SIAM J. COMPUT.
Vol. 36, No. 4, pp. 889–974
© 2006 Society for Industrial and Applied Mathematics

ROBUST PCPS OF PROXIMITY, SHORTER PCPS, AND APPLICATIONS TO CODING*

ELI BEN-SASSON†, ODED GOLDREICH‡, PRAHLADH HARSHA§, MADHU SUDAN¶, AND SALIL VADHAN‖

Abstract. We continue the study of the trade-off between the length of probabilistically checkable proofs (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 $\exp(o(\log\log n)^2)\cdot n$ 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 is $2^{(\log n)^{\varepsilon}}$, for any $\varepsilon > 0$, in the second case (of constant query complexity). 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 the 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 used in PCP constructions, it does differ from previous constructions in fundamental ways, and in particular does not use the "parallelization" step of Arora et al. [*J. ACM,* 45 (1998), pp. 501–555]. 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 $\varepsilon > 0$.

Key words. probabilistically checkable proofs, PCP, locally testable codes, locally decodable codes, property testing

AMS subject classifications. 68Q17, 68P30, 94B60

DOI. 10.1137/S0097539705446810

---

*Received by the editors December 27, 2004; accepted for publication (in revised form) January 24, 2006; published electronically December 15, 2006. Preliminary versions of this paper have appeared in *Proceedings of the 36th ACM Symposium on Theory of Computing* [BGH+04a] and *Electronic Colloquium on Computational Complexity* [BGH+04b]. Part 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.

http://www.siam.org/journals/sicomp/36-4/44681.html

†Department of Computer Science, Technion – Israel Institute of Technology, Haifa, Israel (eli@cs.technion.ac.il).

‡Department of Computer Science, Weizmann Institute of Science, Rehovot, Israel (oded.goldreich@weizmann.ac.il). The research of this author was partially supported by the Israel Science Foundation (grant 460/05).

§Toyota Technological Institute, Chicago, IL 60637 (prahladh@tti-c.org). This work was done while the author was at the Massachusetts Institute of Technology. The work of this author was supported in part by NSF Award CCR-0312575.

¶Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 (madhu@mit.edu). The work of this author was supported in part by NSF Award CCR-0312575.

‖Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138 (salil@eecs.harvard.edu). This author's work was supported in part by NSF grant CCR-0133096, ONR grant N00014-04-1-0478, and by a Sloan Research Fellowship.

1. Introduction. Probabilistically checkable proofs (PCPs) [FGL+96, AS98, ALM+98] (aka holographic proofs [BFLS91]) are NP witnesses that allow efficient probabilistic verification based on probing few bits of the NP witness. The celebrated PCP theorem [AS98, ALM+98] 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. [Hås01, GLST98]).

Optimizing the query complexity of PCPs has attracted a lot of attention, motivated in part by the significance of query complexity for nonapproximability results (see, for example, [BGLR93, BGS95, Hås01, GLST98, ST00]). However, these works guarantee only that the new NP witness (i.e., the PCP) is of a 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 [BFLS91, PS94, HS00, GS02, BSVW03], and in this work we continue the latter research direction.

In our view, the significance of PCPs extends far beyond their applicability to deriving nonapproximability results. The mere fact that NP witnesses can be transformed into a format that supports superfast probabilistic verification is remarkable. From this perspective, the question of how much redundancy is introduced by such a transformation is fundamental. Furthermore, PCPs have been used not only to derive nonapproximability results but also for obtaining positive results (e.g., computationally sound (CS) proofs [Kil92, Mic00] and their applications [Bar01, CGH04]), and the length of the PCP affects the complexity of those applications.

In any case, the length of PCPs is also relevant to nonapproximability results; specifically, it affects their *tightness with respect to the running time* (as noted in [Sze99]). For example, suppose (exact) satisfiability (SAT) has complexity $2^{\Omega(n)}$. The original PCP theorem [AS98, ALM+98] implies only that approximating maximum satisfiability (MaxSAT) requires time 2^{n^α} for some (small) $\alpha > 0$. The work of [PS94] makes α arbitrarily close to 1, whereas the results of [GS02, BSVW03] 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 versus 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 [PS94] 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 $\widetilde{O}(n)$. That is, stretching the NP witness by only a polylogarithmic amount, allows us to dramatically reduce the number of bits read (from n to $n^{0.01}$). More precisely, see the following theorem.[3]

THEOREM 1.1 (implicit in [PS94]). *Satisfiability of circuits of size n can be probabilistically verified by probing an NP witness of length* $\mathrm{poly}(\log n) \cdot n$ *in* $n^{o(1)}$ *bit locations. In fact, for any integer 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* $\mathrm{poly}(\log n) \cdot n^{1/m}$.

