysis and circuit lower bounds, via PCPs of proximity. In *CCC*, volume 137 of *LIPICs*, pages 19:1–19:43. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.CCC.2019.19. (cit. on p. 2, 7, 9, 22, 32, 35, 80, 83, 98)
- [CW21] Timothy M. Chan and R. Ryan Williams. Deterministic APSP, orthogonal vectors, and more: Quickly derandomizing Razborov-Smolensky. *ACM Trans. Algorithms*, 17(1):2:1–2:14, 2021. doi:10.1145/3402926. (cit. on p. 7, 11, 53)
- [CZ19] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. *Annals of Mathematics*, 189(3):653–705, 2019. doi:10.4007/annals.2019.189.3.1. (cit. on p. 1)
- [Erd59] Paul Erdős. Graph theory and probability. *Canadian Journal of Mathematics*, 11:34–38, 1959. doi:10.4153/CJM-1959-003-9. (cit. on p. 1)
- [FGHK16] Magnus Gausdal Find, Alexander Golovnev, Edward A. Hirsch, and Alexander S. Kulikov. A better-than- $3n$ lower bound for the circuit complexity of an explicit function. In *FOCS*, pages 89–98. IEEE Computer Society, 2016. doi:10.1109/FOCS.2016.19. (cit. on p. 1)
- [FS16] Lance Fortnow and Rahul Santhanam. New non-uniform lower bounds for uniform classes. In *CCC*, volume 50 of *LIPICs*, pages 19:1–19:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICs.CCC.2016.19. (cit. on p. 13, 14)
- [FSS84] Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Math. Syst. Theory*, 17(1):13–27, 1984. doi:10.1007/BF01744431. (cit. on p. 4)
- [GGH⁺07] Shafi Goldwasser, Dan Gutfreund, Alexander Healy, Tali Kaufman, and Guy N. Rothblum. Verifying and decoding in constant depth. In *STOC*, pages 440–449. ACM, 2007. doi:10.1145/1250790.1250855. (cit. on p. 9)
- [GGNS23] Karthik Gajulapalli, Alexander Golovnev, Satyajeet Nagargoje, and Sidhant Saraogi. Range avoidance for constant-depth circuits: Hardness and algorithms. *CoRR*, 2023. doi:10.48550/arXiv.2303.05044. (cit. on p. 2)
- [GJS76] M. R. Garey, David S. Johnson, and Larry J. Stockmeyer. Some simplified NP-complete graph problems. *Theor. Comput. Sci.*, 1(3):237–267, 1976. doi:10.1016/0304-3975(76)90059-1. (cit. on p. 83)

- [GLW22] Venkatesan Guruswami, Xin Lyu, and Xiuhan Wang. Range avoidance for low-depth circuits and connections to pseudorandomness. In *APPROX/RANDOM*, volume 245 of *LIPICs*, pages 20:1–20:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.APPROX/RANDOM.2022.20. (cit. on p. 2, 10, 11)
- [GNW11] Oded Goldreich, Noam Nisan, and Avi Wigderson. On Yao’s XOR-lemma. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, volume 6650 of *Lecture Notes in Computer Science*, pages 273–301. Springer, 2011. doi:10.1007/978-3-642-22670-0_23. (cit. on p. 15)
- [Gol11] Oded Goldreich. *A Sample of Samplers: A Computational Perspective on Sampling*, pages 302–332. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011. doi:10.1007/978-3-642-22670-0_24. (cit. on p. 60)
- [GR08] Dan Gutfreund and Guy N. Rothblum. The complexity of local list decoding. In *APPROX-RANDOM*, volume 5171 of *Lecture Notes in Computer Science*, pages 455–468. Springer, 2008. doi:10.1007/978-3-540-85363-3_36. (cit. on p. 10)
- [Hås89] Johan Håstad. Almost optimal lower bounds for small depth circuits. *Adv. Comput. Res.*, 5:143–170, 1989. (cit. on p. 4)
- [HV21] Xuanguai Huang and Emanuele Viola. Average-case rigidity lower bounds. In *CSR*, volume 12730 of *Lecture Notes in Computer Science*, pages 186–205. Springer, 2021. doi:10.1007/978-3-030-79416-3_11. (cit. on p. 54)
- [KKMP21] Robert Kleinberg, Oliver Korten, Daniel Mitropolsky, and Christos H. Papadimitriou. Total functions in the polynomial hierarchy. In *ITCS*, volume 185 of *LIPICs*, pages 44:1–44:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ITCS.2021.44. (cit. on p. 1, 3)
- [Kor21] Oliver Korten. The hardest explicit construction. In *FOCS*, pages 433–444. IEEE, 2021. doi:10.1109/FOCS52979.2021.00051. (cit. on p. 1, 10)
- [Lev87] Leonid A. Levin. One-way functions and pseudorandom generators. *Comb.*, 7(4):357–363, 1987. doi:10.1007/BF02579323. (cit. on p. 15)
- [Li23] Xin Li. Two source extractors for asymptotically optimal entropy, and (many) more. *CoRR*, abs/2303.06802, 2023. arXiv:2303.06802, doi:10.48550/arXiv.2303.06802. (cit. on p. 1)
- [LY22] Jiayu Li and Tianqi Yang. $3.1n - o(n)$ circuit lower bounds for explicit functions. In *STOC*, pages 1180–1193. ACM, 2022. doi:10.1145/3519935.3519976. (cit. on p. 1)
- [Mie09] Thilo Mie. Short PCPPs verifiable in polylogarithmic time with $O(1)$ queries. *Ann. Math. Artif. Intell.*, 56(3-4):313–338, 2009. doi:10.1007/s10472-009-9169-y. (cit. on p. 76)
- [Mul11] Ketan Mulmuley. On P vs. NP and geometric complexity theory: Dedicated to Sri Ramakrishna. *J. ACM*, 58(2):5:1–5:26, 2011. doi:10.1145/1944345.1944346. (cit. on p. 12)
- [MW20] Cody D. Murray and R. Ryan Williams. Circuit lower bounds for nondeterministic quasipolytime from a new easy witness lemma. *SIAM J. Comput.*, 49(5), 2020. doi:10.1137/18M1195887. (cit. on p. 2)
- [Par21] Orr Paradise. Smooth and strong PCPs. *Comput. Complex.*, 30(1):1, 2021. doi:10.1007/s00037-020-00199-3. (cit. on p. 10, 18)
- [Păt08] Mihai Pătraşcu. Succincter. In *FOCS*, pages 305–313. IEEE Computer Society, 2008. doi:10.1109/FOCS.2008.83. (cit. on p. 11)
- [Pip79] Nicholas Pippenger. On simultaneous resource bounds (preliminary version). In *FOCS*, pages 307–311. IEEE Computer Society, 1979. doi:10.1109/SFCS.1979.29. (cit. on p. 4)
- [Ram20] C. Ramya. Recent progress on matrix rigidity - A survey. *CoRR*, abs/2009.09460, 2020. doi:10.48550/arXiv.2009.09460. (cit. on p. 1)

