Robust PCPs of Proximity, Shorter PCPs, and Applications to Coding

Citation

Ben#Sasson, Eli, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. 2006. "Robust PCPs of Proximity, Shorter PCPs, and Applications to Coding." SIAM Journal on Computing 36 (4): 889–974. <https://doi.org/10.1137/s0097539705446810>.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:41467483>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Robust PCPs of Proximity, Shorter PCPs and Applications to Coding

Eli Ben-Sasson ^{*} Oded Goldreich [†] Prahladh Harsha [‡] Madhu Sudan [§]
Salil Vadhan [¶]

March 14, 2005

Abstract

We continue the study of the trade-off between the length of PCPs and their query complexity, establishing the following main results (which refer to proofs of satisfiability of circuits of size n):

1. We present PCPs of length $\exp(\tilde{O}(\log \log n)^2) \cdot n$ that can be verified by making $o(\log \log n)$ Boolean queries.
2. For every $\varepsilon > 0$, we present PCPs of length $\exp(\log^\varepsilon n) \cdot n$ that can be verified by making a constant number of Boolean queries.

In both cases, false assertions are rejected with constant probability (which may be set to be arbitrarily close to 1). The multiplicative overhead on the length of the proof, introduced by transforming a proof into a probabilistically checkable one, is just quasi-polylogarithmic in the first case (of query complexity $o(\log \log n)$), and $2^{(\log n)^\varepsilon}$, for any $\varepsilon > 0$, in the second case (of constant query complexity). In contrast, previous results required at least $2^{\sqrt{\log n}}$ overhead in the length, even to get query complexity $2^{\sqrt{\log n}}$.

Our techniques include the introduction of a new variant of PCPs that we call “Robust PCPs of proximity”. These new PCPs facilitate proof composition, which is a central ingredient in construction of PCP systems. (A related notion and its composition properties were discovered independently by Dinur and Reingold.) Our main technical contribution is a construction of a “length-efficient” Robust PCP of proximity. While the new construction uses many of the standard techniques in PCPs, it does differ from previous constructions in fundamental ways, and in particular does not use the “parallelization” step of Arora *et al.* The alternative approach may be of independent interest.

We also obtain analogous quantitative results for locally testable codes. In addition, we introduce a relaxed notion of locally decodable codes, and present such codes mapping k information bits to codewords of length $k^{1+\varepsilon}$, for any $\varepsilon > 0$.

^{*}Radcliffe Institute for Advanced Study, Cambridge, MA 02139. Email: eli@eecs.harvard.edu.

[†]Department of Computer Science, Weizmann Institute of Science, Rehovot, ISRAEL. Email: oded.goldreich@weizmann.ac.il.

[‡]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139. Email: prahladh@mit.edu. Supported in part by NSF Award CCR-0312575.

[§]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139. Email: madhu@mit.edu. Supported in part by NSF Award CCR-0312575.

[¶]Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138. Email: salil@eecs.harvard.edu. Supported in part by NSF grant CCR-0133096 and a Sloan Research Fellowship.

⁰Part of the work was done while the first, second, fourth and fifth authors were fellows at the Radcliffe Institute for Advanced Study of Harvard University.

Contents

1	Introduction	2
1.1	PCPs with better length vs query trade-off	2
1.2	New notions and main techniques	3
1.3	Related work	5
1.4	Applications to coding problems	6
1.5	Relation to Property Testing	6
1.6	Organization	7
I	All but the main construct	7
2	PCPs and variants: definitions, observations and transformations	8
2.1	Standard PCPs	8
2.2	PCPs of Proximity	9
2.3	Robust Soundness	12
2.4	Composition	13
2.5	Various observations and transformations	15
3	Very short PCPs with very few queries	19
4	Applications to coding problems	25
4.1	Locally Testable Codes	25
4.2	Relaxed Locally Decodable codes	27
4.3	Linearity of the codes	31
II	The main construct: A short, robust PCP of proximity	32
5	Overview of our main construct	32
6	A randomness-efficient PCP	35
6.1	Well-structured Boolean circuits	36
6.2	Arithmetization	40
6.3	The PCP verifier	43
6.4	Analysis of the PCP verifier	46
7	A randomness-efficient PCP of proximity	48
8	A randomness-efficient robust PCP of proximity	53
8.1	Robustness of individual tests	54
8.2	Bundling	58
8.3	Robustness over the binary alphabet	65
8.4	Linearity of encoding	66
	Bibliography	71
III	Appendices	75
A	Hadamard-code-based PCP of proximity	75
B	Randomness-efficient low-degree tests and the sampling lemma	77

1 Introduction

Probabilistically Checkable Proofs [22, 5, 4] (a.k.a. Holographic Proofs [6]) are NP witnesses that allow efficient probabilistic verification based on probing few bits of the NP witness. The celebrated PCP Theorem [5, 4] asserts that probing a constant number of bits suffices, and it turned out that three bits suffice for rejecting false assertions with probability almost 1/2 (cf. [32, 30]).

Optimizing the query complexity of PCPs has attracted a lot of attention, motivated in part by the significance of query complexity for non-approximability results (see, for example, [10, 9, 32, 30, 48]). However, these works only guarantee that the new NP witness (i.e., the PCP) is of length that is upper-bounded by a polynomial in the length of the original NP witness.¹ Optimizing the length of the new NP witness was the focus of [6, 44, 31, 28, 12], and in this work we continue the latter research direction.

In our view, the significance of PCPs extends far beyond their applicability to deriving non-approximability results. The mere fact that NP-witnesses can be transformed into a format that supports super-fast probabilistic verification is remarkable. From this perspective, the question of how much redundancy is introduced by such a transformation is a fundamental one. Furthermore, PCPs have been used not only to derive non-approximability results but also for obtaining positive results (e.g., CS-proofs [37, 41] and their applications [7, 16]), and the length of the PCP affects the complexity of those applications.

In any case, the length of PCPs does also have relevance to non-approximability results; specifically, it affects their *tightness with respect to the running time*. For example, suppose (exact) SAT has complexity $2^{\Omega(n)}$. The original PCP theorem [5, 4] only implies that approximating MaxSAT requires time 2^{n^α} , for some (small) $\alpha > 0$. The work of Polishchuk and Spielman [44] makes α arbitrarily close to 1, whereas the results of [28, 12] further improve the lower-bound to $2^{n^{1-o(1)}}$. Our results reduce the $o(1)$ term.²

1.1 PCPs with better length vs query trade-off

How short can a PCP be? The answer may depend on the number of bits we are willing to read in order to reject false assertions (say) with probability at least 1/2. It is implicit in the work of Polishchuk and Spielman [44] that, for proofs of satisfiability of circuits of size n , if we are willing to read $n^{0.01}$ bits then the length of the new NP witness may be $\tilde{O}(n)$. That is, stretching the NP witness by only a poly-logarithmic amount, allows to dramatically reduce the number of bits read (from n to $n^{0.01}$). More precisely:³

Theorem 1.1 (implicit in [44]) *Satisfiability of circuits of size n can be probabilistically verified by probing an NP witness of length $\text{poly}(\log n) \cdot n$ in $n^{o(1)}$ bit locations. In fact, for any value of a parameter $m \leq \log n$, there is a PCP having randomness complexity $(1 - m^{-1}) \cdot \log_2 n + O(\log \log n) + O(m \log m)$ and query complexity $\text{poly}(\log n) \cdot n^{1/m}$.*

Recall that the proof length of a PCP is at most $2^r \cdot q$, where r is the randomness complexity and q is the query complexity of the PCP. Thus, the first part of the above theorem follows by setting $m = \log \log n / \log \log \log n$ in the second part.

¹We stress that in all the above works as well as in the current work, the new NP witness can be computed in polynomial-time from the original NP witness.

²A caveat: it is currently not known whether these improved lower-bounds can be achieved simultaneously with optimal approximation ratios, but the hope is that this can eventually be done.

³All logarithms in this work are to base 2, but in some places we choose to emphasize this fact by using the notation $\log_2$ rather than $\log$.

Our results show that the query complexity can be reduced dramatically if we are willing to increase the length of the proof slightly. First, with a quasi-polylogarithmic stretch, the query complexity can be made double-logarithmic:

Theorem 1.2 *Satisfiability of circuits of size n can be probabilistically verified by probing an NP witness of length $\exp(\tilde{O}(\log \log n)^2) \cdot n$ in $o(\log \log n)$ bit-locations. In fact, it has a PCP having randomness complexity $\log_2 n + (\log \log n)^2 \cdot \text{poly}(\log \log \log n)$ and query complexity $O(\log \log n / \log \log \log n)$.*

Prior to our work, reducing the query complexity below $\exp(\sqrt{\log n})$ required stretching the NP witness by at least a $\exp(\sqrt{\log n})$ factor. With approximately such a stretch factor, previous works actually achieved constant query complexity (cf. [28, 12]). Thus, Theorem 1.2 represents a vast improvement in the query complexity of PCPs that use very short proofs (i.e., in the range between $\exp(\tilde{O}(\log \log n)^2) \cdot n$ and $\exp(\sqrt{\log n}) \cdot n$). On the other hand, considering NP witnesses that allow probabilistic verification by a *constant* number of queries, we reduce the best known stretch factor from $\exp(\log^{0.5+\varepsilon} n)$ (established in [28, 12]) to $\exp(\log^\varepsilon n)$, for any $\varepsilon > 0$. That is:

Theorem 1.3 *For every constant $\varepsilon > 0$, satisfiability of circuits of size n can be probabilistically verified by probing an NP witness of length $\exp(\log^\varepsilon n) \cdot n$ in a constant number of bit-locations. In fact, it has a PCP having randomness complexity $\log_2 n + \log^\varepsilon n$ and query complexity $O(1/\varepsilon)$.*

It may indeed be the case that the trade-off (between length blow-up factors and query complexity) offered by Theorems 1.1–1.3 merely reflects our (incomplete) state of knowledge. In particular, we wonder whether circuit satisfiability can be probabilistically verified by a PCP having proof-length $n \cdot \text{poly}(\log n)$ and constant query complexity.

1.2 New notions and main techniques

A natural approach to reducing the query complexity in Theorem 1.1 is via the “proof composition” paradigm of [5]. However, that PCP, as constructed in [44], does not seem amenable to composition (when the parameter m is non-constant). Thus, we begin by giving a new PCP construction whose parameters match those in Theorem 1.1, but is suitable for composition. As we will see, we cannot afford the standard proof composition techniques, and thus also introduce a new, more efficient composition paradigm.

The initial PCP. Our new proof of Theorem 1.1 modifies the constructions of Polishchuk and Spielman [44] and Harsha and Sudan [31]. The latter construction was already improved in [28, 12] to reduce the length of PCPs to $n \cdot 2^{\tilde{O}(\sqrt{\log n})}$. Our results go further by re-examining the “low-degree test” (query-efficient tests that verify if a given function is close to being a low-degree polynomial) and first observing that the small-bias sample sets of [12] give an even more significant savings on the randomness complexity of low-degree tests than noticed in their work. However, exploiting this advantage takes a significant effort in modifying known PCP modules, and redefining the ingredients in “proof composition”.

For starters, PCP constructions tend to use many (i.e., a super-constant number of) functions and need to test if each is a low-degree polynomial. In prior results, this was performed efficiently by combining the many different functions on, say m variables, into a single new one on $m + 1$ variables, where the extra variable provides an index into the many different old functions. Testing if the new function is of low-degree, implicitly tests all the old functions. Such tricks, which involve introducing a few extra variables, turn out to be too expensive in our context. Furthermore, for similar reasons, we can not use other “parallelization” techniques [23, 38, 4, 27, 45], which were

instrumental to the proof composition technique of [5]. In turn, this forces us to introduce a new variant of the proof composition method, which is much more flexible than the one of [5]. Going back to the PCP derived in Theorem 1.1, we adapt it for our new composition method by introducing a “bundling” technique that offers a randomness efficient alternative to parallelization.

Our new “proof composition” method refers to two new notions: the notion of a *PCP of proximity* and the notion of a *robust PCP*. Our method is related to the method discovered independently by Dinur and Reingold [19]. (There are significant differences between the two methods; as explained in Section 1.3.)

PCPs of Proximity. Recall that a standard PCP is given an explicit input (which is supposedly in some NP language) as well as access to an oracle that is supposed to encode a “probabilistically verifiable” NP witness. The PCP verifier uses oracle queries (which are counted) in order to probabilistically verify whether the input, which is explicitly given to it, is in the language. In contrast, a PCP of proximity is given access to two oracles, one representing an input (supposedly in the language) and the other being a redundant encoding of an NP-witness (as in a PCP). Indeed, the verifier may query both the input oracle and the proof oracle, but *its queries to the input oracle are also counted in its query complexity*. As usual we focus on verifiers having very low query complexity, certainly smaller than the length of the input. Needless to say, such a constrained verifier cannot hope to distinguish inputs in the language from inputs out of the language, but it is not required to do so. A verifier for a PCP of proximity is only required to accept inputs that are in the language and reject inputs that are *far* from the language (i.e., far in Hamming distance from any input in the language). (PCPs of proximity are related to holographic proofs [6] and to “PCP spot-checkers” [20]; see further discussion in Section 1.3.)

Robust PCPs. To discuss robust PCPs, let us recall the soundness guarantee of standard (non-adaptive) PCPs. The corresponding verifier can be thought of as determining, based on its coin tosses, a sequence of oracle positions and a predicate such that evaluating this predicate on the indicated oracle bits always accepts if the input is in the language and rejects with high probability otherwise. That is, in the latter case, we require that the assignment of oracle bits to the predicate does satisfy the predicate. In a **robust PCP** we strengthen the latter requirement. We require that the said assignment (of oracle bits) not only fails to satisfy the predicate but rather is *far* from any assignment that does satisfy the predicate.

Proof Composition. The key observation is that “proof composition” works very smoothly when we compose an outer “robust PCP” with an inner “PCP of proximity”. We need neither worry about how many queries the outer “robust PCP” makes nor care about what coding the inner “PCP of proximity” uses in its proof oracle (much less apply the same encoding to the outer answers). All that we should make sure is that the lengths of the objects match and that the distance parameter in the robustness condition (of the outer verifier) is at least as big as the distance parameter in the proximity condition (of the inner verifier).

Indeed, Theorems 1.2 and 1.3 are proved by first extending Theorem 1.1 to provide a robust PCP of proximity of similar complexities, and then applying the new “proof composition” method. We stress that our contribution is in providing a proof of Theorem 1.1 that lends itself to a modification that satisfies the robustness property, and in establishing the latter property. In particular, the aforementioned “bundling” is applied in order to establish the robustness property. Some care is also due when deriving Theorem 1.2 using a non-constant number of “proof compositions”. In particular, Theorem 1.2 (resp., Theorem 1.3) is derived in a way that guarantees that the query complexity is linear rather than exponential in the number of “proof compositions”, where the latter is $o(\log \log n)$ (resp., $1/\varepsilon$).

We stress that the flexibility in composing robust PCPs of proximity plays an important role in our ability to derive quantitatively stronger results regarding PCPs. We believe that robust PCPs of proximity may play a similar role in other quantitative studies of PCPs. We note that the standard PCP Theorem of [5, 4] can be easily derived using a much weaker and simpler variant of our basic robust PCP of proximity, and the said construction seems easier than the basic PCPs used in the proof composition of [5, 4].

In addition to their role in our “proof composition” method, PCPs of proximity provide also a good starting point for deriving improved locally testable codes (see discussion in Section 1.4). The relation of PCPs of proximity to “property testing” is further discussed in Section 1.5.

1.3 Related work

As mentioned above, the notion of a PCP of proximity is related to notions that have appeared in the literature. Firstly, the notion of a PCP of proximity generalizes the notion of holographic proofs set forward by Babai, Fortnow, Levin, and Szegedy [6]. In both cases, the verifier is given oracle access to the input, and we count its probes to the input in its query complexity. The key issue is that holographic proofs refer to inputs that are presented in an error-correcting format (e.g., one aims to verify that a graph that is *represented by an error-correcting encoding of its adjacency matrix* (or incidence list) is 3-colorable). In contrast, a PCP of proximity refers to inputs that are presented in any format but makes assertions only about their proximity to acceptable inputs (e.g., one is interested in whether a graph, represented by its adjacency matrix (or incidence list), *is 3-colorable or is far from being 3-colorable*).

PCP of proximity are implicit in the low-degree testers that utilize auxiliary oracles (e.g., an oracle that provides the polynomial representing the value of the function restricted to a queried line); cf. [5, 4].

PCPs of proximity are a special case of the “PCP spot-checkers” defined by Ergün, Kumar and Rubinfeld [20]. On the other hand, PCPs of proximity extend “property testing” [47, 25] by providing the tester with oracle access to a proof (on top of the ordinary input-oracle to which it has access). Thus, the relation of PCPs of proximity to property testing is analogous to the relation of NP to BPP (or RP). Put differently, while property testing provides a notion of approximation for decision procedures, PCP of proximity provides a notion of approximation for (probabilistic) verification procedures. In both cases, approximation means that inputs in the language should be accepted (when accompanied with suitable proofs) while inputs that are far from the language should be rejected (no matter what false proof is provided).

As stated above, our “proof composition” method is related to the method discovered independently by Dinur and Reingold [19]. Both methods use the same notion of PCPs of proximity, but while our method refers to the new notion of robustness (i.e., to the robustness of the outer verifier) the method of Dinur and Reingold refers to the number of (non-Boolean) queries (made by the outer verifier). Indeed, the method of Dinur and Reingold uses a (new) parallelization procedure (which reduces the number of queries by a constant factor), whereas we avoid parallelization altogether (but rather use a related “bundling” of queries into a non-constant number of “bundles” such that robustness is satisfied at the bundle-level).⁴ We stress that we cannot afford the cost

⁴The main part of the bundling technique takes place at the level of analysis, without modifying the proof system at all. Specifically, we show that the answers read by the verifier can be partitioned into a non-constant number of (a-priori fixed) “bundles” so that on any no instance, with high probability a constant fraction of the bundles read should be modified to make the verifier accept. We stress that the fact that certain sets of queries (namely those in each bundle) are always made together is a feature that our particular proof system happens to have (or rather it was somewhat massaged to have). Once “robust soundness” is established at the “bundle level,” we may just modify the

of any known parallelization procedure, because at the very least these procedures increase the length of the proof by a factor related to the answer length, which is far too large in the context of Theorem 1.1 (which in turn serves as the starting point for all the other results in this work). We comment that the parallelization procedure of [19] is combinatorial (albeit inapplicable in our context), whereas our “bundling” relies on the algebraic structure of our proof system.

1.4 Applications to coding problems

The flexibility of PCPs of proximity makes them relatively easy to use towards obtaining results regarding locally testable and decodable codes. In particular, using a suitable PCP of proximity, we obtain an improvement in the rate of locally testable codes (improving over the results of [28, 12]). Loosely speaking, a codeword test (for a code C) is a randomized oracle machine that is given oracle access to a string. The tester may query the oracle at a constant number of bit-locations and is required to (always) accept every codeword and reject with (relatively) high probability every string that is “far” from the code. The locally testable codes of [28, 12] used codewords of length $\exp(\log^{0.5+\varepsilon} k) \cdot k$ in order to encode k bits of information, for any constant $\varepsilon > 0$. Here we reduce the length of the codewords to $\exp(\log^\varepsilon k) \cdot k$. That is:

Theorem 1.4 (loosely stated, see Section 4.1 for details): *For every constant $\varepsilon > 0$, there exists locally testable codes that use codewords of length $\exp(\log^\varepsilon k) \cdot k$ in order to encode k bits of information.*

We also introduce a relaxed notion of locally decodable codes, and show how to construct such codes using any PCP of proximity (and ours in particular). Loosely speaking, a code is said to be locally decodable if whenever relatively few location are corrupted, the decoder is able to recover each information-bit, with high probability, based on a constant number of queries to the (corrupted) codeword. This notion was formally defined by Katz and Trevisan [34] and the best known locally decodable code has codeword of length that is sub-exponential in the number of information bits. We relax the definition of locally decodable codes by requiring that, whenever few location are corrupted, the decoder should be able to recover most of the individual information-bits (based on few queries) and for the rest of the locations, the decoder may output a fail symbol (but not the wrong value). That is, the decoder must still avoid errors (with high probability), but is allowed to say “don’t know” on a few bit-locations. We show that this relaxed notion of local decodability can be supported by codes that have codewords of length that is almost-linear in the number of information bits. That is:

Theorem 1.5 (loosely stated, see Section 4.2 for details): *For every $\varepsilon > 0$, there exists relaxed locally decodable codes that use codewords of length $k^{1+\varepsilon}$ in order to encode k bits of information.*

1.5 Relation to Property Testing

Following Ergün *et al.* [20], we view PCPs of proximity as an extension of property testing [47, 25]. Loosely speaking, a **property tester** is given oracle access to an input and is required to distinguish the case in which the input has the property from the case in which it is far (say in Hamming distance) from any input having the property. Typically, the interest is in testers that query their

proof system so that the bundles become queries and the answers are placed in (any) good error-correcting format, which implies robustness at the bit level.

input on few bit-locations (or at the very least on a sub-linear number of such locations). In a PCP of proximity such a tester (now called a verifier) is also given oracle access to an alleged proof.

We comment that PCPs of proximity are provably stronger than property testers; that is, there are (natural) separations between property testers and PCPs of proximity (which may be viewed as the “approximation” versions of BPP and NP). For further discussions, refer to Section 2.2

1.6 Organization

Theorems 1.2 and 1.3, which are the work’s main results, are proved by constructing and using a Robust PCP of Proximity that achieves a very good trade-off between randomness and query complexity. Thus, this Robust PCP of Proximity is the main building block that underlies our work. Unfortunately, the construction of a very efficient Robust PCP of Proximity is quite involved, and is thus deferred to the second part of this work (which starts with an overview). In the first part of this work we show how the aforementioned Robust PCP of Proximity can be used to derive all the results mentioned in the Introduction (and, in particular, Theorems 1.2 and 1.3). Thus, the overall structure of this work is as follows:

Part I: Using the main building block. We start by providing a definitional treatment of PCPs of proximity and robust PCPs. The basic definitions as well as some observations and useful transformations are presented in Section 2. Most importantly, we analyze the natural composition of an outer robust PCP with an inner PCP of proximity.

In Section 3, we state the properties of our main building block (i.e., a highly efficient Robust PCP of proximity), and show how to derive Theorems 1.2 and 1.3, by composing this Robust PCP of proximity with itself multiple times. Specifically, $o(\log \log n)$ compositions are used to derive Theorem 1.2 and $1/\varepsilon$ compositions are used to derive Theorem 1.3. The coding applications stated in Theorems 1.4 and 1.5 are presented in Section 4.

Part II: Constructing the main building block. We start this part by providing an overview of the construction. This overview (i.e., Section 5) can be read before reading Part I, provided that the reader is comfortable with the notion of a Robust PCP of proximity.

The construction itself is presented in Sections 6–8. We start by presenting a (highly efficient) ordinary PCP (establishing Theorem 1.1), which lends itself to the subsequent modifications. In Section 7, we augment this PCP with a test of proximity, deriving an analogous PCP of proximity. In Section 8 we present a robust version of the PCP of proximity derived in the previous sections.

Part III: Appendices. The construction presented in Section 3 also uses a PCP of proximity of polynomial randomness complexity and constant query complexity. Such a PCP of proximity can be derived by a simple augmentation of the Hadamard-based PCP of [4], which we present in Appendix A.

In Appendix B, we recall results regarding random-efficient low-degree tests and a related sampling lemma, which are used in Part II.

Part I

All but the main construct

2 PCPs and variants: definitions, observations and transformations

Notation: Except when otherwise noted, all *circuits* in this paper have fan-in 2 and fan-out 2, and we allow arbitrary unary and binary Boolean operations as internal gates. The *size* of a circuit is the number of gates. We will refer to the following languages associated with circuits: the **P**-complete language CIRCUIT VALUE, defined as $\text{CKTVAL} = \{(C, w) : C(w) = 1\}$; the **NP**-complete CIRCUIT SATISFIABILITY, defined as $\text{CKTSAT} = \{C : \exists w C(w) = 1\}$; and the also **NP**-complete NONDETERMINISTIC CIRCUIT VALUE, defined as $\text{NCKTVAL} = \{C : (C, w) : \exists z C(w, z) = 1\}$. (In the latter, we assume that the partition of the variables of C into w -variables and z -variables is explicit in the encoding of C .)

We will extensively refer to the *relative* distance between strings/sequences over some alphabet Σ : For $u, v \in \Sigma^\ell$, we denote by $\Delta(u, v)$ the fraction of locations on which u and v differ (i.e., $\Delta(u, v) \triangleq |\{i : u_i \neq v_i\}|/\ell$, where $u = u_1 \cdots u_\ell \in \Sigma^\ell$ and $v = v_1 \cdots v_\ell \in \Sigma^\ell$). We say that u is δ -close to v (resp., δ -far from v) if $\Delta(u, v) \leq \delta$ (resp., $\Delta(u, v) > \delta$). The relative distance of a string to a set of strings is defined in the natural manner; that is, $\Delta(u, S) \triangleq \min_{v \in S} \{\Delta(u, v)\}$. Occasionally, we will refer to the absolute Hamming distance, which we will denote by $\bar{\Delta}(u, v) \triangleq |\{i : u_i \neq v_i\}|$.

Organization of this section: After recalling the standard definition of PCP (in Section 2.1), we present the definitions of PCPs of Proximity and Robust PCPs (in Sections 2.2 and 2.3, respectively). We then turn to discuss (in Section 2.4) the composition of a Robust PCP with a PCP of Proximity. Various observations and transformations regarding the new notions are presented in Section 2.5.

2.1 Standard PCPs

We begin by recalling the formalism of a PCP verifier. Throughout this work, we restrict our attention to *nonadaptive* verifiers, both for simplicity and because one of our variants (namely robust PCPs) only makes sense for nonadaptive verifiers.

Definition 2.1 (PCP verifiers)

- A verifier is a probabilistic polynomial-time algorithm V that, on an input x of length n , tosses $r = r(n)$ random coins R and generates a sequence of $q = q(n)$ queries $I = (i_1, \dots, i_q)$ and a circuit $D : \{0, 1\}^q \rightarrow \{0, 1\}$ of size at most $d(n)$.

We think of V as representing a probabilistic oracle machine that queries its oracle π for the positions in I , receives the q answer bits $\pi|_I \triangleq (\pi_{i_1}, \dots, \pi_{i_q})$, and accepts iff $D(\pi|_I) = 1$.

- We write $(I, D) \stackrel{R}{\leftarrow} V(x)$ to denote the queries and circuit generated by V on input x and random coin tosses, and $(I, D) = V(x; R)$ if we wish to specify the coin tosses R .
- We call r the randomness complexity, q the query complexity, and d the decision complexity of V .

For simplicity in these definitions, we treat the parameters r , q , and d above (and other parameters below) as functions of only the input length n . However, at times we may also allow them to depend on other parameters, which should be understood as being given to the verifier together with the input. We now present the standard notion of PCPs, restricted to perfect completeness for simplicity.

Definition 2.2 (standard PCPs) *For a function $s : \mathbb{Z}^+ \rightarrow [0, 1]$, a verifier V is a probabilistically checkable proof system for a language L with soundness error s if the following two conditions hold for every string x :*

Completeness: *If $x \in L$ then there exists π such that $V(x)$ accepts oracle π with probability 1. Formally,*

$$\exists \pi \quad \Pr_{(I,D) \stackrel{R}{\leftarrow} V(x)} [D(\pi|_I) = 1] = 1.$$

Soundness: *If $x \notin L$ then for every oracle π , the verifier $V(x)$ accepts π with probability strictly less than s . Formally,*

$$\forall \pi \quad \Pr_{(I,D) \stackrel{R}{\leftarrow} V(x)} [D(\pi|_I) = 1] < s(|x|).$$

If s is not specified, then it is assumed to be a constant in $(0, 1)$.

Our main goal in this work is to construct *short* PCPs that use *very few queries*. Recalling that the length of a (nonadaptive) PCP is upper-bounded by $2^{r(n)} \cdot q(n)$, we focus on optimizing the (trade-off between) randomness and query complexities.

We will focus on constructing PCPs for the **NP**-complete problem CIRCUIT SATISFIABILITY, defined as $\text{CKTSAT} = \{C : \exists w \ C(w) = 1\}$. Recall that every language in **NTIME** $(t(n))$ reduces to CKTSAT in time $O(t(n) \log t(n))$ (cf. [33, 43, 17]), and so a nearly linear-sized PCP for CKTSAT implies PCPs for **NTIME** $(t(n))$ of size nearly linear in $t(n)$ for every polynomial $t(n)$.

2.2 PCPs of Proximity

We now present a relaxation of PCPs that only verify that the input is *close* to an element of the language. The advantage of this relaxation is that it allows the possibility that the verifier may read only a small number of bits from the input. Actually, for greater generality, we will divide the input into two parts (x, y) , giving the verifier x explicitly and y as an oracle, and we only count the verifier's queries to the latter. Thus we consider languages consisting of pairs of strings, which we refer to as a **pair language**. One pair language to keep in mind is the CIRCUIT VALUE problem: $\text{CKTVAL} = \{(C, w) : C(w) = 1\}$. For a pair language L , we define $L(x) = \{y : (x, y) \in L\}$. For example, $\text{CKTVAL}(C)$ is the set of satisfying assignments to C . It will be useful below to treat the two oracles to which the verifier has access as a single oracle, thus for oracles π^0 and π^1 , we define the concatenated oracle $\pi = \pi^0 \circ \pi^1$ as $\pi_{b,i} = \pi_i^b$.

Definition 2.3 (PCPs of proximity (PCPPs)) *For functions $s, \delta : \mathbb{Z}^+ \rightarrow [0, 1]$, a verifier V is a probabilistically checkable proof of proximity (PCPP) system for a pair language L with proximity parameter δ and soundness error s if the following two conditions hold for every pair of strings (x, y) :*

Completeness: *If $(x, y) \in L$, then there exists π such that $V(x)$ accepts oracle $y \circ \pi$ with probability 1. Formally,*

$$\exists \pi \quad \Pr_{(I,D) \stackrel{R}{\leftarrow} V(x)} [D((y \circ \pi)|_I) = 1] = 1.$$

Soundness: If y is $\delta(|x|)$ -far from $L(x)$, then for every π , the verifier $V(x)$ accepts oracle $y \circ \pi$ with probability strictly less than $s(|x|)$. Formally,

$$\forall \pi \quad \Pr_{(I,D) \stackrel{R}{\leftarrow} V(x)} [D((y \circ \pi)|_I) = 1] < s(|x|).$$

If s and δ are not specified, then both are assumed to be constants in $(0, 1)$.

Note that the parameters (soundness, randomness, etc.) of a PCPP are measured as a function of the length of x , the explicit portion of the input.

In comparing PCPPs and PCPs, one should note two differences that have conflicting effects. On one hand, the soundness criterion of PCPPs is a relaxation of the soundness of PCPs. Whereas, a PCP is required to reject (with high probability) every input that is not in the language, a PCPP is only required to reject input pairs (x, y) in which the second element (i.e., y) is far from being suitable for the first element (i.e., y is far from $L(x)$). That is, in a PCPP, nothing is required in the case that y is close to $L(x)$ and yet $y \notin L(x)$. On the other hand, the query complexity of a PCPP is measured more stringently, as it accounts also for the queries to the input-part y (on top of the standard queries to the proof π). This should be contrasted with a standard PCP that has free access to all its input, and is only charged for access to an auxiliary proof. To summarize, PCPPs are required to do less (i.e., their performance requirements are more relaxed), but they are charged for more things (i.e., their complexity is evaluated more stringently). Although it may not be a priori clear, the stringent complexity requirement prevails. That is, PCPPs tend to be more difficult to construct than PCPs of the same parameters. For example, while CIRCUIT VALUE has a trivial PCP (since it is in **P**), a PCPP for it implies a PCP for CIRCUIT SATISFIABILITY:

Proposition 2.4 *If CIRCUIT VALUE has a PCPP, then CIRCUIT SATISFIABILITY has a PCP with identical parameters (randomness, query complexity, decision complexity, and soundness).*

An analogous statement holds for any pair language L and the corresponding projection on first element $L_1 \triangleq \{x : \exists y \text{ s.t. } (x, y) \in L\}$; that is, if L has a PCPP then L_1 has a PCP with identical parameters.

Proof: A PCP π that C is satisfiable can be taken to be $w \circ \pi'$, where w is a satisfying assignment to C and π' is a PCPP that $(C, w) \in \text{CKTVAL}$. This proof π can be verified using the PCPP verifier. The key observation is that if $C \notin \text{CIRCUIT SATISFIABILITY}$ then there exists no w that is 1-close to $\text{CIRCUIT VALUE}(C)$, because the latter set is empty. ■

Note that we only obtain a standard PCP for CIRCUIT SATISFIABILITY, rather than a PCP of proximity. Indeed, CIRCUIT SATISFIABILITY is not a pair language, so it does not even fit syntactically into the definition of a PCPP. However, we can give a PCPP for the closely related (and also **NP**-complete) pair language NONDETERMINISTIC CIRCUIT VALUE. Recall that is the language $\text{NCKTVAL} = \{(C, w) : \exists z C(w, z) = 1\}$ (where the variables of C are explicitly partitioned into w -variables and z -variables).

Proposition 2.5 *If CIRCUIT VALUE has a PCPP with proximity parameter $\delta(n)$, soundness $s(n)$, randomness $r(n)$, query complexity $q(n)$, and decision complexity $d(n)$, then NONDETERMINISTIC CIRCUIT VALUE has a PCPP with proximity parameter $2\delta(4n)$, soundness $s(4n)$, randomness $r(4n)$, query complexity $q(4n)$, and decision complexity $d(4n)$.*

Proof: Given a circuit $C(\cdot, \cdot)$ of size n whose variables are partitioned into one group of size k and another of size ℓ , we transform it into a new circuit $C'(\cdot, \cdot)$ of size $n' = 4n$ in which the first group has size $k' \geq \ell$ and the second group has size ℓ . Specifically, we set $t = \lceil \ell/k \rceil$ and $k' = t \cdot k$, and define $C'(x', y)$ to be a circuit that checks whether $x' = x^t$ for some x such that $C(x, y) = 1$. It can be verified that this can be done in size $n + 3tk \leq 4n$ (over the full binary basis). In addition, if w is δ -far from being extendable to a satisfying assignment of C , then w^t is δ -far from being extendable to a satisfying assignment of C' .

Now, the NCKTVAL-verifier, on explicit input C and input oracle $w \in \{0, 1\}^k$, will construct C' as above and expect a proof oracle of the form $z \circ \pi$, where $z \in \{0, 1\}^m$ and π is a PCPP that $w^t \circ z$ satisfies C' as constructed above. That is, the NCKTVAL-verifier will simulate the CKTVAL-verifier on explicit input C' , input oracle $w^t \circ z$ (which can easily be simulated given oracle access to w and z), and proof oracle π . Completeness can be verified by inspection. For soundness, suppose that w is 2δ -far from being extendable to a satisfying assignment of C . Then w^t is 2δ -far from being extendable to a satisfying assignment of C' , which implies that, for any z , $w^t \circ z$ is δ -far from satisfying C' . Thus, by the soundness of the CKTVAL-verifier, the acceptance probability is at most $s(n') = s(4n)$, for any proof oracle π . ■

Relation to property testing: Actually, the requirements from a PCPP for a pair language L refer only to its performance on the (“gap”) promise problem $\Pi = (\Pi_Y, \Pi_N)$, where $\Pi_Y = L$ and $\Pi_N = \{(x, y) : y \text{ is } \delta\text{-far from } L(x)\}$. That is, this PCPP is only required to (always) accept inputs in Π_Y and reject (with high probability) inputs in Π_N (whereas nothing is required with respect to inputs not in $\Pi_Y \cup \Pi_N$). Such a gap problem corresponds to the notion of approximation in *property testing* [47, 25].⁵ Indeed, property testers are equivalent to PCPP verifiers that have no access to an auxiliary proof π . Thus the relation between property testing and PCPPs is analogous to the relation between **BPP** and **NP** (or **MA**). For example, the problem of testing Bipartiteness can be cast by the pair language $L = \{(n, G) : \text{the } n\text{-vertex graph } G \text{ is bipartite}\}$, where the first (i.e., explicit) input is only used to specify the length of the second (i.e., non-explicit) input G , to which the tester has oracle access (measured in its query complexity). We comment that the formulation of pair languages allows to capture more general property testing problems where more information about the property (to be tested) itself is specified as part of the input (e.g., by a circuit, as in CKTVAL).

In both property testers and PCPs of proximity, the interest is in testers/verifiers that query their input (and proof oracle) in only a small (preferably constant, and certainly sublinear) number of bit-locations. It turns out that PCPPs are provably stronger than property testers; that is, there are (natural) separations between property testers and PCPs of proximity. (Some of the following examples were pointed out in [20].) In the adjacency matrix model (cf. [25]), Bipartiteness has a PCP of proximity in which the verifier makes only $O(1/\delta)$ queries and rejects any graph that is δ -far from being bipartite with probability at least $2/3$. (The proof-oracle consists of an assignment of vertices to the two parts, and the verifier queries the assignment of the end-points of $O(1/\delta)$ random edges. This construction also generalizes to k -colorability, and in fact any generalized graph partition property (cf. [25]) with an efficient one-sided tester.) In contrast, Bogdanov and Trevisan [15] showed that any tester for Bipartiteness that rejects graphs that are δ -far from being bipartite must make $\Omega(\delta^{-3/2})$ queries. More drastic separations are known in the incidence-lists (bounded-degree) model (of [26]): testing Bipartiteness (resp., 3-colorability) of n -vertex graphs has query complexity $\Omega(\sqrt{n})$ [26] (resp., $\Omega(n)$ [14]), but again a PCP of proximity will only use

⁵This notion of approximation (of decision problems) should not be confused with the approximation of (search) optimization problems, which is also closely related to PCPs [22, 4].

$O(1/\delta)$ queries.

Another example comes from the domain of codes. For any good code (or “even” any code of linear distance), there exists a PCP of proximity for the property of being a codeword that makes a constant number of queries.⁶ This stands in contrast to the linear lower-bound on the query-complexity of codeword testing for some (good) linear codes, proved by Ben-Sasson *et al.* [11].

Needless to say, there may be interesting cases in which PCPs of proximity do not out-perform property testers.

Queries vs. proximity: Intuitively, the query complexity of a PCPP should depend on the proximity parameter δ . Proposition 2.8 (in Section 2.5) confirms this intuition.