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.

Our results show that the query complexity can be reduced dramatically if we are

[1] We 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.

[2] A 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.

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

willing to increase the length of the proof slightly. First, with a quasi-polylogarithmic stretch, the query complexity can be made double-logarithmic as follows.

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

We mention that the only prior work claiming query complexity below $\exp(\sqrt{\log n})$ (cf. [GS02, BSVW03]) required stretching the NP witness by at least an $\exp(\sqrt{\log n})$ factor. With approximately such a stretch factor, these works actually achieved constant query complexity (cf. [GS02, BSVW03]). 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 $\exp(o(\log\log n)^2)\cdot n$ 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 [GS02, BSVW03]) to $\exp(\log^{\varepsilon} n)$, for any $\varepsilon > 0$ as follows.

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 $n\cdot\mathrm{poly}(\log n)$ and constant query complexity.

1.2. New notions and main techniques. A natural approach to reducing the query complexity of the PCP provided by Theorem 1.1 is via the "proof composition" paradigm of [AS98]. However, that PCP (as constructed in [PS94]) does not seem amenable to composition when the parameter m is nonconstant.[4] The reason is that standard proof composition requires the "outer" proof system to make a constant number of multivalue oracle queries (or be converted to such), whereas this specific PCP does not have this property and we cannot afford the standard parallelization involved in a suitable conversion. 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 [PS94] and [HS00]. The latter construction was already improved in [GS02, BSVW03] to reduce the length of PCPs to $n\cdot 2^{\widetilde{O}(\sqrt{\log n})}$. 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 observing that the small-biased sample sets of [BSVW03] give an even more significant savings on the randomness complexity of low-degree tests than noticed in that 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 superconstant number of) functions and need to test if each is a low-degree polynomial. In prior results,

[4] Also for constant m, we get stronger quantitative results by using our new PCP construction as a starting point.

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 cannot use other "parallelization" techniques [FRS94, LS97, ALM$^+$98, GS00, Raz98], which were instrumental to the proof composition technique of [AS98]. In turn, this forces us to introduce a new variant of the proof composition method, which is much more flexible than the one of [AS98]. 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 [DR04]. (There are significant differences between the two methods, as explained in section 1.3, where we also discuss our method in relation to Szegedy's work [Sze99].)

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 NP 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 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 required only 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). Indeed, PCPs of proximity are related to holographic proofs [BFLS91] and to "PCP spot-checkers" [EKR04]; see a further discussion in section 1.3.

Robust PCPs. To discuss robust PCPs, let us review the soundness guarantee of standard (nonadaptive) 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 satisfies the predicate. In a *robust PCP* we strengthen the latter requirement. We require that the said assignment (of oracle bits) not only fail to satisfy the predicate but also be *far* from any assignment that satisfies the predicate.

Proof composition. The key observation is that our 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 of 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 required when deriving Theorem 1.2 using a nonconstant 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., $1/\varepsilon$). This, in turn, requires obtaining strong bounds on the robustness property of the ("robust") extension of Theorem 1.1.

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 [AS98, ALM$^+$98] can be easily derived using a much weaker and simpler variant of our basic robust PCP of proximity, and the said construction seems easier than the basic PCPs used in the proof composition of [AS98, ALM$^+$98].

In addition to their role in our proof composition method, PCPs of proximity also provide a good starting point for deriving improved locally testable codes (see the discussion in section 1.4). The relation of PCPs of proximity to "property testing" is further discussed in section 1.3.

1.3. Related work. As mentioned above, the notion of a PCP of proximity is related to notions that have appeared in the literature.

Relation to holographic proofs. First, the notion of a PCP of proximity generalizes the notion of holographic proofs set forth by Babai et al. [BFLS91]. 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*).