- [Raz87] Alexander A Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987. doi:10.1007/BF01137685. (cit. on p. 11)
- [RSW22] Hanlin Ren, Rahul Santhanam, and Zhikun Wang. On the range avoidance problem for circuits. In *FOCS*, pages 640–650. IEEE, 2022. doi:10.1109/FOCS54457.2022.00067. (cit. on p. 2, 3, 4, 5, 9, 10, 11, 14, 16, 17, 57, 58, 63, 64, 65, 66, 68, 69, 70, 73, 76, 98)
- [Sha49] Claude E. Shannon. The synthesis of two-terminal switching circuits. *Bell System technical journal*, 28(1):59–98, 1949. doi:10.1002/j.1538-7305.1949.tb03624.x. (cit. on p. 1)
- [Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *STOC*, pages 77–82. ACM, 1987. doi:10.1145/28395.28404. (cit. on p. 11)
- [Smo93] Roman Smolensky. On representations by low-degree polynomials. In *FOCS*, pages 130–138. IEEE Computer Society, 1993. doi:10.1109/SFCS.1993.366874. (cit. on p. 11)
- [Spi96] Daniel A. Spielman. Linear-time encodable and decodable error-correcting codes. *IEEE Trans. Inf. Theory*, 42(6):1723–1731, 1996. doi:10.1109/18.556668. (cit. on p. 13)
- [Vad12] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012. doi:10.1561/0400000010. (cit. on p. 74)
- [Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In *MFCS*, volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1977. doi:10.1007/3-540-08353-7_135. (cit. on p. 1, 9)
- [VW18] Emanuele Viola and Avi Wigderson. Local expanders. *Comput. Complex.*, 27(2):225–244, 2018. doi:10.1007/s00037-017-0155-1. (cit. on p. 73)
- [Wil13] R. Ryan Williams. Improving exhaustive search implies superpolynomial lower bounds. *SIAM J. Comput.*, 42(3):1218–1244, 2013. doi:10.1137/10080703X. (cit. on p. 2)
- [Wil14a] R. Ryan Williams. Nonuniform ACC circuit lower bounds. *J. ACM*, 61(1):2:1–2:32, 2014. doi:10.1145/2559903. (cit. on p. 2, 52)
- [Wil14b] R. Ryan Williams. The polynomial method in circuit complexity applied to algorithm design (invited talk). In *FSTTCS*, volume 29 of *LIPICs*, pages 47–60. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014. doi:10.4230/LIPICs.FSTTCS.2014.47. (cit. on p. 11)
- [Wil18a] R. Ryan Williams. Faster all-pairs shortest paths via circuit complexity. *SIAM J. Comput.*, 47(5):1965–1985, 2018. doi:10.1137/15M1024524. (cit. on p. 11, 52)
- [Wil18b] R. Ryan Williams. Limits on representing Boolean functions by linear combinations of simple functions: Thresholds, ReLUs, and low-degree polynomials. In *CCC*, volume 102 of *LIPICs*, pages 6:1–6:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICs.CCC.2018.6. (cit. on p. 32, 35)
- [Wil18c] R. Ryan Williams. New algorithms and lower bounds for circuits with linear threshold gates. *Theory Comput.*, 14(1):1–25, 2018. doi:10.4086/toc.2018.v014a017. (cit. on p. 7, 12, 53, 86)
- [Yao85] Andrew Chi-Chih Yao. Separating the polynomial-time hierarchy by oracles (preliminary version). In *FOCS*, pages 1–10. IEEE Computer Society, 1985. doi:10.1109/SFCS.1985.49. (cit. on p. 4)

A Missing Proofs in Section 3 and 4

In this section, we complete the missing proofs in the framework of solving range avoidance, remote point, and hard partial truth table via non-trivial satisfying pairs algorithms.

A.1 Satisfying Pairs for $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ Circuits

Theorem 3.7. *Let \mathcal{C} be a typical circuit class, $M' \geq 1$ and $\eta \in (0, 1)$ be parameters. Suppose there is a deterministic algorithm running in time $T^{\text{alg}} = T^{\text{alg}}(N, M)$ that, given as input a list of $\hat{N} \leq N$ $\text{AND}_d \circ \mathcal{C}$ circuits $\{C_i\}$ and a list of $\hat{M} \leq M$ inputs $\{x_j\}$, estimates the following quantity with additive error η :*

$$\Pr_{i \leftarrow [\hat{N}], j \leftarrow [\hat{M}]} [C_i(x_j)].$$

Then, there is a deterministic algorithm running in time $A^d(2^{dl} + M'/M) \cdot (\ell_{\text{Prod}}/N) \cdot O(T^{\text{alg}})$ that, given as input a $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuit C^{Prod} with parameters specified in Remark 3.6, and a list of M' inputs $\{(x_j, \alpha_j)\}$, estimates the following quantity with additive error $\eta \cdot U^d$:

$$\mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow [M']} [C_i^{\text{Prod}}(x_j, \alpha_j)].$$

Proof. For any fixed i and j , we know that

$$\begin{aligned} C_i^{\text{Prod}}(x_j, \alpha_j) &= \prod_{t=1}^d \text{Sum}_{q_t(i)}(x_j, \alpha_j) \\ &= \prod_{t=1}^d \sum_{k \in [A]} \text{coeff}_k(\alpha_j) \cdot C_{\text{id}_{x_k}(\alpha_j, q_t(i))}(x_j) \\ &= \sum_{k_1 \in [A]} \sum_{k_2 \in [A]} \cdots \sum_{k_d \in [A]} \prod_{t=1}^d \left(\text{coeff}_{k_t}(\alpha_j) \cdot C_{\text{id}_{x_{k_t}}(\alpha_j, q_t(i))}(x_j) \right). \end{aligned} \quad (15)$$

As we can enumerate $k_1, k_2, \dots, k_d \in [A]$ in A^d time, it suffices to estimate

$$\mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow [M']} \left[\prod_{t=1}^d \left(\text{coeff}_{k_t}(\alpha_j) \cdot C_{\text{id}_{x_{k_t}}(\alpha_j, q_t(i))}(x_j) \right) \right]. \quad (16)$$

Fix $k_1, k_2, \dots, k_d \in [A]$. Since C^{Prod} is of locality l , we can see that Eq. (16) only depends on dl bits of α_j . We partition $j \in [M']$ into 2^{dl} groups as follows: For each $\alpha \in \{0, 1\}^{dl}$, let \mathcal{J}_α be the set of $j \in [M']$ such that the dl bits of α_j (that Eq. (16) for this j depends on) equals to α . We will estimate Eq. (16) by enumerating $\alpha \in \{0, 1\}^{dl}$, estimating it for $j \leftarrow \mathcal{J}_\alpha$ (instead of $j \leftarrow [M']$), and then taking the average over all possible α .

Now we fix any $\alpha \in \{0, 1\}^{dl}$. We can rephrase the following two items as they no longer depend on α_j :

$$\begin{aligned} \text{coeff}_{k_t}(\alpha_j) &=: \text{coeff}'_t; \\ C_{\text{id}_{x_{k_t}}(\alpha_j, q_t(i))}(x_j) &=: C_{\text{id}_{x'_t}(i)}(x_j). \end{aligned}$$

It then suffices to estimate

$$\mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow \mathcal{J}_\alpha} \left[\prod_{t=1}^d \text{coeff}'_t \cdot C_{\text{id}_{x'_t}(i)}(x_j) \right] = \left(\prod_{t=1}^d \text{coeff}'_t \right) \cdot \mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow \mathcal{J}_\alpha} \left[\bigwedge_{t=1}^d C_{\text{id}_{x'_t}(i)}(x_j) \right]. \quad (17)$$

Each expression of the form $\mathbb{E}_{i,j} \left[\bigwedge_{t=1}^d C_{\text{id}_{x'_t}(i)}(x_j) \right]$ can be reduced to the SATISFYING-PAIRS problem for $\text{AND}_d \circ \mathcal{C}$ circuits. More precisely, we split \mathcal{J}_α into blocks of size M , split $[\ell_{\text{Prod}}]$ into

blocks of size N , and use the assumed algorithm (which works for N $\text{AND}_d \circ \mathcal{C}$ circuits and M inputs) to estimate Eq. (17) within additive error $\left(\eta \cdot \prod_{t=1}^d |\text{coeff}'_t|\right)$ in $T^{\text{alg}} \cdot \lceil |\mathcal{J}_\alpha|/M \rceil \cdot \lceil \ell_{\text{Prod}}/N \rceil$ time.²⁹ We substitute this estimation in Eq. (16) and then in Eq. (15) to obtain our final algorithm.

It remains to upper bound the running time and the additive error of our algorithm.

Running time. Consider the subroutine for estimating Eq. (16). This subroutine itself is reduced to subroutines for each \mathcal{J}_α , which takes $T^{\text{alg}} \cdot \lceil |\mathcal{J}_\alpha|/M \rceil \cdot \lceil \ell_{\text{Prod}}/N \rceil$ time. The time complexity of this subroutine is

$$O(T^{\text{alg}}) \cdot \lceil \ell_{\text{Prod}}/N \rceil \cdot \sum_{\alpha} \lceil |\mathcal{J}_\alpha|/M \rceil \leq O(T^{\text{alg}}) \cdot (2^{dl} + M'/M) \cdot (\ell_{\text{Prod}}/N).$$

We invoked this subroutine A^d times by enumerating $k_1, k_2, \dots, k_d \in [A]$ to estimate Eq. (15), so the total time complexity of our algorithm is $A^d(2^{dl} + M'/M)(\ell_{\text{Prod}}/N)O(T^{\text{alg}})$.