The relation of PCPP to other works: As discussed in the introduction (see Section 1.3), notions related to (and equivalent to) PCPPs have appeared in the literature before [6, 20]. In particular, holographic proofs are a special case of PCPPs (which refer to pair languages $L = \{(n, \mathcal{C}(x)) : x \in L' \cap \{0, 1\}^n\}$, where $\mathcal{C}$ is an error-correcting code and $L' \in \mathbf{NP}$), whereas PCPPs are a special case of “PCP spot-checkers” (when viewing decision problems as a special case of search problems). In addition, PCPPs play an important role also in the work of Dinur and Reingold [19]; again, see Section 1.3. Recall that both our use and their use of PCPPs is for facilitating “proof composition” (of PCP-type constructs). Finally, existing PCP constructions (such as [4]) can be modified to yield PCPPs.

2.3 Robust Soundness

In this section, we present a *strengthening* of the standard PCP soundness condition. Instead of asking that the bits that the verifier reads from the oracle are merely rejected with high probability, we ask that the bits that the verifier reads are *far* from being accepted with high probability. The main motivation for this notion is that, in conjunction with PCPPs, it allows for a very simple composition without the usual costs of “parallelization”.

Definition 2.6 (robust soundness) *For functions $s, \rho : \mathbb{Z}^+ \rightarrow [0, 1]$, a PCP verifier V for a language L has robust-soundness error s with robustness parameter ρ if the following holds for every $x \notin L$: For every oracle π , the bits read by the verifier V are ρ -close to being accepted with probability strictly less than s . Formally,*

$$\forall \pi \quad \Pr_{(I, D) \stackrel{R}{\leftarrow} V(x)} [\exists a \text{ s.t. } D(a) = 1 \text{ and } \Delta(a, \pi|_I) \leq \rho] < s(|x|).$$

If s and ρ are not specified, then they are assumed to be constants in $(0, 1)$. PCPPs with robust-soundness are defined analogously, with the $\pi|_I$ being replaced by $(y \circ \pi)|_I$.

Note that for PCPs with query complexity q , robust-soundness with any robustness parameter $\rho < 1/q$ is equivalent to standard PCP soundness. However, there can be robust PCPs with large query complexity (e.g. $q = n^{\Omega(1)}$) yet constant robustness, and indeed such robust PCPs will be the main building block for our construction.

Various observations regarding robust PCPs are presented in Section 2.5. We briefly mention here the relation of robustness to parallelization; specifically, when applied to a robust PCP, the simple query-reduction technique of Fortnow *et al.* [23] performs less poorly than usual (i.e., the

⁶Indeed, this is a special case of our extension of the result of Babai *et al.* [6], discussed in Section 1.3. On the other hand, this result is simpler than the locally testable code mentioned in Section 1.4, because here the PCP of proximity is not part of the codeword.

resulting soundness is determined by the robustness parameter rather than by the number of queries).

2.4 Composition

As promised, a robust “outer” PCP composes very easily with an “inner” PCPPs. Loosely speaking, we can compose such schemes provided that the decision complexity of the outer verifier matches the input length of the inner verifier, and soundness holds provided that the robustness parameter of the outer verifier upper-bounds the proximity parameter of the inner verifier. Note that composition does not refer to the query complexity of the outer verifier, which is always upper-bounded by its decision complexity.

Theorem 2.7 (Composition Theorem) *Suppose that for functions $r_{\text{out}}, r_{\text{in}}, d_{\text{out}}, d_{\text{in}}, q_{\text{in}} : \mathbb{N} \rightarrow \mathbb{N}$, and $\varepsilon_{\text{out}}, \varepsilon_{\text{in}}, \rho_{\text{out}}, \delta_{\text{in}} : \mathbb{N} \rightarrow [0, 1]$, the following hold:*

- *Language L has a robust PCP verifier V_{out} with randomness complexity r_{out} , decision complexity d_{out} , robust-soundness error $1 - \varepsilon_{\text{out}}$, and robustness parameter ρ_{out} .*
- *CIRCUIT VALUE has a PCPP verifier V_{in} with randomness complexity r_{in} , query complexity q_{in} , decision complexity d_{in} , proximity parameter δ_{in} , and soundness error $1 - \varepsilon_{\text{in}}$.*
- *$\delta_{\text{in}}(d_{\text{out}}(n)) \leq \rho_{\text{out}}(n)$, for every n .*

Then, L has a (standard) PCP, denoted 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 $d_{\text{in}}(d_{\text{out}}(n))$, and*
- *soundness error $1 - \varepsilon_{\text{out}}(n) \cdot \varepsilon_{\text{in}}(d_{\text{out}}(n))$.*

Furthermore, the computation of V_{comp} (i.e. evaluating $(I, D) \leftarrow V_{\text{comp}}(x; R)$) can be performed by some universal algorithm with black-box access to V_{out} and V_{in} . On inputs of length n , this algorithm runs in time n^c for some universal constant c , with one call to V_{out} on an input of length n and one call to V_{in} on an input of length $d_{\text{out}}(n)$. In addition:

- *If (instead of being a PCP) the verifier V_{out} is a PCPP with proximity parameter $\delta_{\text{out}}(n)$ then V_{comp} is a PCPP with proximity parameter $\delta_{\text{out}}(n)$.*
- *If V_{in} has robust-soundness with robustness parameter $\rho_{\text{in}}(n)$, then V_{comp} has robust-soundness with robustness parameter $\rho_{\text{in}}(d_{\text{out}}(n))$.*

Proof: We will use the inner PCPP to verify that the oracle positions selected by the (robust) outer-verifier are close to being accepted by the outer-verifier’s decision circuit. Thus, the new proof will consist of a proof for the outer verifier as well as proofs for the inner verifier, where each of the latter corresponds to a possible setting of the outer verifier’s coin tosses (and is intended to prove that the bits that should have been read by the outer-verifier satisfy its decision circuit). We will index the positions of the new (combined) oracle by pairs such that (out, i) denotes the i ’th position in the part of the oracle that represents the outer-verifier’s proof oracle, and (R, j) denotes the j ’th position in the R ’th auxiliary block (which represents the R -th possible proof oracle (for

the inner verifier's), which in turn is associated with the outer-verifier's coins $R \in \{0, 1\}^{r_{\text{out}}}$. For notational convenience, we drop the input length n from the notation below; all parameters of V_{out} are with respect to length n and all parameters of V_{in} with respect to length $d_{\text{out}}(n)$. With these conventions, here is the description of the composed verifier, $V_{\text{comp}}(x)$:

1. Choose $R \xleftarrow{R} \{0, 1\}^{r_{\text{out}}}$.
2. Run $V_{\text{out}}(x; R)$ to obtain $I_{\text{out}} = (i_1, \dots, i_{q_{\text{out}}})$ and D_{out} .
3. Run $V_{\text{in}}(D_{\text{out}})$ (on random coin tosses) to obtain $I_{\text{in}} = ((b_1, j_1), \dots, (b_{q_{\text{in}}}, j_{q_{\text{in}}}))$ and D_{in} .
(Recall that V_{in} , as a PCPP verifier, expects two oracles, an input oracle and a proof oracle, and thus makes queries of the form (b, j) , where $b \in \{0, 1\}$ indicates which oracle it wishes to query.)
4. For each $\ell = 1, \dots, q_{\text{in}}$, determine the queries of the composed verifier:
 - (a) If $b_\ell = 0$, set $k_\ell = (\text{out}, i_{j_\ell})$; that is, V_{in} 's queries to its input oracle are directed to the corresponding locations in V_{out} 's proof oracle. Recall that the j -th bit in V_{in} 's input oracle is the j -th bit in the input to D_{out} , which in turn is the i_j -th bit in the proof oracle of V_{out} .
 - (b) If $b_\ell = 1$, set $k_\ell = (R, j_\ell)$; that is, V_{in} 's queries to its R 'th possible proof oracle are directed to the corresponding locations in the auxiliary proof. Recall that the j -th bit in the proof oracle that V_{in} is using to verify the claim referring to the outer-verifier coins R is the j -th bit in the R -th block of the auxiliary proof.
5. Output $I_{\text{comp}} = (k_1, \dots, k_{q_{\text{in}}})$ and D_{in} .

The claims about V_{comp} 's randomness, query, decision, and computational complexities can be verified by inspection. Thus we proceed to check completeness and soundness.

Suppose that $x \in L$. Then, by completeness of the outer verifier, there exists a proof π_{out} making V_{out} accept with probability 1. In other words, for every $R \in \{0, 1\}^{r_{\text{out}}}$, if we set $(I_{\text{out}}, D_{\text{out}}) = V_{\text{out}}(x; R)$, we have $D_{\text{out}}(\pi_{\text{out}}|_{I_{\text{out}}}) = 1$. By completeness of the inner verifier, there exists a proof π_R such that $V_{\text{in}}(D_{\text{out}})$ accepts the oracle $\pi_{\text{out}}|_{I_{\text{out}}} \circ \pi_R$ with probability 1. If we set $\pi(t, \cdot) = \pi_t(\cdot)$ for all $t \in \{\text{out}\} \cup \{0, 1\}^{r_{\text{out}}}$, then V_{comp} accepts π with probability 1.

Suppose that $x \notin L$, and let π be any oracle. Define oracles $\pi_t(\cdot) = \pi(t, \cdot)$. By the robust-soundness (of V_{out}), with probability greater than ε_{out} over the choices of $R \in \{0, 1\}^{r_{\text{out}}}$, if we set $(I_{\text{out}}, D_{\text{out}}) = V_{\text{out}}(x; R)$, then $\pi_{\text{out}}|_{I_{\text{out}}}$ is ρ_{out} -far from satisfying D_{out} . Fixing such an R , by the PCPP-soundness of V_{in} (and $\delta_{\text{in}} \leq \rho_{\text{out}}$), it holds that $V_{\text{in}}(D_{\text{out}})$ rejects the oracle $\pi_{\text{out}}|_{I_{\text{out}}} \circ \pi_R$ (or, actually, any proof oracle augmenting the input oracle $\pi_{\text{out}}|_{I_{\text{out}}}$) with probability greater than ε_{in} . Therefore, $V_{\text{comp}}(x)$ rejects oracle π with probability at least $\varepsilon_{\text{out}} \cdot \varepsilon_{\text{in}}$.

The additional items follow by similar arguments. If V_{out} is a PCPP verifier, then the input is of the form (x, y) , where y is given via oracle access. In this case, throughout the proof above we should replace references to the oracle π_{out} with the oracle $y \circ \pi_{\text{out}}$, and for soundness we should consider the case that y is δ_{out} -far from $L(x)$. If V_{in} has robust-soundness, then at the end of the soundness analysis, we note that not only is $\pi_{\text{out}}|_{I_{\text{out}}} \circ \pi_R$ rejected with probability greater than ε_{in} but rather it is ρ_{in} -far from being accepted by V_{in} (and hence also by V_{comp}). ■

The above theorem can serve as a substitute for the original Composition Theorem in the derivation of the original PCP Theorem [4]. Specifically, one simply needs to modify the (pre-composition) verifiers of [4] to both test proximity and have robust soundness. As we shall see

in the next section, robust soundness can be obtained automatically from “parallelized PCPs” (as already constructed in [4]). And the PCPs [4] can easily be made PCPs of proximity by augmenting them with appropriate “proximity tests”. Thus, all the technical work in Part II is not forced by the new notion of robust PCPPs, but rather is aimed at constructing ones which have *nearly linear length*.

2.5 Various observations and transformations

Most of this subsection refers to robust PCPs, but we start with an observation regarding PCPs of proximity.

Queries vs. proximity: Intuitively, the query complexity of a PCPP should depend on the proximity parameter δ . The following proposition confirms this intuition.

Proposition 2.8 (queries vs. proximity) *Suppose pair-language L has a PCPP with proximity parameter δ , soundness error $1 - \varepsilon$, and query complexity q . Suppose further that there exists $(x, y) \in L$ such that $|x| = n$ and $|y| = m$, such that if we let $z \in \{0, 1\}^m$ be a random string of relative Hamming distance $\delta' \triangleq \delta'(x)$ from y , we have*

$$\Pr_z[z \text{ is } \delta\text{-far from } L(x)] \geq \gamma \triangleq \gamma(x).$$

Then

$$q > \frac{\varepsilon \cdot \gamma}{\delta'}$$

In particular, if $L = \text{CKTVAL}$, then $q \geq \varepsilon/(\delta + O(1/n))$.

The first part of Proposition 2.8 does not specify the relation of δ' to δ (although, surely, $\delta' > \delta$ must hold for any $\gamma > 0$, because $\Delta(z, L(x)) \leq \Delta(z, y) = \delta'$). The second part relies on the fact that, for CKTVAL , one may set δ' as low as $\delta + O(1/n)$.

Proof: By completeness, there exists an oracle π such that the PCPP verifier $V(x)$ accepts oracle $y \circ \pi$ with probability 1. Consider $z = y \oplus \eta$, where $\eta \in \{0, 1\}^m$ is a uniformly distributed string with relative Hamming weight δ' . If we invoke $V(x)$ with oracle $z \circ \pi$, then the probability (over the choice of η) that any of the positions read by V has been changed is at most $q \cdot \delta'$. Thus, $V(x)$ rejects oracle $(y \oplus \eta) \circ \pi$ with probability at most $q \cdot \delta'$.

On the other hand, by assumption z is δ -far from $L(x)$ with probability at least γ , in which case $V(x)$ should reject oracle $z \circ \pi$ with probability greater than ε , by the PCPP soundness. Thus $V(x)$ should reject with probability greater than $\gamma \cdot \varepsilon$ (over the choice of z and the coin tosses of V), and we conclude that $q \cdot \delta' > \gamma \cdot \varepsilon$, as desired.

For the application to CKTVAL , let $C : \{0, 1\}^m \rightarrow \{0, 1\}$ be a circuit of size n that accepts only the all-zeroes string 0^m , for $m = \Omega(n)$. Then we have $(C, 0^m) \in \text{CKTVAL}$, but for every $\delta' > \delta$ and every string z of relative Hamming weight δ' , we see that (C, z) is δ -far from satisfying C . Setting $\gamma = 1$ and δ' such that $\delta'm$ is the least integer greater than δm completes the proof. ■

Expected robustness: Occasionally, we will be interested in a variant of robust-soundness, which refers to distance on average rather than with high probability.

Definition 2.9 (expected robustness) For a function $\rho : \mathbb{Z}^+ \rightarrow [0, 1]$, a PCP has expected robustness ρ if for every $x \notin L$, we have

$$\forall \pi, \mathbb{E}_{(I,D) \stackrel{R}{\leftarrow} V(x)} [\Delta(\pi|_I, D^{-1}(1))] > \rho(|x|).$$

Expected robustness for PCPPs is defined analogously.

We now present several generic transformations regarding robustness and soundness. Although we only state them for PCPs, all of these results also hold for PCPPs, with no change in the proximity parameter. The following proposition relates robust-soundness to expected robustness.

Proposition 2.10 (robust-soundness vs. expected robustness) If a PCP has robust-soundness error $1 - \varepsilon$ with robustness ρ , then it has expected robustness $\varepsilon \cdot \rho$. On the other hand, if a PCP has expected robustness ρ , then for every $\varepsilon \leq \rho$, it has robust-soundness error $1 - \varepsilon$ with robustness parameter $\rho - \varepsilon$.

Expected robustness can easily be amplified to standard robustness *with low robust-soundness error*, using any averaging (a.k.a. oblivious) sampler (cf., [24]). Combined with Proposition 2.10, we get a (soundness) error reduction for robust PCPs. For example, using the expander-neighborhood sampler of [29], we have:

Lemma 2.11 (error reduction via expander neighborhoods) If a language L has a PCP with expected robustness ρ , randomness complexity r , query complexity q , and decision complexity d , then for every two functions $s, \gamma : \mathbb{Z}^+ \rightarrow [0, 1]$, then L has PCP having

- robust-soundness error s with robustness parameter $\rho - \gamma$,
- randomness complexity $r + O(\log(1/s) + \log(1/\gamma))$,
- query complexity $O(1/(s\gamma^2)) \cdot q$, and
- decision complexity $O(1/(s\gamma^2)) \cdot d$

An alternative error-reduction procedure that will also be used is given by pairwise independent samples:

Lemma 2.12 (error reduction via pairwise independence) If a language L has a PCP with expected robustness ρ , randomness complexity r , query complexity q , and decision complexity d such that $\rho \cdot 2^r \geq 2$, then L has PCP having

- robust-soundness error $1/2$ with robustness parameter $\rho/2$,
- randomness complexity $2r$,
- query complexity $2q/\rho$, and
- decision complexity $2d/\rho$

Non-Boolean PCPs: The next few transformations involve non-Boolean PCPs. That is, PCPs where the oracle returns symbols over some larger alphabet $\Sigma = \{0, 1\}^a$ rather than bits; we refer to $a = a(n)$ as the **answer length** of the PCP. (Often non-Boolean PCPs are discussed in the language of multi-prover interactive proofs, but it is simpler for us to work with the PCP formulation.)

Robust-soundness of a non-Boolean PCP is defined in the natural way, using Hamming distance over the alphabet Σ . (In the case of a robust non-Boolean PCPP, we still treat the input oracle as binary.)

The first transformation provides a way of converting non-Boolean PCPs to Boolean PCPs in a way that preserves robust-soundness.

Lemma 2.13 (alphabet reduction) *If a language L has a non-Boolean PCP with answer length a , query complexity q , randomness complexity r , decision complexity d , and robust-soundness error s with robustness parameter ρ , then L has a Boolean PCP with query complexity $O(a \cdot q)$, randomness complexity r , decision complexity $d + O(a \cdot q)$, and robust-soundness error s with robustness parameter $\Omega(\rho)$. If, instead of robust-soundness, the non-Boolean PCP has expected robustness ρ , then the Boolean PCP has expected robustness $\Omega(\rho)$.*

The proof uses a good error-correcting code (i.e., constant relative distance and rate). Furthermore, to obtain decision complexity $d + O(a \cdot q)$ we should use a code having linear-size circuits for encoding (cf. [51]). Using more classical codes would only give decision complexity $d + \tilde{O}(a \cdot q)$, which is actually sufficient for our purposes.

Proof: This transformation is analogous to converting non-Boolean error-correcting codes to Boolean ones via “code concatenation”. Let V be the given non-Boolean PCP verifier, with answer length a . Let $\text{ECC} : \{0, 1\}^a \rightarrow \{0, 1\}^b$ for $b = O(a)$ a binary error-correcting code of constant relative minimum distance, which can be computed by an explicit circuit of size $O(a)$. We will augment the original oracle π having a -bit entries with an additional oracle τ having b -bit entries, where τ_i is supposed to be $\text{ECC}(\pi_i)$. (We note that including the original oracle simplifies the argument as well as frees us from assuming a noiseless decoding algorithm.)

Our new verifier $V'(x)$, on oracle access to $\pi \circ \tau$, will simulate $V(x)$, and for each query i made by V , will query the a bits in π_i and the b bits in τ_i , for a total of $q \cdot (a + b)$ binary queries. That is, if V queries positions $I = (i_1, \dots, i_q)$, V' will query positions $I' = ((0, i_1), \dots, (0, i_q), (1, i_1), \dots, (1, i_q))$. If V outputs a decision circuit $D : (\{0, 1\}^a)^q \rightarrow \{0, 1\}$, V' will output the circuit $D' : (\{0, 1\}^a)^q \times (\{0, 1\}^b)^q \rightarrow \{0, 1\}$ defined by

$$D'(x_1, \dots, x_q, y_1, \dots, y_q) = D(x_1, \dots, x_q) \wedge C(x_1, \dots, x_q, y_1, \dots, y_q),$$

where

$$C(x_1, \dots, x_q, y_1, \dots, y_q) = \bigwedge_{i=1}^q (y_i = \text{ECC}(x_i)).$$

Since ECC can be evaluated by a circuit of size $O(a)$, we see that $|D'| = |D| + O(a \cdot q)$, as desired.

For completeness of V' , we note that any accepting oracle π for V can be augmented to an accepting oracle for V' by setting $\tau_i = \text{ECC}(\pi_i)$ for all i . For soundness of V' , suppose $x \notin L$ and let (π, τ) be any pair of oracles. Define a “decoded” oracle $\hat{\pi}$ by setting $\hat{\pi}_i$ to be the string $x \in \{0, 1\}^a$ which minimizes the distance between $\text{ECC}(x)$ and τ_i . We will relate the robustness of V on oracle $\hat{\pi}$ to the robustness of V' on oracles π and τ . Specifically, let $\beta > 0$ be a constant such that the (absolute) minimum distance of ECC is greater than $2\beta \cdot (a + b)$. Then we will show that

for every sequence R of coin tosses and for every $\alpha > 0$, if the bits read by $V'(x; R)$ from $\pi \circ \tau$ are $\alpha\beta$ -close to being accepted, then the bits read by V from $\hat{\pi}$ are α -close to being accepted. Thus, both robustness parameters (standard and expected) decrease by at most a factor of β .

Consider any sequence R of coin tosses, let $(I, D) = V(x; R)$, and write $I = (i_1, \dots, i_q)$. Suppose that $(\pi_{i_1}, \dots, \pi_{i_q}, \tau_{i_1}, \dots, \tau_{i_q})$ is $\alpha\beta$ -close to some $(\pi'_{i_1}, \dots, \pi'_{i_q}, \tau'_{i_1}, \dots, \tau'_{i_q})$ that satisfies $D' = D \wedge C$. Then, for at least a $1 - \alpha$ fraction of $j \in [q]$, the pair (π_{i_j}, τ_{i_j}) is β -close to $(\pi'_{i_j}, \tau'_{i_j}) = (\pi'_{i_j}, \text{ECC}(\pi'_{i_j}))$. For such j , the choice of β implies that $\text{ECC}(\pi'_{i_j})$ is the closest codeword to τ_{i_j} and hence $\hat{\pi}_{i_j} = \pi'_{i_j}$. Since the π' 's satisfy D , we conclude that $\hat{\pi}$'s are α -close to satisfying D , as desired. ■

The usual “parallelization” paradigm of PCPs [38, 4] converts a Boolean PCP with many queries into a non-Boolean PCP with a constant number of queries, where this is typically the first step in PCP composition. As mentioned in the introduction, we cannot afford parallelization, and robust-soundness will be our substitute. Nevertheless, there is a close (but not close enough for us) connection between parallelized PCPs and PCPs with robust-soundness:

Proposition 2.14 (parallelization vs. robustness)

1. *If a language L has a non-Boolean PCP with answer length a , query complexity q , randomness complexity r , decision complexity d , and soundness error s , then L has a (Boolean) PCP with query complexity $O(a \cdot q)$, randomness complexity r , decision complexity $d + O(a \cdot q)$, and robust-soundness error s with robustness parameter $\rho = \Omega(1/q)$.*
2. *If a language L has a (Boolean) PCP with query complexity q , randomness complexity r , decision complexity d , and expected robustness ρ , then L has a 2-query non-Boolean PCP with answer length q , randomness complexity $r + \log q$, decision complexity $d + O(1)$, and soundness error $1 - \rho$.*

Thus, for constant soundness and constant robustness parameter, q -query robust (Boolean) PCPs are essentially equivalent to constant-query non-Boolean PCPs with answer length $\Theta(q)$. However, note that in passing from robust-soundness to a 2-query non-Boolean PCP, the randomness complexity increases by $\log q$. It is precisely this cost that we cannot afford, and hence we work with robust-soundness in the rest of the paper.

Proof: For Part 1, note that any non-Boolean PCP with query complexity q and soundness error s has robust-soundness error s for any robustness parameter $\rho < 1/q$. Thus, the claim follows from Lemma 2.13.

Turning to Part 2, let V be a robust PCP verifier for L with the stated parameters. We use the usual query-reduction technique for PCPs [23], and observe that when applied to a robust PCP, the detection probability (i.e., one minus the soundness error) does not deteriorate by a factor of q as usual. Instead, the detection probability of the resulting 2-query (non-Boolean) PCP equals the expected robustness of V .⁷ Specifically, the 2-query non-Boolean PCP verifier V' is defined as follows:

⁷It may be more instructive (alas more cumbersome) to discuss what is happening in terms of ordinary robustness. Suppose that V has robust-soundness error $s = 1 - d$ with respect to robustness ρ . The standard analysis ignores the robustness and asserts that the 2-query (non-Boolean) PCP has soundness error $s' = 1 - d'$, where $d' = d/q$. This crude analysis implicitly assumes the trivial bound (i.e., $1/q$) of the robustness parameter. A more refined analysis takes advantage of the actual bound of the robustness parameter, and asserts that the 2-query (non-Boolean) PCP has soundness error $s' = 1 - \rho \cdot d$.

- V' expects two oracles, one Boolean oracle π corresponding to the oracle for V , and a second oracle τ with answer length q , indexed by random strings of V .
- On input x , the verifier V' selects a random string R for V and $j \xleftarrow{R} [q]$, and computes $(I, D) = V(x; R)$, where $I = (i_1, \dots, i_q)$. It sets $I' = (R, i_j)$ (which means that the queries for the values τ_R and π_{i_j}) and $D'(a, b) = [(D(a) = 1) \wedge (a_j = b)]$; that is, it accepts if and only if $[D(\tau_R) = 1] \wedge [(\tau_R)_j = \pi_{i_j}]$.

It can be verified that the probability that V' rejects a false assertion is precisely the expected robustness of V . In particular, suppose that $V'(x)$ accepts the oracle pair (π, τ) with probability p . We may assume, without loss of generality, that $D(\tau_R) = 1$ for any R , where $(\cdot, D) = V(x; R)$. Then, it follows that the expected (relative) distance of $\pi|_I$ from τ_R , where $(I, D) = V(x; R)$ for a random R , equals $1 - p$ (because $1 - p = \Pr_{R,j}[(\tau_R)_j \neq \pi_{i_j}]$, which in turn equals $\mathbb{E}_R[\Delta(\tau_R, \pi|_I)]$). This means that on the average, π is $(1 - p)$ -close to assignments that satisfy the corresponding decision circuits. Thus, if $x \notin L$ then $1 - p > \rho$, and $p < 1 - \rho$ follows. ■

Robustness vs. proximity: Finally, for PCPPs, we prove that the robustness parameter is upper-bounded by the proximity parameter.

Proposition 2.15 (robustness vs. proximity) *Suppose a pair-language L has a PCPP with proximity parameter δ and expected robustness ρ . Suppose further that there exists $(x, y) \in L$ such that $|x| = n$ and $|y| = m$, such that if we let $z \in \{0, 1\}^m$ be a random string at relative Hamming distance $\delta' \triangleq \delta'(x)$ from y , we have*

$$\Pr_z[z \text{ is } \delta\text{-far to } L(x)] \geq \gamma \triangleq \gamma(x).$$

Then

$$\rho \leq \delta' / \gamma.$$

In particular, if $L = \text{CKTVAL}$, then $\rho \leq \delta + O(1/n)$.

Proof: The proof is similar to that of Proposition 2.8. By completeness, there exists an oracle π such that the PCPP verifier $V(x)$ accepts oracle $y \circ \pi$ with probability 1. If we run $V(x)$ with oracle $z \circ \pi$ instead, then bits read by V have expected distance at most δ' from being accepted, where the expectation is over the choices of z (even when fixing the coins of V).

On the other hand, z is δ -far from $L(x)$ with probability at least γ , and for any such fixed z the bits read by V from $z \circ \pi$ should have expected distance greater than ρ from being accepted (over the coin tosses of V). Thus, the expected distance of $z \circ \pi$ from being accepted is greater than $\gamma \cdot \rho$, where here the expectation is over the choice of z and the coin tosses of V . We conclude that $\delta' > \gamma \cdot \rho$, as desired.

Recall that in the proof of Proposition 2.8, we have demonstrated the existence of a pair (C, w) such that any string z at distance $\delta' = \delta + O(1/n)$ from w it holds that w is δ -far from satisfying C . Setting $\gamma = 1$, the second part follows. ■

3 Very short PCPs with very few queries

In this section we prove the main results of this work; that is, we establish Theorem 1.2 and 1.3. Our starting point is the following Robust PCP of proximity, which is constructed in the second part of this work (cf. Case I of Theorem 8.1).

Theorem 3.1 (Main Construct) *There exists a universal constant c such for all $n, m \in \mathbb{Z}^+$, $0 < \delta, \gamma < 1/2$ satisfying $n^{1/m} \geq m^{cm}/(\gamma\delta)^3$ and $\delta \leq \gamma/c$, CIRCUIT VALUE has a robust PCP of proximity (for circuits of size n) with the following parameters*

- randomness $(1 - \frac{1}{m}) \log n + O(m \log m) + O(\log \log n) + O(\log(1/\delta))$,
- decision complexity $n^{1/m} \cdot \text{poly}(\log n, 1/\delta)$, which also upper-bounds the query complexity.⁸
- perfect completeness, and
- for proximity parameter δ , the verifier has robust-soundness error γ with robustness parameter $(1 - \gamma)\delta$.

We comment that the condition $\delta < \gamma/c$ merely means that we present robust PCPs of proximity only for the more difficult cases (when δ is small), and our robustness parameter does not improve for larger values of δ . We call the reader's attention to the typically small value of the query and randomness complexities, which yield a proof length that is upper-bounded by $\text{poly}(m^m \log n) \cdot n$ (for δ and γ as small as $1/\text{poly}(m^m, \log n)$), as well as to the small values of the soundness error and the the small deterioration of robustness wrt proximity.

We also need the following robust PCP of proximity which has parameters similar to the PCP constructed by Arora *et al.* [4]. In comparison to the main construct above, this PCPP is not very efficient in randomness. However, as we plan to use this robust PCPP only towards the final stages of composition, we can afford to pay this cost in randomness. This robust PCP of proximity is obtained by applying the error-reduction lemma using pairwise independence (Lemma 2.12) to a variant of the robust PCP of proximity constructed in Theorem 3.1 (cf. Case II of Theorem 8.1). An alternate construction of this robust PCPP can be obtained by adding a suitable proximity test to the PCP constructed by Arora *et al.* [4].

Theorem 3.2 (ALMSS-type Robust PCP of proximity) *For all $n \in \mathbb{Z}^+$ and $\delta \in (0, 1)$, CIRCUIT VALUE has a robust PCP of proximity (for circuits of size n) with the following parameters*

- randomness $O(\log n + \log(1/\delta))$,
- decision complexity $\text{poly}(\log n)/\delta$, which also upper-bounds the query complexity.
- perfect completeness, and
- for proximity parameter δ , the verifier has robust-soundness error $1/2$ with robustness parameter $\Omega(\delta)$.

Using Theorems 3.1 and 3.2, we derive the general trade-off (captured by the following Theorem 3.3) between the length of PCPs and their query complexity.

Theorem 3.3 (Randomness vs. query complexity trade-off for PCPs of proximity) *For every parameter $n, t \in \mathbb{N}$ such that $2 \leq t \leq \frac{2 \log \log n}{\log \log \log n}$ there exists a PCP of proximity for CIRCUIT VALUE (for circuits of size n) with randomness complexity $\log_2 n + O(t + (\log n)^{\frac{1}{t}}) \log \log n + O((\log n)^{\frac{2}{t}}) + t^2 \cdot \text{poly} \log \log \log n$, query complexity $O(1)$, perfect completeness, and soundness error $1 - \Omega(1/t)$ with respect to proximity parameter $1/t$. Alternatively, we can have query complexity $O(t)$ and soundness error $1/2$ maintaining all other parameters the same.*

⁸In fact, we will upper-bound the query complexity by $q = n^{1/m} \cdot \text{poly}(\log n, 1/\delta)$ and show that the verifier's decision can be implemented by a circuit of size $\tilde{O}(q)$, which can also be bounded by $n^{1/m} \cdot \text{poly}(\log n, 1/\delta)$ with a slightly larger unspecified polynomial.

Theorem 3.3 actually asserts a PCP of proximity (for CIRCUIT VALUE), but a PCP for CIRCUIT SATISFIABILITY and a PCPP for NONDETERMINISTIC CIRCUIT VALUE (of the same complexity) follow; see Propositions 2.4 and 2.5.

Theorem 3.3 is proved by using the robust PCP of proximity described in Theorem 3.1. Specifically, this robust PCP of proximity is composed with itself several times (using the Composition Theorem from Section 2). Each such composition drastically reduces the query complexity of the resulting PCP, while only increasing very moderately its randomness complexity. The deterioration of the soundness error and the robustness is also very moderate. After composing the robust PCP of proximity with itself $O(t(n))$ times, we compose the resulting robust PCP with the ALMSS-type robust PCP of proximity thrice to reduce the query complexity to $\text{poly log log log } n$. Finally we compose this resultant robust PCP of proximity with a PCPP of proximity parameter roughly $\Omega(1/t)$ that has query complexity $O(1)$ and exponential length. The latter PCP of proximity can be obtained by a suitable modification of the Hadamard-based PCP of [4], as shown in Appendix A.

Proof: We construct the PCP of proximity of Theorem 3.3 by composing the robust PCP of proximity described in Theorem 3.1 with itself several times. Each such composition reduces the query complexity from n to approximately $n^{1/m}$. Ideally, we would like to do the following: Set $m = (\log n)^{\frac{1}{t}}$ and compose the robust PCPP of Theorem 3.1 with parameter m with itself $t-1$ times. This would result in a robust PCPP of query complexity roughly $n^{1/m^t} = n^{1/\log n} = O(1)$ giving us the desired result. However, we cannot continue this repeated composition for all the $t-1$ steps as the requirements of Theorem 3.1 (namely, $n^{1/m} \geq m^{cm}/(\delta\gamma)^3$) are violated in the penultimate two steps of the repeated composition. So we instead do the following: In the first stage, we compose the (new and) highly efficient verifier from Theorem 3.1 with itself $t-3$ times. This yields a verifier with query complexity roughly $2^{m^2} = \exp(\log^{2/t} n) < n$, while the soundness error is bounded away from 1 and robustness $\Omega(1/t)$. In the second stage, we compose the resultant robust PCPP a constant number of times with the ALMSS-type robust PCPP described in Theorem 3.2 to reduce the query complexity to $\text{poly log log log } n$ (and keeping the other parameters essentially the same). The ALMSS-type PCPP is (relatively) poor in terms of randomness, however the input size to the ALMSS-type PCPP is too small to affect the randomness of the resultant PCPP. Finally, we compose with the Hadamard-based verifier of Theorem A.1 to bring the query complexity down to $O(1)$. In all stages, we invoke the Composition Theorem (Theorem 2.7).

Throughout the proof, n denotes the size of the circuit that is given as explicit input to the PCPP verifier that we construct. We shall actually construct a sequence of such verifiers. Each verifier in the sequence will be obtained by composing the prior verifier (used as the outer verifier in the composition) with an adequate inner verifier. In the first stage, the inner verifier will be the verifier obtained from Theorem 3.1, whereas in the second and third stages it will be the one obtained from Theorem 3.2 and Theorem A.1, respectively. Either way, the inner verifier will operate on circuits of much smaller size (than n) and will use a proximity parameter that is upper-bounded by the robustness parameter of the corresponding outer verifier.

Stage I: Let $m = (\log n)^{\frac{1}{t}} \geq 2$ and $\gamma = \frac{1}{t}$. For this choice of m and γ , let V_0 be the verifier obtained from Theorem 3.1. We recall that the parameters of this verifier: For circuits of size ℓ and any proximity parameter $\delta_0 \in (\gamma/3c, \gamma/c)\mathbb{S}$, its randomness complexity is $r_0(\ell) \triangleq (1 - \frac{1}{m}) \cdot \log_2 \ell + O(\log \log \ell) + O(m \log m) + O(\log t)$, its decision (and query) complexity is $d_0(\ell) \triangleq \ell^{\frac{1}{m}} \cdot \text{poly}(\log \ell, t)$, its soundness error is $s_0 \triangleq \gamma$ and its robustness is $\rho_0 \geq (1 - \gamma)\delta_0$.

We compose V_0 with itself $t-3$ times for the same fixed choice of m and γ to obtain a sequence of

verifiers of increasingly smaller query complexity.⁹ While doing so, we will use the largest possible proximity parameter for the inner verifier (V_0) in each stage; that is, in the i th composition, we set the proximity parameter of the inner verifier to equal the robustness of the outer verifier, where the latter is the result of $i - 1$ compositions of V_0 with itself. We get a sequence of verifiers $V_1, \dots, V_{t-2}$ such that $V_1 = V_0$ and the verifier V_i is obtained by composing (the outer verifier) V_{i-1} with (the inner verifier) V_0 , where the proximity parameter of the latter is set to equal the robustness of the former. Unlike V_0 , which is invoked on different circuit sizes and (slightly) different values of the proximity parameter, all the V_i 's ($i \in [t - 2]$) refer to circuit size n and proximity parameter $\delta \triangleq \gamma/c < 1/t$.

Let r_i, d_i, δ_i, s_i and ρ_i denote the randomness complexity, decision (and query) complexity, proximity parameter, soundness error, and the robustness parameter of the verifier V_i . (Recall that V_i will be composed with the inner-verifier V_0 , where in this composition the input size and proximity parameter of the latter will be set to d_i and ρ_i respectively, and so we will need to verify that $d_i^{1/m} \geq m^{cm}/(\gamma\rho_i)^3$ and $\rho_i < \gamma/c$ for $i < t - 2$).¹⁰ We first claim that the decision complexity, proximity, soundness-error, robustness, and proof size parameters satisfy the following conditions:

1. Decision complexity: $d_i(n) \leq a(n, m)^2 \cdot n^{1/m^i}$, where $a(\ell, m) \triangleq d_0(\ell)/\ell^{1/m} = \text{poly}(\log \ell, t)$.
On the other hand, $d_i(n) \geq n^{1/m^i}$.
2. Proximity: $\delta_i = \delta$.
3. Soundness error: $s_i \leq 1 - (1 - \gamma)^i$. (In particular, $s_i < i\gamma$.)
4. Robustness: $\rho_i \geq (1 - \gamma)^i \cdot \delta$. On the other hand, $\rho_i \leq \rho_0 < \gamma/c$.
5. Proof length: $2^{r_i(n)} d_i(n) \leq b(n, m)^i \cdot n$, where $b(\ell, m) \triangleq 2^{r_0(\ell)} \cdot d_0(\ell)/\ell = \text{poly}(m^m, \log \ell, t)$.

We prove this claim by induction on i . For starters, note that the base case (i.e., $i = 1$) follows from the properties of V_0 : In particular, $d_1(n) \leq \text{poly}(\log n, t) \cdot n^{1/m}$ and $2^{r_1(n)} d_1(n) \leq \text{poly}(m^m, \log n, t) \cdot n$. Turning to the induction step, assuming that these claims hold for V_i , we prove that they hold also for V_{i+1} . For (1), note that

$$\begin{aligned}
d_{i+1}(n) &= d_0(d_i(n)) && \text{[By the Composition Theorem]} \\
&= a(d_i(n), m) \cdot d_i(n)^{1/m} && \text{[By the properties of } V_0\text{]} \\
&\leq a(n, m) \cdot d_i(n)^{1/m} && \text{[By monotonicity of } a(\cdot, \cdot) \text{ and } d_i(n) \leq n\text{]} \\
&\leq a(n, m) \cdot \left(a(n, m)^2 \cdot n^{1/m^i} \right)^{1/m} && \text{[By induction]} \\
&\leq a(n, m)^2 \cdot n^{1/m^{i+1}} && \text{[Using } m \geq 2\text{]}
\end{aligned}$$

and $d_{i+1}(n) \geq d_i(n)^{1/m} \geq n^{1/m^{i+1}}$ also holds. Clearly $\delta_i = \delta$ and the bound on s_i is straightforward from the Composition Theorem. Recalling that the proximity parameter for V_0 in this composition is set to ρ_i , the robustness of the composed verifier V_{i+1} is $\rho_{i+1} = (1 - \gamma)\rho_i = (1 - \gamma)^{i+1}\delta$ as desired. Furthermore, $\rho_i = (1 - \gamma)^i \delta \geq (1 - \frac{1}{t})^t \delta \geq e^{-1}\delta = \gamma/O(1)$. We now move to the last condition

⁹We assume, for simplicity, that $t \geq 3$. Note that it suffices to establish the claimed result for t that is greater than any universal constant.

¹⁰We also need to verify that $n^{1/m} \geq m^{cm}/(\gamma\delta_0)^3$ and $\delta_0 < \gamma/c$ for the initial verifier $V_1 = V_0$ but this is true for our choice of parameters. Furthermore, as ρ_i can only deteriorate with each composition, we have that $\rho_i \leq \rho_0 \leq \gamma/c$. Thus, the only condition that needs to be verified is $d_i^{1/m} \geq m^{cm}/(\gamma\rho_i)^3$ for $i < t - 2$.

(essentially bounding the randomness). Notice first that $r_{i+1}(n) = r_i(n) + r_0(d_i(n))$ and thus

$$\begin{aligned}
2^{r_{i+1}(n)} \cdot d_{i+1}(n) &= 2^{r_i(n)} \cdot 2^{r_0(d_i(n))} \cdot d_0(d_i(n)) && [\text{By the Composition Theorem}] \\
&\leq 2^{r_i(n)} \cdot d_i(n) \cdot b(d_i(n), m) && [\text{By the properties of } V_0] \\
&\leq n \cdot b(n, m)^i \cdot b(n, m) && [\text{By induction and monotonicity of } b(\cdot, \cdot)] \\
&\leq n \cdot b(n, m)^{i+1}
\end{aligned}$$

Thus, Part (5) is verified. Recall that we have to verify that $d_i^{1/m} \geq m^{cm}/(\gamma\rho_i)^3$ for $i < t-2$ as promised before. We have $d_i^{1/m} \geq (n^{1/m^i})^{1/m} = n^{1/m^{i+1}} \geq n^{1/m^{t-2}}$ (since $i < t-2$). Since $m = (\log n)^{\frac{1}{t}}$, we have $n^{1/m^t} = 2$. Hence, $d_i^{1/m} \geq (n^{1/m^t})^{m^2} = 2^{m^2}$. On the other hand, $m^{cm}/(\gamma\rho_i)^3 \leq m^{cm}/(e^{-1}\gamma\delta)^3 = m^{cm} \cdot \text{poly}(t)$. Thus it suffices to verify that $2^{m^2}/m^{cm} \geq \text{poly}(t)$, for $2 \leq t \leq 2 \log \log n / \log \log \log n$, which is straightforward.¹¹

Lastly, we consider the running-time of V_i , denoted T_i . A careful use of the Composition Theorem (Theorem 2.7) indicates that $T_i(n) = \text{poly}(n) + T_{i-1}(n)$, for every $i = 2, \dots, t-2$, where $T_1(n) = \text{poly}(n)$ (since $V_1 = V_0$). Alternatively, unraveling the inductive composition, we note that V_i consists of invoking V_0 for i times, where in the first invocation V_0 is invoked on V_i 's input and in later invocations V_0 is invoked on an input obtained from the previous invocation. Furthermore, the output of V_i is obtained by combining the inputs obtained in these $i \leq t-2 < n$ invocations.

We now conclude the first stage by showing that the final verifier $V_c = V_{t-2}$ has the desired properties. By Part (5) above (and the fact that $d_{t-2} \geq 1$), we have $r_c(n) = r_{t-2}(n) \leq \log n + (t-2) \cdot \log b(n, m) \leq \log n + t \log b(n, m)$. By the definition of $b(n, m)$, we have $\log b(n, m) = O(\log \log n) + O(m \log m) + O(\log t) = O(\log \log n + m \log m)$, whereas $m \log m = (\log n)^{\frac{1}{t}} \cdot \frac{1}{t} \log \log n$. Thus $r_c(n) \leq \log_2 n + O(t \cdot \log \log n) + t \cdot O(m \log m) = \log_2 n + O(t + (\log n)^{\frac{1}{t}}) \cdot \log \log n$. The decision complexity of V_c is $d_c(n) = d_{t-2}(n) \leq a(n, m)^2 \cdot n^{1/m^{t-2}} = a(n, m)^2 \cdot 2^{m^2}$, because $n^{1/m^t} = 2$. Using $a(n, m) = \text{poly}(\log n, t)$, it follows that $d_c(n) \leq 2^{m^2} \cdot \text{poly}(\log n)$. The proximity of V_c equals δ , its soundness error is $s_c = s_{t-2} = 1 - (1 - \gamma)^{t-2} = 1 - (1 - (1/t))^{t-2} < 1/2$, and its robustness is $\rho_c = \rho_{t-2} \geq (1 - \gamma)^{t-2} \delta = \delta/e = \Omega(1/t)$.

Stage II: We now compose the verifier V_c with the ALMSS-type verifier V_a described in Theorem 3.2 thrice to obtain the verifiers V' , V'' , and V''' respectively (i.e., V' equals V_c composed with V_a , V'' equals V' composed with V_a , and V''' equals V'' composed with V_a). We compose as before setting the proximity parameter of the inner verifier equal to the robustness parameter of the outer verifier. Recall from Theorem 3.2 that the ALMSS-type verifier V_a has the following parameters: randomness $r_a(\ell, \delta) = O(\log \ell + \log(1/\delta))$, decision complexity $d_a(\ell, \delta) = \text{poly} \log \ell / \delta$, soundness error $s_a(\ell, \delta) = 1/2$ and robustness $\rho_a(\ell, \delta) = \Omega(\delta)$ for input size ℓ and proximity parameter δ . Thus each composition with the inner verifier V_a adds $O(\log q + \log(1/\rho))$ to the randomness while reducing the query complexity to $(\text{poly} \log q)/\rho$ where q and ρ are the decision complexity and robustness parameter respectively of the outer verifier. Furthermore, the robustness parameter deteriorates by at most a constant factor while the soundness error remains a constant (provided that the soundness error of the outer verifier is a constant). Hence, the parameters of the verifiers V' , V'' and V''' are as follows:

Parameters of V' (recall that $d_c = 2^{m^2} \cdot \text{poly}(\log n)$ and $\delta = \Omega(1/t)$):

$$r' = r_c + O(m^2 + \log \log n + \log t) \qquad d' = \text{poly}(m, \log \log n) \cdot t$$

¹¹Note that as t varies from 2 to $2 \log \log n / \log \log \log n$, the value of m varies from $\sqrt{\log n}$ to $\sqrt{\log \log n}$. For $t \in [2, 2 \log \log n / \log \log \log n]$, the maximum value of $\text{poly}(t)$ is $\text{poly}(\log \log n / \log \log \log n) = \text{poly}(\log \log n)$. On the other hand, for $m \in [\sqrt{\log \log n}, \sqrt{\log n}]$, the minimum value of $2^{m^2}/m^{cm} > 2^{m^2/2}$ is $2^{\sqrt{\log \log n^2/2}} = \sqrt{\log n} > \text{poly}(\log \log n)$.

$$\begin{array}{ll}
s' = 1 - \Omega(1) & \rho' = \Omega(1/t) \\
\text{Parameters of } V'': & \\
r'' = r' + O(\log m + \log \log \log n + \log t) & d'' = \text{poly}(\log m, \log \log \log n, \log t) \cdot t \\
s'' = 1 - \Omega(1) & \rho'' = \Omega(1/t) \\
\text{Parameters of } V''': & \\
r''' = r'' + O(\log \log m + \log \log \log \log n + \log t) & d''' = \text{poly}(\log \log m, \log \log \log \log n, \log t) \cdot t \\
s''' = 1 - \Omega(1) & \rho''' = \Omega(1/t)
\end{array}$$

while the proximity parameter for all three verifiers is that of V_c (i.e., δ). We have that

$$\begin{aligned}
r''' &= \log_2 n + O(t + (\log n)^{1/t}) \cdot \log \log n + O(m^2), \\
q''' < d''' &= \text{poly}(\log \log \log n, \log \log m, \log t) \cdot t,
\end{aligned}$$

whereas $\delta''' = \delta = 1/(ct)$, $s''' = 1 - \Omega(1)$ and $\rho''' = \Omega(1/t)$. Substituting $m = (\log n)^{\frac{1}{t}}$, we get $r''' = \log_2 n + O(t + (\log n)^{1/t}) \cdot \log \log n + O((\log n)^{\frac{2}{t}})$ and $q''' = t \cdot \text{poly}(\log \log \log n)$.

Stage III: Finally, we compose V''' with the Hadamard-based inner verifier V_h of Theorem A.1 to obtain our final verifier V_f . The query complexity of V_h and hence that of V_f is constant. The randomness complexity of V_f is $r_f(n) \triangleq r'''(n) + r_h(q'''(n)) = r'''(n) + t^2 \cdot \text{poly}(\log \log \log n)$, because $r_h(\ell) = O(\ell^2)$. Thus, $r_f(n) = \log_2 n + O(t + (\log n)^{\frac{1}{t}}) \cdot \log \log n + O((\log n)^{\frac{2}{t}}) + t^2 \cdot \text{poly}(\log \log \log n)$. On proximity parameter δ_h , the soundness error of V_h is $s_h = 1 - \Omega(\delta_h)$. Setting $\delta_h = \rho''' = \Omega(1/t)$, we conclude that the soundness error of V_f on proximity parameter δ is $1 - \Omega(1/t)$.

To obtain soundness error $1/2$, we repeat perform $O(t)$ repetitions of V_h , yielding a query complexity of $O(t)$. This can be done without increasing the randomness complexity by using “recycled randomness” (specifically, the neighbors of a uniformly selected vertex in a Ramanujan expander graph; see [24, Apdx. C.4]). ■

Corollaries: Recall that Theorem 3.3 asserts a PCP of proximity with randomness complexity $\log_2 n + A_t(n)$, where $A_t(n) \triangleq O(t + (\log n)^{\frac{1}{t}}) \log \log n + O((\log n)^{\frac{2}{t}}) + t^2 \cdot \text{poly} \log \log \log n$ and query complexity $O(t)$ (for soundness error $1/2$). For $t \in [2, \dots, \frac{0.99 \log \log n}{\log \log \log n}]$, we have $(\log n)^{\frac{1}{t}} > (\log \log n)^{1/0.99}$ and so $A_t(n) = O((\log n)^{\frac{2}{t}})$. On the other hand, for $t \geq \frac{1.01 \log \log n}{\log \log \log n}$, we have $(\log n)^{\frac{1}{t}} \leq (\log \log n)^{1/1.01}$ and so $A_t(n) = t^2 \cdot \text{poly} \log \log \log n$.

- Two extreme choices of $t(n)$ are when $t(n) = \frac{2}{\varepsilon}$, for some $\varepsilon > 0$ (which maintains a constant query complexity), and $t(n) = \frac{2 \log \log n}{\log \log \log n}$ (which minimizes the randomness complexity of the verifier). Setting $t(n) = \frac{2}{\varepsilon}$ yields Theorem 1.3 (i.e., constant query complexity $O(1/\varepsilon)$ and randomness $\log_2 n + O(\log^\varepsilon n)$), whereas setting $t(n) = \frac{2 \log \log n}{\log \log \log n}$ yields Theorem 1.2 (i.e., query complexity $O((\log \log n)/\log \log \log n)$ and randomness $\log_2 n + (\log \log n)^2 \cdot \text{poly}(\log \log \log n)$).

Thus, both Theorems 1.2 and 1.3 follow from Theorem 3.3.

- By Proposition 2.5, we conclude that for every $2 \leq t(n) \leq \frac{2 \log \log n}{\log \log \log n}$, there exists a PCP of proximity for NONDETERMINISTIC CIRCUIT VALUE of the same complexities (i.e., randomness complexity $\log_2 n + A_t(n)$, query complexity $O(t(n))$, perfect completeness, and soundness error $1/2$ with respect to proximity $\delta = \Omega(1/t(n))$).

Comment: We note that the tight bound on the robustness (as a function of the proximity parameter) in our main construct (Theorem 3.1) plays an important role in the proof of Theorem 3.3.

The reason is that when we compose two robust PCPs of proximity, the proximity parameter of the second must be upper-bounded by the robustness parameter of the first. Thus, when we compose many robust PCPs of proximity, the robustness parameter deteriorates exponentially in the number of composed systems where the base of the exponent is determined by the tightness of the robustness (of the second verifier). That is, let $\tau \triangleq \rho/\delta$, where δ and ρ are the proximity and robustness parameters of the system. Then composing this system t times with itself, means that at the lowest PCP-instance we need to set the proximity parameter to be τ^{t-1} times the initial proximity. This requires the lowest PCP-instance to make at least $1/\tau^{t-1}$ queries (or be composed with a PCP of proximity that can handle proximity parameter τ^t , which again lower-bounds the number of queries). For a constant $\tau < 1$, we get $\exp(t)$ query complexity, whereas for $\tau = 1 - \gamma = (1 - (1/t))$ we get query complexity that is linear in $1/((1 - \gamma)^t \cdot \gamma) = O(t)$. Finally, we argue that in the context of such an application, setting $\gamma = 1/t$ is actually the “natural” choice. Such a choice, assigns each proof-oracle encountered in the composition almost equal weight (of $1/t$); that is, such a proof oracle is assigned weight $1/t$ when it appears as the current proof-oracle and maintains its weight when it appears as part of the input-oracle in subsequent compositions.

4 Applications to coding problems

In this section we show that, combined with any good code, any PCP of proximity yields a Locally Testable Code (LTC). Using our PCPs of proximity, we obtain an improvement in the rate of LTCs (improving over the results of [28, 12]). We also introduce a relaxed notion of Locally Decodable Codes, and show how to construct such codes using any PCP of proximity (and ours in particular).

Preliminaries: We consider codes mapping sequences of k (input) bits into sequences of $n \geq k$ (output) bits. Such a generic code is denoted by $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$, and the elements of $\{C(x) : x \in \{0, 1\}^k\} \subseteq \{0, 1\}^n$ are called **codewords** (of C). Throughout this section, *the integers k and n are to be thought of as parameters*, and we are typically interested in the relation of n to k (i.e., how n grows as a function of k). Thus, we actually discuss infinite families of codes (which are associated with infinite sets of possible k ’s), and whenever we say that some quantity of the code is a constant we mean that this quantity is constant for the entire family (of codes).

The **distance** of a code $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$ is the minimum (Hamming) distance between its codewords; that is, $\min_{x \neq y} \{\overline{\Delta}(C(x), C(y))\}$, where $\overline{\Delta}(u, v)$ denotes the number of bit-locations on which u and v differ. *Throughout this work, we focus on codes of “linear distance”*; that is, codes $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$ of distance $\Omega(n)$. The **distance** of $w \in \{0, 1\}^n$ from a code $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$, denoted $\overline{\Delta}_C(w)$, is the minimum distance between w and the codewords; that is, $\overline{\Delta}_C(w) \triangleq \min_x \{\overline{\Delta}(w, C(x))\}$. Recall that, for $\delta \in [0, 1]$, the n -bit long strings u and v are said to be δ -far (resp., δ -close) if $\overline{\Delta}(u, v) \geq \delta \cdot n$ (resp., $\overline{\Delta}(u, v) \leq \delta \cdot n$).

4.1 Locally Testable Codes

Loosely speaking, by a **codeword test** (for the code $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$) we mean a randomized (non-adaptive) oracle machine, called a **tester**, that is given oracle access to $w \in \{0, 1\}^n$ (viewed as a function $w : \{1, \dots, n\} \rightarrow \{0, 1\}$). The tester may query the oracle at a constant number of bit-locations and is required to (always) accept every codeword and reject with (relatively) high probability every oracle that is “far” from the code. Indeed, since our focus is on positive results, we use a strict formulation in which the tester is required to accept each codeword with probability 1. (This corresponds to “perfect completeness” in the PCP setting.) Specifically, following [28], we require the codeword test to reject any oracle $w \in \{0, 1\}^n$ with probability $\Omega(\overline{\Delta}_C(w)/n) - o(1)$. In

particular, for every $\delta > 0$, there exists an $\varepsilon > 0$ such that strings that are δ -far from the code are rejected with probability at least ε . Note that this definition requires nothing with respect to non-codewords that are relatively close to the code (i.e., are at relative distance at most $o(1)$ from C).

Using an adequate PCP of proximity, we can transform any code to a related code that has a codeword tester. This is done by appending each codeword with a PCP of proximity proving the codeword is indeed the encoding of a message. One technical problem that arises is that the PCP of proximity constitutes most of the length of the new encoding and we cannot assume much about the Hamming distance between different proofs of the same statement, thus the distance of the new code may deteriorate. But this is easily fixed by repeating the codeword many times, so that the PCP of proximity constitutes only a small fraction of the total length.¹² Specifically, given a code $C_0 : \{0, 1\}^k \rightarrow \{0, 1\}^m$ of minimal relative distance δ_0 , we consider the code $C(x) \triangleq (C_0(x)^t, \pi(x))$, where $t = d(k) \cdot |\pi(x)|/|C_0(x)|$ such that (say) $d(k) = \log k$, and $\pi(x)$ is a PCP of proximity with a proximity parameter $\delta_{\text{pcpp}} \leq \delta_0/10$ and constant soundness error $\varepsilon > 0$ (we make no attempt to optimize constants). We first note that the replicated versions of $C_0(x)$ dominate the length of $C(x)$. As for $\pi(x)$, it is a PCP of proximity that refers to an input of the form $z \in \{0, 1\}^m$ and asserts that there exists an $x = x_1 \cdots x_k$ (indeed the one that is a parameter to π) such that $z = C_0(x)$.

Observe that the code C maps k -bit long strings to codewords of length $n \triangleq (d(k) + 1) \cdot \ell$, where $\ell = s_0(m)^{1+o(1)}$ denotes the length of the PCP of proximity and $s_0(m)$ denotes the size of the circuit for encoding relative to C_0 . In addition, C approximately preserves the relative distance of C_0 , because a codeword of C is dominated by repetitions of C_0 . Using a good code (with $s_0(m) = \tilde{O}(m)$), we may set $m = O(k)$, and obtain $n = k^{1+o(1)}$. Specifically, using $\ell = s_0(m) \cdot \exp(\log^\varepsilon s_0(m))$, we obtain $n < k \cdot \exp(\log^{2\varepsilon} k)$. The codeword test emulates the PCP-verifier in the natural way. Specifically, given oracle access to $w = (w_1, \dots, w_t, \pi) \in \{0, 1\}^{t \cdot m + \ell}$, the codeword tester selects uniformly $i \in [t]$, and emulates the PCP-verifier providing it oracle access to the input-oracle w_i and to the proof-oracle π . In addition, the tester checks that the repetitions are valid (by inspecting a randomly selected position in two randomly selected m -bit long blocks). Let us denote this tester by T .

Proposition 4.1 *The oracle machine T is a codeword tester for C .*

Theorem 1.4 follows.

Proof: Completeness follows immediately from the completeness of the PCP of proximity, so we focus on the soundness condition. We consider an arbitrary $w = (w_1, \dots, w_t, \pi) \in \{0, 1\}^{t \cdot m + \ell}$. Assuming that w is δ -far from C , it follows that $w' = (w_1, \dots, w_t)$ is $(\delta - o(1))$ -far from the code $C' \triangleq \{C_0(x)^t : x \in \{0, 1\}^k\}$ (because the repetition of C_0 constitutes an $1 - o(1)$ fraction of the length of w). We consider two cases. The first case is that $\sum_{i=1}^t \overline{\Delta}_{C_0}(w_i) \geq \delta' \cdot tm$, where $\delta' = (\delta - o(1))/4$. In this case, at least $\delta'/2$ fraction of w_i are $\delta'/2$ -far from C_0 , so the PCP-verifier will reject each of them with probability at least $1 - \varepsilon$ (recall ε is the soundness error of the PCP of proximity). Thus, w will be rejected with probability $\geq \varepsilon \cdot \delta'/2$. The second case is that $\sum_{i=1}^t \overline{\Delta}_{C_0}(w_i) \leq \delta' \cdot tm$. Consider codewords c_i 's that are closest to the w_i 's (i.e., $\sum_{i=1}^t \overline{\Delta}(w_i, c_i) \leq \delta' \cdot tm$). Then it holds that $c' = (c_1, \dots, c_t)$ is $3\delta'$ -far from C' , because $\overline{\Delta}_{C'}(c') \geq \overline{\Delta}_{C'}(w') - \overline{\Delta}(w', c')$. Let c_0 be a codeword

¹²Throughout this section we will use repetitions to adjust the “weights” of various parts of our codes. An alternative method would be to work with weighted Hamming distance (i.e. where different coordinates of a codeword receive different weights), and indeed these two methods (weighting and repeating) are essentially equivalent. For the sake of explicitness we work only with repetitions.

of C_0 such that c' is closest to c_0^t . Then $3\delta' \cdot tm \leq \sum_{i=1}^t \overline{\Delta}(c_i, c_0) \leq \mathbb{E}_{r \in [t]} [\sum_{i=1}^t \overline{\Delta}(c_i, c_r)]$, where r is uniformly distributed in $[t]$. In this case, the “replication test” that checks if two random blocks w_r, w_s are equivalent, will reject with probability

$$\begin{aligned} \mathbb{E}_{r,s \in [t]} [\overline{\Delta}(w_r, w_s)/m] &\geq \frac{1}{m} \cdot \mathbb{E}_{r,s \in [t]} [\overline{\Delta}(c_r, c_s) - \overline{\Delta}(c_r, w_r) - \overline{\Delta}(c_s, w_s)] \\ &= \frac{1}{tm} \cdot \left(\mathbb{E}_{r \in [t]} \left[\sum_{i=1}^t \overline{\Delta}(c_r, c_i) \right] - 2 \cdot \sum_{i=1}^t \overline{\Delta}(c_i, w_i) \right) \\ &\geq 3\delta' - 2 \cdot \delta' \end{aligned}$$

We conclude that, in both cases, the tester rejects with probability at least $\Omega(\delta - o(1))$. $\blacksquare$

4.2 Relaxed Locally Decodable codes

We first recall the definition of Locally decodable Codes (LDCs), as formally defined by Katz and Trevisan [34]. A code $C : \{0,1\}^k \rightarrow \{0,1\}^n$ is **locally decodable** if for some constant $\delta > 0$ (which is independent of k) there exists an efficient oracle machine M that, on input any index $i \in [k]$ and access to any oracle $w \in \{0,1\}^n$ such that $\overline{\Delta}(w, C(x)) \leq \delta n$, recovers the i -th bit of x with probability at least $2/3$ while making a constant number of queries to w . That is, whenever relatively few location are corrupted, the decoder should be able to recover each information-bit, with high probability, based on a constant number of queries to the (corrupted) codeword.

Katz and Trevisan showed that if M makes q queries then $n = \Omega(k^{1+1/(q-1)})$ must hold [34].¹³ This lower-bound is quite far from the best known upper-bound, due to Beimal *et al.* [8], that asserts $n \leq \exp(k^{O((\log \log q)/(q \log q))})$. It has been conjectured that, for a constant number of queries, n should be exponential in k ; that is, for every constant q there exists a constant $\varepsilon > 0$ such that $n > \exp(k^\varepsilon)$ must hold. In view of this state of affairs, it is natural to relax the definition of Locally decodable Codes, with the hope of obtaining more efficient constructions (e.g., $n = \text{poly}(k)$).

We relax the definition of Locally decodable Codes by requiring that, whenever few location are corrupted, the decoder should be able to recover most (or almost all) of the individual information-bits (based on few queries) and for the rest (of the locations) the decoder may output a fail symbol (but not the wrong value). That is, the decoder must still avoid errors (with high probability), but is allowed to say “don’t know” on a few bit-locations. The following definition is actually weaker; yet, the (aforementioned) stronger formulation is obtained when considering $\rho \approx 1$ (and using standard amplification to reduce the error from $1/3$ to any desired constant).

Definition 4.2 (Relaxed LDC) *A code $C : \{0,1\}^k \rightarrow \{0,1\}^n$ is relaxed locally decodable if for some constants $\delta, \rho > 0$ there exists an efficient probabilistic oracle machine M that makes a constant number of queries and satisfies the following two conditions with respect to any $w \in \{0,1\}^n$ such that $\overline{\Delta}(w, C(x)) \leq \delta n$:*

1. *On input any index $i \in [k]$ and given access to the oracle w , with probability at least $2/3$, machine M outputs either the i -th bit of x or a special failure symbol, denoted $\perp$.*
2. *For at least a ρ fraction of the indices $i \in [k]$, on input i and oracle access to $w \in \{0,1\}^n$, with probability at least $2/3$, machine M outputs the i -th bit of x .*

¹³Their lower-bound refers to non-adaptive decoders, and yields a lower-bound of $n = \Omega(k^{1+1/(2^q-1)})$ for adaptive decoders. A lower-bound of $n = \Omega(k^{1+1/O(q)})$ for adaptive decoders was presented in [18], and lower-bound of $n = \Omega(k^{1+1/(q/2-1)})$ for non-adaptive decoders was presented in [36]. (We note that below we use a non-adaptive (relaxed) decoder.)

Furthermore, if $w = C(x)$ is a codeword then the decoder correctly recovers every bit of x with probability at least $2/3$.

One may strengthen the definition by requiring that ρ be greater than $1/2$ or any other favorite constant smaller than 1 (but probably refrain from setting $\rho > 1 - \delta$ or so). In fact, we achieve the latter stronger form.

Note that it is very easy to come up with constructions that satisfy each one of the two main conditions of Definition 4.2. For example, Condition 1 can be satisfied by (any code and) a trivial decoder that always returns $\perp$. On the other hand, the identity encoding (combined with a trivial decoder) satisfies Condition 2. Our aim, however, is to have a construction that satisfies both conditions (as well as the furthermore clause) and beats the performance of the known locally decodable codes.

The above definition only considers strings w that are δ -close to the code. However, using our construction from the previous section, any relaxed LDC can be augmented so that strings that are δ -far from the code are rejected with high probability (i.e. for every index i , the decoder outputs $\perp$ with high probability), and this can be done with only a nearly linear ($n \cdot \exp(\log^\varepsilon n)$) increase in size.

Motivation to our construction: We seek a code of linear distance that has some weak “local decodability” properties. One idea is to separate the codeword into two parts, the first providing the distance property (eg using any code of linear distance) and the second allowing for “local decodability” (eg using the identity map). It is obvious that a third part that guarantees the consistency of the first part should be added, and it is natural to try to use a PCP of proximity in the latter part. The natural decoder will check consistency (via the PCPP), and in case it detects no error will decode according to the second part. Indeed, the second part may not be “robust to corruption” but the first part is “robust to corruption” and consistency means that both parts encode the same information. Considering this vague idea, we encounter two problems. First, a PCP of proximity is unlikely to detect a small change in the second part. Thus, if we use the identity map in the second part then the decoder may output the wrong value of some (although few) bits. Put in other words, the “proximity relaxation” in PCPPs makes sense for the first part of the codewords but not for the second. Our solution is to provide proofs of the consistency *for each bit* in the second part with the entire first part. The second problem is that the PCPPs are much longer than the first two parts, whereas the corruption rate is measured in terms of the entire codeword. This is easy to fix by repeating the first two parts sufficiently many times. However, it is important not to “overdo” this repetition, because if the third part is too short, then corrupting it may prevent meaningful decoding even at low corruption rates (measured in terms of the entire codeword).

The actual construction. Let $C_0 : \{0, 1\}^k \rightarrow \{0, 1\}^m$ be a good code of relative distance δ_0 , then we encode $x \in \{0, 1\}^k$ by $C(x) \triangleq (x^t, C_0(x)^{t'}, \pi_1(x), \dots, \pi_k(x))$, where $t = |\pi_1(x), \dots, \pi_k(x)|/|x|$ (resp., $t' = |\pi_1(x), \dots, \pi_k(x)|/|C_0(x)|$), and $\pi_i(x)$ is a PCP of proximity to be further discussed. We first note that the replicated versions of x (resp., $C_0(x)$) takes a third of the total length of $C(x)$. As for $\pi_i(x)$, it is a PCP of proximity that refers to an input of the form $(z_1, z_2) \in \{0, 1\}^{m+m}$ and asserts that there exists an $x = x_1 \cdots x_k$ (indeed the one that is a parameter to π_i) such that $z_1 = x_i^m$ and $z_2 = C_0(x)$.¹⁴ We use our PCP of proximity from Theorem 3.3 (with $t(n)$ set to be a sufficiently large constant to be determined later). Let $\delta_{\text{pcpp}} > 0$ be the proximity parameter of the

¹⁴Indeed z_1 is merely the bit x_i repeated $|C_0(x)|$ times in order to give equal weight to each part in measuring proximity.

PCP of proximity and let q denote the number of queries it makes in order to support a soundness error of $1/6$. We note that, in our PCP of proximity, the queries to the proof oracle (resp., input oracle) are uniformly distributed. It follows from this uniformity and perfect completeness of the PCP of proximity, that if we flip at most a δ' fraction of the bit-locations of the input-oracle and the proof-oracle then the verifier will accept with probability at least $1 - q \cdot \delta'$, (Following the alternative argument, we obtain the weaker lower-bound of $1 - (2/\delta_{\text{pcpp}}) \cdot q\delta'$.)

Observe that the code C maps k -bit long strings to codewords of length $n \triangleq 3 \cdot k \cdot \ell$, where $\ell = s_0(m)^{1+o(1)}$ denotes the length of the PCP of proximity and $s_0(m)$ denotes the size of the circuit for encoding relative to C_0 . Using a good code — of constant relative distance δ_0 , $m = O(k)$, and $s_0(m) = \tilde{O}(m)$ — we obtain $n = k^{2+o(1)}$. The relative distance of C is at least $\delta_0/3$.

We now turn to the description of the decoder D . On input $i \in [k]$ and oracle access to $w = (w_1, w_2, w_3) \in \{0, 1\}^n$, the decoder invokes the PCPP-verifier providing it with access to an input-oracle (z_1, z_2) and a proof oracle π that are defined and emulated as follows: The decoder selects uniformly $r \in [t]$ and $r' \in [t']$, and defines each bit of z_1 to equal the $((r-1)k + i)$ -th bit of w_1 , the string z_2 is defined to equal the r' -th (m -bit long) block of w_2 , and π is defined to equal the i -th block (ℓ -bit long) of w_3 . That is, when the verifier asks to access the j -th bit of z_1 (resp., z_2) [resp., π], the decoder answers with the $((r-1)k + i)$ -th bit of w_1 (resp., $((r'-1)m + j)$ -th bit of w_2) [resp., the $((i-1)\ell + j)$ -th bit of w_3]. If the verifier rejects then the decoder outputs a special (failure) symbol. Otherwise, it outputs the $((r-1)k + i)$ -th bit of w_1 .

Proposition 4.3 *There exist universal constants $\varepsilon > 0$ and c such that the code C is relaxed locally decodable for any $\delta \leq \varepsilon$ and with $\rho = 1 - c\delta$.*

Proof: We prove that the decoder D satisfies the two main conditions of Definition 4.2 (for $\rho = 1 - O(\delta)$). Fixing any $x \in \{0, 1\}^k$, we consider an arbitrary oracle $w = (w_1, w_2, w_3)$ that is δ -close to $C(x)$, where w_1 (resp., w_2) denotes the alleged replication of x (resp., $C_0(x)$) and $w_3 = (w_{3,1}, \dots, w_{3,k})$ denotes the part of the PCPs of proximity. Note that w_1 (resp., w_2) is 3δ -close to x^t (resp., to $C_0(x)^t$). To analyze the performance of $D^w(i)$, we define random variables Z_1 and Z_2 that correspond to the input-oracles to which the PCP-verifier is given access. Specifically, $Z_1 = \sigma^m$, where σ is set to equal the $((r-1)k + i)$ -th bit of w_1 , when r is uniformly distributed in $[t]$. Likewise, Z_2 is determined to be the r' -th block of w_2 , where r' is uniformly distributed in $[t']$. Finally, we set the proof-oracle, π , to equal the i -th block of w_3 .

We start with Condition 1. We bound the probability that the decoder outputs $\neg x_i$ by considering three cases:

Case 1: $\sigma = x_i$. Then Condition 1 is satisfied regardless of whether the decoder outputs σ or $\perp$ (which are the only two possibilities).

Case 2: Z_2 is 18δ -far from $C_0(x)$. This case occurs with probability at most $1/6$, since the expected distance of Z_2 from $C_0(x)$ is at most 3δ .

Case 3: Z_2 is 18δ -close to $C_0(x)$ and $\sigma \neq x_i$. Then, (Z_1, Z_2) is $\min\{1/2, \delta_0 - 18\delta\}$ -far from any string of the form $(y_i^m, C_0(y))$. (Either all of Z_1 must be changed, or Z_2 must be changed to a codeword other than $C_0(x)$.) Choosing $\delta_{\text{pcpp}} < \min\{1/2, \delta_0 - 18\delta\}$, we conclude that the PCPP verifier accepts (and the decoder outputs $\neg x_i$) with probability at most $1/6$.

Thus, in total, the decoder outputs $\neg x_i$ with probability at most $1/6 + 1/6 = 1/3$.

We now turn to Condition 2. Recall that $w_1 = w_{1,1} \cdots w_{1,t\ell}$ (resp., w_3) is 3δ -close to x^t (resp., to $\pi_1(x), \dots, \pi_k(x)$). It follows that for all but at most an $\alpha/2$ fraction of the i 's, $w_{1,i}w_{1,k+i} \cdots w_{1,(t-1)k+i}$

is $6\delta/\alpha$ -close to x_i^t (resp., the i -th block of w_3 is $6\delta/\alpha$ -close to $\pi_i(x)$). Thus, for at least a $1 - \alpha$ fraction of the i 's, Z_1 equals x_i^m with probability at least $1 - 6\delta/\alpha$ and the i -th block of w_3 is $6\delta/\alpha$ -close to $\pi_i(x)$. Let us fix such a good i . Recall that the expected (fractional) distance of Z_2 from $C_0(x)$ is at most 3δ and that the decoder D outputs σ (where $Z_1 = \sigma^m$) if and only if the PCP-verifier accepts. We need to lower-bound the probability that the PCP-verifier accepts the oracle-input (Z_1, Z_2) when accessing the proof oracle π (i.e., the i -th block of w_3). Recall that the PCP-verifier makes at most q queries, where each query is uniformly distributed in one of the oracles, and that the PCP-verifier would have accepted $(x_i^m, C_0(x))$ with probability 1 when accessing the proof $\pi_i(x)$. It follows that this verifier rejects (Z_1, Z_2) when accessing the (corrupted) proof π with probability at most

$$\Pr[Z_1 \neq x_i^m] + q \cdot \mathbb{E} \left[\frac{\overline{\Delta}(Z_2, C_0(x))}{m} \right] + q \cdot \frac{\overline{\Delta}(\pi, \pi_i(x))}{\ell} < \frac{6\delta}{\alpha} - q \cdot 3\delta + q \cdot \frac{6\delta}{\alpha} < \frac{7q\delta}{\alpha}$$

Now, $7q\delta/\alpha < 1/3$, provided that $\alpha > 21q\delta$. That is, we may use $\alpha = 22q\delta$. Thus, Condition 2 holds for $\rho = 1 - \alpha = 1 - 22q\delta$. ■

Improving the rate: The reason that our code has quadratic length codewords (i.e., $n = \Omega(k^2)$) is that we augmented a standard code with proofs regarding the relation of the standard codeword to the value of *each* information bit. Thus, we had k proofs each relating to a statement of length $\Omega(k)$. Consider the following improvement. Partition the message into $\sqrt{k}$ blocks, each of length $\sqrt{k}$. Encode the original message and each one of the smaller blocks, via good error correcting codes. Let w be the encoding of the entire message, and $w_i, i = 1, \dots, \sqrt{k}$ be the encodings of the blocks. Now, for every $i = 1, \dots, \sqrt{k}$ append a PCP of proximity for the claim “ w_i is the encoding of the i th block of a message encoded by w ”. In addition, for each message bit x_j residing in block i , append a PCP of proximity of the statement “ x_j is the j th bit of the $\sqrt{k}$ -bit long string encoded in w_i ”. The total encoding length has decreased, because we have $\sqrt{k}$ proofs of statements of length $O(k)$ and k proofs of statements of length $O(\sqrt{k})$, leading to a total length that is almost linear in $k^{3/2}$.

In general, for any constant ℓ , we consider ℓ successively finer partitions of the message into blocks, where the $(i+1)$ -st partition is obtained by breaking each block of the previous partition into $k^{1/\ell}$ equally sized pieces. Thus, the i -th partition uses $k^{i/\ell}$ blocks, each of length $k^{1-i/\ell}$. Encoding is done by providing, for each $i = 0, 1, \dots, \ell$, encodings of all of the blocks in the i -th partition in a good error-correcting code. (Thus, for $i = 0$ we provide the encoding of the entire messages, whereas for $i = \ell$ we provide an “encoding” of individual bits.) Each of these $\ell + 1$ levels of encodings will be assigned equal weight (via repetitions) in the new codeword. In addition, the new codeword will contain PCPs of proximity that assert the consistency of “directly related” blocks (i.e., blocks of consecutive levels that contain one another). That is, for every $i = 1, \dots, \ell$ and $j \in [k^{i/\ell}]$, we place a proof that the encoding of the j -th block in the i -th level is consistent with the encoding of the $\lceil j/k^{1/\ell} \rceil$ -th block in the $(i-1)$ -st level. The i -th such sequence of proofs contains $k^{i/\ell}$ proofs, where each such proof refers to statements of length $O(k^{1-i/\ell} + k^{1-(i-1)/\ell}) = O(k^{1-(i-1)/\ell})$, which yields a total length of proofs that is upper-bounded by $k^{i/\ell} \cdot (k^{1-(i-1)/\ell})^{1+o(1)} = k^{1+1/\ell+o(1)}$. Each of these sequences will be assigned equal weight in the new codeword, and the total weight of all the encodings will equal the total weight of all proofs. The new decoder will just check the consistency of the ℓ relevant proofs and act accordingly. (Thus, the decoder uses $O(\ell)$ queries.) We stress that, as before, the proofs in use are PCPs of proximity. In the current context these proofs refer to two input-oracles of vastly different length, and so the bit-positions of the shorter input-oracle are

given higher “weight” (by repetition) such that both input-oracles are assigned the same weight.¹⁵ The proof of Proposition 4.3 can be extended with ε and c now depending on ℓ (i.e., $\varepsilon = \Omega(1/\ell)$ and $c = O(\ell)$). That is:

Proposition 4.4 *There exist universal constants $\varepsilon' > 0$ and c' such that, for every ℓ , the aforementioned code is relaxed locally decodable (by $O(\ell)$ queries) for any $\delta \leq \varepsilon'/\ell$ and with $\rho = 1 - c'\ell\delta$.*

Theorem 1.5 follows.