Additive error. Our estimation of Eq. (17) is within additive error $\left(\eta \cdot \prod_{t=1}^d |\text{coeff}'_t|\right)$. Thus our estimation of Eq. (16) is within additive error of

$$\eta \cdot \sum_{\alpha \in \{0,1\}^{dl}} \frac{|\mathcal{J}_\alpha|}{M'} \cdot \prod_{i=1}^d |\text{coeff}'_i| = \eta \cdot \mathbb{E}_j \left[\prod_{i=1}^d |\text{coeff}_{k_i}(\alpha_j)| \right].$$

It follows that our estimation of Eq. (15) is within additive error of

$$\begin{aligned} & \eta \cdot \sum_{k_1 \in [A]} \sum_{k_2 \in [A]} \cdots \sum_{k_d \in [A]} \mathbb{E}_j \left[\prod_{i=1}^d |\text{coeff}_{k_i}(\alpha_j)| \right] \\ &= \eta \cdot \mathbb{E}_j \left[\left(\sum_{k \in [A]} |\text{coeff}_k(\alpha_j)| \right)^d \right] \\ &\leq \eta \cdot U^d. \end{aligned} \quad \square$$

Lemma 4.5. *Let \mathcal{C} be a typical circuit class, $M' \geq 1$ and $\eta \in (0, 1)$ be parameters. Suppose there is a deterministic algorithm running in time $T^{\text{alg}} = T^{\text{alg}}(N, M)$ that, given as input a list of $\hat{M} \leq M$ $\text{AND}_d \circ \mathcal{C}$ circuits $\{C_i\}$ and a list of $\hat{N} \leq N$ inputs $\{x_j\}$ of length $n \cdot \text{polylog}(\ell)$, estimates the following quantity with additive error η :*

$$\Pr_{i \leftarrow [\hat{M}], j \leftarrow [\hat{N}]} [C_i(x_j)].$$

Then, for any constant $\ell_{\mathcal{C}} > 0$, there is a deterministic algorithm running in time $A^d \cdot (\ell_{\mathcal{C}}^d + \ell_{\text{Prod}}/N) \cdot (2^{dl} + M'/M) \cdot O(T^{\text{alg}})$ that, given as input a $\text{Prod}_d \circ \text{Sum}$ circuit C^{Prod} with parameters specified in Remark 4.4, a list of ℓ_x strings $\{x_j\}$, a list of M' inputs $\{\alpha_j\}$, and a list of M' \mathcal{C} circuits $\{C_j\}$ from $\{0, 1\}^{|x|}$ to $\{0, 1\}^{\ell_{\mathcal{C}}}$, estimates the following quantity with additive error $\eta \cdot U^d$:

$$\mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow [M']} [C_i^{\text{Prod}}(C_j(x), \alpha_j)].$$

Recall here that $C_j(x) = C_j(x_1) \circ C_j(x_2) \circ \cdots \circ C_j(x_{\ell_x})$.

²⁹Note that at most one of the block may contain less than M inputs. However, the assumed algorithm works for input number $\leq M$ as well, and this will not have any blow-up on the error factor.

Proof Sketch of Lemma 4.5. We identify $\text{id}_{x_k}(\alpha, i) \in [\ell_y]$ with $(\text{id}_{x_k^x}(\alpha, i), \text{id}_{x_k^{\mathcal{C}}}(\alpha, i)) \in [\ell_x] \times [\ell_{\mathcal{C}}]$ (note that $\ell_y = \ell_x \ell_{\mathcal{C}}$). Then,

$$\begin{aligned} C_i^{\text{Prod}}(C_j(x), \alpha_j) &= \prod_{t=1}^d \text{Sum}_{q_t(i)}(C_j(x), \alpha_j) \\ &= \prod_{t=1}^d \sum_{k \in [A]} \text{coeff}_k(\alpha_j) \cdot (C_j)_{\text{id}_{x_k^{\mathcal{C}}}(\alpha_j, q_t(i))} (x_{\text{id}_{x_k^x}(\alpha_j, q_t(i))}) \\ &= \sum_{k_1 \in [A]} \sum_{k_2 \in [A]} \cdots \sum_{k_d \in [A]} \prod_{t=1}^d \left(\text{coeff}_{k_t}(\alpha_j) \cdot (C_j)_{\text{id}_{x_{k_t}^{\mathcal{C}}}(\alpha_j, q_t(i))} (x_{\text{id}_{x_{k_t}^x}(\alpha_j, q_t(i))}) \right). \end{aligned} \quad (18)$$

As we can enumerate $k_1, k_2, \dots, k_d \in [A]$ in A^d time, it suffices to estimate

$$\mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow [M']} \left[\prod_{t=1}^d \left(\text{coeff}_{k_t}(\alpha_j) \cdot (C_j)_{\text{id}_{x_{k_t}^{\mathcal{C}}}(\alpha_j, q_t(i))} (x_{\text{id}_{x_{k_t}^x}(\alpha_j, q_t(i))}) \right) \right]. \quad (19)$$

Fix $k_1, k_2, \dots, k_d \in [A]$. Since C^{Prod} is of locality l , we can see that Eq. (19) only depends on dl bits of α_j . We partition $j \in [M']$ into 2^{dl} groups as follows: For each $\alpha \in \{0, 1\}^{dl}$, let \mathcal{J}_α be the set of $j \in [M']$ such that the dl bits of α_j (that Eq. (19) for this j depends on) equals to α . We will estimate Eq. (19) by enumerating $\alpha \in \{0, 1\}^{dl}$, estimating it for $j \leftarrow \mathcal{J}_\alpha$ (instead of $j \leftarrow [M']$), and then taking the average over all possible α .

Now we fix any $\alpha \in \{0, 1\}^{dl}$. We can rephrase the following items as they no longer depend on α_j :

$$\begin{aligned} \text{coeff}_{k_t}(\alpha_j) &=: \text{coeff}'_t; \\ \text{id}_{x_{k_t}^x}(\alpha_j, q_t(i)) &=: \text{id}_{x_t^{x'}}(i); \\ \text{id}_{x_{k_t}^{\mathcal{C}}}(\alpha_j, q_t(i)) &=: \text{id}_{x_t^{\mathcal{C}'}}(i). \end{aligned}$$

It then suffices to estimate

$$\begin{aligned} &\mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow \mathcal{J}_\alpha} \left[\prod_{t=1}^d \text{coeff}'_t \cdot (C_j)_{\text{id}_{x_t^{\mathcal{C}'}}(i)} (x_{\text{id}_{x_t^{x'}}(i)}) \right] \\ &= \left(\prod_{t=1}^d \text{coeff}'_t \right) \cdot \mathbb{E}_{i \leftarrow [\ell_{\text{Prod}}], j \leftarrow \mathcal{J}_\alpha} \left[\bigwedge_{t=1}^d (C_j)_{\text{id}_{x_t^{\mathcal{C}'}}(i)} (x_{\text{id}_{x_t^{x'}}(i)}) \right]. \end{aligned} \quad (20)$$

Now for $\beta \in [\ell_{\mathcal{C}'}]^d$, let $\mathcal{I}_\beta := \{i \in [\ell_{\text{Prod}}] : \forall t \in [d], \text{id}_{x_t^{\mathcal{C}'}}(i) = \beta_t\}$. We enumerate over β , and now it suffices to estimate

$$\mathbb{E}_{i \leftarrow \mathcal{I}_\beta, j \leftarrow \mathcal{J}_\alpha} \left[\bigwedge_{t=1}^d (C_j)_{\beta_t} (x_{\text{id}_{x_t^{x'}}(i)}) \right]. \quad (21)$$

Each expression of the form $\mathbb{E}_{i,j} \left[\bigwedge_{t=1}^d (C_j)_{\beta_t} (x_{\text{id}_{x_t^{x'}}(i)}) \right]$ can be reduced to the SATISFYING-PAIRS problem for $\text{AND}_d \circ \mathcal{C}$ circuits. More precisely, we split \mathcal{I}_β into blocks of size N and \mathcal{J}_α into blocks of size M , and use the assumed algorithm to estimate Eq. (21). By similar argument as in Theorem 3.7, the additive error of our algorithm is bounded by $\eta \cdot U^d$.

Complexity. The subroutine for estimating Eq. (21) takes $O(T^{\text{alg}}) \cdot [|\mathcal{J}_\alpha|/M] \cdot [|\mathcal{I}_\beta/N|]$ time. Therefore, the subroutine for estimating Eq. (20) takes

$$\sum_{\beta \in [\ell_{\mathcal{C}}]^d} O(T^{\text{alg}}) \cdot [|\mathcal{J}_\alpha|/M] \cdot [|\mathcal{I}_\beta|/N] = O(T^{\text{alg}}) \cdot [|\mathcal{J}_\alpha|/M] \cdot (\ell_{\mathcal{C}}^d + \ell_{\text{Prod}}/N)$$

time. It then follows that the subroutine for estimating Eq. (19) takes

$$\sum_{\alpha \in \{0,1\}^{dl}} O(T^{\text{alg}}) \cdot [|\mathcal{J}_\alpha|/M] \cdot (\ell_{\mathcal{C}}^d + \ell_{\text{Prod}}/N) = O(T^{\text{alg}}) \cdot (2^{dl} + M'/M) \cdot (\ell_{\mathcal{C}}^d + \ell_{\text{Prod}}/N)$$

time, and finally, estimating Eq. (18) takes

$$O(T^{\text{alg}}) \cdot (2^{dl} + M'/M) \cdot (\ell_{\mathcal{C}}^d + \ell_{\text{Prod}}/N) \cdot A^d$$

time, which is the total time complexity of our algorithm. \diamond

A.2 Verifying PCPP with Satisfying Pairs

Lemma 3.8. *Let \mathcal{C} be a typical circuit class and $d \geq 2$ be an even number. Suppose there is an algorithm that takes as input a list of $N = 2^{r_{\text{col}}}$ AND $_{2d} \circ \mathcal{C}$ circuits $\{C_i\}$ and a list of N inputs $\{x_j\}$, runs in deterministic T^{alg} time, and estimates the following quantity with additive error η :*

$$\Pr_{i,j \leftarrow [2^{r_{\text{col}}}]}[C_i(x_j)].$$

Then there is an algorithm that takes the circuit C , $(w_1, w_2, \dots, w_{\hat{H}_{\text{proof}}})$, and $(\alpha_1, \alpha_2, \dots, \alpha_{\hat{H}_{\text{proof}}})$ as input, runs in deterministic $O((3A)^{2d} T^{\text{alg}}) \cdot (2^{2dl+r_{\text{shared}}} + T \log^{O(m)} T / 2^{2r_{\text{col}}})$ time, and satisfies the following:

(Completeness) *If for every $i \in [\hat{H}_{\text{proof}}]$, it holds that (1) for every $j \in [W_{\text{proof}}]$, $\pi_{i,j}^{\text{Real}} \in [0, 1]$; (2) $\|\pi_i^{\text{Real}} - \pi_i^{\text{Bool}}\|_1 \leq \delta$, then the algorithm accepts.*

(Soundness) *If the algorithm accepts, then it holds that*

1. *for every $\text{seed.shared} \in \{0, 1\}^{r_{\text{shared}}}$ and $\iota \in [q]$, $\|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d^d \leq 1 + 2\eta \cdot U^d$;*
2. $\mathbb{E}_{i \leftarrow [\hat{H}_{\text{proof}}], j \leftarrow [W_{\text{proof}}]} [|\pi_{i,j}^{\text{Real}} - \pi_{i,j}^{\text{Bool}}|^d] \leq 4^d \cdot \delta + 2^{d+1} \eta (2U + 1)^{2d}$.

Proof. Fix seed.shared and ι , we first estimate

$$\|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d^d := \mathbb{E}_{\text{seed.row}, \text{seed.col}} \left[|\pi_{\text{irow}[\iota], \text{icol}[\iota]}^{\text{Real}}|^d \right].$$

Recall that

$$\pi_{i,j}^{\text{Real}} = \sum_{k \in [A]} \text{coeff}_k(\alpha_i) \cdot C_{\text{id}_{\text{x}}k(\alpha_i, j)}(w_i).$$

Therefore, we can build a $\text{Prod}_d \circ \text{Sum} \circ \mathcal{C}$ circuit $C_{\text{norm}} := C_{\text{norm}}(\text{seed.shared}, \iota)$ as follows.

Circuit C_{norm}

(Inputs) The input consists of (w, α) with the intended meaning that $w = w_{\text{irow}[\iota]}$ and $\alpha = \alpha_{\text{irow}[\iota]}$.

(Bottom circuits) The bottom circuit is exactly C (taking input w). Thus, there are ℓ output gates of \mathcal{C} circuits with the i -th one being precisely the i -th output gate of C .

(Intermediate linear sum gates) There are $2^{r_{\text{col}}}$ intermediate linear sum gates. For each seed.col ,

$$\text{Sum}_{\text{seed.col}}(w, \alpha) = \sum_{k \in [A]} \text{coeff}_k(\alpha) \cdot C_{\text{id}_{x_k}(\alpha, i_{\text{col}[l]})}(w).$$

(Output product gates) There are $2^{r_{\text{col}}}$ product gates. For each seed.col , the seed.col -th output gate is simply

$$(C_{\text{norm}})_{\text{seed.col}}(w, \alpha) = (\text{Sum}_{\text{seed.col}}(w, \alpha))^d.$$

Recall that this circuit C_{norm} has parameters as follows:

- the number of gates in each layer: $\ell_{\mathcal{L}} = \ell$, $\ell_{\text{Sum}} = 2^{r_{\text{col}}}$, $\ell_{\text{Prod}} = 2^{r_{\text{col}}}$;
- the fan-in of the top Prod gates d ;
- the fan-in A , coefficient sum U , and locality l of the linear sum layer.

We invoke Theorem 3.7 on the circuit C_{norm} and $M' := 2^{r_{\text{row}}}$ inputs $\{(w_{i_{\text{row}[l]}}, \alpha_{i_{\text{row}[l]}})\}_{\text{seed.row}}$. We obtain an estimation $\text{EST}_{\text{norm}} = \text{EST}_{\text{norm}}(\text{seed.shared}, \iota)$ where

$$\left| \text{EST}_{\text{norm}} - \|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d^d \right| \leq \eta \cdot U^d.$$

If $\text{EST}_{\text{norm}} > 1 + \eta \cdot U^d$, then we reject the input. Otherwise, we proceed to verify that π^{Real} and π^{Bool} are close. Consider the polynomial $P(z) := z^d(1-z)^d$, which intuitively measures how close z is to Boolean. We will estimate

$$\mathbb{E}_{i \leftarrow [\hat{H}_{\text{proof}}], j \leftarrow [W_{\text{proof}}]} \left[P(\pi_{i,j}^{\text{Real}}) \right]. \quad (22)$$

Similarly, we estimate Eq. (22) by building a $\text{Prod}_{2d} \circ \text{Sum} \circ \mathcal{C}$ circuit C_{diff} .

Circuit C_{diff}

(Inputs and bottom circuits) The inputs and bottom circuits of C_{diff} are exactly the same as C_{norm} .

(Intermediate linear sum gates) There are $2W_{\text{proof}}$ intermediate linear sum gates. Let $j \in [W_{\text{proof}}]$, then the $2j$ -th linear sum gate computes $(\pi^{\text{Real}})_j$, and the $(2j+1)$ -th one computes $1 - (\pi^{\text{Real}})_j$. That is,

$$\text{Sum}_{2j}(w, \alpha) = \sum_{k \in [A]} \text{coeff}_k(\alpha) \cdot C_{\text{id}_{x_k}(\alpha, j)}(w); \quad \text{Sum}_{2j+1}(w, \alpha) = 1 - \text{Sum}_{2j}(w, \alpha).$$

Implementation of the linear sum layer: Since we did not allow $\text{coeff}_k(\alpha)$ to depend on i (the output index in $[2W_{\text{proof}}]$), we need to be careful when implementing the linear sum layer. The fan-in of this layer will be $2A+1$ (instead of A). We identify $[2A+1]$ with the disjoint union of $[A] \times \{0, 1\}$ and $\{\star\}$ (where \star denotes the constant term 1 in $\text{Sum}_{2j+1}(w, \alpha)$). Let $\text{id}_{x'}$ and coeff' be the id_{x} and coeff functions of the intermediate linear sum gates of C_{diff} :

(Function $\text{id}_{x'_k}(\alpha, i)$) We write $i = 2j + b$ where $j \in [W_{\text{proof}}]$ and $b \in \{0, 1\}$. If $k = (k', b') \in [A] \times \{0, 1\}$, then $\text{id}_{x'_k}(\alpha, i)$ returns $\text{id}_{x_{k'}}(\alpha, j)$ if $b = b'$ and returns **ZERO** if $b \neq b'$. If $k = \star$ then $\text{id}_{x'_k}(\alpha, i)$ returns **ZERO** if $b = 0$ and returns **ONE** if $b = 1$.

(Function $\text{coeff}'_k(\alpha)$) If $k = (k', b') \in [A] \times \{0, 1\}$, then $\text{coeff}'_k(\alpha) = (-1)^{b'} \cdot \text{coeff}_{k'}(\alpha)$. If $k = \star$ then $\text{coeff}'_k(\alpha) = 1$.