Proof Sketch: Following the outline of the proof of Proposition 4.3, we consider an oracle $(w_0, w_1, \dots, w_\ell, \pi_1, \dots, \pi_\ell)$ that is δ close to an encoding of $x \in \{0, 1\}^k$, where each w_i is supposed to consist of encodings of the $k^{i/\ell}$ (non-overlapping) $k^{1-i/\ell}$ -bit long blocks of x , and π_i consists of the corresponding proofs of consistency. A key observation is that each of these sequences is sufficiently close to what it is suppose to be (because $(2\ell + 1) \cdot \delta < 3\varepsilon'$ is a sufficiently small constant). Condition 1 is thus established as in the proof of Proposition 4.3, while noting that if the encoding of some bit does not fit the encoding of the entire x then a mismatch must occur between the encoding of some block (in some w_i) and its immediate super-block (in w_{i-1}). To establish Condition 2, we first upper-bound the number of $i \in [k]$ for which one of the relevant blocks or proofs has been “drastically corrupted” (i.e., corrupted too much). Corrupting a δ fraction of the codeword, only allows to “drastically corrupt” a δ fraction of the blocks and proofs at each level, which amounts to disrupting the decoding of $O(\ell\delta)$ of the bit positions. ■

Open Problem: We wonder whether one can obtain relaxed-LDC that can be decoded using q queries while having length $n = o(k^{q/(q-1)})$. The existence of such relaxed-LDC will imply that our relaxation (i.e., relaxed-LDC) is actually strict, because such codes will beat the lower-bound currently known for LDC (cf. [34]). Alternatively, it may be possible to improve the lower-bound for (q -query) LDC to $n > k^{1+(c/q)}$, for any constant c and every sufficiently large constant q (where, as usual, k is a parameter whereas q is a fixed constant). (In fact, some conjecture that n must be superpolynomial in k , for any constant q .)

4.3 Linearity of the codes

We note that the codes presented above (establishing both Theorems 1.4 and 1.5) are actually GF(2)-linear codes, whenever the base code C_0 is also GF(2)-linear. Proving this assertion reduces to proving that the PCPs of proximity used (in the aforementioned constructions) have proof-oracles in which each bit is a linear functions of the bits to which the proof refers. The main part of the latter task is undertaken in Section 8.4, where we show the the main construct (i.e., the PCPs of proximity stated in Theorems 3.1 and 3.2) when applied to a linear circuit yields a an GF(2)-linear transformation of assignments (satisfying the circuit) to proof-oracles (accepted by the verifier). In addition, we need to show that also the construction underlying the proof of Theorem 3.3 satisfy this property. This is done next, and consequently we get:

Proposition 4.5 *If C is a linear circuit (see Definition 8.13), then there is a linear transformation T mapping satisfying assignments w of C to proof oracles $T(w)$ such that the PCPP verifier of Theorem 3.3 will, on input C , accept oracle $(w, T(w))$ with probability 1.*

¹⁵In a sense, this was also done in the simpler code analyzed in Proposition 4.3.

Proof Sketch: In Section 8.4, we establish a corresponding result for the main construct (i.e., Proposition 8.14 refers to the linearity of the construction used in the proof of Theorem 8.1, which in turn underlies Theorems 3.1 and 3.2). Here we show that linearity is preserved in composition as well as by the most inner (or bottom) verifier.

In each composition step, we append the proof-oracle with new (inner) PCPs of proximity per each test of the (outer) verifier. Since all these tests are linear, we can apply Proposition 8.14 and infer that the new appended information is a linear transformation of the input-oracle and the outer proof-oracle (where, by induction, the latter is a linear transformation of the input).

At the bottom level of composition we apply a Hadamard based PCP (Section A). The encoding defined there is not $GF(2)$ -linear (rather it is quadratic), but this was necessary for dealing with non-linear gates. It can be verified that for a linear circuit, one can perform all necessary tests of Section A with the Hadamard encoding of the input. Thus, we conclude this final phase of the encoding is also linear, and this completes the proof of Proposition 4.5. ■

Part II

The main construct: A short, robust PCP of proximity

5 Overview of our main construct

Throughout this section, n denotes the length of the explicit input given to the PCPP verifier, which in case of CIRCUIT VALUE is defined as the size of the circuit (given as explicit input). As stated in the introduction, our main results rely on the following highly efficient robust PCP of proximity.

Theorem 3.1 (Main Construct - restated): *There exists a universal constant c such for all $n, m \in \mathbb{Z}^+$, $0 < \delta, \gamma < 1/2$ satisfying $n^{1/m} \geq m^{cm}/(\gamma\delta)^3$ and $\delta \leq \gamma/c$, CIRCUIT VALUE has a robust PCP of proximity (for circuits of size n) with the following parameters*

- *randomness* $(1 - \frac{1}{m}) \log n + O(m \log m) + O(\log \log n) + O(\log(1/\delta))$,
- *decision complexity* $n^{1/m} \cdot \text{poly}(\log n, 1/\delta)$, which also upper-bounds the query complexity.¹⁶
- *perfect completeness*, and
- *for proximity parameter δ , the verifier has robust-soundness error γ with robustness parameter $(1 - \gamma)\delta$.*

We comment that the condition $\delta < \gamma/c$ merely means that we present robust PCPs of proximity only for the more difficult cases (when δ is small). A (simplified) variant of Theorem 3.1 also yields the ALMSS-type Robust PCP of proximity (of Theorem 3.2). Following is an overview of the proof of Theorem 3.1; the actual proof is given in the subsequent three sections.

Theorem 3.1 is proved by modifying a construction that establishes Theorem 1.1. We follow [31] and modify their construction. (An alternative approach would be to start from [44], but that

¹⁶In fact, we will upper-bound the query complexity by $q = n^{1/m} \cdot \text{poly}(\log n, 1/\delta)$ and show that the verifier's decision can be implemented by a circuit of size $\tilde{O}(q)$, which can also be bounded by $n^{1/m} \cdot \text{poly}(\log n, 1/\delta)$ with a slightly larger unspecified polynomial.

construction does not seem amenable to achieving robust soundness.) The construction of [31] may be abstracted as follows: To verify the satisfiability of a circuit of size n , a verifier expects oracles $F_i : F^m \rightarrow F$, $i \in \{1, \dots, t = \text{poly log } n\}$, where F is a field and m is a parameter such that $F^m \approx m^m \cdot n$. The verifier then needs to test that (1) each of the F_i 's is close to a m -variate polynomial of low degree and (2) the polynomials satisfy some consistency properties which verify that F_i is locally consistent with F_{i-1} .¹⁷ (These consistency checks include tests which depend on the input circuit and verify that F_i 's actually encode a satisfying assignment to the circuit.)

We work within this framework — namely our verifier will also try to access oracles for F_i 's and test low-degreeness and consistency. Our key modification to this construction is a randomness-reduction in the low-degree test obtained by using the small collection of (small-biased) lines of [12], while using only the “canonical” representations of these lines (and avoiding any complication that was introduced towards “proof composition”). In particular, unlike in [31, 28, 12], we cannot afford to pack the polynomials $F_1, \dots, F_t$ into a single polynomial (by using an auxiliary variable that blows-up the proof length by a factor of the size of the field in use). Instead, we just maintain all these t polynomials separately and test them separately to obtain Theorem 1.1. (In the traditional framework of parallelized PCPs, this would give an unaffordable increase in the number of (non-Boolean) queries. However, we will later ameliorate this loss by a “bundling technique” that will yield robust-soundness.)

The resulting PCP is converted into a PCP of proximity by comparing the input-oracle (i.e. supposed satisfying assignment to the circuit) to the proof-oracle (which is supposed to include an encoding of the said assignment). That is, we read a random location of the input and the corresponding location of the proof oracle, and test for equality. Actually, these locations of the proof-oracle must be accessed via a self-correction mechanism (rather than merely probing at the desired points of comparison), since they constitute only a small part of the proof oracle (and thus corruptions there may not be detected). (This technique was already suggested in [6].)

The most complex and subtle part of the proof of Theorem 3.1 is establishing the robust-soundness property. We sketch how we do this below, first dealing with the low-degree test and the consistency tests separately, and then showing how to reconcile the two “different” fixes.

Low-degree tests of $F_1, \dots, F_t$: Selecting a random line $\ell : F \rightarrow F^m$ (from the aforementioned sample space), we can check that (for each i) the restriction of F_i to the line ℓ (i.e., the function $f_i(j) \triangleq F_i(\ell(j))$) is a low-degree (univariate) polynomial. Each of these tests is *individually robust*; that is, if F_i is far from being a low-degree polynomial then with high probability the restriction of F_i to a random line ℓ (in the sample space) is far from being a low-degree polynomial. The problem is that the conjunction of the t tests is not sufficiently robust; that is, if one of the F_i 's is δ -far from being a low-degree polynomial then it is only guaranteed that the sequence of t restrictions (i.e., the sequence of the f_i 's) is (δ/t) -far from being a sequence of t low-degree (univariate) polynomials. Thus robustness decreases by a factor of t , which we cannot afford for nonconstant t .

Our solution is to observe that we can “bundle” the t functions together into a function $\bar{F} : F^m \rightarrow F^t$ such that if one of the F_i 's is far from being a low-degree polynomial then the restriction of $\bar{F}$ to a random line will be far from being a bundling of t low-degree univariate polynomials. Specifically, for every $x \in F^m$, define $\bar{F}(x) \triangleq (F_1(x), \dots, F_t(x))$. To test that $\bar{F}$ is a bundling of low-degree polynomials, select a random line ℓ (as above), and check that $\bar{f}_\ell(j) = \bar{F}(\ell(j))$ is a bundling of low-degree univariate polynomials. Thus, we establish *robustness at the bundle level*; that is, if one of the F_i 's is far from being low degree then, with high probability, one must modify $\bar{f}_\ell$ on a

¹⁷Strictly speaking, the consistency checks are a little more complicated, with the functions really being indexed by two subscripts and consistency tests being between $F_{i,j}$ and $F_{i,j-1}$, as well as between $F_{i,0}$ and $F_{i+1,0}$. However, these differences don't alter our task significantly — we ignore them in this section to simplify our notation.

constant fraction of values in order to make the test accept. The point is that this robustness refers to Hamming distance over the alphabet F^t , rather than alphabet F as before. We can afford this increase in alphabet size, as we later encode the values of $\bar{F}$ using an error-correcting code in order to derive *robustness at the bit level*.

We wish to highlight a key point that makes the above approach work: when we look at the values of $\bar{F}$ restricted to a random line, we get the values of the individual F_i 's restricted to some random line, which is exactly what a low-degree test of each F_i needs. This fact is not very surprising, given that we are subjecting all F_i 's to the same test. *But what happens when we need to make two different types of tests?* This question is not academic and does come up in the consistency tests.

Consistency tests: To bundle the t consistency tests between F_i and F_{i+1} we need to look into the structure of these tests. We note that for every i , a random test essentially refers to the values of F_i and F_{i+1} on (random) i -th axis-parallel lines. That is, for every i , and a random $x' = (x_1, \dots, x_{i-1}) \in F^{i-1}$ and $x'' = (x_{i+1}, \dots, x_m) \in F^{m-i}$, we need to check some relation between $F_i(x', \cdot, x'')$ and $F_{i+1}(x', \cdot, x'')$.¹⁸ Clearly, querying $\bar{F}$ as above on the i th axis-parallel line, we can obtain the relevant values from $\bar{F}(x', \cdot, x'')$, but this works only for one specific value of i , and other values of i will require us to make other queries. The end result would be that we'll gain nothing from the bundling (i.e., from $\bar{F}$) over using the individual F_i 's, which yields a factor of t loss in the robustness.¹⁹ Fortunately, a different bundling works in this case.

Consider $\bar{F}'$ such that $\bar{F}'(x) \triangleq (F_1(x), GF(2)(S(x)), \dots, F_t(S^{t-1}(x)))$, for every $x \in F^m$, where S denotes a (right) cyclic-shift (i.e., $S(x_1, \dots, x_m) = (x_m, x_1 \dots x_{m-1})$ and $S^i(x_1, \dots, x_m) = (x_{m-(i-1)}, \dots, x_m, x_1, x_2, \dots, x_{m-i})$). Now, if we ask for the value of $\bar{F}'$ on the first and last axis-parallel lines (i.e., on $(\cdot, x_2, \dots, x_m)$ and $(x_2, \dots, x_m, \cdot) = S^{-1}(\cdot, x_2, \dots, x_m)$), then we get all we need for all the m tests. Specifically, for every i , the i -th component in the bundled function $\bar{F}'(\cdot, x_2, \dots, x_m)$ is $F_i(S^{i-1}(\cdot, x_2, \dots, x_m)) = F_i(x_{m-i+2}, \dots, x_m, \cdot, x_2, \dots, x_{m-i+1})$, whereas the $(i+1)$ -st component in $\bar{F}'(S^{-1}(\cdot, x_2, \dots, x_m))$ is $F_{i+1}(S^i(S^{-1}(\cdot, x_2, \dots, x_m))) = F_{i+1}(x_{m-i+2}, \dots, x_m, \cdot, x_2, \dots, x_{m-i+1})$. Thus, we need only to query two bundles (rather than t), and robustness only drops by a constant factor.

Reconciling the two bundlings: But what happens with the low-degree tests that we need to do (which were “served” nicely by the original bundling $\bar{F}$)? Note that we cannot use both $\bar{F}$ and $\bar{F}'$, because this will require testing consistency between them, which will introduce new problems as well as a cost in randomness that we cannot afford. Fortunately, the new bundling (i.e., $\bar{F}'$), designed to serve the axis-parallel line comparisons, can also serve the low-degree tests. Indeed, the various F_i 's will not be inspected on the same lines, but this does not matter, because the property of being a low-degree polynomial is preserved when “shifted” (under S).

Tightening the gap between robustness and proximity: The above description suffices for deriving a weaker version of Theorem 3.1 in which the robustness is only (say) $\delta/3$ rather than $(1 - \gamma)\delta$ for a parameter γ that may be set as low as $1/\text{poly}(\log n)$. Such a weaker result yields a weaker version of Theorem 3.3 in which the query complexity is exponentially larger (e.g., for proof-length $\exp(o(\log \log n)^2) \cdot n$, we would have obtained query complexity $\exp(o(\log \log n)) = \log^{o(1)} n$ rather than $o(\log \log n)$); see comment at the end of Section 3. To obtain the stronger bound on the robustness parameter, we take a closer look at the conjunction of the standard PCP test and

¹⁸Again, this is an oversimplification, but suffices to convey the main idea of our solution.

¹⁹It turns out that for constant m (e.g., $m = 2$) this does not pose a problem. However, a constant m would suffice only for proving a slightly weaker version of Theorem 1.2 (where $o(\log \log n)$ is replaced by $\log \log n$). but not for proving Theorem 1.3, which requires setting $m = \log^\varepsilon n$, for constant $\varepsilon > 0$.

the proximity test. The PCP test can be shown to have constant robustness $c > 0$, whereas the proximity test can be shown to have robustness $\delta' \triangleq (1 - \gamma)\delta$. When combining the two tests, we obtain robustness equal to $\min(\alpha c, (1 - \alpha)\delta')$, where α is the relative length of queries used in the PCP test (as a fraction of the total number of queries). A natural choice, which yields the weaker result, is to weight the queries (or replicate the smaller part) so that $\alpha = 1/2$. (This yields robustness of approximately $\min(c, \delta')/2$.) In order to obtain the stronger bound, we assign weights such that $\alpha = \gamma$, and obtain robustness $\min(\gamma c, (1 - \gamma)\delta') > \min(\Omega(\gamma), (1 - 2\gamma)\delta)$, which simplifies to $(1 - 2\gamma)\delta$ for $\delta < \gamma/O(1)$. (The above description avoids the fact that the PCP test has constant soundness error, but the soundness error can be decreased to γ by using sequential repetitions while paying a minor cost in randomness and *while approximately preserving the robustness*. We comment that the proximity test, as is, has soundness error γ .)

6 A randomness-efficient PCP

In this section, we present a vanilla version (Theorem 6.1) of Theorem 3.1. More specifically, we construct a regular PCP for CIRCUIT SATISFIABILITY (i.e., a robust PCP of proximity without either the robustness or proximity properties). This construction favors over earlier PCP constructions in the fact that it is very efficient in randomness. As mentioned earlier, this theorem suffices to prove Theorem 1.1.

Theorem 6.1 *There exists a universal constant $0 < \varepsilon < 1$ such that the following holds. Suppose $m \in \mathbb{Z}^+$ satisfies $m \leq \log n / \log \log n$. Then there exists a PCP for CIRCUIT SATISFIABILITY (for circuits of size n) with the following parameters*

- *randomness* $(1 - \frac{1}{m}) \log n + O(m \log m) + O(\log \log n)$,
- *query complexity* $q = O(m^2 n^{1/m} \log^2 n)$ and *decision complexity* $\tilde{O}(q)$,
- *perfect completeness*,
- *and soundness error* $1 - \varepsilon$.

The construction of the PCP for CIRCUIT SATISFIABILITY proceeds in three steps. First, we transform the input circuit φ to a well-structured circuit φ' along the lines of Polishchuk and Spielman [44, 52] (Section 6.1). φ' is only slightly larger than φ , but has an algebraic structure that will be crucial to our verification process. Any legal assignment to the gates of φ (i.e. one that preserves the functionality of the gates of φ) can be transformed to a legal assignment to φ' . The important property of φ' is the following: If we encode an assignment to the gates of φ' using a specific sequence of Reed-Muller codewords (i.e. low degree polynomials), then the legality of the assignment can be locally verified (by reading a small random portion of the encoding). The encoding via low degree polynomials (and resulting local tests) is as in Harsha and Sudan [31] and is described in Section 6.2. Thus, our PCP verifier will essentially test (i) that the encoding of the purported satisfying assignment to φ' is formed of low degree polynomials, (this part will be done using the randomness-efficient low degree test of Ben Sasson *et al.* [12]); and (ii) that the assignment is legal. Section 6.3 describes the construction of the PCP verifier and Section 6.4 analyzes its properties. Most of the above results are implicit in the literature, but carefully abstracting the results and putting them together helps us in significantly reducing the randomness of the PCP verifier.

6.1 Well-structured Boolean circuits

The main problem with designing a randomness-efficient PCP verifier directly for CIRCUIT SATISFIABILITY is that we need to encode the assignment to all gates of the input circuit using certain Reed-Muller based codes, in such a way that will allow us to locally verify the legality of all gates of the circuit, using only the encoded assignment. In order to do this, we require the circuit to have a well-behaved structure (amenable to our specific encoding and verification demands). Of course, an arbitrary circuit does not necessarily have this structure, but luckily we have the technology to overcome this. More to the point, we can restructure any circuit into a well-behaved circuit that will suit our needs. The natural encoding (used e.g. in the Hadamard based PCP, Section A) incurs a quadratic blowup in size. To get over this problem, Polishchuk and Spielman [44, 52] introduced a different, more efficient restructuring process that embeds the input circuit into well-structured graphs known as de Bruijn graphs. Indeed, the blowup in circuit size using these circuits is merely by a logarithmic multiplicative factor, and their usefulness for the local verification of legal assignments will become evident later (in Section 6.2). As in Polishchuk and Spielman [44, 52], we embed the input circuit into wrapped de Bruijn graphs (see Definition 6.2). We use a slightly different definition of de Bruijn graphs, more convenient for our purposes, than that used in [44, 52]. However it can easily be checked that these two definitions yield isomorphic graphs. The main advantage with the de Bruijn graphs is that the neighborhood relations can be expressed very easily using simple bit-operations like cyclic-shifts and bit-flips. In [44, 52] the vertex set of these graphs is *identified* with a vector space. We instead work with a strict embedding of these graphs in a vector space where the vertices are a *strict* subset of the vector space. The benefit of both approaches is that the neighborhood functions can be expressed as affine functions (see Section 6.2 for more details). The reason for our approach will be explained at the end of Section 6.2.

Definition 6.2 *The wrapped de Bruijn graph $\mathcal{G}_{N,l}$ is a directed graph with l layers each with 2^N nodes which are represented by N -bit strings. The layers are numbered $0, 1, \dots, l-1$. The node represented by $v = (b_0, \dots, b_{i^*}, \dots, b_{N-1})$ in layer i has edges pointing to the nodes represented by $\Gamma_{i,0}(v) = (b_0, \dots, b_{i^*}, \dots, b_{N-1})$ and $\Gamma_{i,1}(v) = (b_0, \dots, b_{i^*} \oplus 1, \dots, b_{N-1})$ in layer $(i+1)$ modulo l , where i^* is i modulo N and $a \oplus b$ denotes the sum of a and b modulo 2.*

See Figure 1 for an example.

We now describe how to embed a circuit into a wrapped de Bruijn graph (see Figure 2 for a simple example). Given a circuit C with n gates (including both input and output gates), we associate with it a wrapped de Bruijn graph $\mathcal{G}_{N,l}$ where $N = \log n$ and $l = 5N = 5 \log n$. We then associate the nodes in layer 0 with the gates of the circuit. Now, we wish to map each wire in the circuit to a path in $\mathcal{G}_{N,l}$ between the corresponding nodes of layer 0. Standard packet-routing techniques (see [39]) can be used to show that if the number of layers l is at least $5N$ then such a routing can be done with edge-disjoint paths. (Recall that we work with circuits whose fan-in and fan-out are 2.)

Thus, we can find “switches” for each of the nodes in layers $1, \dots, l-1$ of the graph such that the output of each gate (i.e., node in layer 0) is routed to the input of the gates that require it. Each node has two inputs and two outputs, and thus there is a constant number of switches routing incoming edges to outgoing ones (See Figure 3). For nodes in layer 0, instead of specifying a switch, we specify the functionality of the Boolean gate associated to that node in the circuit (e.g. AND, OR, PARITY, NOT, INPUT, OUTPUT). Actually unary gates (such as NOT and OUTPUT) have two forms (NOT, NOT’, OUTPUT, OUTPUT’) in order to specify which of the two incoming edges in the de Bruijn graph to use.

Figure 1: The wrapped de Bruijn graph $\mathcal{G}_{3,3}$. Notice the first and last layer are the same.

This specifies the embedding of the input circuit into a well-structured circuit (based on a de Bruijn graph). More precisely, let $\mathcal{C} = \{\text{Type of switching actions}\} \cup \{\text{Type of Boolean gates}\}$ be the set of allowable gates of the well-structured circuit (See Figure 3). Given a circuit on n gates, we can construct, in polynomial time, a wrapped de Bruijn graph $\mathcal{G}_{N,l}$ (where $N = \log n$ and $l = 5 \log N$) and l functions $T_0, T_1, \dots, T_{l-1} : \{0, 1\}^N \rightarrow \mathcal{C}$ where each function T_i is a specification of the gates of layer i (i.e. a specification of the switching action or Boolean functionality).

We now demonstrate how to translate a proof that a circuit is satisfiable into an assignment that satisfies the embedded circuit. A proof that a circuit is satisfiable consists of an assignment of 0's and 1's to the inputs and the gates of the circuit such that each gate's output is consistent with its inputs and the output gate evaluates to 1. The corresponding assignment to the embedded circuit consists of an assignment of 0's and 1's to the edges entering and leaving the nodes of the wrapped de Bruijn graph that is consistent with the functionality of the gates (in layer 0) and the switching actions of the nodes (in the other layers). Since we are assigning values to nodes of the embedded graph (and not their edges), the assignment actually associates a 4-tuple of 0's and 1's to each of the nodes in the graph indicating the value carried by the four edges incident at that node (two incoming and two outgoing). More formally, the embedded assignment is given by a set of l functions $A_0, A_1, \dots, A_{l-1}$ where each function $A_i : \{0, 1\}^N \rightarrow \{0, 1\}^4$ specifies the values carried by the 4 edges incident at that vertex.

Figure 2: Embedding of a circuit into $\mathcal{G}_{3,3}$. In this example all paths between nodes at 0 layer are vertex disjoint. For general circuits we merely need edge disjoint paths.

We now list the constraints on the embedded circuit that assure us that the only legal assignments are ones that correspond to legal satisfying assignments of the original circuit, i.e. assignments that correctly propagate along the edges of the circuit, correctly compute the value of every gate and produce a 1 at the output gate.

Definition 6.3 *The assignment constraints for each node of the well-structured circuit require:*

- the two outgoing values at the node are propagated correctly to the incoming values of its neighbors at the next level,
- for nodes at layers $\neq 0$, the two outgoing values have the unique values dictated by the incoming values and the switching action,
- for non-OUTPUT nodes in layer 0, both outgoing values equal the unique value dictated by the gate functionality and the incoming values (the INPUT functionality merely requires that the two outgoing values are equal to each other)
- for nodes in layer 0 with an OUTPUT functionality, the appropriate incoming value equals 1

Let $\psi : \mathcal{C} \times (\{0, 1\}^4)^3 \rightarrow \{0, 1\}$ be the boolean function such that $\psi(t, a, a_0, a_1) = 0$ iff a node whose T -gate is t , A -assignment is a , and whose neighbors in the next layer have A -assignments a_0 and a_1 respectively, satisfies the aforementioned assignment constraints.

Observe that the definition of ψ is independent of N , the assignments A_i and gates T_i . By definition, the assignment $A = (A_0, \dots, A_{l-1})$ is legal for an embedded circuit defined by $T_0, \dots, T_{l-1}$ if and only if for every layer i and every node v in layer i ,

$$\psi\left(T(v), A(v), A(\Gamma_{i,0}(v)), A(\Gamma_{i,1}(v))\right) = 0.$$

We are now ready to formally define the well-structured circuit satisfiability problem.

Figure 3: Some gates of a well-structured circuit. Gates 1–2 are switching gates, and gate 3 sits in layer 0 and computes the parity (xor) function.

Definition 6.4 *The problem STRUCTURED-CKTSAT has as its instances $\langle \mathcal{G}_{N,l}, \{T_0, T_1, \dots, T_{l-1}\} \rangle$ where $\mathcal{G}_{N,l}$ is a wrapped de Bruijn graph with l layers and $T_i : \{0, 1\}^N \rightarrow \mathcal{C}$ is a specification of the node types of layer i of the graph (T_i 's are specified by a table of values).*

$\langle \mathcal{G}_{N,l}, \{T_0, \dots, T_{l-1}\} \rangle \in \text{STRUCTURED-CKTSAT}$ if there exists a set of assignments $A_0, A_1, \dots, A_{l-1}$ where $A_i : \{0, 1\}^N \rightarrow \{0, 1\}^4$ is an assignment to the nodes of layer i of $\mathcal{G}_N$ such that for all layers i and all nodes v in layer i ,

$$\psi \left(T(v), A(v), A(\Gamma_{i,0}(v)), A(\Gamma_{i,1}(v)) \right) = 0.$$

The above discussion also demonstrates the existence of a reduction from CKTSAT to STRUCTURED-CKTSAT which does not blow up the length of the target instance by more than a logarithmic multiplicative factor.

Proposition 6.5 *There exists a polynomial time reduction $\mathcal{R}$ from CKTSAT to STRUCTURED-CKTSAT such that for any circuit C , it holds that $C \in \text{CKTSAT}$ if and only if $\mathcal{R}(C) \in \text{STRUCTURED-CKTSAT}$. Moreover, if C is a circuit of size n , then $\mathcal{R}(C) = \langle \mathcal{G}_{N,l}, \{T_0, \dots, T_{l-1}\} \rangle$ where $N = \lceil \log n \rceil$ and $l = 5N$.*

Remark 6.6 *The above reduction though known to take polynomial time (via routing techniques) is not known to be of almost linear time.*

Remark 6.7 *We observe that if C is a satisfiable circuit, then any set of assignments $A_0, \dots, A_l$ proving that the reduced instance $\mathcal{R}(C) = \langle \mathcal{G}_{N,l}, \{T_0, \dots, T_{l-1}\} \rangle$ is a YES instance of STRUCTURED-CKTSAT contains within it a satisfying assignment to the circuit C . Specifically, let I be the set of nodes in layer 0 that have gate functionality INPUT associated with them. Then the assignment A_0 restricted to the set of nodes I (i.e., $A_0|_I$) contains a satisfying assignment. More precisely, the satisfying assignment is obtained by concatenating the third bit (i.e., first outgoing bit) of $A_0|_i \in \{0, 1\}^4$ for all $i \in I$. Conversely, every satisfying assignment w to C can be extended to $A_0, \dots, A_{l-1}$ such that $A_0|_I$ contains w . This is done by computing the values of all gates in the computation of $C(w)$, setting the outgoing bits of A_0 according to these values, and routing them throughout*

Figure 4: Example of legal and illegal assignments. The two vertices on the left are the inputs (at layer $i - 1$) to a gate at layer i . Recall that assignments evaluate each incoming and outgoing edge of a gate.

$\mathcal{G}_{N,l}$ according to the switching actions to obtain $A_1, \dots, A_{l-1}$ and the incoming bits of A_0 . This observation will be vital while constructing PCPs of proximity (see Section 7).

Remark 6.8 Suppose the input circuit C is a linear circuit, in the sense that all gates are INPUT, OUTPUT, or PARITY gates, and the OUTPUT gates test for 0 rather than 1 (See Definition 8.13). Then it can be verified that the transformation mapping satisfying assignments w of C to legal assignments $A_0, \dots, A_{l-1}$ of $\mathcal{R}(C)$ is $\text{GF}(2)$ -linear. The reason is that each gate in the computation of $C(w)$ is a $\text{GF}(2)$ -linear function of w . These remarks will be used in the coding applications, to obtain linear codes (see Section 8.4 for more information).

6.2 Arithmetization

In this section, we construct an algebraic version of STRUCTURED-CKTSAT by arithmetizing it along the lines of Harsha and Sudan [31]. The broad overview of the arithmetization is as follows: We embed the nodes in each layer of the wrapped de Bruijn graph $\mathcal{G}_{N,l}$ in a vector space and extend the gate specifications and assignments to low-degree polynomials over this space. Finally, we express the assignment constraints (Definition 6.3) as a pair of polynomial identities satisfied by these polynomials.

First for some notation. Let m be a parameter. Set h such that $h = N/m$ where 2^N is the number of nodes in each layer of the de Bruijn graph. Choose a finite extension field F of $\text{GF}(2)$ of size roughly $c_F m^2 2^h = c_F m^2 2^{N/m}$ where c_F is a suitably large constant to be specified later. Specifically, take $F = \text{GF}(2)^f$ for $f = h + 2 \log m + \log c_F$. Let $\{e_0, e_1, \dots, e_{f-1}\}$ be a basis of F over $\text{GF}(2)$. Set H to be a subspace of $\text{GF}(2)^f$ (and hence a subset of F) spanned by $\{e_0, \dots, e_{h-1}\}$. Note that H^m is a subset of the space F^m . Furthermore, $|H^m| = 2^N$. Hence, we can embed each layer of the graph $\mathcal{G}_{N,l}$ in F^m by identifying the node $v = (b_0, \dots, b_{N-1}) \in \{0, 1\}^N$ with the element $(b_0 e_0 + \dots + b_{h-1} e_{h-1}, b_h e_0 + \dots + b_{2h-1} e_{h-1}, \dots, b_{(m-1)h} e_0 + \dots + b_{mh-1} e_{h-1})$ of H^m . Henceforth, we use both representations (N -bit string and element of H^m) interchangeably. The representation will be clear from the context.

Any assignment $S : H^m \rightarrow F$ can be interpolated to obtain a polynomial $\tilde{S} : F^m \rightarrow F$ of degree at most $|H|$ in each variable (and hence a total degree of at most $m|H|$) such that $\tilde{S}|_{H^m} = S$ (i.e.,

the restriction of $\tilde{S}$ to H^m coincides with the function S). Conversely, any polynomial $\tilde{S} : F^m \rightarrow F$ can be interpreted as an assignment from H^m to F by considering the function restricted to the sub-domain H^m .

Recall that $\mathcal{C}$ and $\{0,1\}^4$ are the set of allowable gates and assignments given by the gate functions T_i and assignments A_i in the STRUCTURED-CKTSAT problem. We identify $\mathcal{C}$ with some fixed subset of F and we identify $\{0,1\}^4$ with the set of elements spanned by $\{e_0, e_1, e_2, e_3\}$ over $\text{GF}(2)$. With this identification, we can view the assignments A_i and gates T_i as functions $A_i : H^m \rightarrow F$ and $T_i : H^m \rightarrow F$ respectively. Furthermore, we can interpolate these functions, as mentioned above, to obtain polynomials $\tilde{A}_i : F^m \rightarrow F$ and $\tilde{T}_i : F^m \rightarrow F$ of degree at most $m|H|$ over F .

We now express the neighborhood functions of the graph in terms of affine functions over F^m . This is where the nice structure of the wrapped de Bruijn graph will be useful. For any positive integer i , define affine transformations $\tilde{\Gamma}_{i,0}, \tilde{\Gamma}_{i,1} : F^m \rightarrow F^m$ as follows: $\tilde{\Gamma}_{i,0}$ is the identity function. For $\tilde{\Gamma}_{i,1}$, first let $t = \lfloor i/h \rfloor \bmod m$ and $u = i \bmod h$. Then $\tilde{\Gamma}_{i,1}(z_0, \dots, z_{m-1}) = (z_0, \dots, z_{t-1}, z_t + e_u, z_{t+1}, \dots, z_{m-1})$.²⁰ It can be checked from the above definition that for any layer i and node x in layer i (which we view as a point in H^m), we have $\tilde{\Gamma}_{i,j}(x) = \Gamma_{i,j}(x)$ for $j = 0, 1$. In other words, $\tilde{\Gamma}$ is an extension of the neighborhood relations of the graph $\mathcal{G}_{N,l}$ over F^m .

Finally, we now express the assignment constraints (Definition 6.3) as polynomial identities. The first of these identities checks that the assignments given by the assignment polynomial $\tilde{A}_i$ are actually elements of $\{0,1\}^4$ for points in H^m . For this purpose, let $\psi_0 : F \rightarrow F$ be the univariate polynomial of degree 2^4 given by

$$\psi_0(z) = \prod_{\alpha \in \{0,1\}^4} (z - \alpha) \quad (1)$$

This polynomial satisfies $\psi_0(z) = 0$ iff $z \in \{0,1\}^4$ (recall we identified $\{0,1\}^4$ with a subset of F spanned by $e_0, \dots, e_3$). We check that $\psi_0(\tilde{A}_i(x)) = 0$ for all $x \in H^m$ and all layers i . We then arithmetize the rule ψ (from Definition 6.3) to obtain a polynomial $\psi_1 : F^4 \rightarrow F$. In other words, $\psi_1 : F^4 \rightarrow F$ is a polynomial such that $\psi_1(t, a, a_0, a_1) = \psi(t, a, a_0, a_1)$ for all $(t, a, a_0, a_1) \in \mathcal{C} \times (\{0,1\}^4)^3$. The degree of ψ_1 can be made constant, because $|\mathcal{C}|$ and $|\{0,1\}^4|$ are constant.²¹ The two polynomial identities we would like to check are $\psi_0(\tilde{A}_i(x)) = 0$ and $\psi_1(\tilde{T}_i(x), \tilde{A}_i(x), \tilde{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \tilde{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))) = 0$ for all $x \in H^m$. For notational convenience, we express these two conditions together as a pair of polynomials $\psi' = (\psi_0, \psi_1) : F^4 \rightarrow F^2$ such that $\psi'(x_1, x_2, x_3, x_4) = (\psi_0(x_2), \psi_1(x_1, x_2, x_3, x_4))$. Let κ be the maximum of the degree of these two polynomials. In order to make these polynomial identities sufficiently redundant, we set c_F to be a sufficiently large constant (say 100) such that $\kappa m^2 2^h / |F|$ is low.

We have thus reduced STRUCTURED-CKTSAT to an algebraic consistency problem, which we shall call the AS-CKTSAT problem (short for ALGEBRAIC-STRUCTURED-CKTSAT)²².

Definition 6.9 *The promise problem $\text{AS-CKTSAT} = (\text{AS-CKTSAT}^{\text{YES}}, \text{AS-CKTSAT}^{\text{NO}})$ has as its instances $\langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle$ where F is a finite extension field of $\text{GF}(2)$ (i.e.,*

²⁰ An alternate description of $\tilde{\Gamma}_{i,1}$ is as follows: Since $F = \text{GF}(2)^f$, we can view F^m as mf -dimensional space over $\text{GF}(2)$. Hence, any vector $(z_0, \dots, z_{m-1})$ can be written as $(b_{0,0}, \dots, b_{0,f-1}, b_{1,0}, \dots, b_{1,f-1}, \dots, b_{m-1,0}, \dots, b_{m-1,f-1})$. Furthermore, we note that for any vector $(z_0, \dots, z_{m-1})$ in H^m , $b_{r,s} = 0$ for all $s \geq h$ and all r . It can now be checked that $\tilde{\Gamma}_{i,1}$ is the affine transformation that flips the bit $b_{t,u}$ where $t = \lfloor i/h \rfloor \bmod m$ and $u = i \bmod h$.

²¹ Notice that we do not specify ψ_1 uniquely at this stage. Any choice of a constant-degree polynomial will work in this section, but to enforce linearity, we will use a somewhat non-standard choice in Section 8.4. Specifically, we argue that if C is a linear circuit, then ψ_1 can be picked to be $\text{GF}(2)$ -linear transformations over $\text{GF}(2)$, and we point out that ψ_0 is a $\text{GF}(2)$ -linear transformation. For more details see Section 8.4.

²² AS-CKTSAT is actually a promise problem.

$F = \text{GF}(2)^f$ for some f), H a $\text{GF}(2)$ -linear subspace of F and $\tilde{T}_i : F^m \rightarrow F$, for $i = 0, \dots, l-1$, a sequence of polynomials of degree d , where $|H| = n^{1/m}$, $d = m \cdot |H|$, and $F = c_F \cdot md$. The field F is specified by an irreducible polynomial $p(x)$ of degree f over $\text{GF}(2)$, H is taken to be spanned by the first $h = \log |H|$ canonical basis elements, and each of the polynomials $\tilde{T}_i$ is specified by a list of coefficients.

- $\langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle \in \text{AS-CKTSAT}^{\text{YES}}$ if there exist a sequence of degree d polynomials $\tilde{A}_i : F^m \rightarrow F$, $i = 0, \dots, l-1$ such that for all $i = 0, \dots, l-1$ and all $x \in H^m$,

$$\psi' \left(\tilde{T}_i(x), \tilde{A}_i(x), \tilde{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \tilde{A}_{i+1}(\tilde{\Gamma}_{i,1}(x)) \right) = (0, 0)$$

- $\langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle \in \text{AS-CKTSAT}^{\text{NO}}$ if for all functions $\tilde{A}_i : F^m \rightarrow F$, $i = 0, \dots, l-1$ there exists an $i = 0, \dots, l-1$ and $x \in H^m$ such that,

$$\psi' \left(\tilde{T}_i(x), \tilde{A}_i(x), \tilde{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \tilde{A}_{i+1}(\tilde{\Gamma}_{i,1}(x)) \right) \neq (0, 0)$$

where $\tilde{\Gamma}_{i,j}$'s and ψ' are as defined earlier. (Recall that the $\tilde{\Gamma}$'s are linear function while ψ' represents a pair of polynomials of degree at most κ .)

From the above discussion we have the following reduction from STRUCTURED-CKTSAT to AS-CKTSAT.

Proposition 6.10 *There exists a polynomial-time computable function $\mathcal{R}$ mapping any instance $\mathcal{I} = \langle \mathcal{G}_{N,l}, \{T_0, T_1, \dots, T_{l-1}\} \rangle$ of STRUCTURED-CKTSAT and parameter $m \leq \log n / \log \log n$ (where $n = |\mathcal{I}|$) to an instance $\mathcal{R}(\mathcal{I}, 1^m)$ of AS-CKTSAT such that*

$$\begin{aligned} \mathcal{I} \in \text{STRUCTURED-CKTSAT} &\implies \mathcal{R}(\mathcal{I}, 1^m) \in \text{AS-CKTSAT}^{\text{YES}} \\ \mathcal{I} \notin \text{STRUCTURED-CKTSAT} &\implies \mathcal{R}(\mathcal{I}, 1^m) \in \text{AS-CKTSAT}^{\text{NO}} \end{aligned}$$

Moreover, if $\mathcal{R}(\mathcal{I}, 1^m) = \langle 1^{n'}, 1^{m'}, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l'-1}\} \rangle$, then $n' = 2^N$ (the number of nodes in each layer of the de Bruijn graph $\mathcal{G}_{N,l}$), $m' = m$, and $l' = l$ (the number of layers in the de Bruijn graph).

Combining Propositions 6.5 and 6.10, we have the following.

Proposition 6.11 *There exists a polynomial-time computable function $\mathcal{R}$ mapping any circuit C and parameter $m \leq \log n / \log \log n$ (where $n = |C|$) to an instance $\mathcal{R}(C, 1^m)$ of AS-CKTSAT such that $C \in \text{CKTSAT} \iff \mathcal{R}(C, 1^m) \in \text{AS-CKTSAT}$.*

Moreover, if C is a circuit of size n then $\mathcal{R}(C, 1^m) = \langle 1^{n'}, 1^{m'}, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l'-1}\} \rangle$, where $n' = \Theta(n)$, $m' = m$, and $l' \leq 5 \log n'$. Thus, $|\mathcal{R}(C, 1^m)| = O((c_F m^2)^m \log n) \cdot |C|$.

Remark 6.12 *Following Remark 6.7, if C is a satisfiable circuit, then any set of polynomials $\tilde{A}_0, \dots, \tilde{A}_{l-1}$ proving that the reduced instance $\mathcal{R}(C, 1^m) = \langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle$ is a YES instance of AS-CKTSAT contain within it a satisfying assignment to the circuit C . Specifically, the set I (of layer-0 nodes with INPUT functionality in $\mathcal{G}_{N,l}$) from Remark 6.7 can now be viewed as a subset $I \subseteq H^m$. Then the polynomial $\tilde{A}_0 : F^m \rightarrow F$ restricted to the set I (i.e., $\tilde{A}_0|_I$) contains a satisfying assignment (again as a concatenation of third-bits). Conversely, every satisfying assignment w to C can be extended to a set of polynomials $\tilde{A}_0, \dots, \tilde{A}_{l-1}$ such that $\tilde{A}_0|_I$ contains w . This is done by taking low-degree extensions of the functions $A_0, \dots, A_{l-1}$ from Remark 6.7.*

Remark 6.13 Following Remark 6.8, if C is a linear circuit, then the mapping of satisfying assignments w of C to polynomials $\tilde{A}_0, \dots, \tilde{A}_{l-1}$ satisfying $\mathcal{R}(C)$ is $\text{GF}(2)$ -linear. This is due to Remark 6.8, the association of $\{0, 1\}^4$ with the linear space spanned by $\{e_0, e_1, e_2, e_3\}$ in F , and from the fact that the interpolation from A_i to $\tilde{A}_i$ is F -linear and hence $\text{GF}(2)$ -linear. For more information see Section 8.4.

Comment: Recall that the arithmetization was obtained by considering low-degree extensions over F^m of functions from H^m to H . If H were a subfield of the field F this step would have caused a quadratic blow-up, and we avoid this problem by not insisting that H be a field. In [44, 52], H is a field and $F = H^2$ is an extension of it, but the PCP system refers only to a $O(|H|)$ -sized of F . We cannot take this approach because we will be using a total low-degree test, which needs to refer to the entire vector space F^m . In contrast, in [44, 52] an individual low-degree test is used, which can work with a subset of F^m .

6.3 The PCP verifier

We design a PCP verifier for CKTSAT via the reduction to AS-CKTSAT based on the randomness-efficient low-degree tests of Ben-Sasson *et al.* [12]. Given a circuit C , the verifier reduces it to an instance of the problem AS-CKTSAT (Proposition 6.11). The proof consists of a sequence of oracles $\tilde{A}_i : F^m \rightarrow F$ for $i = 0, \dots, l-1$ and an auxiliary sequence of oracles $P_{i,j} : F^m \rightarrow F^2$ for $i = 0, \dots, l-1$ and $j = 0, \dots, m$. For each i and j , we view the auxiliary oracle $P_{i,j} : F^m \rightarrow F^2$ as a pair of functions $P_{i,j}^{(0)} : F^m \rightarrow F$ and $P_{i,j}^{(1)} : F^m \rightarrow F$ (i.e., $P_{i,j}(x) = (P_{i,j}^{(0)}(x), P_{i,j}^{(1)}(x))$). This auxiliary sequence of oracles helps the verifier to check that the functions $\tilde{A}_i$ satisfy condition ψ' (see Definition 6.9).

The verifier expects the first subsequence of auxiliary oracles $P_{i,0}(\cdot)$ for $i = 0, \dots, l-1$, to satisfy the following relation:

$$P_{i,0}(x) = \psi' \left(\tilde{T}_i(x), \tilde{A}_i(x), \tilde{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \tilde{A}_{i+1}(\tilde{\Gamma}_{i,1}(x)) \right) \quad \forall x \in F^m \quad (2)$$

Furthermore, it expects $P_{i,0}(x) = 0$ for every $x \in H^m$. Indeed, by Definition 6.9, we have:

- Lemma 6.14** 1. If $\langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle$ is a YES instance of AS-CKTSAT, satisfied by polynomials $\tilde{A}_0, \dots, \tilde{A}_{l-1}$, and $P_{0,0}, \dots, P_{l-1,0}$ are defined according to Equation 2, then $P_{i,0}(x) = (0, 0)$ for all $x \in H^m$.
2. If $\langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle$ is a NO instance of AS-CKTSAT, then for any sequences of functions $\tilde{A}_0, \dots, \tilde{A}_{l-1}$, $P_{0,0}, \dots, P_{l-1,0}$, either Equation 2 fails to hold for some i or $P_{i,0}(x) \neq (0, 0)$ for some i and some $x \in H^m$.

Recalling that the degree of the constraint ψ' (see Definition 6.9) is at most κ and that the $\tilde{A}_i$'s are of degree at most $d = m \cdot |H|$, we observe that the $P_{i,0}$'s can be taken to be of degree at most κd in Part 1.

As mentioned above, the verifier now needs to check that the functions $P_{i,0}$ vanish on the set H^m . For this we use a “zero-propagation test”, based on the sum-check protocol of Lund *et al.* [40]. Specifically, the verifier expects the remaining set of auxiliary oracles $P_{i,j} = (P_{i,j}^{(0)}, P_{i,j}^{(1)})$ ($i = 0, \dots, l-1$ and $j = 1, \dots, m$) to satisfy the following relations: Let $H = \{h_0, \dots, h_{|H|-1}\}$ be

some fixed enumeration of the elements in H . For all $b \in \{0, 1\}$,

$$P_{i,j}^{(b)} \left(\underbrace{x_1, \dots, x_{j-1}}_{\substack{\text{---} \\ \text{---} \\ \text{---}}} , x_j, \underbrace{x_{j+1}, \dots, x_m}_{\substack{\text{---} \\ \text{---} \\ \text{---}}} \right) = \sum_{k=0}^{|H|-1} P_{i,j-1}^{(b)} \left(\underbrace{x_1, \dots, x_{j-1}}_{\substack{\text{---} \\ \text{---} \\ \text{---}}} , h_k, \underbrace{x_{j+1}, \dots, x_m}_{\substack{\text{---} \\ \text{---} \\ \text{---}}} \right) x_j^k, \quad (3)$$

$\forall (x_1, \dots, x_m) \in F^m$

These relations ensure that for all i and $j \geq 1$, $P_{i,j}(\cdot)$ vanishes on $F^j \times H^{m-j}$ iff the function $P_{i,j-1}(\cdot)$ vanishes on $F^{j-1} \times H^{m-j+1}$. In other words:

Lemma 6.15 $P_{i,j}^{(b)}|_{F^j \times H^{m-j}} \equiv 0 \iff P_{i,j-1}^{(b)}|_{F^{j-1} \times H^{m-j+1}} \equiv 0$.

Thus, for all i , $P_{i,m}$ vanishes on the entire space F^m iff $P_{i,0}$ vanishes on H^m . Also, as $P_{i,0}^{(b)}$ has degree at most κd in each variable, so does $P_{i,j}^{(b)}$ for each i and j . Hence, the degree of $P_{i,j}^{(b)}$ is at most κd .

Thus, the verifier needs to make the following checks

- **LOW-DEGREE TEST**

For $i = 0, \dots, l-1$ and $j = 0, \dots, m$, the sequence of functions $\tilde{A}_i$ are polynomials of degree at most $d = m \cdot |H|$ and the sequence of functions $P_{i,j}$ are pairs of polynomials of degree at most κd ,

- **EDGE-CONSISTENCY TEST**

For $i = 0, \dots, l-1$, the functions $P_{i,0}$ obey Equation (2),

- **ZERO PROPAGATION TEST**

For $i = 0, \dots, l-1$ and $j = 1, \dots, m$, the functions $P_{i,j}$ satisfy Equation (3),

- **IDENTITY TEST**

For $i = 0, \dots, l-1$, the functions $P_{i,m}$ are identically zero on the entire domain F^m .

The Low-Degree test in most earlier construction of PCP verifiers is performed using the “line-point” test. The “line-point” low degree test first chooses a random line, a random point on this line and checks if the restriction of the function to the line (given by a univariate polynomial) agrees with the value of the function at the point. A random line l is typically chosen by choosing two random points $x, y \in F^m$ and setting $l = l_{x,y} = \{x + ty | t \in F\}$. However, this requires $2m \log |F|$ bits of randomness which is too expensive for our purposes. We save on randomness by using the low-degree test of Ben-Sasson *et al.* [12] based on small-biased spaces (see Section B for more details). The low-degree test of [12] uses pseudorandom lines instead of totally random lines in the following sense: The pseudorandom line $l = l_{x,y}$ is chosen by choosing the first point x at random from F^m , while the second point y is chosen from a λ -biased subset S_λ of F^m . This needs only $\log |S_\lambda| + \log |F|^m$ bits of randomness. We further save on randomness by the use of *canonical lines*²³. Consider any pseudorandom line $l = l_{x,y}$ where $x \in F^m$ and $y \in S_\lambda$. We observe that for every $x' \in l$, we have $l_{x',y} = l_{x,y}$. In other words, $|F|$ different choices of random bits leads to the same line $l_{x,y}$. We prevent this redundancy by representing each line in a canonical manner. A canonical line is chosen by first choosing a random point y from the λ -biased set S_λ . We view this

²³It is to be noted that the canonical representation of lines has been used either implicitly or explicitly in the soundness analysis of all earlier uses of the Low-Degree Test. However, this is the first time that the canonical representation is used to save on the number of random bits.

y as specifying the direction (i.e., slope) of the line. This direction partitions the space F^m into $|F|^{m-1}$ parallel lines (each with direction y). We enumerate these lines arbitrarily and select one of them uniformly at random. Thus, choosing a random canonical line costs only $\log |S_\lambda| + \log |F|^{m-1}$ bits of randomness. A further point to be noted is that we perform a “line” test instead of the regular “line-point” test: The test queries the function for all points along the canonical line $l_{x,y}$ and verifies that the restriction of the function to this line is a low-degree polynomial.

Having performed the low-degree test (i.e., verified that the polynomials $\tilde{A}_i$ ’s and $P_{i,j}$ ’s are close to low-degree polynomials), the verifier then performs each of the NODE-CONSISTENCY TEST, ZERO PROPAGATION TEST, and IDENTITY TESTS by choosing a suitable small-sized sample in the entire space and checking if the corresponding condition is satisfied on that sample. For the ZERO PROPAGATION TEST indeed the natural sample is an axis-parallel line. For the EDGE-CONSISTENCY TEST and IDENTITY TEST, the sample we use is any set of $|F|$ points selected from a partition of F^m into $|F|^{m-1}$ equal sets.

We are now ready to formally describe the PCP verifier for CKTSAT. We parameterize the PCP verifier in terms of m , the number of dimensions in our intermediate problem AS-CKTSAT, and λ , the parameter of the λ -biased sets of F^m required for the low-degree tests of Ben-Sasson *et al.* [12]. We rely on the fact that λ -biased subsets of F^m of size at most $\text{poly}(\log |F|^m, 1/\lambda)$ can be constructed efficiently [42, 2].

PCP-VERIFIER $_{m,\lambda}^{\tilde{A}_i, P_{i,j}; i=0,\dots,l-1; j=0,\dots,m}(C)$.

1. Use Proposition 6.11 to reduce the instance C of CKTSAT, using parameter m , to an instance $\langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle$ of AS-CKTSAT, and set $d = m \cdot |H|$.

Notation: We let $S_\lambda \subset F^m$ be a λ -biased set of size at most $\left(\frac{\log |F|^m}{\lambda}\right)^2$ [2]. Let $F^m = \bigsqcup_{\eta=1}^{|F|^{m-1}} U_\eta$ and $F^m = \bigsqcup_{\eta=1}^{|F|^{m-1}} V_\eta$ be two arbitrary partitions of the space F^m into $|F|$ -sized sets each.

2. Choose a random string R of length $\log(|S_\lambda| \cdot |F|^{m-1})$. [Note: We reuse R in all tests, but only the LOW-DEGREE TEST utilizes the full length of R .]

3. LOW-DEGREE TEST

Use random string R to determine a random canonical line $\mathcal{L}$ in F^m using the λ -biased set S_λ .

For $i = 0, \dots, l-1$,

Query oracle $\tilde{A}_i$ on all points along the line $\mathcal{L}$ and reject if the restriction $\tilde{A}_i$ to $\mathcal{L}$ is not a (univariate) polynomial of degree at most d .

For $i = 0, \dots, l-1$, $j = 0, \dots, m$, and $b \in \{0, 1\}$,

Query oracle $P_{i,j}^{(b)}$ on all points along the line $\mathcal{L}$ and reject if the restriction of $P_{i,j}^{(b)}$ to $\mathcal{L}$ is not a (univariate) polynomial of degree at most κd .

4. EDGE-CONSISTENCY TEST

Use the random string R to determine a random set U_η of the partition $F^m = \bigsqcup_{\eta=1}^{|F|^{m-1}} U_\eta$.

For $i = 0, \dots, l-1$,

For all $x \in U_\eta$, query $P_{i,0}(x)$, $\tilde{A}_i(x)$, $\tilde{A}_{i+1}(\tilde{\Gamma}_{i,0}(x))$ and $\tilde{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))$ and reject if Equation (2) is not satisfied.

5. ZERO PROPAGATION TEST

For $i = 0, \dots, l-1$, $j = 1, \dots, m$, and $b \in \{0, 1\}$,

Use random string R to determine a random j th axis-parallel line in F^m of the form $\mathcal{L} = \{(a_1, \dots, a_{j-1}, X, a_{j+1}, \dots, a_m) : X \in F\}$. Query $P_{i,j-1}^{(b)}$ and $P_{i,j}^{(b)}$ along all the points in $\mathcal{L}$. Reject if either the restriction of $P_{i,j-1}^{(b)}$ or $P_{i,j}^{(b)}$ to $\mathcal{L}$ is not a univariate polynomial of degree at most κd or if any of the points on the line $\mathcal{L}$ violate Equation (3).

6. IDENTITY TEST

Use the random string R to determine a random set V_η of the partition $F^m = \bigsqcup_{\eta=1}^{|F|^{m-1}} V_\eta$. For $i = 0, \dots, l-1$,

For all $x \in V_\eta$, query $P_{i,m}(x)$. Reject if any of these $P_{i,m}(x)$ is not $(0, 0)$.

Accept if none of the above tests reject.

Remark 6.16

1. The LOW-DEGREE TEST requires $\log(|S_\lambda| \cdot |F|^{m-1})$ random bits to generate a canonical line in F^m using the λ -biased set, while each of the other tests require at most $\log(|F|^{m-1})$ bits of randomness. Hence, the string R suffices for each of the tests. For the settings of parameters we use, $\log(|S_\lambda| \cdot |F|^{m-1})$ is typically significantly smaller than $\log(|F|^m)$, which we would not be able to afford.
2. The EDGE-CONSISTENCY TEST and IDENTITY TEST in the “standard” sense are usually performed by selecting a random point in the space F^m and checking whether the corresponding condition is satisfied. However, we state these tests in a “non-standard” manner using partitions of the space F^m into $|F|$ sized tests so that these tests can easily be adapted when we construct the robust PCP verifier (see Section 8). The non-standard tests are performed in the following manner: Choose a random set in the partition and perform the standard test for each point in the set. At present, we can work with any partition of F^m , however we will later need specific partitions to get “robustness”.

6.4 Analysis of the PCP verifier

We now analyze the PCP verifier above. The analysis below assumes that the parameters satisfy $m \leq \log n / \log \log n$ and $\lambda \leq 1/c \log n$ for a sufficiently large constant c . Theorem 6.1 can be deduced by setting $\lambda = 1/c \log n$.

Complexity: The PCP VERIFIER makes $O(lm|F|) = O(m^3 n^{1/m} \log n)$ queries each of which expects as an answer an element of F or F^2 (i.e., a string of length $O(\log |F|)$). Hence, the total (bit) query complexity is $O(lm|F| \log |F|) = O(lm \cdot c_F m^2 n^{1/m} \log(c_F m^2 n^{1/m}))$. Recalling that $l = 5 \log n$, this quantity is at most $O(m^2 n^{1/m} \log^2 n)$ for $m \leq \log n$. For the decision complexity, we note that the main computations required are (a) testing whether a function is a low-degree univariate polynomial over F (for LOW-DEGREE TEST and ZERO PROPAGATION TEST), (b) evaluating ψ' on $|F|$ quadruples of points (for EDGE-CONSISTENCY TEST), and (c) univariate polynomial interpolation and evaluation (for testing (3) in ZERO PROPAGATION TEST). We now argue that each of these can be done with a nearly linear ($\tilde{O}(|F|)$) number of operations over F , yielding a nearly linear ($\tilde{O}(q)$) decision complexity overall. Each evaluation of ψ' can be done with a constant number of F -operations because ψ' is of constant degree. Polynomial interpolation and evaluation can be done with a nearly linear number of F -operations by [50, 49], and testing whether a function is of low degree reduces to polynomial interpolation (interpolate to represent as a polynomial of degree $|F| - 1$ and check that the high-degree coefficients are zero). Each F -operations can be

done with $\tilde{O}(\log |F|)$ bit-operations, using the polynomial multiplication algorithm of [50, 49] (over $\text{GF}(2)$).

The number of random bits used by the verifier is exactly $\log(|S_\lambda| \cdot |F|^{m-1})$. Let $n' = |F|^m$. Then $\log(|S_\lambda| \cdot |F|^{m-1}) = (1 - \frac{1}{m}) \log n' + \log \left(\text{poly} \left(\frac{\log n'}{\lambda} \right) \right) = (1 - \frac{1}{m}) \log n' + O(\log \log n') + O(\log(\frac{1}{\lambda}))$. Now, $n' = (c_F m^2)^m n$. Hence, $\log n' = \log n + 2m \log m + O(m)$ and $\log \log n' = \log \log n + O(\log m)$. Thus, the total randomness is at most $(1 - \frac{1}{m}) \log n + O(m \log m) + O(\log \log n) + O(\log(\frac{1}{\lambda}))$.

We summarize the above observations in the following proposition for future reference.

Proposition 6.17 *The randomness, query and decision complexities of the PCP-VERIFIER are $r = (1 - \frac{1}{m}) \log n + O(m \log m) + O(\log \log n) + O(\log(\frac{1}{\lambda}))$, $q = O(m^2 n^{1/m} \log^2 n)$ and $d = \tilde{O}(q)$ respectively.*

Completeness: If C is satisfiable, then the reduction reduces it to an YES instance of AS-CKTSAT. Then by definition there exist polynomials $\tilde{A}_i$ that satisfy constraint ψ' . Setting $P_{i,j}$ according to Equations (2) and (3), we notice that the verifier accepts with probability one.

Soundness: To prove the soundness, we need to prove that if C is not satisfiable then the verifier accepts with probability bounded away from 1. We will prove a stronger statement. Recall from Remark 6.12 that the function $\tilde{A}_0 : F^m \rightarrow F$ supposedly has the satisfying assignment embedded within it. Let $I \subset F^m$ be the set of locations in F^m that contains the assignment (i.e., $\tilde{A}_0|_I$ is supposedly the assignment).

Lemma 6.18 *There exists a constant c and a constant $0 < \varepsilon_0 < 1$ such that for all ε, m, λ satisfying $\varepsilon \leq \varepsilon_0$, $m \leq \log n / \log \log n$ and $\lambda \leq 1/c \log n$, the following holds. If the verifier accepts proof oracles $\{\tilde{A}_i\}$ and $\{P_{i,j}\}$ with probability greater than $1 - \varepsilon$, then $\tilde{A}_0$ is 4ε -close to some polynomial $\hat{A}_0$ of degree md such that $C(\hat{A}_0|_I) = 1$.*

Proof: Let α be the universal constant from Theorem B.4. Set $\varepsilon_0 = \min\{\alpha, \frac{1}{22}\}$. Let $d = m2^h$, and choose c_F to be a large enough constant such that $\kappa md / |F| = \kappa / c_F \leq \varepsilon_0$. Suppose each of the functions $\tilde{A}_i$ are 4ε -close to some polynomial of degree md and each of the functions $P_{i,j}^{(b)}$ is 4ε -close to some polynomial of degree κmd . If this were not the case, then by Theorem B.4 the LOW-DEGREE TEST accepts with probability at most $1 - \varepsilon$ for the polynomial that is 4ε -far. It can be verified that the parameters satisfy the requirements of Theorem B.4, for sufficiently large choices of the constants c_F and c and sufficiently small ε .

For each $i = 0, \dots, l-1$, let $\hat{A}_i : F^m \rightarrow F$ be the polynomial of degree at most md that is 4ε -close to $\tilde{A}_i$. Similarly, for each $i = 0, \dots, l-1, j = 0, \dots, m$ and $b \in \{0, 1\}$, let $\hat{P}_{i,j}^{(b)}$ be the polynomial of degree at most κmd that is 4ε -close to $P_{i,j}^{(b)}$. Such polynomials are uniquely defined since every two polynomials of degree κmd disagree in at least a $1 - \frac{\kappa md}{|F|} \geq 1 - \varepsilon_0 > 8\varepsilon$ fraction of points. As in the case of the $P_{i,j}$'s, let $\hat{P}_{i,j} : F^m \rightarrow F^2$ be the function given by $\hat{P}_{i,j}(x) = (\hat{P}_{i,j}^{(0)}(x), \hat{P}_{i,j}^{(1)}(x))$.

By hypothesis, $\tilde{A}_0|_I$ does not satisfy C . Then, by Lemmas 6.14 and 6.15, at least one of the following must hold.

(a) **There exists $i = 0, \dots, l-1$ and $b \in \{0, 1\}$ such that $\hat{P}_{i,m}^{(b)} \not\equiv 0$.**

Then for this i , the IDENTITY TEST fails unless a random set V_η is chosen such that for all $x \in V_\eta$, $P_{i,m}^{(b)}(x) = 0$. Hence, it must be the case that for all $x \in V_\eta$, either $P_{i,m}^{(b)}(x) \neq \hat{P}_{i,m}^{(b)}(x)$ or $\hat{P}_{i,m}^{(b)}(x) = 0$. Since the V_η 's form a partition of F^m , the probability of this occurring is

upper-bounded by the probability that a random $x \in F^m$ satisfies either $P_{i,m}^{(b)}(x) \neq \hat{P}_{i,m}^{(b)}(x)$ or $\hat{P}_{i,m}^{(b)}(x) = 0$. This probability is at most $4\varepsilon + \frac{\kappa md}{|F|} = 4\varepsilon + \frac{\kappa}{c_F} \leq 5\varepsilon_0$, where we use the fact that $\hat{P}_{i,m}^{(b)}$ is 4ε -close to $P_{i,m}^{(b)}$ and that a nonzero polynomial of degree κmd vanishes on at most a $\kappa md/|F|$ fraction of points.

(b) **There exists $i = 0, \dots, l-1$ such that $\hat{P}_{i,0}$, $\hat{A}_i$, and $\hat{A}_{i+1}$ do not obey Equation (2).**

In other words, $\hat{P}_{i,0}(x) \neq \psi'((\tilde{T}_i(x), \hat{A}_i(x), \hat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \hat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))))$. Then for this i , the EDGE-CONSISTENCY TEST fails unless a random partition U_η is chosen such that for all $x \in U_\eta$, $P_{i,0}(x) = \psi'((\tilde{T}_i(x), \hat{A}_i(x), \hat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \hat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))))$. Hence, it must be the case that for every $x \in U_\eta$, that one of the following holds:

$$\begin{aligned} P_{i,0}^{(0)}(x) \neq \hat{P}_{i,0}^{(0)}(x); \quad & P_{i,0}^{(1)}(x) \neq \hat{P}_{i,0}^{(1)}(x); \quad & \tilde{A}_i(x) \neq \hat{A}_i(x); \quad & \tilde{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)) \neq \hat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)); \\ \tilde{A}_{i+1}(\tilde{\Gamma}_{i,1}(x)) \neq \hat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x)); \quad & \hat{P}_{i,0}(x) = \psi'((\tilde{T}_i(x), \hat{A}_i(x), \hat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \hat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x)))) \end{aligned}$$

The probability of this happening is at most the probability that a random $x \in F^m$ satisfies these conditions, which is at most $5 \cdot 4\varepsilon + \frac{\kappa md}{|F|} \leq 21\varepsilon_0$.

(c) **For some $i = 0, \dots, l-1$, $j = 1, \dots, m$, and $b \in \{0, 1\}$, $\hat{P}_{i,j}^{(b)}$ does not obey Equation (3).**

In other words, $\hat{P}_{i,j}^{(b)}(\dots, x_j, \dots) \neq \sum_{k=1}^{|H|} \hat{P}_{i,j-1}^{(b)}(\dots, h_j, \dots) x_i^k$. Then, for this i, j , the ZERO PROPAGATION TEST rejects unless a random axis parallel line $\mathcal{L}$ is chosen such that both $P_{i,j}^{(b)}|_{\mathcal{L}}$ and $P_{i,j-1}^{(b)}|_{\mathcal{L}}$ are polynomials of degree at most κd and for every $x \in \mathcal{L}$, $P_{i,j}^{(b)}(\dots, x, \dots) = \sum_{k=0}^{|H|-1} P_{i,j-1}^{(b)}(\dots, h_k, \dots) x^k$. Suppose we have that for all $x \in \mathcal{L}$, $P_{i,j}^{(b)}(x) = \hat{P}_{i,j}^{(b)}(x)$ and $P_{i,j-1}^{(b)}(x) = \hat{P}_{i,j-1}^{(b)}(x)$. Then, it must be the case that for all $x \in \mathcal{L}$, $\hat{P}_{i,j}^{(b)}(\dots, x, \dots) = \sum_{k=0}^{|H|-1} \hat{P}_{i,j-1}^{(b)}(\dots, h_k, \dots) x^k$. Since the axis-parallel lines cover F^m uniformly, the probability of this occurring is at most the probability of a random $x \in F^m$ satisfying this condition which is at most $\frac{\kappa md}{c_F} \leq \varepsilon$. The probability that both $P_{i,j}^{(b)}|_{\mathcal{L}}$ and $P_{i,j-1}^{(b)}|_{\mathcal{L}}$ are polynomials of degree κd and either $P_{i,j}^{(b)}|_{\mathcal{L}} \neq \hat{P}_{i,j}^{(b)}|_{\mathcal{L}}$ or $P_{i,j-1}^{(b)}|_{\mathcal{L}} \neq \hat{P}_{i,j-1}^{(b)}|_{\mathcal{L}}$ is at $2 \cdot 4\varepsilon/(1 - \varepsilon_0) \leq 9\varepsilon_0$, since $P_{i,j}^{(b)}$ and $P_{i,j-1}^{(b)}$ are 4ε -close to $\hat{P}_{i,j}^{(b)}$ and $\hat{P}_{i,j-1}^{(b)}$ respectively, and any two distinct polynomials of degree κmd disagree on at least a $1 - \kappa md/|F| \geq 1 - \varepsilon_0$ fraction of points. Hence, the ZERO PROPAGATION TEST accepts with probability at most $10\varepsilon_0$.

We thus have that the verifier accepts with probability at most $\max\{1 - \varepsilon, 5\varepsilon_0, 21\varepsilon_0, 10\varepsilon_0\} = 1 - \varepsilon$. ■

Proof (of Theorem 6.1): Theorem 6.1 is proved using the PCP-VERIFIER defined in this section setting $\lambda = 1/c \log n$. Step 1 of the verifier reduces the instance C of CKTSAT to an instance $\langle 1^{n'}, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle$ of AS-CKTSAT. We have from Proposition 6.11 that $n' = \Theta(n)$ and $l = O(\log n)$ where n is the size of the input circuit C . Setting $n = n'$ in Proposition 6.17, we have that the randomness, query and decision complexity of the verifier are as claimed in Theorem 6.1. The soundness of the verifier follows from Lemma 6.18. ■

7 A randomness-efficient PCP of proximity

In this section, we modify the PCP for CIRCUIT SATISFIABILITY and construct a PCP of proximity for CIRCUIT VALUE while maintaining all the complexities. (Recall that the latter is stronger than

the former, via Proposition 2.4.) We do so by adding a proximity test to the PCP-VERIFIER defined in Section 6.3. This new proximity test, as the name suggests, checks the closeness of the input to the satisfying assignment that is supposed to be encoded in the proof oracle (see Remark 6.12). This check is done by locally decoding a bit (or several bits) of the input from its encoding and comparing it with the actual input oracle.

Theorem 7.1 *There exists universal constants c and $0 < \varepsilon < 1$ such that the following holds for all $n, m \in \mathbb{Z}^+$ and $0 < \delta < 1$ such that $m \leq \log n / \log \log n$. There exists a PCP of proximity for CIRCUIT VALUE (for circuits of size n) with the following parameters*

- *randomness* $(1 - \frac{1}{m}) \log n + O(m \log m) + O(\log \log n) + O(\log(1/\delta))$,
- *query complexity* $q = O(m^2 n^{1/m} \log^2 n)$ and *decision complexity* $d = \tilde{O}(q)$,
- *perfect completeness*,
- *Case I: if $n^{1/m} \geq m^{cm}/\delta^3$, then the verifier has soundness error $1 - \varepsilon$ for proximity parameter δ .*
- *Case II: if $n^{1/m} < m^{cm}/\delta^3$, then the verifier has soundness error $1 - \delta/c$ for proximity parameter δ .*

Recall that the PCPP-VERIFIER is supposed to work as follows: The verifier is given explicit access to a circuit C with n gates on k input bits and oracle access to the input w in the form of an input oracle $W : [k] \rightarrow \{0, 1\}$. The verifier should accept W with probability 1 if it is a satisfying assignment and accept it with probability at most $1 - \varepsilon$ if it is δ -far from any satisfying assignment.

For starters, we assume that $k \geq n/5$. In other words, the size of the input w is linear in the size of the circuit C . The reason we need this assumption is that we wish to verify the proximity of w to a satisfying assignment, but our proofs encode the assignment to all n gates of the circuit, thus it better be the case that w is a non-negligible fraction of the circuit. This assumption is not a major restriction, because if this is not the case we work with the circuit C' and input w' which are as follows: For $t = \lceil n/k \rceil$, C' is a circuit with $n' = n + 3tk$ gates on $k' = tk$ input bits such that $C'(w') = 1$ iff $w' = w^t$ for some w such that $C(w) = 1$; that is, C' checks if its input is t copies of some satisfying assignment of C . (It can be verified that C' can indeed be implemented on a circuit of size $n + 3tk$ over the full binary basis.) We choose t such that $k' \geq n'/10$. However, note that the input oracle W cannot be altered. So the verifier emulates the input w' using the original input oracle $W : [k] \rightarrow \{0, 1\}$ in the straight-forward manner. Whenever it wants to read the i -th bit of the w' , it queries the $((i-1) \bmod k) + 1$ -th bit of w .

Remark 7.2 *The above transformation from (C, w) to (C', w') is a generic one that increases the length of the input oracle compared to the proof oracle. The circuit C' checks that w' is a repetition codeword in order to maintain the distance features of C ; that is, if w is δ -far from the set of satisfying assignments of C then $w' = w^t$ is also δ -far from the satisfying assignments of C' .*

As in the case of the PCP-VERIFIER described in Section 6.3, the PCPP-VERIFIER is constructed by reducing the input circuit C , an instance of CKTSAT, using parameter m , to an instance $\langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle$ of AS-CKTSAT. The proof oracle for the PCPP-VERIFIER is the same as that of the PCP-VERIFIER (i.e., the proof oracle consists of a sequence of functions $\tilde{A}_i : F^m \rightarrow F, i = 0, \dots, l-1$ and $P_{i,j} : F^m \rightarrow F^2, i = 0, \dots, l-1, j = 0, \dots, m$ where $l = 5 \log n$).

Recall that the function $\tilde{A}_0 : F^m \rightarrow F$ is supposed to contain within it an assignment (See Remarks 6.7, 6.12). Let $I \subseteq H^m \subset F^m$ be the set of locations in F^m that contain the assignment. The PCPP-VERIFIER in addition to the tests of the PCP-VERIFIER performs the following PROXIMITY TEST to check if the assignment given by $\tilde{A}_0|_I$ matches with the input oracle W . Specifically:

$$\text{PCPP-VERIFIER}_{m,\lambda,\delta}^{W; \tilde{A}_i, P_{i,j}; i=0,\dots,l-1; j=0,\dots,m}(C).$$

1. Run $\text{PCP-VERIFIER}_{m,\lambda}^{W; \tilde{A}_i, P_{i,j}}(C)$ and reject if it rejects.
Let R be the random string generated during the execution of this step.
2. PROXIMITY TEST

Case I: $n \geq 8000|F|^{m-1}/\delta^3$.

Use random string R to determine a random canonical line $\mathcal{L}$ in F^m using the λ -biased set S_λ . Query oracle $\tilde{A}_0$ on all points along the line $\mathcal{L}$ and reject if the restriction $\tilde{A}_0$ to $\mathcal{L}$ is not a polynomial of degree at most $d = m \cdot |H|$. Query the input oracle W on all locations corresponding to those in $I \cap \mathcal{L}$ and reject if W disagrees with $\tilde{A}_0$ on any of the locations in $I \cap \mathcal{L}$.

Case II: $n < 8000|F|^{m-1}/\delta^3$.

Use random string R to determine a position $i \xleftarrow{R} \{1, \dots, k\}$ in the input and a direction $y \xleftarrow{R} S_\lambda$. Let $x \in I$ be the point corresponding to i in H^m , and let $\mathcal{L}$ be the canonical line through x in direction y . Query oracle $\tilde{A}_0$ on all points along the line $\mathcal{L}$ and reject if the restriction $\tilde{A}_0$ to $\mathcal{L}$ is not a polynomial of degree at most $d = m \cdot |H|$. Query the input oracle W at location i and reject if $W[i] \neq \tilde{A}_0(x)$.

By inspection, the proximity test increases the query and decision complexity by (even less than) a constant factor. For the randomness complexity, we first consider Case I, the randomness is used only to generate a random canonical line (as in the PCP verifier), so the randomness complexity is $\log(|F|^{m-1} \cdot |S_\lambda|)$ as before. However, in order to prove soundness, we will need to assume not only that $\lambda \leq 1/c \log n$ for some constant c (as before), but also that $\lambda \leq \delta^3/m^{cm}$.²⁴ Thus, setting $\lambda = \min\{1/c \log n, \delta^3/m^{cm}\}$, the randomness complexity increases by at most $O(m \log m) + O(\log(1/\delta))$, as claimed in Theorem 7.1. For Case II, the randomness complexity is $\log(n \cdot |S_\lambda|) \leq \log(8000|F|^{m-1} \cdot |S_\lambda|/\delta^3)$, which is the same as in Case I plus an additional $O(\log(1/\delta)) + O(1)$ random bits. Summarizing the above observations for future reference, we have the following proposition.

Proposition 7.3 *The randomness, query and decision complexities of the PCPP-VERIFIER are $r = (1 - \frac{1}{m}) \log n + O(m \log m) + O(\log \log n) + O(\log(1/\delta))$, $q = O(m^2 n^{1/m} \log^2 n)$ and $d = \tilde{O}(q)$ respectively.*

It is straightforward to check perfect completeness of this verifier. To prove soundness, we observe that if the input W is δ -far from satisfying the circuit, then one of the following must happen: (1) the verifier detects an inconsistency in the proof oracle or (2) the input oracle does not match with the encoding of the input in the proof oracle. In case of the former, we prove soundness by invoking Lemma 6.18 while in the latter case, we prove soundness by analyzing the proximity test. These ideas are explained in detail in the following two lemmas which prove the soundness of the verifier.