The locality of $(\text{id}_{x'}, \text{coeff}')$ is still l . The coefficient sum becomes $2U+1$.

(Output product gates) There are W_{proof} product gates. For each $j \in [W_{\text{proof}}]$, the j -th output gate is

$$C_{\text{diff}}(w, \alpha) = (\text{Sum}_{2^j}(w, \alpha) \cdot \text{Sum}_{2^{j+1}}(w, \alpha))^d.$$

The parameters of the circuit C_{diff} are as follows:

- the number of gates in each layer: $\ell_{\mathcal{L}} = \ell$, $\ell_{\text{Sum}} = 2W_{\text{proof}}$, $\ell_{\text{Prod}} = W_{\text{proof}}$;
- the fan-in of the top Prod gates $2d$;
- the fan-in $2A + 1$, coefficient sum $2U + 1$, and locality l of the linear sum layer.

We invoke Theorem 3.7 on the circuit C_{diff} and the set of $M' := \hat{H}_{\text{proof}}$ inputs $\{(w_i, \alpha_i)\}_{i \in [\hat{H}_{\text{proof}}]}$, and obtain an estimation EST_{diff} where

$$|\text{EST}_{\text{diff}} - (22)| \leq \eta \cdot (2U + 1)^{2d}.$$

We then accept if and only if $\text{EST}_{\text{diff}} \leq 2^d \cdot \delta + \eta(2U + 1)^{2d}$.

Complexity. Our algorithm calls the algorithm in Theorem 3.7 as a subroutine on the circuits C_{norm} and C_{diff} . It is easy to see that $A^d(2^{dl} + 2^{r_{\text{row}}}/2^{r_{\text{col}}}) \cdot O(T^{\text{alg}})$ time is spent on each C_{norm} . Similarly, it takes $(2A + 1)^{2d}(2^{2dl} + \hat{H}_{\text{proof}}/2^{r_{\text{col}}}) \cdot (W_{\text{proof}}/2^{r_{\text{col}}}) \cdot O(T^{\text{alg}})$ time to process C_{diff} . It follows that our algorithm runs in deterministic time

$$\begin{aligned} & O(2^{r_{\text{shared}}}) \cdot A^d(2^{dl} + 2^{r_{\text{row}}}/2^{r_{\text{col}}}) \cdot O(T^{\text{alg}}) + (2A + 1)^{2d}(2^{2dl} + \hat{H}_{\text{proof}}/2^{r_{\text{col}}}) \cdot (W_{\text{proof}}/2^{r_{\text{col}}}) \cdot O(T^{\text{alg}}) \\ &= O((3A)^{2d}T^{\text{alg}}) \cdot \left(2^{dl+r_{\text{shared}}} + 2^{r-2r_{\text{col}}} + 2^{2dl} \cdot W_{\text{proof}}/2^{r_{\text{col}}} + \hat{H}_{\text{proof}} \cdot W_{\text{proof}}/2^{2r_{\text{col}}}\right) \\ &= O((3A)^{2d}T^{\text{alg}}) \cdot \left(2^{2dl+r_{\text{shared}}} + T \log^{O(m)} T / 2^{2r_{\text{col}}}\right). \end{aligned}$$

(Recall that since $\log(W_{\text{proof}}/2^{r_{\text{col}}}) = (5/m) \log T$, $r_{\text{shared}} \geq (10/m) \log T$, we have $W_{\text{proof}}/2^{r_{\text{col}}} \leq 2^{r_{\text{shared}}}$. Also, from $r := \log T + O(\log \log T + m \log m)$ and $m = \Theta(\log n/\delta)$ we know that $2^{r-2r_{\text{col}}} \leq T \log^{O(m)} T / 2^{2r_{\text{col}}}$.)

Now it suffices to prove the completeness and soundness requirements. Before that, we need the following fact regarding the polynomial P . Let $z \in \mathbb{R}$, $d_{\text{bin}}(z)$ be the distance between z to the closest Boolean value; namely $d_{\text{bin}}(z) := \min\{|z|, |1 - z|\}$.

Fact A.1. For every $z \in \mathbb{R}$, $d_{\text{bin}}(z)^d \cdot 2^{-d} \leq P(z) \leq d_{\text{bin}}(z)^d \cdot (1 + d_{\text{bin}}(z))^d$. \diamond

Completeness. Suppose that for every $i \in [\hat{H}_{\text{proof}}]$ and $j \in [W_{\text{proof}}]$, $\pi_{i,j}^{\text{Real}} \in [0, 1]$. Then for every $\iota \in [q]$ and $\text{seed.shared} \in \{0, 1\}^{r_{\text{shared}}}$, $\|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d^d \leq 1$, and thus $\text{EST}_{\text{norm}} \leq 1 + \eta \cdot U^d$.

Suppose in addition that for every $i \in [\hat{H}_{\text{proof}}]$, $\|\pi_i^{\text{Real}} - \pi_i^{\text{Bool}}\|_1 \leq \delta$. Then:

$$\begin{aligned} (22) &\leq \mathbb{E}_{i,j} \left[d_{\text{bin}}(\pi_{i,j}^{\text{Real}})^d \cdot (1 + d_{\text{bin}}(\pi_{i,j}^{\text{Real}}))^d \right] \\ &\leq 2^d \cdot \mathbb{E}_{i,j} \left[d_{\text{bin}}(\pi_{i,j}^{\text{Real}})^d \right] \\ &\leq 2^d \cdot \delta. \end{aligned}$$

Therefore, $\text{EST}_{\text{diff}} \leq 2^d \cdot \delta + \eta(2U + 1)^{2d}$ and our algorithm accepts.

Soundness. Suppose our algorithm accepts.

1. For every `seed.shared` and ι , we have $\text{EST}_{\text{norm}} \leq 1 + \eta \cdot U^d$ and thus $\|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d^d \leq 1 + 2\eta \cdot U^d$.
2. We have (22) $\leq 2^d \cdot \delta + 2\eta(2U + 1)^{2d}$ and

$$\begin{aligned} \mathbb{E}_{i \leftarrow [\hat{H}_{\text{proof}}], j \leftarrow [W_{\text{proof}}]} \left[|\pi_{i,j}^{\text{Real}} - \pi_{i,j}^{\text{Bool}}|^d \right] &= \mathbb{E}_{i,j} \left[d_{\text{bin}}(\pi_{i,j}^{\text{Real}})^d \right] \\ &\leq 2^d \cdot (22) && \text{(Fact A.1)} \\ &\leq 4^d \cdot \delta + 2^{d+1} \eta (2U + 1)^{2d}. && \square \end{aligned}$$

Lemma 4.6. Let \mathcal{C} be a typical circuit class and $d \geq 2$ be an even number. Suppose there is an algorithm that takes as inputs a list of $2^{r_{\text{col}}}$ $\text{AND}_{2d} \circ \mathcal{C}$ circuits $\{C_i\}$ and a list of $2^{r_{\text{col}}}$ inputs $\{x_j\}$ of length $n \cdot \text{polylog}(\ell)$, runs in deterministic T^{alg} time, and estimates the following quantity with additive error η :

$$\Pr_{i,j \leftarrow [2^{r_{\text{col}}}] } [C_i(x_j)].$$

Then there is an algorithm that takes the strings w_1, w_2, \dots, w_ℓ , circuits $(C_1, C_2, \dots, C_{\hat{H}_{\text{proof}}})$, and $(\alpha_1, \alpha_2, \dots, \alpha_{\hat{H}_{\text{proof}}})$ as inputs, runs in deterministic $O((3A)^{2d} T^{\text{alg}}) \cdot (2^{2d\ell + r_{\text{shared}}} + T \log^{O(m)} T / 2^{2r_{\text{col}}})$ time, and satisfies the following:

(Completeness) If for every $i \in [\hat{H}_{\text{proof}}]$, it holds that (1) for every $j \in [W_{\text{proof}}]$, $\pi_{i,j}^{\text{Real}} \in [0, 1]$; (2) $\|\pi_i^{\text{Real}} - \pi_i^{\text{Bool}}\|_1 \leq \delta$, then the algorithm accepts.