²⁴Actually, for the proximity test we only need $\lambda \leq \delta/m^{cm}$, however to prove robustness of the proximity test (see Section 8.1) we require $\lambda \leq \delta^3/m^{cm}$.

Lemma 7.4 (Case I) *There exists a constant c and a constant $\varepsilon > 0$ such that for all m, λ, δ satisfying $n \geq 8000|F|^{m-1}/\delta^3$, $\lambda \leq 1/c \log n$, and $\lambda \leq \delta/m^{cm}$, the following holds. If the input w given by the input oracle $W : [k] \rightarrow \{0, 1\}$ is δ -far from satisfying the circuit, then for any proof oracle the verifier rejects W with probability at least ε .*

Proof: Set ε to be the constant ε_0 in Lemma 6.18.

Case (i): $\tilde{A}_0$ is not 4ε -close to any polynomial $\hat{A}_0$ of degree md such that $C(\hat{A}_0|_I) = 1$. Then by Lemma 6.18, we conclude that the verifier rejects with probability at least ε .

Case (ii): $\tilde{A}_0$ is 4ε -close to some polynomial $\hat{A}_0$ of degree md such that $C(\hat{A}_0|_I) = 1$. Since w is δ -far from any satisfying assignment, the assignment given by $\hat{A}_0|_I$ must be at least δ -far from w . Let $B \subset F^m$ denote the set of locations in I where the assignment given by $\hat{A}_0$ disagrees with w (i.e., $B = \{x \in I \mid \hat{A}_0(x) \text{ disagrees with } w \text{ at } x\}$). Hence, $|B|/|I| \geq \delta$. Since $|I| = k \geq n/5$, we have $|B| \geq \delta n/5$. Consider the following 2 events.

[Event I]: $\tilde{A}_0|_{\mathcal{L}}$ is 5ε -far from $\hat{A}_0|_{\mathcal{L}}$.

By the Sampling Lemma (Lemma B.3) with $\mu = 4\varepsilon$ and $\zeta = \varepsilon$, this event occurs with probability at most $\left(\frac{1}{|F|} + \lambda\right) \cdot \frac{4\varepsilon}{\varepsilon^2} \leq \frac{1}{4}$ since $|F|, \frac{1}{\lambda} \geq 32/\varepsilon$.

[Event II]: $B \cap \mathcal{L} = \emptyset$.

Again by the Sampling Lemma (Lemma B.3) with $\mu = \zeta = \frac{|B|}{|F^m|}$, this event occurs with probability at most $\left(\frac{1}{|F|} + \lambda\right) \cdot \frac{|F^m|}{|B|} = \left(\frac{1}{|F|} + \lambda\right) \cdot \frac{5|F^m|}{\delta n} \leq \frac{1}{4}$, where the last inequality follows because $n \geq 8000|F|^{m-1}/\delta^3 \geq 40|F|^{m-1}/\delta$ and $\lambda \leq \delta/(40(c_F m^2)^m)$.

Suppose Event I does not occur. Then, if $\hat{A}_0|_{\mathcal{L}} \neq \tilde{A}_0|_{\mathcal{L}}$, the PROXIMITY TEST rejects since then $\tilde{A}_0|_{\mathcal{L}}$ cannot be a polynomial of degree at most d as it is 5ε -close to the polynomial $\hat{A}_0$ and hence cannot be closer to any other polynomial (as $5\varepsilon \leq \frac{1}{2}(1 - \frac{d}{|F|}) = \frac{1}{2}(1 - \frac{1}{c_F})$). Now if $\hat{A}_0|_{\mathcal{L}} = \tilde{A}_0|_{\mathcal{L}}$ and Event II does not occur, then the PROXIMITY TEST detects a mismatch between the input oracle W and $\tilde{A}_0|_{\mathcal{L}}$. Hence, if both Event I and Event II do not occur, then the test rejects. Thus, the probability of the test accepting in this case is at most the probability of either Event I or Event II occurring which is at most $1/2$.

Thus, the probability that the verifier accepts is at most $\max\{1 - \varepsilon, \frac{1}{2}\} = 1 - \varepsilon$. This completes the proof of the lemma. ■

Lemma 7.5 (Case II) *There exists a constant c such that for all m, λ, δ satisfying $n < 8000|F|^{m-1}/\delta^3$, $\lambda \leq 1/c \log n$, and $\lambda \leq \delta/m^{cm}$, the following holds. If the input w given by the input oracle $W : [k] \rightarrow \{0, 1\}$ is δ -far from satisfying the circuit, then for any proof oracle the verifier rejects W with probability δ/c .*

Proof: Set $\mu = 4\delta/c \leq \delta/16$.

Case (i): $\tilde{A}_0$ is not μ -close to any polynomial $\hat{A}_0$ of degree md such that $C(\hat{A}_0|_I) = 1$. Then by Lemma 6.18, we conclude that the verifier rejects with probability at least $\mu/4 = \delta/c$.

Case (ii): $\tilde{A}_0$ is μ -close to some polynomial $\hat{A}_0$ of degree md such that $C(\hat{A}_0|_I) = 1$. Since w is δ -far from any satisfying assignment, the assignment given by $\hat{A}_0|_I$ must be δ -far from w . With

probability greater than δ over the choice of $i \in \{1, \dots, k\}$ (and the corresponding point $x \in I$ in H^m), we have $W[i] \neq \hat{A}_0(x)$. If this occurs, the only way the verifier can accept is if $\tilde{A}_0|_{\mathcal{L}}$ is a degree md polynomial other than $\hat{A}_0|_{\mathcal{L}}$. Below, we will bound the probability of this latter event by $\delta/2$, and thereby conclude that the verifier rejects with probability at least $\delta - \delta/2 > \delta/c$.

Recall that $\tilde{A}_0$ and $\hat{A}_0$ are μ -close. Thus, if we show that the points on $\mathcal{L}$ are almost-uniformly distributed, it will follow that the expected distance between $\tilde{A}_0|_{\mathcal{L}}$ and $\hat{A}_0|_{\mathcal{L}}$ is small, and hence they are distinct degree md polynomials with small probability. Let $B \subseteq F^m$ be the set of points where $\tilde{A}_0$ and $\hat{A}_0$ disagree. Consider any point z on $\mathcal{L}$ other than x . That is, $z = x + ty$, where $t \in F \setminus \{0\}$ is fixed, x is chosen uniformly in I , and y is chosen uniformly in S_λ . The probability (over x, y) that z lands in B is precisely $1/(|I| \cdot |S_\lambda|)$ times the number of (directed) edges between I and B in the (expander) graph on vertex set F^m with edge set $E = \{(a, b) : (a - b) \in tS_\lambda\}$. This graph has second eigenvalue λ , so by the Expander Mixing Lemma, the fraction of edges between I and B satisfies

$$\frac{|E \cap (I \times B)|}{|E|} \leq \frac{|I|}{|F^m|} \cdot \frac{|B|}{|F^m|} + \lambda \cdot \sqrt{\frac{|I|}{|F^m|} \cdot \frac{|B|}{|F^m|}}.$$

Multiplying by $|E|/(|I||S_\lambda|)$, and using $|B|/|F^m| \leq \mu$ and $|E| = |F^m| \cdot |S_\lambda|$, we obtain

$$\begin{aligned} \frac{|E \cap (I \times B)|}{|I| \cdot |S_\lambda|} &\leq \mu + \lambda \cdot \sqrt{\frac{|F^m|}{|I|}} \cdot \mu \\ &\leq \frac{\delta}{16} + \lambda \cdot \sqrt{\frac{(c_F m^2 n^{1/m})^m}{n}} \\ &\leq \frac{\delta}{8}, \end{aligned}$$

where in the last inequality we use $\lambda \leq \delta/m^{cm}$. Thus, the expected fraction of $\mathcal{L} \setminus \{x\}$ that lands in B is at most $\delta/8$. By Markov's inequality, with probability at least $1 - \delta/2$, $\tilde{A}_0|_{\mathcal{L} \setminus \{x\}}$ and $\hat{A}_0|_{\mathcal{L} \setminus \{x\}}$ are $1/4$ -close. This implies that $\tilde{A}_0|_{\mathcal{L}}$ cannot be a polynomial of degree md other than $\hat{A}_0|_{\mathcal{L}}$ (since two distinct polynomials agree in at most md points, and $(md-1)/|F| < 1/4$). By Markov's Inequality, this occurs with probability at most $\delta/2$, as claimed.

■

Proof (of Theorem 7.1): Theorem 7.1 is proved using the PCPP-VERIFIER defined in this section setting $\lambda = \min\{1/c \log n, \delta^3/m^{cm}\}$. The randomness and decision (query) complexity follow from Proposition 7.3. The only fact to be verified is the soundness of the verifier. The soundness requirements of Theorem 7.1 are different for the two cases when $n^{1/m} \geq m^{cm}/\delta^3$ and $n^{1/m} < m^{cm}/\delta^3$. On the other hand, the PCPP-VERIFIER's actions are different depending on whether $n \geq 8000|F|^{m-1}/\delta^3$ or $n < 8000|F|^{m-1}/\delta^3$. These two cases of the soundness requirements and that of the PCPP-VERIFIER do not directly correspond to each other. We first note that the condition $n \geq 8000|F|^{m-1}/\delta^3$ is equivalent to $n^{1/m} \geq 8000(c_F m^2)^{m-1}/\delta^3$. Suppose $n^{1/m} \geq m^{cm}/\delta^3$ for a suitably large constant c (i.e., Case I of Theorem 7.1). Then, $n^{1/m} \geq 8000(c_F m^2)^{m-1}/\delta^3$ or equivalently $n \geq 8000|F|^{m-1}/\delta^3$. Hence, Lemma 7.4 applies and we have that the verifier has soundness error $1 - \varepsilon$ for proximity parameter δ . Now, suppose $n^{1/m} < m^{cm}/\delta^3$ (i.e., Case II of Theorem 7.1). In this case, we do not know which of $n^{1/m}$ or $8000(c_F m^2)^{m-1}/\delta^3$ is greater. Hence, we do not know which one of Lemma 7.4 or Lemma 7.5 applies. Thus, we can only guarantee the

weaker of the two soundness claims. Hence, the soundness-error for proximity parameter δ is at most $\max\{1 - \varepsilon, 1 - \delta/c\} = 1 - \delta/c$. This proves Theorem 7.1. $\blacksquare$

8 A randomness-efficient robust PCP of proximity

In this section, we modify the PCP of proximity for CIRCUIT VALUE constructed in Section 7 to design a robust PCP of proximity, while essentially maintaining all complexities. Recall the definition of robustness: If the input oracle W is δ -far from a satisfying assignment, a “regular” PCPP verifier for most choices of its random coins rejects the input; that is, it observes an inconsistency in the input. A robust PCPP verifier, on the other hand, for most choices of its random coins not only notices an inconsistency in the input but also observes that a considerable portion of the input read by it has to be modified to remove this inconsistency.

Theorem 8.1 *There exists a universal constant c such that the following holds for all $n, m \in \mathbb{Z}^+$, $\delta, \gamma > 0$ satisfying $m \leq \log n / \log \log n$: There exists a robust PCP of proximity for CIRCUIT VALUE (for circuits of size n) with the following parameters*

- *randomness* $(1 - \frac{1}{m}) \log n + O(m \log m) + O(\log \log n) + O(\log(1/\delta))$,
- *query complexity* $q = O((m^2 n^{1/m} \log^2 n)/\gamma)$ and *decision complexity* $d = \tilde{O}(q)$,
- *perfect completeness*, and
- *Case I: if $n^{1/m} \geq m^{cm}/\delta^3$, then for every $\delta' > \delta$, the verifier has expected robustness*

$$\min\{\Omega(\gamma), (1 - \gamma) \cdot (\delta' - (\delta/2))\}$$

for proximity parameter δ' .
- *Case II: if $n^{1/m} < m^{cm}/\delta^3$, then the verifier has expected robustness $\Omega(\delta)$ for proximity parameter δ .*

Note that the expected robustness in Case I can be very close to the proximity parameter δ' (provided $\delta \ll \delta' \ll \gamma \ll 1$), whereas in Case II it is always a constant factor smaller. Thus, Case I is suitable for a large number of proof composition operations, whereas Case II is useful when the query complexity of the outer verifier is already very small (and Case I can no longer be applied). Indeed, this is how these cases (which yield Theorems 3.1 and 3.2, respectively) are used in the proof of Theorem 3.3.

How Theorem 8.1 implies Theorems 3.1 and 3.2: Our main construct (i.e., Theorem 3.1) follows from Case I of Theorem 8.1 by using the error-reduction lemma (Lemma 2.11). Specifically, replacing δ by $\delta'\gamma$ and using $\delta' < \gamma/c$, Case I yields expected robustness of $\min\{\Omega(\gamma), (1 - \gamma) \cdot (\delta' - \delta'\gamma/2)\}$, which is lower-bounded by $\rho' \triangleq (1 - \gamma)^2 \cdot \delta'$. Applying Lemma 2.11 with a slackness parameter of $\gamma' \triangleq \gamma\rho'$ and $s = \gamma$, yields robust-soundness error of γ with robustness parameter of $\rho' - \gamma' = (1 - \gamma)^3 \cdot \delta'$ for proximity parameter δ' . Using $\gamma \leq 1/2$, note that the randomness increases by an additive term of $O(\log(1/\gamma')) + O(\log(1/\delta')) = O(\log(1/\delta'))$, and the decision complexity increases by a multiplicative factor of $1/(\gamma \cdot (\gamma\rho')^2) = \text{poly}(1/\delta')$. These modifications do not affect the general form of the corresponding complexities claimed in Theorem 3.1, and the latter follows (when substituting δ' for δ and γ by $\gamma/3$). We now turn to establishing Theorem 3.2: Setting $m = (\log n)/\log \log n$ (which satisfies the condition of Case II) and $\gamma = 1/2$, we obtain randomness $O(\log n/\delta)$, decision complexity $\tilde{O}((\log n)^{3+1+2}) = \text{poly}(\log n)$ and expected robustness $\Omega(\delta)$ for proximity parameter δ . Applying Lemma 2.12, Theorem 3.2 follows.

Overview of the proof of Theorem 8.1: We “robustify” our PCPP-VERIFIER in 3 steps. First we observe that a single execution of the verifier actually involves several tests (in fact $lm + 2l$ LOW-DEGREE TESTS, l EDGE-CONSISTENCY TESTS, lm ZERO PROPAGATION TESTS, l IDENTITY TESTS and a single PROXIMITY TEST). In the first step (Section 8.1), we observe that each of these tests is robust individually. In the second step (Section 8.2), we perform a “bundling” of the queries so that a certain set of queries can always be asked together. This bundling achieves robustness, albeit over a much a larger alphabet. In the final step (Section 8.3), we use a good error-correcting code to transform the “bundles” into regular bit-queries such that robustness over the binary alphabet is achieved.

8.1 Robustness of individual tests

For each possible random string R , the PCPP-VERIFIER performs several tests. More precisely, it performs $l(m + 2)$ LOW-DEGREE TESTS, l EDGE-CONSISTENCY TESTS, lm ZERO PROPAGATION TESTS, l IDENTITY TESTS and a single PROXIMITY TEST. In this section, we prove that each of these test are robust individually. In other words, we show that when one of these tests fail, it fails in a “robust” manner; that is, a considerable portion of the input read by the test has to be modified for the test to pass.

First, some notation. We view functions $g, g' : F^m \rightarrow F$ as strings of length $|F|^m$ over the alphabet F , so their relative Hamming distance $\Delta(g, g')$ is simply $\Pr_x[g(x) \neq g'(x)]$. As before, let $I \subseteq H^m \subset F^m$ be the set of locations in F^m that contains the assignment.

Let $0 < \varepsilon < 1$ be a small constant to be specified later. As before, for $i = 0, \dots, l - 1$, $j = 0, \dots, m$ and $b \in \{0, 1\}$, let $\hat{A}_i$ (resp., $\hat{P}_{i,j}^{(b)}$) be the closest polynomials of degree md (resp., κmd) to $\tilde{A}_i$ and $P_{i,j}$ respectively. (If there is more than one polynomial, choose one arbitrarily.) The proof of the soundness of the PCPP-VERIFIER (see Sections 6 and 7) was along the following lines: If the input oracle $W : [k] \rightarrow \{0, 1\}$ is δ -far from satisfying the circuit, then one of the following must happen (changing ε by a factor of 2).

1. There exists a $i = 0, \dots, l - 1$ such that $\tilde{A}_i$ is 8ε -far from every degree md polynomial or there exists a $i = 0, \dots, l - 1$, $j = 0, \dots, m$ and $b \in \{0, 1\}$ such that $P_{i,j}^{(b)}$ is 8ε -far from every degree κmd polynomial. In this case, the LOW-DEGREE TEST detects the error with probability at least 2ε .
2. There exists $i = 0, \dots, l - 1$ and $b \in \{0, 1\}$, such that $\Delta(P_{i,m}^{(b)}, \hat{P}_{i,m}^{(b)}) \leq 8\varepsilon$ and $\hat{P}_{i,m} \neq 0$. In this case, the IDENTITY TEST detects the error with probability at least $1 - 10\varepsilon$.
3. There exists $i = 0, \dots, l - 1$, $j = 1, \dots, m$ and $b \in \{0, 1\}$ such that $\Delta(P_{i,j}, \hat{P}_{i,j}) \leq 8\varepsilon$, $\Delta(P_{i,j-1}, \hat{P}_{i,j-1}) \leq 8\varepsilon$, and $\hat{P}_{i,j}(\dots, x_j, \dots) \neq \sum_{k=0}^{|H|-1} \hat{P}_{i,j-1}(\dots, h_k, \dots) x_j^k$. In this case, the ZERO PROPAGATION TEST detects the error with probability at least $1 - 20\varepsilon$.
4. There exists a $i = 0, \dots, l - 1$ such that $\Delta(P_{i,0}^{(0)}, \hat{P}_{i,0}^{(0)}) \leq 8\varepsilon$, $\Delta(P_{i,0}^{(1)}, \hat{P}_{i,0}^{(1)}) \leq 8\varepsilon$, $\Delta(\tilde{A}_i, \hat{A}_i) \leq 8\varepsilon$, $\Delta(\tilde{A}_{i+1}, \hat{A}_{i+1}) \leq 8\varepsilon$, and $\hat{P}_{i,0}(x) \neq \psi'((\tilde{T}_i(x), \hat{A}_i(x), \hat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \hat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))))$. In this case, the EDGE-CONSISTENCY TEST detects the error with probability at least $1 - 42\varepsilon$.
5. $\Delta(\tilde{A}_0, \hat{A}_0) \leq 8\varepsilon$ but W and $\hat{A}_0|_I$ disagree on at least δ fraction of the points. In this case, the PROXIMITY TEST detects the error with probability at least $1/2$ (in Case I).

Claims 8.2 to 8.7 below strengthen the above analysis and show that one of the tests not only detects the error, but a significant portion of the input read by that test needs to be modified in

order to make the test accept. More formally, recall that each of our tests T (randomly) generates a pair (I, D) where I is a set of queries to make to its oracle and D is the predicate to apply to the answers. For such a pair $(I, D) \leftarrow T$ and an oracle π , we define the *distance of $\pi|_I$ to T* to be the relative Hamming distance between $\pi|_I$ and the nearest satisfying assignment of D . Similarly, we say that π has *expected distance ρ from satisfying T* if the expectation of the distance of $\pi|_I$ to T over $(I, D) \xleftarrow{R} T$ equals ρ .

We then have the following claims about the robustness of the individual tests.

The robustness of the LOW-DEGREE TEST can be easily be inferred from the analysis of the λ -biased low-degree test due to Ben-Sasson *et al.* [12] as shown below.

Claim 8.2 *The following holds for all sufficiently small $\varepsilon > 0$. If $A : F^m \rightarrow F$ (resp., $P : F^m \rightarrow F$) is 8ε -far from every polynomial of degree md (resp., degree κmd), then the expected distance of A (resp. P) from satisfying the LOW-DEGREE TEST with degree parameter d (resp., κd) is at least 2ε .*

Proof: Recall that the LOW-DEGREE TEST chooses a random canonical line $\mathcal{L}$ and checks if $A|_{\mathcal{L}}$ is a univariate polynomial of degree d . For each canonical line $\mathcal{L}$, define $A_{\text{lines}}(\mathcal{L})$ to be the degree d univariate polynomial mapping $\mathcal{L} \rightarrow F$ having maximum agreement with A on $\mathcal{L}$, breaking ties arbitrarily. The distance of $A|_{\mathcal{L}}$ to satisfying LOW-DEGREE TEST is precisely $\Delta(A|_{\mathcal{L}}, A_{\text{lines}}(\mathcal{L}))$.

The low-degree test LDT of Ben-Sasson *et al.* [12] works as follows (see Section B for more details): The test LDT has oracle access to a points-oracle $f : F^m \rightarrow F$ and a lines oracle g . It chooses a random canonical line $\mathcal{L}$ using the λ -biased set, queries the lines-oracle g on the line $\mathcal{L}$ and the points-oracle f on a random point x on $\mathcal{L}$. It accepts iff $g(\mathcal{L})$ agrees with f at x .

By inspection, the probability that $\text{LDT}^{A, A_{\text{lines}}}$ rejects the points-oracle A and lines-oracle A_{lines} as defined above equals $\mathbb{E}_{\mathcal{L}}[\Delta(A|_{\mathcal{L}}, A_{\text{lines}}(\mathcal{L}))]$. By Theorem B.4, if A is 8ε -far from every degree md polynomial, then $\text{LDT}^{A, A_{\text{lines}}}$ rejects with probability at least 2ε (for sufficiently small ε). (Recall that our parameters satisfy the conditions of Theorem B.4 for sufficiently large choices of the constants c and c_F .) Thus, A has expected distance 2ε from satisfying our LOW-DEGREE TEST, as desired.

The intuition behind the proofs of robustness of IDENTITY TEST, ZERO PROPAGATION TEST, and EDGE-CONSISTENCY TEST is as follows. The key point to be noted is that the checks made by each of these tests are in the form of polynomial identities. Hence, if the functions read by these tests are close to being polynomials, then it follows from the Schwartz-Zippel Lemma that the inputs read by these tests either satisfy these polynomial identities or are in fact far from satisfying them. We formalize this intuition and prove the robustness of IDENTITY TEST, EDGE-CONSISTENCY TEST, and ZERO PROPAGATION TEST in Claims 8.3, 8.4, and 8.5 respectively.

Claim 8.3 *The following holds for all sufficiently small $\varepsilon > 0$. If for some $i = 0, \dots, l-1$ and $b \in \{0, 1\}$, $\Delta(P_{i,m}^{(b)}, \hat{P}_{i,m}^{(b)}) \leq 8\varepsilon$ and $\hat{P}_{i,m}^{(b)}(\cdot) \not\equiv 0$, then $P_{i,m}$ has expected distance at least $1 - 9\varepsilon$ from satisfying the IDENTITY TEST.*

Proof: The expected distance of $P_{i,m}$ from satisfying the IDENTITY TEST equals

$$\begin{aligned} \mathbb{E}_{V_\eta}[\Delta(P_{i,m}|_{V_\eta}, 0)] &= \Delta(P_{i,m}, 0) \text{ (since the } \{V_\eta\} \text{ are a partition)} \\ &\geq \Delta(\hat{P}_{i,m}, 0) - \Delta(P_{i,m}, \hat{P}_{i,m}) \\ &\geq \left(1 - \frac{\kappa md}{|F|}\right) - 8\varepsilon \text{ (by Schwartz-Zippel and hypothesis)} \\ &\geq 1 - 9\varepsilon \end{aligned}$$

Claim 8.4 *The following holds for all sufficiently small $\varepsilon > 0$. Suppose for some $i = 0, \dots, l-1$, we have $\Delta(P_{i,0}^{(0)}, \widehat{P}_{i,0}^{(0)}) \leq 8\varepsilon$, $\Delta(P_{i,0}^{(1)}, \widehat{P}_{i,0}^{(1)}) \leq 8\varepsilon$, $\Delta(\tilde{A}_i, \widehat{A}_i) \leq 8\varepsilon$, $\Delta(\tilde{A}_{i+1}, \widehat{A}_{i+1}) \leq 8\varepsilon$, and $\widehat{P}_{i,0}(\cdot) \neq \psi'(\tilde{T}_i(\cdot), \widehat{A}_i(\cdot), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,0}(\cdot)), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,1}(\cdot)))$. Then $\{P_{i,0}(\cdot), A_i(\cdot), A_{i+1}(\tilde{\Gamma}_{i,0}(\cdot)), A_{i+1}(\tilde{\Gamma}_{i,1}(\cdot))\}$ has expected distance at least $(1 - 41\varepsilon)/5$ from satisfying the EDGE-CONSISTENCY TEST.*

Proof: Note that the distance of $\{P_{i,0}(\cdot), A_i(\cdot), A_{i+1}(\tilde{\Gamma}_{i,0}(\cdot)), A_{i+1}(\tilde{\Gamma}_{i,1}(\cdot))\}|_{U_\eta}$ from satisfying the EDGE-CONSISTENCY TEST is at least $1/5$ times the distance of $P_{i,0}(\cdot)|_{U_\eta}$ to the function $\psi'(\tilde{T}_i(\cdot), A_i(\cdot), A_{i+1}(\tilde{\Gamma}_{i,0}(\cdot)), A_{i+1}(\tilde{\Gamma}_{i,1}(\cdot)))|_{U_\eta}$ (Since for each point $x \in U_\eta$ where the latter two functions disagree, at least one of $P_{i,0}, A_i, A_{i+1} \circ \tilde{\Gamma}_{i,0}, A_{i+1} \circ \tilde{\Gamma}_{i,1}$ needs to be changed at x to make the test accept). As in the proof of Claim 8.3, we have:

$$E_{U_\eta}[\Delta(P_{i,0}(\cdot)|_{U_\eta}, \psi'(\tilde{T}_i(\cdot), A_i(\cdot), A_{i+1}(\tilde{\Gamma}_{i,0}(\cdot)), A_{i+1}(\tilde{\Gamma}_{i,1}(\cdot)))|_{U_\eta})] \geq \left(1 - \frac{\kappa md}{|F|}\right) - 5 \cdot 8\varepsilon \geq 1 - 41\varepsilon,$$

where the $(1 - \kappa md/|F|)$ term corresponds to the distance if we replace all five functions with their corrected polynomials (e.g., $\widehat{P}_{i,0}, \widehat{A}_i, \widehat{A}_{i+1} \circ \tilde{\Gamma}_{i,0}, \widehat{A}_{i+1} \circ \tilde{\Gamma}_{i,1}$) and the $-5 \cdot 8\varepsilon$ accounts for the distance between each of the five functions and their corrected polynomials. Thus, the overall expected distance to satisfying the EDGE-CONSISTENCY TEST is at least $(1 - 41\varepsilon)/5$.

Claim 8.5 *The following holds for all sufficiently small $\varepsilon > 0$. Suppose for some $i = 0, \dots, l-1$, $j = 1, \dots, m$, and $b \in \{0, 1\}$, we have $\Delta(P_{i,j}^{(b)}, \widehat{P}_{i,j}^{(b)}) \leq 8\varepsilon$, $\Delta(P_{i,j-1}^{(b)}, \widehat{P}_{i,j-1}^{(b)}) \leq 8\varepsilon$, and $\widehat{P}_{i,j}^{(b)}(\dots, x_j, \dots) \neq \sum_{k=0}^{|H|-1} \widehat{P}_{i,j-1}^{(b)}(\dots, h_k, \dots) x_j^k$. Then $(P_{i,j}^{(b)}, P_{i,j-1}^{(b)})$ has expected distance at least $(1 - 19\varepsilon)/2$ from satisfying the ZERO PROPAGATION TEST.*

Proof: Suppose that $\mathcal{L}$ is a j th axis-parallel line such that

$$\widehat{P}_{i,j}^{(b)}(\dots, x_j, \dots)|_{\mathcal{L}} \neq \sum_{k=0}^{|H|-1} \widehat{P}_{i,j-1}^{(b)}(\dots, h_k, \dots) x_j^k|_{\mathcal{L}},$$

Then in order for the ZERO PROPAGATION TEST to accept, either $P_{i,j}^{(b)}|_{\mathcal{L}}$ must be modified to equal a degree κd polynomial other than $\widehat{P}_{i,j-1}^{(b)}(\dots, x_j, \dots)|_{\mathcal{L}}$ or $P_{i,j-1}^{(b)}|_{\mathcal{L}}$ must be modified to equal a degree κd polynomial other than $\widehat{P}_{i,j-1}^{(b)}(\dots, x_j, \dots)|_{\mathcal{L}}$. (Recall that the ZERO PROPAGATION TEST checks that the said restrictions are in fact polynomials of degree κd .) This would require modifying $P_{i,j}^{(b)}|_{\mathcal{L}}$ (resp., $P_{i,j-1}^{(b)}|_{\mathcal{L}}$) in at least a $1 - \kappa d/|F| - \Delta(P_{i,j}^{(b)}|_{\mathcal{L}}, \widehat{P}_{i,j}^{(b)}|_{\mathcal{L}})$ (resp., $1 - \kappa d/|F| - \Delta(P_{i,j-1}^{(b)}|_{\mathcal{L}}, \widehat{P}_{i,j-1}^{(b)}|_{\mathcal{L}})$) fraction of points. This implies that the pair $(P_{i,j}^{(b)}|_{\mathcal{L}}, P_{i,j-1}^{(b)}|_{\mathcal{L}})$ would have to be modified in at least a

$$\frac{1}{2} \cdot \left(1 - \frac{\kappa d}{|F|} - \Delta(P_{i,j}^{(b)}|_{\mathcal{L}}, \widehat{P}_{i,j}^{(b)}|_{\mathcal{L}}) - \Delta(P_{i,j-1}^{(b)}|_{\mathcal{L}}, \widehat{P}_{i,j-1}^{(b)}|_{\mathcal{L}})\right)$$

fraction of points.

Thus the expected distance of $(P_{i,j}^{(b)}, P_{i,j-1}^{(b)})$ to satisfying the ZERO PROPAGATION TEST is at least

$$\begin{aligned}
& \frac{1}{2} \cdot \mathbb{E}_{\mathcal{L}} \left[1 - \frac{\kappa d}{|F|} - \Delta(P_{i,j}^{(b)}|_{\mathcal{L}}, \widehat{P}_{i,j}^{(b)}|_{\mathcal{L}}) - \Delta(P_{i,j-1}^{(b)}|_{\mathcal{L}}, \widehat{P}_{i,j-1}^{(b)}|_{\mathcal{L}}) \right] \\
& - \Pr_{\mathcal{L}} \left[\widehat{P}_{i,j}^{(b)}(\dots, x_j, \dots)|_{\mathcal{L}} \equiv \sum_{k=0}^{|H|-1} \widehat{P}_{i,j-1}^{(b)}(\dots, h_k, \dots) x_j^k |_{\mathcal{L}} \right] \\
& \geq \frac{1}{2} (1 - \varepsilon - 8\varepsilon - 8\varepsilon) - \frac{\kappa d}{|F|} \\
& \geq \frac{1}{2} (1 - 19\varepsilon).
\end{aligned}$$

We are now left with analyzing the robustness the PROXIMITY TEST. Note that the input of the PROXIMITY TEST comes in two parts: (a) the restriction of A_0 to the line $\mathcal{L}$ and (b) the input W restricted to the line $\mathcal{L}$. Unlike earlier tests, we do not collate the robustness of these two parts of the input but express them separately. The robustness of the PROXIMITY TEST is proved by repeated applications of the Sampling Lemma (Lemma B.3).

Let $B \subset F^m$ denote the set of locations in I where the assignment given by $\widehat{A}_0$ disagrees with W (i.e., $B = \{x \in I \mid \widehat{A}_0(x) \text{ disagrees with } W \text{ at } x\}$). Recall that $|I| = k \geq n/5$.

Claim 8.6 (Case I) *There exists a constant c and a constant $\varepsilon > 0$ such that for all $m, \lambda, \delta, \delta'$ satisfying $n \geq 8000|F|^{m-1}/\delta^3$, $\lambda \leq 1/c \log n$, $\lambda \leq \delta^3/m^{cm}$, $\delta' > \delta$, the following holds. Suppose $\Delta(\widehat{A}_0, \widehat{A}_0) \leq 1/4$ and the input oracle W is δ' -far from $\widehat{A}_0|_I$ (i.e., $|B|/|I| \geq \delta'$), then with probability at least $1 - \delta/4$ (over the choice of the canonical line $\mathcal{L}$) either at least a ε -fraction of $A_0|_{\mathcal{L}}$ or at least a $(\delta' - \delta/4)$ -fraction of $W|_{\mathcal{L}}$ needs to be changed to make the PROXIMITY TEST accept.*

This claim is the robust analogue of Lemma 7.4. Observe that the robustness of the verifier is expressed separately for the proof and input oracles. As expected, the robustness of the input oracle depends on the proximity parameter δ while that of the proof oracle is independent of δ .

Proof: Consider the following three events.

Event 1: $\Delta(\tilde{A}_0|_{\mathcal{L}}, \widehat{A}_0|_{\mathcal{L}}) \geq 1/3$.

By the Sampling Lemma (Lemma B.3) with $\mu = 1/4$ and $\zeta = 1/12$, this event occurs with probability at most $\left(\frac{1}{|F|} + \lambda\right) \cdot \frac{1/4}{(1/12)^2} \leq \frac{\delta}{12}$ since $|F| \geq (8000|F|^m)/(\delta^3 n) > (12^3/2)/\delta$ and $\lambda < 2\delta/12^3$.

Event 2: $\frac{|I \cap \mathcal{L}|}{|\mathcal{L}|} > \left(1 + \frac{\delta}{8}\right) \cdot \frac{|I|}{|F|^m}$.

Again by the Sampling Lemma (Lemma B.3) with $\mu = |I|/|F|^m \geq \frac{n}{5|F|^m}$ and $\zeta = \frac{\delta\mu}{8}$, this event occurs with probability at most

$$\left(\frac{1}{|F|} + \lambda\right) \cdot \frac{8^2}{\delta^2 \mu} = \left(\frac{1}{|F|} + \lambda\right) \cdot \frac{320|F|^m}{\delta^2 n} \leq \frac{\delta}{12},$$

where the last inequality follows from the fact that $n \geq 24 \cdot 320 \cdot |F|^{m-1}/\delta^3$ and $\lambda \leq \delta^3/(24 \cdot 320(c_F m^2)^m)$.

Event 3: $\frac{|B \cap \mathcal{L}|}{|\mathcal{L}|} < \frac{|B|}{|F^m|} - \frac{\delta}{8} \cdot \frac{|I|}{|F^m|}$.

Again by the Sampling Lemma (Lemma B.3) with $\mu = |B|/|F^m| = \frac{\delta' n}{5|F|^m}$ and $\zeta = \frac{\delta n}{40|F|^m}$, this event occurs with probability at most

$$\left(\frac{1}{|F|} + \lambda\right) \cdot \frac{\mu}{\zeta^2} \leq \left(\frac{1}{|F|} + \lambda\right) \cdot \frac{320|F|^m}{\delta^2 n} \leq \frac{\delta}{12}.$$

Hence, the probability that at least one of the three events occurs is at most $\delta/4$.

Now, suppose none of the three events occur. We then get that

$$\frac{|B \cap \mathcal{L}|}{|I \cap \mathcal{L}|} \geq \frac{|B| - \delta|I|/8}{(1 + \delta/8)|I|} = \frac{\delta' - \delta/8}{1 + \delta/8} \geq \delta' - \delta/4.$$

Now for the PROXIMITY TEST to accept the pair $(\tilde{A}_0|_{\mathcal{L}}, W \cap \mathcal{L})$, either we must change $\tilde{A}_0|_{\mathcal{L}}$ to a polynomial other than $\hat{A}_0|_{\mathcal{L}}$ or correct the input for all $x \in B \cap \mathcal{L}$. The former requires us to change at least $(1 - \frac{d}{|F|} - 1/3) \geq 1/2$ fraction of the points of $A_0|_{\mathcal{L}}$ while the latter requires us to change at least $\delta' - \delta/4$ -fraction of the input read (i.e., the input oracle W restricted to the line $\mathcal{L}$). This proves the claim. ■

Claim 8.7 (Case II) *There exists a constant c such that for all m, λ, δ satisfying $n < 8000|F|^{m-1}/\delta^3$, $\lambda \leq 1/c \log n$, $\lambda \leq \delta^3/m^{cm}$, the following holds. Suppose $\Delta(\tilde{A}_0, \hat{A}_0) \leq \delta/16$ and the input oracle W is δ -far from $\hat{A}_0|_I$ (i.e., $|B|/|L| \geq \delta$), then with probability at least $\delta/2$ (over the choice of index i and direction $y \xleftarrow{R} S_\lambda$), either at least a $1/2$ -fraction of $A_0|_{\mathcal{L}}$ or $W[i]$ (i.e., the entire portion of the input oracle read by the verifier) needs to be changed to make the PROXIMITY TEST accept.*

This claim is the robust analogue of Lemma 7.5. As before, the robustness of the verifier is expressed separately for the proof and input oracles.

Proof: Since w is δ -far from any satisfying assignment, the assignment given by $\hat{A}_0|_I$ must be δ -far from w . Thus with probability greater than δ over the choice of $i \in \{1, \dots, k\}$ (and the corresponding point $x \in I$), we have $W[i] \neq \hat{A}_0(x)$. If this occurs, the only way to make the verifier accept is to either change $W[i]$ or change $\tilde{A}_0|_{\mathcal{L}}$ to a degree md polynomial other than $\hat{A}_0|_{\mathcal{L}}$. As in the proof of Claim 8.6, with probability at least $1 - \delta/2$, $\tilde{A}_0|_{\mathcal{L} \setminus \{x\}}$ and $\hat{A}_0|_{\mathcal{L} \setminus \{x\}}$ have distance at most $1/4$, and hence $\tilde{A}_0|_{\mathcal{L}}$ would have to be changed in at least $1 - ((md - 1)/|F|) - 1/4 \geq 1/2$ points to be a degree md polynomial other than $\hat{A}_0|_{\mathcal{L}}$. Thus, with probability at least $\delta - \delta/2 = \delta/2$, either $W[i]$ would have to change or at least half of $\tilde{A}_0|_{\mathcal{L}}$ would have to change to make the verifier accept. ■

8.2 Bundling

In Section 8.1, we showed that each of the tests performed by the PCPP verifier is individually robust. However, we need to show that the conjunction of all these tests is also robust. This is not true for the PCPP verifier in its present form for the following reason: Suppose the input oracle W is δ -far from satisfying the circuit. We then know that one of the tests detects this fact with non-negligible probability. Moreover as seen in Section 8.1, this test is robust. However, since this test is only one of the $O(lm)$ tests being performed by the verifier, the oracle bits read by this test comprise a small fraction of the total query complexity of the verifier. For instance, the number of

bits read by a single LOW-DEGREE TEST is about $1/lm$ times the query complexity. This causes the robustness of the verifier to drop by a factor of at least lm . Note that the issue here is not the fact that the verifier performs different types of tests (i.e., LOW-DEGREE TEST, IDENTITY TEST, ZERO PROPAGATION TEST, etc) but rather that it performs many instances of each test and that the soundness analysis only guarantees that one of these test instances rejects (robustly). This is not sufficient to make the verifier robust.