(Soundness) If the algorithm accepts, then it holds that

1. for every `seed.shared` $\in \{0, 1\}^{r_{\text{shared}}}$ and $\iota \in [q]$, $\|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d^d \leq 1 + 2\eta \cdot U^d$;
2. $\mathbb{E}_{i \leftarrow [\hat{H}_{\text{proof}}], j \leftarrow [W_{\text{proof}}]} \left[|\pi_{i,j}^{\text{Real}} - \pi_{i,j}^{\text{Bool}}|^d \right] \leq 4^d \cdot \delta + 2^{d+1} \eta (2U + 1)^{2d}$.

Proof Sketch. We first estimate $\|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d^d$ for fixed `seed.shared` and ι . Recall that

$$\pi_{i,j}^{\text{Real}} = \sum_{k \in [A]} \text{coeff}_k(\alpha_i) \cdot C_i(w_{\text{id}_{X_k}(\alpha_i, j)}).$$

We build a $\text{Prod}_d \circ \text{Sum}$ circuit $C_{\text{norm}} := C_{\text{norm}}(\text{seed.shared}, \iota)$ as follows.

Circuit C_{norm}

(Inputs) The input consists of (y, α) with the intended meaning that $y = (y_1, y_2, \dots, y_\ell)$ where $y_i = C_{i_{\text{row}[\iota]}(w_i)}$, and $\alpha = \alpha_{i_{\text{row}[\iota]}}$.

(Linear sum gates) There are $2^{r_{\text{col}}}$ linear sum gates. For each `seed.col`,

$$\text{Sum}_{\text{seed.col}}(y, \alpha) = \sum_{k \in [A]} \text{coeff}_k(\alpha) \cdot y_{\text{id}_{X_k}(\alpha, i_{\text{col}[\iota]})}.$$

(Output product gates) There are $2^{r_{\text{col}}}$ product gates. For each `seed.col`, the `seed.col`-th output gate is simply

$$(C_{\text{norm}})_{\text{seed.col}}(y, \alpha) = (\text{Sum}_{\text{seed.col}}(y, \alpha))^d.$$

Recall that this circuit C_{norm} has parameters as follows:

- the number of gates in each layer: $\ell_{\text{Sum}} = 2^{r_{\text{col}}}$, $\ell_{\text{Prod}} = 2^{r_{\text{col}}}$;
- the fan-in of the top Prod gates d ;
- the fan-in A , coefficient sum U , and locality l of the linear sum layer.

We invoke Lemma 4.5 on the circuit C_{norm} , strings w_1, w_2, \dots, w_ℓ , a list of $2^{r_{\text{row}}}$ inputs $\{\alpha_{i_{\text{row}[l]}}\}_{\text{seed.row}}$, and a list of $2^{r_{\text{row}}}$ size- s \mathcal{C} circuits $\{C_{i_{\text{row}[l]}}\}_{\text{seed.row}}$. Here $\ell_{\mathcal{C}} = 1$. We thus obtain an estimation $\text{EST}_{\text{norm}} = \text{EST}_{\text{norm}}(\text{seed.shared}, \iota)$ where

$$\left| \text{EST}_{\text{norm}} - \|f_{\text{seed.shared}, \iota}^{\text{Real}}\|_d^d \right| \leq \eta \cdot U^d.$$

If $\text{EST}_{\text{norm}} > 1 + \eta \cdot U^d$, then we reject the input. Otherwise, we proceed to verify that π^{Real} and π^{Bool} are close. Consider the polynomial $P(z) := z^d(1-z)^d$. we will estimate

$$\mathbb{E}_{i \leftarrow [\hat{H}_{\text{proof}}], j \leftarrow [W_{\text{proof}}]} \left[P(\pi_{i,j}^{\text{Real}}) \right]. \quad (23)$$

Similarly, we estimate Eq. (23) by building a $\text{Prod}_{2d} \circ \text{Sum}$ circuit C_{diff} .

Circuit C_{diff}

(Inputs) The inputs are exactly the same as C_{norm} .

(Linear sum gates) There are $2W_{\text{proof}}$ linear sum gates. Let $j \in [W_{\text{proof}}]$, then the $2j$ -th linear sum gate computes $(\pi^{\text{Real}})_j$, and the $(2j+1)$ -th one computes $1 - (\pi^{\text{Real}})_j$. That is,

$$\text{Sum}_{2j}(y, \alpha) = \sum_{k \in [A]} \text{coeff}_k(\alpha) \cdot y_{\text{id}_{\times k}(\alpha, j)}; \quad \text{Sum}_{2j+1}(y, \alpha) = 1 - \text{Sum}_{2j}(y, \alpha).$$

The implementation of the linear sum layer is the same as in Lemma 3.8, and we omit it here.

(Output product gates) There are W_{proof} product gates. For each $j \in [W_{\text{proof}}]$, the j -th output gate is

$$C_{\text{diff}}(y, \alpha) = (\text{Sum}_{2j}(y, \alpha) \cdot \text{Sum}_{2j+1}(y, \alpha))^d.$$

The parameters of the circuit C_{diff} are as follows:

- the number of gates in each layer: $\ell_{\text{Sum}} = 2W_{\text{proof}}$, $\ell_{\text{Prod}} = W_{\text{proof}}$;
- the fan-in of the top Prod gates $2d$;
- the fan-in $2A+1$, coefficient sum $2U+1$, and locality l of the linear sum layer.

We invoke Lemma 4.5 on the circuit C_{diff} , strings w_1, w_2, \dots, w_ℓ , a list of \hat{H}_{proof} inputs $\{\alpha_i\}$, and a list of \hat{H}_{proof} size- s \mathcal{C} circuits $\{C_i\}$. Here $\ell_{\mathcal{C}} = 1$. We obtain an estimation EST_{diff} where

$$\left| \text{EST}_{\text{diff}} - (23) \right| \leq \eta \cdot (2U+1)^{2d}.$$

We accept if and only if $\text{EST}_{\text{diff}} \leq 2^d \cdot \delta + \eta(2U+1)^{2d}$.

The correctness and complexity are analysed the same as in Lemma 3.8, so we omit it here. \diamond

A.3 Proof of Claim 3.9

We need the following technical lemma (see [CW19b, Lemma 28] and [RSW22, Lemma 4.9]):

Lemma A.2. *Let $d \geq 2$ be an integer, $f_1, f_2, \dots, f_d, g_1, g_2, \dots, g_d : [N] \rightarrow \mathbb{R}$ be functions. For all $i \in [d]$, suppose that $\|f_i\|_d \leq 1$, and define $\varepsilon := \sum_{i=1}^d \|f_i - g_i\|_d$. Then*

$$\left| \mathbb{E}_{x \leftarrow [N]} \left[\prod_{i=1}^d f_i(x) - \prod_{i=1}^d g_i(x) \right] \right| \leq (1 + \varepsilon)^{d-1} \cdot \varepsilon.$$

The above lemma is a consequence of the following generalisation of Hölder's inequality:

Fact A.3. Let $f_1, f_2, \dots, f_d : [N] \rightarrow \mathbb{R}$ be functions, $f : [N] \rightarrow \mathbb{R}$ be their product, i.e., $f(x) = \prod_{i=1}^d f_i(x)$. Then $\|f\|_1 \leq \prod_{i=1}^d \|f_i\|_d$.

Proof of Lemma A.2. Let $\varepsilon_i := \|f_i - g_i\|_d$, then $\varepsilon = \sum_{i=1}^d \varepsilon_i$. Define

$$\text{Hyb}_i := \mathbb{E}_{x \leftarrow [N]} \left[\prod_{j=1}^i f_j(x) \cdot \prod_{j=i+1}^d g_j(x) \right].$$

Then, for every $1 \leq i \leq d$,

$$\begin{aligned} |\text{Hyb}_i - \text{Hyb}_{i-1}| &\leq \mathbb{E}_{x \leftarrow [N]} \left[\left| \prod_{j=1}^{i-1} f_j(x) \cdot \prod_{j=i+1}^d g_j(x) \cdot (f_i(x) - g_i(x)) \right| \right] \\ &\leq \prod_{j=1}^{i-1} \|f_j\|_d \cdot \prod_{j=i+1}^d \|g_j\|_d \cdot \|f_i - g_i\|_d && \text{(Fact A.3)} \\ &\leq \prod_{j=2}^d (1 + \varepsilon_j) \cdot \varepsilon_i \\ &\leq (1 + \varepsilon)^{d-1} \cdot \varepsilon_i. \end{aligned}$$

It follows that

$$\left| \mathbb{E}_{x \leftarrow [N]} \left[\prod_{i=1}^d f_i(x) - \prod_{i=1}^d g_i(x) \right] \right| = |\text{Hyb}_d - \text{Hyb}_0| \leq \sum_{i=1}^d |\text{Hyb}_i - \text{Hyb}_{i-1}| \leq (1 + \varepsilon)^{d-1} \cdot \varepsilon. \quad \square$$

Recall that for $S, S' \subseteq [q]$, we define

$$\begin{aligned} p_{\text{acc}}(\text{seed.shared}, S, S') &:= \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\prod_{\ell \in S} (\text{row}_{\ell}^{\text{Bool}})_{\text{icol}[\ell]} \cdot \prod_{\ell \in S'} p_{c_{\ell}} \right], \\ p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S') &:= \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\prod_{\ell \in S} (\text{row}_{\ell}^{\text{Real}})_{\text{icol}[\ell]} \cdot \prod_{\ell \in S'} p_{c_{\ell}} \right], \\ \delta_{\text{seed.shared}} &:= \sum_{\ell: \text{itype}[\ell] = \text{proof}} \|f_{\text{seed.shared}, \ell}^{\text{Bool}} - f_{\text{seed.shared}, \ell}^{\text{Real}}\|_{2q}. \end{aligned}$$

Claim 3.9. For every $S, S' \subseteq [q]$,

$$|p_{\text{acc}}(\text{seed.shared}, S, S') - p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S')| \leq (1 + \delta_{\text{seed.shared}})^{2q-1} \cdot \delta_{\text{seed.shared}}.$$

Proof. Define the following $2(|S| + |S'|)$ functions $f_i^{\text{Bool}}, g_j^{\text{Bool}}, f_i^{\text{Real}}, g_j^{\text{Real}}$, where $i \in S$ and $j \in S'$. Each function takes $(\text{seed.row}, \text{seed.col})$ as inputs, and:

$$f_i^{\text{Bool}} := (\text{row}_i^{\text{Bool}})_{\text{icol}[i]}; \quad f_i^{\text{Real}} := (\text{row}_i^{\text{Real}})_{\text{icol}[i]}; \quad \text{and} \quad g_j^{\text{Bool}} = g_j^{\text{Real}} := p_{c_j}.$$

(Note: for convenience, we omit the input $(\text{seed.row}, \text{seed.col})$.) It follows that $\|f_i^{\text{Bool}}\|_{2q}, \|g_j^{\text{Bool}}\|_{2q} \leq 1$; for every $j \in S'$, $\|g_j^{\text{Bool}} - g_j^{\text{Real}}\|_{2q} = 0$; and for every $i \in S$,

$$\|f_i^{\text{Bool}} - f_i^{\text{Real}}\|_{2q} = \begin{cases} 0 & \text{if } \text{itype}[i] = \text{input}; \\ \|f_{\text{seed.shared}, i}^{\text{Bool}} - f_{\text{seed.shared}, i}^{\text{Real}}\|_{2q} & \text{if } \text{itype}[i] = \text{proof}. \end{cases}$$

Therefore, by Lemma A.2,

$$\begin{aligned}
& |p_{\text{acc}}(\text{seed.shared}, S, S') - p_{\text{acc}}^{\text{Real}}(\text{seed.shared}, S, S')| \\
&= \left| \mathbb{E}_{\substack{\text{seed.row} \leftarrow \{0,1\}^{r_{\text{row}}} \\ \text{seed.col} \leftarrow \{0,1\}^{r_{\text{col}}}}} \left[\prod_{i \in S} f_i^{\text{Bool}} \cdot \prod_{j \in S'} g_j^{\text{Bool}} - \prod_{i \in S} f_i^{\text{Real}} \cdot \prod_{j \in S'} g_j^{\text{Real}} \right] \right| \\
&\leq (1 + \delta_{\text{seed.shared}})^{2q-1} \cdot \delta_{\text{seed.shared}}. \quad \square
\end{aligned}$$

A.4 Proof of Lemma 3.10

Lemma 3.10. *Let $f : [N] \times [q] \rightarrow \mathbb{R}_{\geq 0}$ be a function and $d \geq 1$ be a constant. Suppose that*

1. *for every $s \in [N]$ and $i \in [q]$, $f(s, i) \leq \alpha$ (where $\alpha \geq 1$);*
2. $\mathbb{E}_{s,i}[f(s, i)^d] \leq \delta$.

Let $f(s) := \sum_{i \in [q]} f(s, i)$. Then

$$\mathbb{E}_s[(1 + f(s))^{d-1} \cdot f(s)] \leq q\delta^{1/d}(2q\alpha)^{d-1}.$$

Proof. By Jensen's inequality,

$$\mathbb{E}_s[f(s)] = q \mathbb{E}_{s,i}[f(s, i)] \leq q\delta^{1/d}.$$

It follows that for every $k \geq 1$,

$$\mathbb{E}_s[f(s)^k] \leq \mathbb{E}_s[f(s)] \cdot \max_s \{f(s)\}^{k-1} \leq q\delta^{1/d} \cdot (q\alpha)^{k-1}.$$

Finally, we have

$$\mathbb{E}_s[(1 + f(s))^{d-1} \cdot f(s)] = \sum_{i=0}^{d-1} \binom{d-1}{i} \cdot \mathbb{E}_s[f(s)^{i+1}] \leq q\delta^{1/d}(2q\alpha)^{d-1}. \quad \square$$

A.5 An XOR Lemma in [CLW20]

Theorem 2.5. *Let $N \in \mathbb{N}$, $0 < \varepsilon, \delta < 1/10$, $k := O(\log(1/\varepsilon)/\delta)$, $\tilde{N} := N^k$, and $a := O(\log^2 N/(\varepsilon\delta)^2)$. There is an algorithm $\text{Amp} : \{0, 1\}^N \rightarrow \{0, 1\}^{\tilde{N}}$ computable in deterministic $\text{poly}(\tilde{N})$ time, and a linear sum circuit $C : \{0, 1\}^{\tilde{N}} \times \{0, 1\}^a \rightarrow \mathbb{R}^N$ such that the following hold.*

(List-decoding) *For every string $\tilde{f} \in \{0, 1\}^{\tilde{N}}$ that is $(1/2 - \varepsilon)$ -close to $\text{Amp}(f)$ for some hidden string f , there is an advice $\alpha \in \{0, 1\}^a$, such that (1) for every $i \in [N]$, $C(\tilde{f}, \alpha)_i \in [0, 1]$; and (2) $\|C(\tilde{f}, \alpha) - f\|_1 \leq \delta$.*

(Complexity) *The fan-in, coefficient sum, and locality of C are at most $O(\log N/(\varepsilon\delta)^2)$, $O(1/\varepsilon)$, and $\log \tilde{N}$ respectively.*

The lemma is implied by the XOR lemma in [CLW20]. For simplicity, we identify a string f of length N and a Boolean function $f : [N] \rightarrow \{0, 1\}$, where $f(x)$ outputs the x -th bit of f . For a string $f \in \{0, 1\}^N$, denote $f^{\oplus k}$ to be the following string of length N^k . For each $(x_1, x_2, \dots, x_k) \in [N]^k$, we have

$$f^{\oplus k}(x_1, x_2, \dots, x_k) := \bigoplus_{i=1}^k f(x_i).$$

And we simply let $\text{Amp}(f) := f^{\oplus k}$.

The decoder. For a length- k vector $\vec{v}^\perp \in ([N] \cup \{\perp\})^k$ and $i \in [N]$, let \vec{v}^i denote the vector where each \perp in \vec{v}^\perp is replaced by i . (In the decoder, we will only need the case where each \vec{v}^\perp contains exactly one \perp , so \vec{v}^i simply replaces that single \perp by i .)

Let $\tilde{f} : [N]^k \rightarrow \{0, 1\}$ be a codeword (treated as a Boolean function). For $A' := O(\log N/(\varepsilon\delta)^2)$ and $r := \frac{(2+\delta)\varepsilon}{1-\delta}$, our decoder will take a list of vectors $\vec{v}_1^\perp, \vec{v}_2^\perp, \dots, \vec{v}_{A'}^\perp \in ([N] \cup \{\perp\})^k$ and a list of signs $\sigma_1, \sigma_2, \dots, \sigma_{A'} \in \{0, 1\}$ as advice. Intuitively, \vec{v}_i^\perp denotes a segment of \tilde{f} that has noticeable correlation with f , and σ_i denotes whether the correlation is positive or negative; our linear sum decoder uses the average of $\tilde{f}(\vec{v}_j^\perp) \oplus \sigma_j$ as a prediction of f_i .