For this purpose, we “bundle” the various functions in the proof oracle so that the inputs required for the several test instances can be read together. This maintains the robustness of the individual tests, albeit over a larger alphabet. To understand this “bundling”, let us assume for the present that the only type of tests that the verifier performs is the LOW-DEGREE TEST. There exists a natural bundling in this case. Instead of $l(m+2)$ different oracles $\{\tilde{A}_i\}$ and $\{P_{i,j}\}$, we have one oracle Π which bundles together the data of all these oracles. The oracle $\Pi : F^m \rightarrow F^{l \cdot (2m+3)}$ is supposed to satisfy $\Pi(x) = (\tilde{A}_0(x), \dots, \tilde{A}_{l-1}(x), P_{0,0}(x), \dots, P_{l-1,m}(x))$ for all $x \in F^m$. It can now be easily checked that over this proof oracle, the conjunction of all the LOW-DEGREE TESTS is robust (over alphabet $F^{l \cdot (2m+3)}$) with the same soundness and robustness parameters as a single LOW-DEGREE TEST (over alphabet F). However, this natural bundling does not lend itself to the other tests performed by the PCPP verifier (namely, ZERO PROPAGATION TEST, and EDGE-CONSISTENCY TEST). Next, we provide an alternate bundling and massage our verifier slightly to work with this bundling.

First for some notation. As mentioned earlier, we will be able to prove robustness of the verifier via bundling, however over a larger alphabet. This large alphabet will be $\Sigma = F^{l+2l \cdot (m+1)}$. Unlike before, the proof oracle for the robust PCPP verifier will consist of only one function $\Pi : F^m \rightarrow \Sigma$. The robust PCPP verifier simulates the PCPP verifier as follows: To answer the queries of the PCPP verifier, the robust verifier bundles several queries together, queries the new proof oracle Π and then unbundles the answer to obtain the answers of the queries of the original PCPP verifier. For convenience, we index the $l + 2l \cdot (m+1)$ coordinates of $\Sigma = F^{l+2l \cdot (m+1)}$ as follows: The first l coordinates are indexed by a single index i ranging from 0 to $l-1$, while the remaining $2l \cdot (m+1)$ are indexed by a triplet of indices (i, j, b) where i ranges over $0, \dots, l-1$, j ranges over $0, \dots, m$ and $b \in \{0, 1\}$. Let $S : F^m \rightarrow F^m$ denote the (linear) transformation that performs one cyclic shift to the right; that is, $S(x_0, \dots, x_{m-1}) = (x_{m-1}, x_0, \dots, x_{m-2})$. The bundling of the proof oracles $\tilde{A}_i$ ’s and $P_{i,j}$ ’s by the modified proof oracle Π is as follows:

$$\forall x \in F^m, \quad \begin{cases} \Pi(x)_i = \tilde{A}_i \left(S^{\lfloor \frac{i}{h} \rfloor} (x) \right) & i = 0, \dots, l-1 \\ \Pi(x)_{(i,j,b)} = P_{i,j}^{(b)} \left(S^{j + \lfloor \frac{i}{h} \rfloor} (x) \right) & i = 0, \dots, l-1; j = 0, \dots, m \text{ and } b \in \{0, 1\} \end{cases} \quad (4)$$

where $h = \log |H| = \log n/m$. Note that the size of the new proof oracle Π is exactly equal to the sum of the size of the oracles $\tilde{A}_i$ ’s and $P_{i,j}$ ’s.

We now state how the robust verifier performs the unbundling and the individual tests. We consider each step of the PCPP verifier and present their robust counterparts.

The first steps of the PCP-VERIFIER (and PCPP-VERIFIER) are independent of the proof oracle and are performed as before. That is, the robust verifier, as before, reduces the CKTSAT instance to an instance $\langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle$ of AS-CKTSAT, sets $d = m \cdot |H|$, and generates a random string R of length $\log(|S_\lambda| \cdot |F|^{m-1})$. The remaining steps are proof-oracle dependent and we will discuss each of them in detail.

Proximity Test. For the proximity test, the only portion of the proof oracle that we require is the portion containing $\tilde{A}_0$. Since $\Pi(x)_0$ is $\tilde{A}_0 \circ S^{\lfloor \frac{0}{h} \rfloor} (x) = \tilde{A}_0(x)$, the ROBUST PROXIMITY TEST can easily be describes as follows:

ROBUST PROXIMITY TEST^{W; \Pi}(R)

Case I: $n \geq 8000|F|^{m-1}/\delta^3$.

Use random string R to determine a random canonical line $\mathcal{L}$ in F^m using the λ -biased set S_λ . Query oracle Π on all points along the line $\mathcal{L}$. Unbundle $\Pi(\mathcal{L})$ to obtain the values of $\tilde{A}_0$ on all points along the line $\mathcal{L}$ and reject if the restriction $\tilde{A}_0$ to $\mathcal{L}$ is not a polynomial of degree at most d . Query the input oracle W on all locations corresponding to those in $I \cap \mathcal{L}$ and reject if W disagrees with $\tilde{A}_0$ on any of the locations in $I \cap \mathcal{L}$.

Case II: $n < 8000|F|^{m-1}/\delta^3$.

Use random string R to determine a position $i \xleftarrow{R} \{1, \dots, k\}$ in the input and a direction $y \xleftarrow{R} S_\lambda$. Let $x \in I$ be the point corresponding to i in H^m , and let $\mathcal{L}$ be the canonical line through x in direction y . Query oracle Π on all points along the line $\mathcal{L}$. Unbundle $\Pi(\mathcal{L})$ to obtain the values of $\tilde{A}_0$ on all points along the line $\mathcal{L}$ and reject if the restriction $\tilde{A}_0$ to $\mathcal{L}$ is not a polynomial of degree at most d . Query the input oracle W at location i and reject if $W[i] \neq \tilde{A}_0(x)$.

Low-Degree Test. We note that the distance of the polynomial $\tilde{A}_i : F^m \rightarrow F$ to being degree k (for any $k \in \mathbb{Z}^+$) is exactly the same as that of $\tilde{A}_i \circ S^{\lfloor \frac{i}{h} \rfloor} : F^m \rightarrow F$ since $S^{\lfloor \frac{i}{h} \rfloor}$ is an invertible linear transformation. Hence, it is sufficient if we check that $\tilde{A}_i \circ S^{\lfloor \frac{i}{h} \rfloor}$ is low-degree. The case with the $P_{i,j}^{(b)}$'s is similar. Thus, the new ROBUST LOW-DEGREE TEST can be described as follows:

ROBUST LOW-DEGREE TEST^{\Pi}(R)

Use random string R to determine a random canonical line $\mathcal{L}$ in F^m using the λ -biased set S_λ .

Query the oracle Π on all points along the line $\mathcal{L}$.

For $i = 0, \dots, l-1$,

Unbundle $\Pi(\mathcal{L})$ to obtain the values of $\tilde{A}_i \circ S^{\lfloor \frac{i}{h} \rfloor}$ on all points along the line $\mathcal{L}$ and reject if the restriction $\tilde{A}_i \circ S^{\lfloor \frac{i}{h} \rfloor}$ to $\mathcal{L}$ is not a polynomial of degree at most d .

For $i = 0, \dots, l-1$, $j = 0, \dots, m$ and $b \in \{0, 1\}$,

Unbundle $\Pi(\mathcal{L})$ to obtain the values of $P_{i,j}^{(b)} \circ S^{j+\lfloor \frac{i}{h} \rfloor}$ on all points along the line $\mathcal{L}$ and reject if the restriction of $P_{i,j}^{(b)} \circ S^{j+\lfloor \frac{i}{h} \rfloor}$ to $\mathcal{L}$ is not a polynomial of degree at most κd .

Thus, effectively we are testing $\tilde{A}_i$ (respectively $P_{i,j}$) using the line space $S^{\lfloor \frac{i}{h} \rfloor} \circ S_\lambda$ (respectively $S^{j+\lfloor \frac{i}{h} \rfloor} \circ S_\lambda$).

Identity Test. In the case of the IDENTITY TEST, we observe that $P_{i,m}^{(b)}$ vanishes on F^m iff $P_{i,m}^{(b)} \circ S^{m+\lfloor \frac{i}{h} \rfloor}$ vanishes on F^m . Recall that we were allowed to use arbitrary partitions of the space F^m . The set of random 1st axis-parallel lines is one such partition and we use this partition.

ROBUST IDENTITY TEST^{\Pi}(R)

Use random string R to determine a random 1st axis-parallel line in F^m of the form $\mathcal{L} = (X, a_1, \dots, a_{m-1})$. Query the oracle Π on all points along the line $\mathcal{L}$.

For $i = 0, \dots, l-1$ and $b \in \{0, 1\}$,

Unbundle $\Pi(\mathcal{L})$ to obtain the values of $P_{i,m}^{(b)} \circ S^{m+\lfloor \frac{i}{h} \rfloor}$ on all points along the line $\mathcal{L}$ and reject if any of these is non-zero.

Edge Consistency Test. For any $x \in F^m$, we say that $P_{i,0}$ is *well-formed* at x if the Equation (2) is satisfied for this x . The EDGE-CONSISTENCY TEST verifies that $P_{i,0}$ is well-formed for all $x \in U_\eta$ and $i = 0, \dots, l-1$. This was done earlier by reading the values of $P_{i,0}, \tilde{A}_i, \tilde{A}_{i+1} \circ \tilde{\Gamma}_{i,0} = \tilde{A}_{i+1}$ and $\tilde{A}_{i+1} \circ \tilde{\Gamma}_{i,1}$ for all $x \in U_\eta$.

Let $\mathcal{L}$ be a random 1st axis-parallel line. The robust version of this test checks that $P_{i,0}$ is well-formed for all points on $S^{\lfloor \frac{i}{h} \rfloor}(\mathcal{L})$. Consider any $x = (x_0, \dots, x_{m-1}) \in \mathcal{L}$. To verify that $P_{i,0}$ is well-formed at $S^{\lfloor \frac{i}{h} \rfloor}(x)$, the verifier needs the values $P_{i,0}(S^{\lfloor \frac{i}{h} \rfloor}(x)), \tilde{A}_i(S^{\lfloor \frac{i}{h} \rfloor}(x)), \tilde{A}_{i+1}(S^{\lfloor \frac{i}{h} \rfloor}(x))$ and $\tilde{A}_{i+1} \circ \tilde{\Gamma}_{i,1}(S^{\lfloor \frac{i}{h} \rfloor}(x))$. We will show that all these values can be obtained from unbundling the value of Π on $\mathcal{L}$ and $S^{-1}(\mathcal{L})$. Clearly, the values $P_{i,0}(S^{\lfloor \frac{i}{h} \rfloor}(x))$ and $\tilde{A}_i(S^{\lfloor \frac{i}{h} \rfloor}(x))$ can be obtained from unbundling the value of Π at x . The other two values that we require are $\tilde{A}_{i+1}(S^{\lfloor \frac{i}{h} \rfloor}(x))$ and $\tilde{A}_{i+1} \circ \tilde{\Gamma}_{i,1}(S^{\lfloor \frac{i}{h} \rfloor}(x))$. We first show that $\tilde{\Gamma}_{i,1}(S^{\lfloor \frac{i}{h} \rfloor}(x)) = S^{\lfloor \frac{i}{h} \rfloor}(x')$ for $x' = (x_0 + e_{(i \bmod h)}, x_1, \dots, x_{m-1}) \in \mathcal{L}$ (recall that $\{e_0, \dots, e_{f-1}\}$ are a basis for F over $\text{GF}(2)$ and $\{e_0, \dots, e_{h-1}\}$ span $H \subset F$). For this purpose, we first recall the definition of $\tilde{\Gamma}_{i,1}$: $\tilde{\Gamma}_{i,1}(z_0, \dots, z_{m-1}) = (z_0, \dots, z_{t-1}, z_t + e_u, z_{t+1}, \dots, z_{m-1})$ where $t = \lfloor i/h \rfloor \bmod m$ and $u = i \bmod h$. Furthermore, since S^m is the identity map, we have that $S^{\lfloor \frac{i}{h} \rfloor \bmod m} = S^{\lfloor \frac{i}{h} \rfloor}$. With these observations, we have the following:

$$\begin{aligned} \tilde{\Gamma}_{i,1}(S^{\lfloor \frac{i}{h} \rfloor}(x)) &= \tilde{\Gamma}_{i,1}(S^{\lfloor i/h \rfloor \bmod m}(x)) \\ &= \tilde{\Gamma}_{i,1}(S^{\lfloor i/h \rfloor \bmod m}(x_0, \dots, x_{m-1})) \\ &= S^{\lfloor i/h \rfloor \bmod m}(x_0 + e_{(i \bmod h)}, x_1, \dots, x_{m-1}) \\ &= S^{\lfloor \frac{i}{h} \rfloor}(x') \end{aligned}$$

Now, $S^{\lfloor \frac{i+1}{h} \rfloor}$ is either $S^{\lfloor \frac{i}{h} \rfloor}$ or $S^{\lfloor \frac{i}{h} \rfloor+1}$ depending on the value of i . Suppose $S^{\lfloor \frac{i+1}{h} \rfloor} = S^{\lfloor \frac{i}{h} \rfloor}$. We then have that $\tilde{A}_{i+1}(S^{\lfloor \frac{i}{h} \rfloor}(x)) = A_{i+1}(S^{\lfloor \frac{i+1}{h} \rfloor}(x))$ and $\tilde{A}_{i+1} \circ \tilde{\Gamma}_{i,1}(S^{\lfloor \frac{i}{h} \rfloor}(x)) = \tilde{A}_{i+1}(S^{\lfloor \frac{i}{h} \rfloor}(x')) = \tilde{A}_{i+1}(S^{\lfloor \frac{i+1}{h} \rfloor}(x'))$. Both these values can be obtained by unbundling the value of Π on $\mathcal{L}$ (since both x and x' lie on $\mathcal{L}$). In the other case, where $S^{\lfloor \frac{i+1}{h} \rfloor} = S^{\lfloor \frac{i}{h} \rfloor+1}$, we have $A_{i+1}(S^{\lfloor \frac{i}{h} \rfloor}(x)) = A_{i+1}(S^{\lfloor \frac{i+1}{h} \rfloor}(S^{-1}x))$ and $\tilde{A}_{i+1} \circ \tilde{\Gamma}_{i,1}(S^{\lfloor \frac{i}{h} \rfloor}(x)) = A_{i+1}(S^{\lfloor \frac{i}{h} \rfloor}(x')) = A_{i+1}(S^{\lfloor \frac{i+1}{h} \rfloor}(S^{-1}x'))$. These values can be obtained by unbundling the value of Π on $S^{-1}(\mathcal{L})$. Thus, to check that $P_{i,0}$ is well-formed for all points on $S^{\lfloor \frac{i}{h} \rfloor}(\mathcal{L})$, it suffices if the verifier queries Π on all points on $\mathcal{L}$ and $S^{-1}(\mathcal{L})$.

ROBUST EDGE-CONSISTENCY TEST $^\Pi(R)$

Use the random string R to determine a random 1st axis-parallel line in F^m of the form $\mathcal{L} = (X, a_2, \dots, a_m)$. Query the oracle Π along all points in the lines $\mathcal{L}$ and $S^{-1}(\mathcal{L})$.

For $i = 0, \dots, l-1$,

For all $x \in S^{\lfloor \frac{i}{h} \rfloor}(\mathcal{L})$, reject if $P_{i,0}$ is not well-formed at x . [Note that all the values required for this verification can be obtained by unbundling $\Pi(\mathcal{L})$ and $\Pi(S^{-1}(\mathcal{L}))$.]

Zero Propagation Test. For each $i = 0, \dots, l-1$ and $b \in \{0, 1\}$, the ZERO PROPAGATION TEST checks that $P_{i,0}^{(b)}$ vanishes on H^m by verifying that Equation (3) is satisfied for all $j = 1, \dots, m-1$

(we also need to check that $P_{i,m}^{(b)} \equiv 0$, however this is taken care by the IDENTITY TEST). Since $S(H^m) = H^m$, checking if $P_{i,0}^{(b)}$ vanishes on H^m is equivalent to checking if $P_{i,0}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor}$ vanishes on H^m . Hence, we can perform the zero propagation on the polynomials $P_{i,0}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor}; i = 0, \dots, l-1, b \in \{0, 1\}$ instead of the polynomials $P_{i,0}^{(b)}; i = 0, \dots, l-1, b \in \{0, 1\}$. In other words, we need to verify the following equation instead of Equation (3).

$$P_{i,j}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor} \left(\underbrace{x_1, \dots, x_{j-1}}_{}, \underbrace{x_j, x_{j+1}, \dots, x_m}_{} \right) = \sum_{k=0}^{|H|-1} P_{i,j-1}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor} \left(\underbrace{x_1, \dots, x_{j-1}}_{}, \underbrace{h_k, x_{j+1}, \dots, x_m}_{} \right) x_j^k, \quad \forall (x_1, \dots, x_m) \in F^m \quad (5)$$

This equation can be further rewritten in terms of the cyclic shift S as follows:

$$P_{i,j}^{(b)} \left(S^{\lfloor \frac{i}{h} \rfloor + j - 1} (x_1, x_2, \dots, x_m) \right) = \sum_{k=0}^{|H|-1} P_{i,j-1}^{(b)} \left(S^{\lfloor \frac{i}{h} \rfloor + j - 1} (h_k, x_2, \dots, x_m) \right) x_1^k, \quad \forall (x_1, \dots, x_m) \in F^m \quad (6)$$

This helps us to rewrite the ZERO PROPAGATION TEST with bundling as follows:

ZERO PROPAGATION TEST ^{$\Pi(R)$}

Use random string R to determine a random 1st axis-parallel line in F^m of the form $\mathcal{L} = (X, a_2, \dots, a_m)$. Query the oracle Π along all points in the lines $\mathcal{L}$ and $S^{-1}(\mathcal{L})$.
For $i = 0, \dots, l-1, j = 1, \dots, m$, and $b \in \{0, 1\}$

Unbundle $\Pi(\mathcal{L})$ to obtain the value of $P_{i,j-1}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor + j - 1}$ on all points along the line $\mathcal{L}$. Similarly, unbundle $\Pi(S^{-1}(\mathcal{L}))$ to obtain the value of $P_{i,j}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor + j}$ on all points along the line $S^{-1}(\mathcal{L})$ (Equivalently, this is the value of $P_{i,j}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor + j - 1}$ on all points along the line $\mathcal{L}$). Reject either if the restriction of $P_{i,j-1}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor + j - 1}$ or $P_{i,j}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor + j - 1}$ to $\mathcal{L}$ is not a polynomial of degree at most κd or if any of the points on the line $\mathcal{L}$ violate Equation (6).

The integrated robust verifier. Having presented the robust version of each of the tests, the integrated robust verifier is as follows.

ROBUST-PCPP-VERIFIER ^{$W; \Pi$} _{m, λ, δ} (C).

1. Using Proposition 6.11, reduce the instance C of CKTSAT, using parameter m , to an instance $\langle 1^n, 1^m, F, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\} \rangle$ of AS-CKTSAT, and set $d = m \cdot |H|$.

We let $S_\lambda \subset F^m$ be a λ -biased set of size at most $\left(\frac{\log |F|^m}{\lambda} \right)^2 [2]$.

2. Choose a random string R of length $\log(|S_\lambda| \cdot |F|^{m-1})$. [Note: We reuse R in all tests, but only the LOW-DEGREE TEST utilizes the full length of R .]
3. Run ROBUST LOW-DEGREE TEST ^{$\Pi(R)$} .
4. Run ROBUST EDGE-CONSISTENCY TEST ^{$\Pi(R)$} .
5. Run ROBUST ZERO PROPAGATION TEST ^{$\Pi(R)$} .

6. Run ROBUST IDENTITY TEST ^{Π} (R).
7. Run ROBUST PROXIMITY TEST ^{$W;\Pi$} (R).

Reject if any of the above tests reject.

The randomness of the ROBUST-PCPP-VERIFIER is exactly the same as before whereas the query complexity and decision complexity increase by a constant factor²⁵.

Proposition 8.8 *The randomness, query and decision complexities of the ROBUST-PCPP-VERIFIER are $r = (1 - \frac{1}{m}) \log n + O(m \log m) + O(\log \log n) + O(\log(1/\delta))$, $q = O(m^2 n^{1/m} \log^2 n)$ and $d = \tilde{O}(q)$ respectively.*

It is straightforward to check perfect completeness of this verifier.

Robustness analysis of the integrated verifier. To state the robust soundness, it is useful for us to separate the robustness wrt the input oracle and wrt the proof oracle. Let $W : [k] \rightarrow \{0, 1\}$ be the input oracle and Π the proof oracle. For every sequence of coin tosses R (and a given setting of parameters), let $\Delta_{\text{inp}}^{W, \Pi}(R)$ (resp., $\Delta_{\text{pf}}^{W, \Pi}(R)$) denote the fraction of the bits read from W (resp. Π) that would need to be changed to make the ROBUST-PCPP-VERIFIER accept on coin tosses R . Then the following lemmas state the robustness property of our verifier.

Lemma 8.9 (Case I) *There are constants $c \in \mathbb{Z}^+$ and $\rho > 0$ such the following holds for every $n, m \in \mathbb{Z}^+$, $\delta, \delta' > 0$ satisfying $m \leq \log n / \log \log n$, $n^{1/m} \geq m^{cm} / \delta^3$, $\lambda \leq \min\{1/c \log n, \delta^3 / m^{cm}\}$, $\delta' > \delta$. If W is δ' -far from satisfying the circuit, then for any proof oracle $\Pi : F^m \rightarrow \Sigma$, either $\mathbb{E}_R[\Delta_{\text{pf}}^{W, \Pi}(R)] \geq \rho$ or $\mathbb{E}_R[\Delta_{\text{inp}}^{W, \Pi}(R)] \geq \delta' - \delta/2$.*

Proof: Unbundle the proof oracle Π to obtain the functions $\tilde{A}_i$ and $P_{i,j}$ using Equation (4). Consider the action of the PCPP-VERIFIER (i.e., the non-robust verifier) on the proof oracles $\tilde{A}_i, P_{i,j}$ and input oracle W .

Let ε be a sufficiently small constant such that the Claims 8.2–8.6 hold. Suppose W is δ' -far from satisfying the circuit. We then know that one of the following holds and that the corresponding test instance of the PCPP-VERIFIER rejects its input robustly (see Claims 8.2 to 8.6).

1. There exists a $i = 0, \dots, l-1$ such that $\tilde{A}_i$ is 8ε -far from every degree md polynomial or there exists a $i = 0, \dots, l-1$, $j = 0, \dots, m$ and $b \in \{0, 1\}$ such that $P_{i,j}^{(b)}$ is 8ε -far from every degree κmd polynomial. In this case, the expected distance of $\tilde{A}_i$ (or resp. $P_{i,j}^{(b)}$) from satisfying the LOW-DEGREE TEST with degree parameter d (resp., κd) is at least 2ε (Claim 8.2).
2. There exists $i = 0, \dots, l-1$ and $b \in \{0, 1\}$, such that $\Delta(P_{i,m}^{(b)}, \hat{P}_{i,m}^{(b)}) \leq 8\varepsilon$ and $\hat{P}_{i,m} \not\equiv 0$. In this case, $P_{i,m}$ has expected distance at least $1 - 9\varepsilon$ from satisfying the IDENTITY TEST (Claim 8.3).
3. There exists a $i = 0, \dots, l-1$ such that $\Delta(P_{i,0}^{(0)}, \hat{P}_{i,0}^{(0)}) \leq 8\varepsilon$, $\Delta(P_{i,0}^{(1)}, \hat{P}_{i,0}^{(1)}) \leq 8\varepsilon$, $\Delta(\tilde{A}_i, \hat{A}_i) \leq 8\varepsilon$, $\Delta(\tilde{A}_{i+1}, \hat{A}_{i+1}) \leq 8\varepsilon$, and $\hat{P}_{i,0}(x) \neq \psi'((\tilde{T}_i(x), \hat{A}_i(x), \hat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \hat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))))$. In this case, $\{P_{i,0}(\cdot), A_i(\cdot), A_{i+1}(\tilde{\Gamma}_{i,0}(\cdot))A_{i+1}(\tilde{\Gamma}_{i,1}(\cdot))\}$ has expected distance at least $(1 - 41\varepsilon)/5$ from satisfying the EDGE-CONSISTENCY TEST (Claim 8.4).

²⁵Though the new proof oracle returns elements of Σ and not bits, we express the query complexity as the number of bits read by the verifier rather than the number of symbols (i.e., elements of $|\Sigma|$) to maintain consistency across calculating the query complexity into the proof and input oracles.

4. There exists $i = 0, \dots, l-1$, $j = 1, \dots, m$ and $b \in \{0, 1\}$ such that $\Delta(P_{i,j}, \hat{P}_{i,j}) \leq 8\varepsilon$, $\Delta(P_{i,j-1}, \hat{P}_{i,j-1}) \leq 8\varepsilon$, and $\hat{P}_{i,j}(\dots, x_j, \dots) \neq \sum_{k=0}^{|H|-1} \hat{P}_{i,j-1}(\dots, h_k, \dots) x_j^k$. In this case, $(P_{i,j}^{(b)}, P_{i,j-1}^{(b)})$ has expected distance at least $(1 - 19\varepsilon)/2$ from satisfying the ZERO PROPAGATION TEST (Claim 8.5).
5. $\Delta(\tilde{A}_0, \hat{A}_0) \leq 8\varepsilon$ but W and $\hat{A}_0|_I$ disagree on at least δ fraction of the points. In this case, with probability at least $1 - \delta/4$ (over the choice of the canonical line $\mathcal{L}$) either at least a ε -fraction of $A_0|_{\mathcal{L}}$ or at least a $(\delta' - \delta/4)$ -fraction of $W|_{\mathcal{L}}$ needs to be changed to make the PROXIMITY TEST accept (Claim 8.6).

This implies that either A_0 has expected distance $(1 - \delta/4)\varepsilon \geq \varepsilon/2$ or W has expected distance $(1 - \delta/4)(\delta' - \delta/4) \geq (\delta' - \delta/2)$ from satisfying the PROXIMITY TEST.

For instance, let's us assume $\tilde{A}_0$ is 8ε -far from being low degree so the LOW-DEGREE TEST rejects it robustly; that is, for a random canonical line $\mathcal{L}$, the expected distance of $\tilde{A}_0|_{\mathcal{L}}$ from satisfying the LOW-DEGREE TEST is at least 2ε . Recall from Equation (4) that $\tilde{A}_0(x)$ is one of the co-ordinates in the bundled $\Pi(x)$. Hence, if $\tilde{A}_0|_{\mathcal{L}}$ is ρ -far from satisfying the LOW-DEGREE TEST, so is $\Pi_{\mathcal{L}}$ from satisfying the ROBUST LOW-DEGREE TEST. Thus, Π has expected distance at least 2ε from satisfying the ROBUST LOW-DEGREE TEST. Now, the oracle positions read by the ROBUST LOW-DEGREE TEST constitute a constant fraction of the oracle positions read by the ROBUST PCPP-VERIFIER, so Π has expected distance $\Omega(\varepsilon)$ from satisfying the ROBUST-PCPP-VERIFIER. Thus, the robustness of the individual test instance is transferred to the combined ROBUST LOW-DEGREE TEST by bundling. The case with the other test types is similar. We thus have that $E_R[\Delta_{\text{pf}}^{W,\Pi}(R)] \geq \Omega(\varepsilon)$ or $E_R[\Delta_{\text{inp}}^{W,\Pi}(R)] \geq \delta' - \delta/2$. The lemma then follows by setting $\rho = \Omega(\varepsilon)$. ■

Lemma 8.10 (Case II) *There is a constant $c \in \mathbb{Z}^+$ such the following holds for every $n, m \in \mathbb{Z}^+$, $\delta > 0$ satisfying $m \leq \log n / \log \log n$, $n^{1/m} < m^{cm}/\delta^3$, $\lambda \leq \min\{1/c \log n, \delta^3/m^{cm}\}$. If W is δ -far from satisfying the circuit, then for any proof oracle $\Pi : F^m \rightarrow \Sigma$, either $E_R[\Delta_{\text{pf}}^{W,\Pi}(R)] \geq \delta/c$ or $E_R[\Delta_{\text{inp}}^{W,\Pi}(R)] \geq \delta/c$.*

Proof: This proof proceeds in the same way as Lemma 8.9 except that in this case we set ε to be $\delta/(8 \cdot 32)$ which is much smaller than the constant required by Claims 8.2–8.5. The reason for this is that the individual robustness of the PROXIMITY TEST - Case II (Claim 8.7) requires that $\Delta(\tilde{A}_0, \hat{A}_0) \leq \delta/16$ unlike in Case I (Claim 8.7) which just required $\Delta(\tilde{A}_0, \hat{A}_0) \leq 1/4$.

Suppose W is δ -far from satisfying C , then as before we have five different cases. The first four cases are exactly as before. We mention below the only case that differs from before (i.e., the fifth one).

5. $\Delta(\tilde{A}_0, \hat{A}_0) \leq 8\varepsilon = \delta/16$ but W and $\hat{A}_0|_I$ disagree on at least δ fraction of the points. In this case, with probability at least $\delta/2$ (over the choice of index i and direction $y \stackrel{R}{\leftarrow} S_\lambda$), either at least a $1/2$ -fraction of $A_0|_{\mathcal{L}}$ or $W[i]$ (i.e., the entire portion of the input oracle read by the verifier) needs to be changed to make the PROXIMITY TEST accept (Claim 8.7).

This implies that either A_0 has expected distance $\frac{\delta}{2} \cdot \frac{1}{2} = \Omega(\varepsilon)$ or W has expected distance $\delta/2$ from satisfying the PROXIMITY TEST.

Arguing as before, we have that $E_R[\Delta_{\text{pf}}^{W,\Pi}(R)] \geq \Omega(\varepsilon) = \Omega(\delta)$ or $E_R[\Delta_{\text{inp}}^{W,\Pi}(R)] \geq \delta/2$. The lemma then follows by setting c to be a large enough constant such that $\delta/c = \min\{\Omega(\delta), \delta/2\}$. ■

8.3 Robustness over the binary alphabet

The transformation from a robust verifier over the alphabet Σ to one over the binary alphabet is analogous to converting non-Boolean error correcting codes to Boolean ones via “code catenation”. This transformation is exactly the same transformation as the one in the proof of Lemma 2.13. However, we cannot directly use Lemma 2.13 as we may apply the “code concatenation” process only to the proof oracle Π and not to the input oracle W . However, this is not a problem, because the input oracle is already binary and has good robustness.

Let $\text{ECC} : \{0, 1\}^{\log |\Sigma|} \rightarrow \{0, 1\}^b$ for $b = O(\log |\Sigma|)$ be a binary error-correcting code of constant relative minimum distance, which can be computed by an explicit circuit of size $O(\log |\Sigma|)[51]$. We augment the original proof oracle Π , viewed now as having $\log |\Sigma|$ -bit long entries (i.e., elements of Σ) with an additional oracle Υ having b -bit long entries, where $\Upsilon(x)$ is supposed to be $\text{ECC}(\Pi(x))$.

Our new verifier V , on oracle access to the input W and proof $\Pi \circ \Upsilon$, will simulate the ROBUST-PCPP-VERIFIER. The queries to the input oracle are performed just as before however, for each query $x \in F^m$ in the proof oracle Π made by ROBUST-PCPP-VERIFIER, V will query the corresponding $\log |\Sigma|$ bits in $\Pi(x)$ and the b bits in $\Upsilon(x)$. Thus, the query complexity of V is at most $\log |\Sigma| + b$ times the number of queries issued by the earlier verifier. Since $b = O(\log |\Sigma|)$, the query complexity of the new verifier V is a constant times that of the previous one. The randomness is exactly the same. The action of the new verifier V is as follows: Suppose ROBUST-PCPP-VERIFIER issues queries $x_1, \dots, x_{q_1}$ to the proof oracle Π , and queries $i_1, \dots, i_{q_2}$ to the input oracle, then V issues queries $x_1, \dots, x_{q_1}$ to the proof oracle Π , a similar set of queries $x_1, \dots, x_{q_1}$ to the proof oracle Υ and $i_1, \dots, i_{q_2}$ to the input oracle. V accepts $(\Pi(x_1), \dots, \Pi(x_{q_1}), \Upsilon(x_1), \dots, \Upsilon(x_{q_1}), W(i_1), \dots, W(i_{q_2}))$ iff the ROBUST-PCPP-VERIFIER accepts $(\Pi(x_1), \dots, \Pi(x_{q_1}), W(i_1), \dots, W(i_{q_2}))$ and $\Upsilon(x_i) = \text{ECC}(\Pi(x_i))$ for all $i = 1, \dots, q_1$. It is straightforward to check that V has perfect completeness if ROBUST-PCPP-VERIFIER has perfect completeness. For the robust soundness, we define $\Delta_{\text{inp}}^{W, \Pi \circ \Upsilon}(R)$ and $\Delta_{\text{pf}}^{W, \Pi \circ \Upsilon}(R)$ wrt V analogously to before (cf. just before Lemma 8.9), but referring to distance over $\{0, 1\}$ (rather than Σ) for the proof oracle. The proofs of the following claims regarding the robust-soundness of V mimics the proof of Lemma 2.13.

Lemma 8.11 (Case I) *There are constants $c \in \mathbb{Z}^+$ and $\rho > 0$ such the following holds for every $n, m \in \mathbb{Z}^+$, $\delta, \delta' > 0$ satisfying $m \leq \log n / \log \log n$, $n^{1/m} \geq m^{cm} / \delta^3$, $\lambda \leq \min\{1/c \log n, \delta^3 / m^{cm}\}$, $\delta' > \delta$. If W is δ' -far from satisfying the circuit, then then for any proof oracles $\Pi : F^m \rightarrow \{0, 1\}^{\log |\Sigma|}$, $\Upsilon : F^m \rightarrow \{0, 1\}^b$, either $\mathbb{E}_R[\Delta_{\text{pf}}^{W, \Pi \circ \Upsilon}(R)] \geq \rho$ or $\mathbb{E}_R[\Delta_{\text{inp}}^{W, \Pi \circ \Upsilon}(R)] \geq \delta' - \delta/2$.*

Lemma 8.12 (Case II) *There is a constant $c \in \mathbb{Z}^+$ such the following holds for every $n, m \in \mathbb{Z}^+$, $\delta > 0$ satisfying $m \leq \log n / \log \log n$, $n^{1/m} < m^{cm} / \delta^3$, $\lambda \leq \min\{1/c \log n, \delta^3 / m^{cm}\}$. If W is δ -far from satisfying the circuit, then for any proof oracle $\Pi : F^m \rightarrow \Sigma$, either $\mathbb{E}_R[\Delta_{\text{pf}}^{W, \Pi}(R)] \geq \delta/c$ or $\mathbb{E}_R[\Delta_{\text{inp}}^{W, \Pi}(R)] \geq \delta/c$.*

It is to be noted that the expected robustness of the proof oracle (ρ in Case I and δ/c in Case II) is not the same as similar parameters in Lemmas 8.9 and 8.10, but weaker by a constant factor as suggested in Lemma 2.13.

Finally, we conclude by proving Theorem 8.1.

Proof (of Theorem 8.1): Theorem 8.1 is proved using the ROBUST-PCPP-VERIFIER defined in this section setting $\lambda = \min\{1/c \log n, \delta^3 / m^{cm}\}$. The randomness, query and decision complexity of the ROBUST-PCPP-VERIFIER (i.e., before the transformation to the binary alphabet) are as mentioned in Proposition 8.8. As mentioned in the earlier paragraph, the transformation from

the alphabet Σ to the binary alphabet maintains the randomness complexity while the query (and decision) complexity increase by at most a constant factor. Hence, the randomness, query and decision complexities of the verifier are as claimed in Theorem 8.1

So far, we have considered the proof and input oracle separately. Hence the expected robustness in Lemmas 8.11 and 8.12 were expressed separately for the proof and input oracles. We can consider them together by giving weights to the two oracle portions in the decision circuits (i.e. repeating queries, see Remark 7.2). Also, the expected robustness mentioned in Theorem 8.1 are different for the two cases when $n^{1/m} \geq m^{cm}/\delta^3$ and $n^{1/m} < m^{cm}/\delta^3$ (and, as in the proof of Theorem 7.1, these two cases do not directly correspond to the two cases expressed in Lemmas 8.11 and 8.12 respectively). We will give weights to the input and proof oracles differently in each of these cases.

Case I: $n^{1/m} \geq m^{cm}/\delta^3$

In this case, we give weight $(1 - \gamma)$ to the input oracle and γ to the proof oracle, where γ is as specified in Theorem 8.1. Recall that these weights mean that each query to the input oracle is repeated several times such that the relative length of the input-part in the decision circuit is $1 - \gamma$. These repeated queries may increase the query (and decision) complexity increase by a factor of at most $1/\gamma$. Note that weighting does not affect the randomness complexity (or any other parameter such as the proximity parameter δ).

Since $n^{1/m} \geq m^{cm}/\delta^3$, we have $n^{1/m} \geq 8000(c_F m^2)^{m-1}/\delta^3$ or equivalently $n \geq 8000|F|^{m-1}/\delta^3$. Hence, Lemma 8.11 applies and we have that either $E_R[\Delta_{\text{pf}}^{W, \Pi \circ \Upsilon}(R)] \geq \rho$ or $E_R[\Delta_{\text{inp}}^{W, \Pi \circ \Upsilon}(R)] \geq \delta' - \delta/2$. Note that the first expression refers to the “expected robustness” of the proof-part whereas the second expression refers to the input-part. The overall expected robustness is obtained by a weighted average of these two expressions, where the weights are with respect to the aforementioned weighting (which assigns weight γ to the input-part). Hence, the expected robustness with respect to the said weighting is

$$\gamma \cdot E_R[\Delta_{\text{pf}}^{W, \Pi \circ \Upsilon}(R)] + (1 - \gamma) \cdot E_R[\Delta_{\text{inp}}^{W, \Pi \circ \Upsilon}(R)] \geq \min\{\gamma \cdot \rho, (1 - \gamma) \cdot (\delta' - \delta/2)\}.$$

Case II: $n^{1/m} < m^{cm}/\delta^3$

In this case, we give equal weight to the input and proof oracles (i.e., we effectively set $\gamma = 1/2$). This weighting may increase the query (and decision) complexity increase by at most a factor of 2, and has no effect on any other parameter.

In this case, we do not know which of $n^{1/m}$ or $8000(c_F m^2)^{m-1}/\delta^3$ is greater. Hence, we do not know which one of Lemma 8.11 or Lemma 8.12 applies. Thus, we can only guarantee the weaker of the two expected robustness claims. Hence, we have that either $E_R[\Delta_{\text{pf}}^{W, \Pi \circ \Upsilon}(R)] \geq \min\{\rho, \delta/c\} = \delta/c$ or $E_R[\Delta_{\text{inp}}^{W, \Pi \circ \Upsilon}(R)] \geq \min\{\delta - \delta/2, \delta/c\} = \delta/c$. Hence, the expected robustness with respect to the said weighting is

$$\frac{1}{2} \cdot E_R[\Delta_{\text{pf}}^{W, \Pi \circ \Upsilon}(R)] + \frac{1}{2} \cdot E_R[\Delta_{\text{inp}}^{W, \Pi \circ \Upsilon}(R)] \geq \frac{1}{2} \cdot \frac{\delta}{c} = \Omega(\delta).$$

Thus, in both cases, the expected robustness is as claimed. Noting that the other parameters (e.g., the randomness and decision complexities) are as claimed, Theorem 8.1 follows. $\blacksquare$

8.4 Linearity of encoding

In this section we point out that, for linear circuits (to be defined below), the mapping from an assignment to the corresponding PCP of proximity is linear. Throughout this section, “linear”

means $GF(2)$ -linear (yet, we will sometimes refer to F -linearity, for an extension field F of $GF(2)$). The main motivation to the current study is to derive linear codes satisfying local-testability and relaxed local-decodability (i.e., Theorems 1.4 and 1.5, respectively). Specifically, the constructions presented in Section 4 yield linear codes provided that the corresponding PCP of proximity is linear in the aforementioned sense.

We call a circuit is *linear* if it is a conjunction of linear constraints. However, instead of representing this conjunction via AND gates, it is more convenient for us to work with circuits that have multiple output gates, one for each linear constraint. That is:

Definition 8.13 *A multi-output circuit is linear if all its internal gates are parity gates and an input is accepted by it if and only if all output gates evaluate to zero.*

Proposition 8.14 *If C is a linear circuit, then there is a linear transformation T mapping satisfying assignments w of C to proof oracles $T(w)$ such that the PCPP verifier of Theorem 3.1 will, on input C , accept oracle $(w, T(w))$ with probability 1. Moreover, all the decision circuits produced by the verifier, on input C , can be made linear (while maintaining the claimed decision complexity).*

In the rest of this section, we provide a proof of Proposition 8.14, starting with an assignment w that satisfies the linear circuit. We prove that the mapping from w to a proof-oracle is linear by reviewing our construction of this mapping and ensuring that all steps in this construction are linear transformations.

Phase I - STRUCTURED-CKTSAT: In this phase (described in Section 6.1) we write down the values to all gates of the circuit and route them along the wrapped de Bruijn graph. Actually, we make a few minor and straightforward modifications to Definition 6.3: we allow multiple output gates (rather than a single output gate) and require that each such gate evaluates to zero (rather than to 1).²⁶ Also, here we deal with gate types that are linear (e.g., XOR), rather than arbitrary (e.g., AND and OR).

Since all the circuit gates are linear functions of the input, the values on the wires leaving the zero-th layer of the well-structured circuit (i.e., the last two bits of the mapping $A_0 : \{0,1\}^N \rightarrow \{0,1\}^4$ in Section 6.1) are linear in the input (i.e., in w). As to A_i , $i > 0$, (and the first two bits of A_0) notice that it is obtained by permuting the values of the previous layer A_{i-1} and setting some wires to zero (if they are not needed in the routing, e.g. gates 3 and 4 in Figure 3). These operations are linear, and so all assignment functions are linear in the input.

Phase II - Arithmetization: In this phase (described in Section 6.2) we extend the values given by A_i to an evaluation of a low-degree multivariate polynomial over some finite field F that is an extension field of $GF(2)$ of degree f . Each value of A_i is four bits long (say b_0, b_1, b_2, b_3) and identified with the element $b_0e_0 + b_1e_1 + b_2e_2 + b_3e_3$, where $e_0, \dots, e_{f-1}$ is a basis for F viewed as a vector space over $GF(2)$. We view A_i as a function $A_i : H^m \rightarrow F$ and construct a low-degree extension $\tilde{A}_i : F^m \rightarrow F$ of A_i by interpolation. on all inputs in H^m and use these values to interpolate and evaluate $\tilde{A}_i$ on all points in F^m . Notice that interpolation is F -linear and hence also $GF(2)$ -linear. We conclude that the values of $\tilde{A}_i$ on all points in F^m is a linear transformation of the values of A_i . Since A_i is linear in the input assignment, so is $\tilde{A}_i$.

Clarification: Many parts of our encoding (starting with $\tilde{A}_i$) consist of evaluations of multivariate polynomials $P(x)$ over F^m . The linearity we claim is not linearity in x (the free variables of the polynomial). Rather, we claim the table of values $\{P(a) : a \in F^m\}$ is linear in the initial assignment

²⁶Recall that an input is accepted by the linear circuit if and only if all output gates evaluate to zero.

w , which may be viewed as the information encoded in this table. In contrast, throughout this section, x is merely an index to this table. For example, in Phase II we showed the table $\{\tilde{A}_i(a) : a \in F^m\}$ is obtained by a linear transformation applied to the table $\{A_i(a') : a' \in H^m\}$ (but we certainly do not claim $\tilde{A}_i(a)$ is linear in a). That is, each $\tilde{A}_i(a)$ is a linear combination of the $A_i(a')$'s.

Phase III - The Constraint Polynomials: We now discuss the polynomials $P_{i,0}^{(0)}$ and $P_{i,1}^{(1)}$ defined in Equation (2), and show their values are a linear transformation of the values of $\tilde{A}_i$. The first polynomial (i.e., $P_{i,0}^{(0)}$) is obtained by applying the univariate polynomial ψ_0 defined in Equation 1 to each value of $\tilde{A}_i$ (i.e., $P_{i,0}^{(0)}(x) = \psi_0(\tilde{A}_i(x))$). By definition, ψ_0 evaluates to zero iff its input, when represented as a vector in $GF(2)^f$, belongs to the linear space spanned by $\{e_0, e_1, e_2, e_3\}$. This polynomial defines a linear transformation, as claimed by the following lemma.

Lemma 8.15 *Let L be a $GF(2)$ -linear subspace of $F = GF(2)^f$ and $\psi_L(t) = \prod_{\alpha \in L} (t - \alpha)$. Then the mapping $\psi_L : F \rightarrow F$ is linear.*

Proof: We use the fact that for any integer i , the transformation $t \mapsto t^{2^i}$ is linear; that is, $(t + t')^{2^i} = t^{2^i} + t'^{2^i}$. Our main claim is that the polynomial $\psi_L(t)$ can be written as $\sum_i c_i t^{2^i}$ and hence is linear (being a sum of linear transformations). We prove this claim by induction on the dimension of $L \subseteq GF(2)^f$. Indeed, for $\dim(L) = 0$ (i.e., $L = \{0^f\}$), it holds that $\psi_L(t) = t$ and our claim follows. In the induction step, write L as $L = L' \cup \{\alpha + L'\}$ where L' is some linear space of dimension $k - 1$ and $\alpha \in L \setminus L'$. Clearly, $\psi_L(t) = \psi_{L'}(t) \cdot \psi_{L'}(t + \alpha)$. Using the inductive hypothesis for L' (and the linearity of $t \mapsto t^{2^j}$), we get

$$\begin{aligned} \psi_L(t) &= \left(\sum_i c_i \cdot t^{2^i} \right) \cdot \left(\sum_j c_j \cdot (t + \alpha)^{2^j} \right) \\ &= \left(\sum_i c_i \cdot t^{2^i} \right) \cdot \left(\sum_j c_j \cdot (t^{2^j} + \alpha^{2^j}) \right) \\ &= \sum_{i,j} c_i c_j t^{2^i} t^{2^j} + \sum_{i,j} c_i c_j t^{2^i} \alpha^{2^j} \\ &= \sum_i c_i^2 t^{2^{i+1}} + \sum_i c'_i t^{2^i} \end{aligned}$$

where $c'_i = \sum_j c_i c_j \alpha^{2^j}$ and $\sum_{i \neq j} c_i c_j t^{2^i} t^{2^j} = 2 \sum_{i < j} c_i c_j t^{2^i} t^{2^j} = 0$ (because F has characteristic 2). This completes the proof of the inductive claim.

We now turn to the second polynomial, $P_{i,0}^{(1)}$. Recall that $P_{i,0}^{(1)}(x) = \psi_1(s, a, a_0, a_1)$, where $s = \tilde{T}_i(x)$, $a = \tilde{A}_i(x)$ and $a_j = \tilde{A}_{i+1}(\tilde{\Gamma}_{i,j}(x))$. It can be verified that $\tilde{T}_i(x)$ (which represents the gate type) is independent of the input w to the circuit, and by our previous discussion a, a_0, a_1 are linear in the input w (to the circuit). Thus, it will suffice to show that ψ' is linear in its last three inputs. When discussing Equation (2) we did not go into the specific construction of the polynomial ψ' because only its functionality mattered, and we showed that there exists some constant-degree polynomial that does the job. But for our current purposes (of showing linearity) we need to present some specific polynomial ψ' that is linear (as an operator over $GF(2)^f$) and has the desired properties needed by the verification process. To do this, recall $\mathcal{C}$ is the set of allowable

gates in the well-structured circuit, and so we define $\delta_{s_0}(z)$ to be the (minimal degree) uni-variate polynomial of degree $|\mathcal{C}|$ that is 1 if $z = s_0$ and is 0 if $z \in \mathcal{C} \setminus \{s_0\}$, and write ψ' as

$$\psi'(s, a, a_0, a_1) = \sum_{s_0 \in \mathcal{C}} \delta_{s_0}(s) \cdot \psi'_{s_0}(a, a_0, a_1) \quad (7)$$

Claim 8.16 *For any $s_0 \in \mathcal{C}$ that can occur as a gate in a well-structured circuit constructed from a linear circuit C , the polynomial $\psi'_{s_0}(a, a_0, a_1)$ of Equation 7 can be written as a linear transformation (of (a, a_0, a_1)).*

Proof: Recall that the value of $\psi'_{s_0}(a, a_0, a_1)$ is supposed to represent whether or not the four least significant bits of the three inputs (denoted a' , a'_0 and a'_1) satisfy some condition. By inspecting Definition 6.3, it can be verified that (in our case) this condition is a linear one. That is, $\psi'_{s_0}(a, a_0, a_1) = 0$ if and only if the triplet (a', a'_0, a'_1) , viewed as a 12-bit vector over $GF(2)$, belongs to some specific linear space $L_{s_0} \subseteq GF(2)^{12}$.

Recall that we may assume that $a = 0^{f-4}a'$ (and similarly for a_0 and a_1), because this condition is imposed by the constraint polynomial $P_{i,0}^{(0)}$. Thus, we seek a polynomial (over F^3) such that if each of its three inputs belongs to $Span(e_0, \dots, e_3)$ then it will output 0 iff the inputs reside in the linear space that is analogous to L_{s_0} ; that is, the input (a, a_0, a_1) should evaluate to 0 iff $a' \circ a'_0 \circ a'_1 \in L_{s_0}$. To obtain this, we assume the existence of $\alpha \in F$ such that multiplying an element by α corresponds to a left cyclic shift by four positions (e.g., $\alpha \cdot \sigma_0 \cdots \sigma_{f-1} = \sigma_4 \cdots \sigma_{f-1} \sigma_0 \cdots \sigma_3$). Such an element exists for the standard representation of F . Using this element we can write $\psi'_{s_0} : F^3 \rightarrow F$ as

$$\psi'_{s_0}(a, a_0, a_1) = \psi_{L_{s_0}}(\alpha^2 a + \alpha a_0 + a_1)$$

where $\psi_{L_{s_0}}$ is the univariate polynomial that is zero iff its input is in L_{s_0} . Note that, for inputs in $Span(e_0, \dots, e_3)$, indeed $\psi'_{s_0}(a, a_0, a_1) = 0$ iff $a' \circ a'_0 \circ a'_1 \in L_{s_0}$. By Lemma 8.15, $\psi_{L_{s_0}}$ is linear. It follows that ψ'_{s_0} is linear, because multiplication by a fixed element of F (i.e., α) is a linear operation. ■

Recall $\delta_{s_0}(s)$ depends only on the circuit and not on its input (i.e., w). Thus, each summand of (7) is linear in w and hence the sum is itself linear in w . We conclude that the table of evaluations of the polynomials given by Equation (2) is obtained by linear transformations applied to the input to the circuit.

Phase IV - The Sum-check Polynomials: In this phase (described by Equation (3)) we apply a sequence of interpolations to previously constructed polynomials $P_{i,j}^{(b)}$. Each such interpolation is an F -linear transformation and hence also a $GF(2)$ -linear one. Thus, the sequence of polynomials $P_{i,j}^{(b)}$ is obtained by a linear transformation applied to the input.

Phase V - Bundling and Encoding: In this phase (described in Sections 8.2 and 8.3) we apply some cyclic shifts to the (values of the) sequence of $l + 2l(m + 1)$ polynomials obtained in the previous phases. Then we bundle the polynomials together, obtaining an alphabet of size $|F|^{l+2l(m+1)}$. This bundling does not change the encoding (only the partitioning of the proof into symbols) and hence is also a linear transformation. Finally, we apply an error correcting code to each symbol in order to reduce the alphabet size (from $|F|^{l+2l(m+1)}$ to binary, and this is also a linear transformation as long as the error correcting code is itself linear.

The result of this shifting, bundling and encoding is the actual proof given to the (outer) verifier of Theorem 8.1. Notice this transformation from $l + 2l(m + 1)$ polynomials (each evaluated in F) to one proof (over the binary alphabet) is linear, because all three parts of it are linear.

Now we argue that all tests performed by the verifier are linear and the decision complexity claimed in Theorem 8.1 can be achieved by using small linear circuits. This can be seen by inspecting the various tests described in Section 6.3, noticing that they all check either linear or F -linear conditions, and applying the general result of Strassen [53] showing that any algebraic circuit that computes a linear function (as a formal polynomial) can be converted into a linear circuit with only a constant-factor increase in size. This completes the proof of Proposition 8.14.

Acknowledgments

We are grateful to Avi Wigderson for collaborating with us at early stages of this research and to Irit Dinur for inspiring discussions at late stages of this research.

References

- [1] AJTAI, M., IWANIEC, H., KOMLOS, J., PINTZ, J., AND SZEMERÉDI, E. Construction of a thin set with small Fourier coefficients. *Bull. London Math. Soc.* 22, (1990), 583–590.
- [2] ALON, N., GOLDBREICH, O., HÅSTAD, J., AND PERALTA, R. Simple constructions of almost k -wise independent random variables. *Journal of Random Structures and Algorithms* 3, 3 (Fall 1992), 289–304.
- [3] ALON, N., AND MANSOUR, Y. ε -discrepancy sets and their applications for interpolation. *Information Processing Letters* 54, 6 (June 1995), 337–342.
- [4] ARORA, S., LUND, C., MOTWANI, R., SUDAN, M., AND SZEGEDY, M. Proof verification and the hardness of approximation problems. *Journal of the ACM* 45, 3 (May 1998), 501–555. (Preliminary Version in *33rd FOCS*, 1992).
- [5] ARORA, S., AND SAFRA, S. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM* 45, 1 (Jan. 1998), 70–122. (Preliminary Version in *33rd FOCS*, 1992).
- [6] BABAI, L., FORTNOW, L., LEVIN, L. A., AND SZEGEDY, M. Checking computations in polylogarithmic time. In *Proc. 23rd ACM Symp. on Theory of Computing* (New Orleans, Louisiana, 6–8 May 1991), pp. 21–31.
- [7] BARAK, B. How to go beyond the black-box simulation barrier. In *Proc. 42nd IEEE Symp. on Foundations of Comp. Science* (Las Vegas, Nevada, 14–17 Oct. 2001), pp. 106–115.
- [8] BEIMEL, A., ISHAI, Y., KUSHILEVITZ, E., AND RAYMOND, J. F. Breaking the $O(n^{1/(2k-1)})$ barrier for information-theoretic private information retrieval. In *Proc. 43rd IEEE Symp. on Foundations of Comp. Science* (Vancouver, British Columbia, Canada, 16–19 Nov. 2002), pp. 261–270.
- [9] BELLARE, M., GOLDBREICH, O., AND SUDAN, M. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM Journal of Computing* 27, 3 (June 1998), 804–915. (Preliminary Version in *36th FOCS*, 1995).
- [10] BELLARE, M., GOLDWASSER, S., LUND, C., AND RUSSELL, A. Efficient probabilistically checkable proofs and applications to approximation. In *Proc. 25th ACM Symp. on Theory of Computing* (San Diego, California, 16–18 May 1993), pp. 294–304.
- [11] BEN-SASSON, E., HARSHA, P., AND RASKHODNIKOVA, S. Some 3CNF properties are hard to test. In *Proc. 35th ACM Symp. on Theory of Computing* (San Diego, California, 9–11 June 2003), pp. 345–354.
- [12] BEN-SASSON, E., SUDAN, M., VADHAN, S., AND WIGDERSON, A. Randomness-efficient low degree tests and short PCPs via epsilon-biased sets. In *Proc. 35th ACM Symp. on Theory of Computing* (San Diego, California, 9–11 June 2003), pp. 612–621.
- [13] BLUM, M., LUBY, M., AND RUBINFELD, R. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences* 47, 3 (Dec. 1993), 549–595. (Preliminary Version in *22nd STOC*, 1990).

- [14] BOGDANOV, A., OBATA, K., AND TREVISAN, L. A lower bound for testing 3-colorability in bounded-degree graphs. In *Proc. 43rd IEEE Symp. on Foundations of Comp. Science* (Vancouver, Canada, 16–19 Nov. 2002), pp. 93–102.
- [15] BOGDANOV, A., AND TREVISAN, L. Lower bounds for testing bipartiteness in dense graphs. Tech. Rep. TR02-064, Electronic Colloquium on Computational Complexity, 2002.
- [16] CANETTI, R., GOLDREICH, O., AND HALEVI, S. The random oracle methodology, revisited. In *Proc. 30th ACM Symp. on Theory of Computing* (Dallas, Texas, 23–26 May 1998), pp. 209–218.
- [17] COOK, S. A. Short propositional formulas represent nondeterministic computations. *Information Processing Letters* 26, 5 (Jan. 1988), 269–270.
- [18] DESHPANDE, A., JAIN, R., KAVITHA, T., RADHAKRISHNAN, J., AND LOKAM, S. V. Better lower bounds for locally decodable codes. In *Proc. 17th IEEE Conference on Computational Complexity* (Montréal, Québec, Canada, 21–24 May 2002), pp. 184–193.
- [19] DINUR, I., AND REINGOLD, O. PCP testers: Towards a more combinatorial proof of PCP theorem. (In Preparation), 2003.
- [20] ERGÜN, F., KUMAR, R., AND RUBINFELD, R. Fast approximate PCPs. In *Proc. 31st ACM Symp. on Theory of Computing* (Atlanta, Georgia, 1–4 May 1999), pp. 41–50.
- [21] EVEN, G., GOLDREICH, O., LUBY, M., NISAN, N., AND VELIČKOVIĆ, B. Efficient Approximations of Product Distributions. *Journal of Random Structures and Algorithms* 13, 1 (Aug. 1998), 1–16. (Preliminary Version in *24th STOC*, 1992).
- [22] FEIGE, U., GOLDWASSER, S., LOVÁSZ, L., SAFRA, S., AND SZEGEDY, M. Interactive proofs and the hardness of approximating cliques. *Journal of the ACM* 43, 2 (Mar. 1996), 268–292. (Preliminary version in *32nd FOCS*, 1991).
- [23] FORTNOW, L., ROMPEL, J., AND SIPSER, M. On the power of multi-prover interactive protocols. *Theoretical Computer Science* 134, 2 (Nov. 1994), 545–557. (Preliminary Version in *3rd IEEE Symp. on Structural Complexity*, 1988).
- [24] GOLDREICH, O. A sample of samplers – a computational perspective on sampling. Tech. Rep. TR97-020, Electronic Colloquium on Computational Complexity, 1997.
- [25] GOLDREICH, O., GOLDWASSER, S., AND RON, D. Property testing and its connection to learning and approximation. *Journal of the ACM* 45, 4 (July 1998), 653–750. (Preliminary Version in *37th FOCS*, 1996).
- [26] GOLDREICH, O., AND RON, D. Property testing in bounded degree graphs. *Algorithmica* 32, 2 (Jan. 2002), 302–343. (Preliminary Version in *29th STOC*, 1997).
- [27] GOLDREICH, O., AND SAFRA, S. A combinatorial consistency lemma with application to proving the PCP theorem. *SIAM Journal of Computing* 29, 4 (2000), 1132–1154. (Preliminary Version in *RANDOM*, 1997).
- [28] GOLDREICH, O., AND SUDAN, M. Locally testable codes and PCPs of almost linear length. In *Proc. 43rd IEEE Symp. on Foundations of Comp. Science* (Vancouver, Canada, 16–19 Nov. 2002), pp. 13–22. (See ECC Report TR02-050, 2002).

- [29] GOLDBREICH, O., AND WIGDERSON, A. Tiny families of functions with random properties: A quality–size trade–off for hashing. *Journal of Random structures and Algorithms* 11, 4 (Dec. 1997), 315–343. (Preliminary Version in *26th STOC*, 1994).
- [30] GURUSWAMI, V., LEWIN, D., SUDAN, M., AND TREVISAN, L. A tight characterization of NP with 3-query PCPs. In *Proc. 39th IEEE Symp. on Foundations of Comp. Science* (Palo Alto, California, 8–11 Nov. 1998), pp. 18–27.
- [31] HARSHA, P., AND SUDAN, M. Small PCPs with low query complexity. *Computational Complexity* 9, 3–4 (Dec. 2000), 157–201. (Preliminary Version in *18th STACS*, 2001).
- [32] HÅSTAD, J. Some optimal inapproximability results. *Journal of the ACM* 48, 4 (July 2001), 798–859. (Preliminary Version in *29th STOC*, 1997).
- [33] HENNIE, F. C., AND STEARNS, R. E. Two-tape simulation of multitape Turing machines. *Journal of the ACM* 13, 4 (Oct. 1966), 533–546.
- [34] KATZ, J., AND TREVISAN, L. On the efficiency of local decoding procedures for error-correcting codes. In *Proc. 32nd ACM Symp. on Theory of Computing* (Portland, Oregon, 21–23 May 2000), pp. 80–86.
- [35] KATZ, N. M. An estimate for character sums. *Journal Amer. Math. Soc.* 2, 2 (1989), 197–200.
- [36] KERENIDIS, I., AND DE WOLF, R. Exponential lower bound for 2-query locally decodable codes via a quantum argument. In *Proc. 35th ACM Symp. on Theory of Computing* (San Diego, California, 9–11 June 2003), pp. 106–115.
- [37] KILIAN, J. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *Proc. 24th ACM Symp. on Theory of Computing* (Victoria, British Columbia, Canada, 4–6 May 1992), pp. 723–732.
- [38] LAPIDOT, D., AND SHAMIR, A. Fully parallelized multi prover protocols for NEXP-time (extended abstract). In *Proc. 32nd IEEE Symp. on Foundations of Comp. Science* (San Juan, Puerto Rico, 1–4 Oct. 1991), pp. 13–18.
- [39] LEIGHTON, F. T. *Introduction to Parallel Algorithms and Architectures*. Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1992.
- [40] LUND, C., FORTNOW, L., KARLOFF, H., AND NISAN, N. Algebraic methods for interactive proof systems. In *Proc. 31st IEEE Symp. on Foundations of Comp. Science* (St. Louis, Missouri, 22–24 Oct. 1990), pp. 2–10.
- [41] MICALI, S. Computationally sound proofs. *SIAM Journal of Computing* 30, 4 (2000), 1253–1298. (Preliminary Version in *35th FOCS*, 1994).
- [42] NAOR, J., AND NAOR, M. Small-bias probability spaces: Efficient constructions and applications. In *Proc. 22nd ACM Symp. on Theory of Computing* (Baltimore, Maryland, 4–16 May 1990), pp. 213–223.
- [43] PIPPENGER, N., AND FISCHER, M. J. Relations among complexity measures. *Journal of the ACM* 26, 2 (Apr. 1979), 361–381.

- [44] POLISHCHUK, A., AND SPIELMAN, D. A. Nearly-linear size holographic proofs. In *Proc. 26th ACM Symp. on Theory of Computing* (Montréal, Québec, Canada, 23–25 May 1994), pp. 194–203.
- [45] RAZ, R. A parallel repetition theorem. *SIAM Journal of Computing* 27, 3 (June 1998), 763–803. (Preliminary Version in *27th STOC*, 1995).
- [46] RAZBOROV, A. A., SZEMERÉDI, E., AND WIGDERSON, A. Constructing small sets that are uniform in arithmetic progressions. *Combinatorics, Probability and Computing* 2 (1993), 513–518.
- [47] RUBINFELD, R., AND SUDAN, M. Robust characterizations of polynomials with applications to program testing. *SIAM Journal of Computing* 25, 2 (Apr. 1996), 252–271. (Preliminary Version in *23rd STOC*, 1991 and *3rd SODA*, 1992).
- [48] SAMORODNITSKY, A., AND TREVISAN, L. A PCP characterization of NP with optimal amortized query complexity. In *Proc. 32nd ACM Symp. on Theory of Computing* (Portland, Oregon, 21–23 May 2000), pp. 191–199.
- [49] SCHÖNHAGE, A. Schnelle multiplikation von polynomen über Körpern der charakteristik 2 (German). *Acta Informatica* 7, 4 (1977), 395–398.
- [50] SCHÖNHAGE, A., AND STRASSEN, V. Schnelle multiplikation großer zahlen (German). *Computing* 7, 3–4 (1971), 281–292.
- [51] SPIELMAN, D. Linear-time encodable and decodable error-correcting codes. *IEEE Transactions on Information Theory* 42, 6 (Nov. 1996), 1723–1732. (Preliminary Version in *27th STOC*, 1995).
- [52] SPIELMAN, D. A. *Computationally Efficient Error-Correcting Codes and Holographic Proofs*. PhD thesis, Massachusetts Institute of Technology, June 1995.
- [53] STRASSEN, V. Vermeidung von Divisionen (German). *J. Reine Angew. Math.*, 264 (1973), 184–202.

Part III

Appendices

A Hadamard-code-based PCP of proximity

In this section we note that the Hadamard-code-based inner verifier from Arora *et al.* [4] can be converted in to a PCP of proximity. Recall that the inner verifier of [4] accesses $O(1)$ input oracles, where the i th oracle is supposed to provide the Hadamard *encoding* of some string w_i , and verifies that their concatenation satisfies some given circuit C .

Here we simplify this verifier to work with a single string w and the verifier accesses a *single input oracle* that represents this string itself (not some encoding of it), and verifies that w is close to an assignment acceptable by the circuit C given as explicit input.

Theorem A.1 *There exists a constant $\delta_0 > 0$ such that there exists a PCP of proximity for CIRCUIT VALUE (for circuits of size n) with randomness complexity $O(n^2)$, query complexity $O(1)$, perfect completeness, soundness error $1 - \delta$, and proximity parameter 5δ for any $\delta \leq \delta_0$. That is, inputs that are δ -far from satisfying the circuit are rejected with probability at least $\min(\delta, \delta_0)/5$.*

Notice that we do not claim robustness of this PCP of proximity. This is because we don't intend to use this verifier (or any verifier derived from it) as the outer verifier during composition. However, this verifier is robust (in a trivial sense). Indeed, any PCP of proximity with $O(1)$ query complexity is trivially ρ -robust for some constant $\rho > 0$ (since the relative distance between two query patterns is lower-bounded by the inverse of number of bits queried).

Proof: Let V denote the claimed verifier. We first list the oracles used by V , then we describe the tests that V performs, and finally we will verify that V 's complexities are as claimed and analyze its performance (most notably its soundness and proximity).

Oracles. Let C be a circuit with n gates on m input bits. The verifier accesses an input oracle $W : [m] \rightarrow \{0, 1\}$ (representing a string $w \in \{0, 1\}^m$), and a proof oracle $\Pi = (A, B)$, with $A : \{0, 1\}^n \rightarrow \{0, 1\}$ and $B : \{0, 1\}^{n \times n} \rightarrow \{0, 1\}$.

To motivate the verifier's tests, we describe what is expected from the oracles in the “completeness” case, i.e., when $C(w) = 1$. The input oracle, by definition, gives the string w , i.e., $W[i] = w_i$. Now let $z \in \{0, 1\}^n$ be the string of values of all the gates of the circuit C (including the input, the internal gates, and the output gate(s)). W.l.o.g., assume $z = w \circ y$, where y represents the values assumed for internal gates. The oracle A is expected to give the values of all linear functions at z (over $\text{GF}(2)$); and the oracle B is supposed to give the value of all quadratic functions at z . More precisely $A = A[x]_{x \in \{0, 1\}^n}$ is expected to be $A[x] = \sum_{i=1}^n x_i z_i = x^T z$ (where x and z are being thought of as column vectors). Similarly, $B = B[M]_{M \in \{0, 1\}^{n \times n}}$ is expected to be $B[M] = \sum_{i,j} M_{ij} z_i z_j = z^T M z$ (where M is an $n \times n$ matrix). In order to verify that w satisfies C , the verifier will verify that A and B have indeed been constructed according to some string z as above, that z represents an accepting computation of the circuit, and finally that A is the encoding of some string $w' \circ y$ where w' is close to the string w given by the input oracle W .

Tests. Given the circuit C , the verifier first constructs polynomials $P_1(z), \dots, P_n(z)$ as follows. Viewing the variables $\{z_i\}$ as representing the values at the individual gates of the circuit C (with $z_1, \dots, z_m$ being the input gates), the polynomial $P_i(z)$ is the quadratic polynomial (over $\text{GF}(2)$)

expressing the constraint imposed by the i 'th gate of the circuit on an accepting computation. For example:

$$P_i(z) = \begin{cases} z_i - z_j z_k & \text{if the } i\text{th gate is an AND gate with inputs from gates } j \text{ and } k. \\ z_i - z_j - z_k + z_j z_k & \text{if the } i\text{th gate is an OR gate with inputs from gates } j \text{ and } k. \\ z_i - (1 - z_j) & \text{if the } i\text{th gate is a NOT gate with input from gate } j. \\ z_i - (z_j + z_k) & \text{if the } i\text{th gate is a PARITY gate with inputs from gates } j \text{ and } k. \\ 1 - z_j & \text{if the } i\text{th gate is an output gate with input from gate } j. \\ 0 & \text{if the } i\text{th gate is an input gate (i.e. } i \leq m). \end{cases}$$

Note that $z = w \circ y$ reflects the computation of C on an acceptable input w iff $P_i(z) = 0$ for every $i \in [n]$. The verifier conducts the following tests:

Codeword tests: These tests refer to (A, B) being a valid encoding of some string $z \in \{0, 1\}^n$. That is, these tests check that both A and B are linear functions, and that B is consistent with A . In the latter check, the verifier employs a self-correction procedure (cf. [13]) to the oracle B . (There is no need to employ self-correction to A , because it is queried at random locations.)

Linearity of A : Pick x_1, x_2 uniformly at random from $\{0, 1\}^n$ and verify that $A[x_1 + x_2] = A[x_1] + A[x_2]$.

Linearity of B : Pick M_1, M_2 uniformly at random from $\{0, 1\}^{n \times n}$ and verify that $B[M_1 + M_2] = B[M_1] + B[M_2]$.

Consistency of A and B : Pick x_1, x_2 uniformly at random from $\{0, 1\}^n$ and M uniformly from $\{0, 1\}^{n \times n}$ and verify that $B[M + x_1 x_2^T] - B[M] = A[x_1]A[x_2]$.

Circuit test: This test checks that the string z encoded in (A, B) represents an accepting computation of C ; that is, that $P_i(z) = 0$ for every $i \in [n]$. The test checks that a random linear combination of the P_i 's evaluates to 0, while employing self-correction to A and B .

Pick $\alpha_1, \dots, \alpha_n \in \{0, 1\}$ uniformly and independently and let $\sum_{k=1}^n \alpha_k P_k(z) = c_0 + \sum_i \ell_i z_i + \sum_{i,j} Q_{i,j} z_i z_j$. Pick $x \in \{0, 1\}^n$ and $M \in \{0, 1\}^{n \times n}$ uniformly at random. Verify that $c_0 + (A[x + \ell] - A[x]) + (B[M + Q] - B[M]) = 0$.

Proximity test: This test checks that the m -bit long prefix of the string z , encoded in A , matches (or is close to) the input oracle W , while employing self-correction to A .

Pick $j \in [m]$ and $x \in \{0, 1\}^n$ uniformly. Let $e_j \in \{0, 1\}^n$ denote the vector that is 1 in the j th coordinate and 0 everywhere else. Verify that $W[j] = A[x + e_j] - A[x]$.

The verifier accepts if all the tests above accept, else it rejects.

Resources. The verifier uses $O(n^2)$ random bits and makes $O(1)$ binary queries.

Completeness. It is straightforward to see that if w , the string given by W satisfies C , then letting z be the set of values of the gates of C and letting $A[x] = x^T z$ and $B[M] = z^T M z$ will satisfy all tests above. Thus the verifier has perfect completeness.

Soundness (with proximity). It follows directly from the analysis of [4] that there exists a $\delta_0 > 0$ such that for every $\delta \leq \delta_0$, if the **Codeword tests** and the **Circuit test** above accept with probability at least $1 - \delta$ then the oracle A is 2δ -close to the Hadamard encoding of some string $z = w' \circ y$ such that $C(w')$ accepts. Now we augment this soundness with a proximity condition. Suppose the verifier also accepts the **Proximity test** with probability at least $1 - \delta$. Then we have that

$w_j \neq A[x+e_j] - A[x]$ with probability at most δ . Furthermore the events $A[x+e_j] \neq (x+e_j)^T z$, and $A[x] \neq x^T z$ happen with probability at most 2δ each. Thus, with probability at least $1 - 5\delta$ (over the possible choices of j and x), both $w_j = A[x+e_j] - A[x]$ and $A[x+e_j] - A[x] = (x+e_j)^T z - x^T z$ hold. Since $(x+e_j)^T z - x^T z = e_j^T z = z_j = w'_j$, it follows that, with probability at least $1 - 5\delta$ (over the choices of j), $w_j = w'_j$. In other words, the string w represented by the oracle W is at distance at most 5δ away from some string w' that is accepted by the circuit C . ■

B Randomness-efficient low-degree tests and the sampling lemma

Following [12], our construction makes heavy use of small-bias spaces [42] to save on randomness when choosing random lines. For a field F and parameters $m \in \mathbb{Z}^+$ and $\lambda > 0$, we require a set $S \subseteq F^m$ that is λ -biased (with respect to the additive group of F^m). Rather than define small-bias spaces here, we simply state the properties we need. (See, e.g., [12] for definitions and background on small-bias spaces.)

Lemma B.1 *For every F of characteristic 2, $m \in \mathbb{Z}^+$, and $\lambda > 0$, there is an explicit construction of a λ -biased set $S \subseteq F^m$ of size at most $(\log |F^m|)/\lambda^2$ [2].*

We now discuss the properties of such sets that we will use.

Expanding Cayley Graphs. λ -biased sets are very useful pseudorandom sets in algebraic applications, and this is due in part to the expansion properties of the Cayley graphs they generate:

Lemma B.2 *If $S \subseteq F^m$ is λ -biased and we let G_S be the graph with vertex set F^m and edge set $\{(x, x+s) : x \in F^m, s \in S\}$, then all the nontrivial eigenvalues of G_S have absolute value at most $\lambda|S|$.*

Randomness-Efficient Line Samplers. In [12], Lemma B.2 was used to prove the following sampling lemma. This lemma says that if one wants to estimate the density of a set $B \subseteq F^m$ using lines in F^m as the sample sets, one does not need to pick a random line in F^m which costs $2 \log |F^m|$ random bits. A pseudorandom line whose slope comes from an λ -biased set will do nearly as well, and the randomness is only $(1 + o(1)) \cdot \log |F^m|$. In what follows $l_{x,y}$ is the line passing through point x in direction y , formally: $l_{x,y} = \{x + ty : t \in F\}$

Lemma B.3 ([12], Sampling Lemma 4.3) *Suppose $S \subseteq F^m$ is λ -biased. Then, for any $B \subseteq F^m$ of density $\mu = |B|/|F^m|$, and any $\zeta > 0$,*

$$\Pr_{x \in F^m, y \in S} \left[\left| \frac{|l_{x,y} \cap B|}{|l_{x,y}|} - \mu \right| > \zeta \right] \leq \left(\frac{1}{|F|} + \lambda \right) \cdot \frac{\mu}{\zeta^2}.$$

Randomness-Efficient Low Degree Tests Ben-Sasson *et al.* [12] use the randomness-efficient Sampling Lemma B.3 to obtain randomness efficient low degree tests, by performing a “line vs. point” test only for pseudorandom lines with a direction y coming from a small λ -biased set. That is for a set $S \subseteq F^m$, we consider lines of the form $l_{x,y}(t) = x + ty$, for $x \in F^m$ and $y \in S$, and let $\mathbb{L}$ be the set of all such lines, where each line is parametrized in a canonical way.

Then for functions $f : F^m \rightarrow F$, and $g : \mathbb{L} \rightarrow P_d$, where P_d is the set of univariate polynomials of degree at most d over F , we let $\text{LDT}_{S,d}^{f,g}$ be the test that uniformly selects $l \xleftarrow{R} \mathbb{L} \triangleq \{l_{x,y} : x \in F^m, y \in S\}$ and $t \in F$, and accepts iff $g(l)(t) = f(l(t))$. That is, the value of the degree d univariate polynomial $g(l)$ at point t equals the value of f at $l(t)$. We quote their main theorem and will use it in our constructions.

Theorem B.4 ([12], Theorem 4.1) *There exists a universal constant $\alpha > 0$ such that the following holds. Let $d \leq |F|/3$, $m \leq \alpha|F|/\log |F|$, $S \subseteq F^m$ be a λ -biased set for $\lambda \leq \alpha/(m \log |F|)$, and $\delta \leq \alpha$. Then, for every $f : F^m \rightarrow F$ and $g : \mathbb{L} \rightarrow P_d$ such that f is at least 4δ -far from, any polynomial of degree at most md , we have the following:*

$$\Pr[\text{LDT}_{S,d}^{f,g} = \text{rej}] > \delta.$$