More formally, given an input $i \in [N]$, the decoder outputs ³⁰

$$\text{dec}(\tilde{f})_i := \frac{1}{r} \mathbb{E}_{j \leftarrow [A']} \left[(\tilde{f}(\vec{v}_j^\perp) \oplus \sigma_j) - 1/2 \right] + 1/2. \quad (24)$$

Correctness. We establish the correctness of this decoder by the following lemma.

Lemma A.4. *Let $k \geq 1$, $\delta \in (0, 1/10)$, $\varepsilon := (1 - \delta)^{k-1}(1/2 - \delta)$, and $A' := O(\log N/(\varepsilon\delta)^2)$. For every string $\tilde{f} \in \{0, 1\}^{N^k}$ that is $(1/2 - \varepsilon)$ -close to $f^{\oplus k}$ for some hidden string $f \in \{0, 1\}^N$, there is a list of A' vectors $\vec{v}_1^\perp, \vec{v}_2^\perp, \dots, \vec{v}_{A'}^\perp \subseteq ([N] \cup \{\perp\})^k$, and a list of signs $\sigma_1, \sigma_2, \dots, \sigma_{A'} \in \{0, 1\}$, such that (1) for every $i \in [N]$, $\text{dec}(\tilde{f})_i \in [0, 1]$; and (2) $\|\text{dec}(\tilde{f}) - f\|_1 \leq \delta$.*

Proof. We use induction on k . Suppose $k = 1$, then one can verify by direct calculation that the lemma holds by setting $\vec{v}_1^\perp = (\perp)$ and $\sigma_1 = 0$. Now suppose $k > 1$ and the lemma holds for $k - 1$.

Fix $i \in [N]$ and let $\vec{v}^\perp \in ([N] \cup \{\perp\})^k$ denote some vector whose first coordinate is \perp and other coordinates are from $[N]$. Think of every coordinate of \vec{v}^\perp , except the first, is drawn independently and uniformly from $[N]$. Define

$$p_i := \Pr_{\vec{v}^\perp \leftarrow \{\perp\} \times [N]^{k-1}} \left[\tilde{f}(\vec{v}^\perp) = f^{\oplus k}(\vec{v}^\perp) \right].$$

Case I: Suppose there is some $i_0 \in [N]$ such that $|p_{i_0} - 1/2| > \varepsilon/(1 - \delta)$. Let $b \in \{0, 1\}$ be a bit, consider the sub-string $\tilde{f}' \in \{0, 1\}^{N^{k-1}}$ such that $\tilde{f}'(\vec{v}^\perp) = \tilde{f}(\vec{v}^\perp) \oplus b$. Then, for some $b \in \{0, 1\}$, \tilde{f}' is $(1/2 - \varepsilon/(1 - \delta))$ -close to $f^{\oplus(k-1)}$.

By the induction hypothesis, there is a list of A' vectors $\vec{u}_1^\perp, \dots, \vec{u}_{A'}^\perp \subseteq ([N - 1] \cup \{\perp\})^k$ and a list of signs $\sigma'_1, \dots, \sigma'_{A'} \in \{0, 1\}$ such that the vector dec' satisfies the conclusion of the lemma, where

$$\text{dec}'_i := \frac{1}{r} \mathbb{E}_{j \leftarrow [A']} \left[(\tilde{f}'(\vec{u}_j^\perp) \oplus \sigma'_j) - 1/2 \right] + 1/2.$$

For each j , let \vec{v}_j^\perp be the concatenation of i_0 and \vec{u}_j^\perp , and let $\sigma_j = \sigma'_j \oplus b$. We have that $\text{dec}(\tilde{f})_i$ is exactly dec'_i and we are done.

³⁰Eq. (24) is perhaps easier to understand when we change the basis from $\{0, 1\}$ to $\{1, -1\}$; we choose the basis $\{0, 1\}$ only to be consistent with other parts of this paper. When we change the basis to $\{1, -1\}$, XOR becomes multiplication and the assertion “ $a = b$ ” becomes simply $a \cdot b$. Thus Eq. (24) becomes

$$\text{dec}(\tilde{f})_i = \frac{1}{r} \mathbb{E}_{j \leftarrow [A']} \left[\tilde{f}(\vec{v}_j^\perp) \cdot \sigma_j \right],$$

which is simply the average of all A' predictions, amplified by a factor of $1/r$.

Case II: Suppose for every $i \in [N]$, we have $|p_i - 1/2| \leq \varepsilon/(1 - \delta)$. Note that, since \tilde{f} is $(1/2 - \varepsilon)$ -close to $f^{\oplus k}$, we have

$$\mathbb{E}_{i \leftarrow [N]} [p_i] \geq 1/2 + \varepsilon.$$

We sample each $\vec{v}_j^\perp \leftarrow \{\perp\} \times [N]^{k-1}$ independently at random. Let

$$\tilde{p}_i := \Pr_{j \leftarrow [A']} [\tilde{f}(\vec{v}_j^i) = f^{\oplus k}(\vec{v}_j^i)].$$

Let $\eta := \frac{\varepsilon\delta}{2(1-\delta)}$. By a Chernoff bound, w.p. $1 - Ne^{-2\eta^2 t} > 0$, for every $i \in [N]$, we have $|p_i - \tilde{p}_i| \leq \eta$.

Let $\sigma_j := f^{\oplus(k-1)}((\vec{v}_j^\perp)_{2 \sim k}) = \bigoplus_{l=2}^k f((\vec{v}_j^\perp)_l)$, then we have

$$\tilde{p}_i = \Pr_{j \leftarrow [t]} [f_i = \tilde{f}(\vec{v}_j^i) \oplus \sigma_j].$$

Note that

$$\begin{aligned} \text{dec}(\tilde{f})_i &= \frac{1}{r} (\tilde{p}_i \cdot (f_i - 1/2) + (1 - \tilde{p}_i) \cdot (1/2 - f_i)) + 1/2 \\ &= \frac{(f_i - 1/2)(2\tilde{p}_i - 1)}{r} + 1/2. \end{aligned}$$

Claim A.5. For every $i \in [N]$, $\text{dec}(\tilde{f})_i \in [0, 1]$.

Proof.

$$\begin{aligned} |\text{dec}(\tilde{f})_i - 1/2| &= \frac{1}{r} |(f_i - 1/2)(2\tilde{p}_i - 1)| \\ &= \frac{1}{r} |\tilde{p}_i - 1/2| \\ &\leq \frac{1 - \delta}{(2 + \delta)\varepsilon} \left(\frac{\varepsilon}{1 - \delta} + \eta \right) \\ &= 1/2. \end{aligned} \quad \diamond$$

Claim A.6. $\|\text{dec}(\tilde{f}) - f\|_1 = \mathbb{E}_{i \leftarrow [N]} [|\text{dec}(\tilde{f})_i - f_i|] \leq \delta$.

Proof.

$$\begin{aligned} &\mathbb{E}_{i \leftarrow [N]} [|\text{dec}(\tilde{f})_i - f_i|] \\ &= \mathbb{E}_{i \leftarrow [N]} \left[\left| \frac{(f_i - 1/2)(2\tilde{p}_i - 1)}{r} + 1/2 - f_i \right| \right] \\ &= \mathbb{E}_{i \leftarrow [N]} \left[\left| \frac{\tilde{p}_i - 1/2}{r} - 1/2 \right| \right] \\ &= 1/2 - \frac{\mathbb{E}_{i \leftarrow [N]} [\tilde{p}_i] - 1/2}{r} \\ &\leq 1/2 - \frac{\varepsilon - \eta}{r} \leq \delta, \end{aligned} \tag{25}$$

where Eq. (25) is because we have shown in the proof of Claim A.5 that $\frac{1}{r} |\tilde{p}_i - 1/2| \leq 1/2$ for every i . \diamond

Combining Claim A.5 and A.6, the lemma is proved. \square

Complexity. It remains to determine the complexity of the decoder defined in Eq. (24). The advice string α contains the vectors $\vec{v}_1^\perp, \vec{v}_2^\perp, \dots, \vec{v}_{A'}^\perp$ and the signs $\sigma_1, \sigma_2, \dots, \sigma_{A'}$. It is clear that the fan-in is at most $A' + 1 = O(\log N / (\varepsilon \delta)^2)$. The coefficient sum is $O(1/r) = O(\varepsilon^{-1})$. Since the k -th term is

$$\text{coeff}_k(\alpha) = (-1)^{\sigma_j} / (A' r), \quad \text{and} \quad \text{idx}_k(\alpha, i) = \vec{v}_k^i,$$

it follows that each term only depends on $\log \tilde{N}$ bits of α .