Relation to property testing. PCPs 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. [AS98, ALM$^+$98]. PCPs of proximity are a natural special case of the PCP spot-checkers defined by Ergün, Kumar, and Rubinfeld [EKR04]. On the other hand, PCPs of proximity extend property testing [RS96, GGR98]. 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 input on few bit locations (or at the very least on a sublinear 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. 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, a PCP of proximity

provides a notion of approximation for (probabilistic) proof-verification procedures. In both cases, approximation means that inputs in the language should be accepted (when accompanied by suitable proofs), while inputs that are far from the language should be rejected (no matter what false proof is provided).

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.

Relation to assignment testers and another proof composition method. As stated above, our proof composition method is related to the method discovered independently by Dinur and Reingold [DR04]. Both methods use the same notion of PCPs of proximity (which are called assignment testers in [DR04]). A key difference between the two methods is that, 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 nonconstant number of "bundles" such that robustness is satisfied at the bundle level).[5] We stress that we cannot afford the cost 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 [DR04] is combinatorial (albeit inapplicable in our context), whereas our bundling relies on the algebraic structure of our proof system.

Relation to Szegedy's work [Sze99]. Some of the ideas presented in the current work are implicit in Szegedy's work [Sze99]. In particular, notions of robustness and proximity are implicit in [Sze99], in which a robust PCP of proximity (attributed to [PS94]) is composed with itself in a way that is similar to our composition theorem. We note that Szegedy does not seek to obtain PCPs with improved parameters, but rather to suggest a framework for deriving nicer proofs of existing results such as those in [PS94]. Actually, he focuses on proving the main result of [PS94] (i.e., a PCP of nearly linear length and constant number of queries) using as a building block a robust PCP of proximity that has length $\widetilde{O}(n)$ and makes $\widetilde{O}(\sqrt{n})$ queries (plus the constant query PCP of [ALM+98]).

We note that the aforementioned robust PCP of proximity is not presented in [Sze99], but is rather attributed to [PS94]. Indeed, observe that Theorem 1.1 above (implicit in [PS94]) achieves $\widetilde{O}(n)$ length and $\widetilde{O}(\sqrt{n})$ queries when the parameter $m = 2$. Thus, Szegedy's assertion is that this PCP can be strengthened to be a robust PCP of proximity, similarly to our main construct (specifically, Theorem 3.1, specialized to $m = 2$). However, our main construct achieves stronger parameters than those claimed in [Sze99], especially with respect to robust soundness. Indeed,

[5] The 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 nonconstant 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 need only modify the 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.

the parameters claimed in [Sze99] allow only for the robust PCP of proximity to be composed with itself a constant number of times.[6] As mentioned above, a significant amount of our effort is aimed at ensuring that our robust PCP of proximity has sufficiently strong parameters to be composed a nonconstant number of times and, moreover, to ensure that the query complexity grows only linearly rather than exponentially with the number of compositions. (See section 3.2 for further explanation.)

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 error-correcting 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 [GS02, BSVW03]). 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 [GS02, BSVW03] 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 $\exp(\log^{\varepsilon} k) \cdot k$ as follows.

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 [KT00] and the best known locally decodable code has codewords of a length that is subexponential in the number of information bits. We relax the definition of locally decodable codes by requiring that, whenever few locations are corrupted, the decoder 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 a length that is almost linear in the number of information bits as follows.

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

1.5. Subsequent work. Since the presentation of our results, there has been considerable progress in the construction of short PCPs.

Ben-Sasson and Sudan [BS05] constructed "shorter" PCPs for NP at the cost of a slightly larger query complexity. More precisely, they construct PCPs of length $n \cdot \mathrm{poly}(\log n)$ (to prove satisfiability of circuits of size n) that can be verified by querying at most $\mathrm{poly}(\log n)$ bits of the proof. They achieve this improvement in

[6] In the language of section 2, his soundness and robustness parameters are unspecified functions of the proximity parameter. In retrospect, it seems that the ideas of [PS94] may lead to a robust PCP of proximity with robustness that is at best linearly related to the proximity parameter; this would make the query complexity increase exponentially with the number of compositions (as discussed in section 3.2).

length by constructing PCPs of proximity for a specific problem, verifying membership in a *Reed–Solomon* code (i.e., verifying if a given function is close to the evaluation of some univariate polynomial of specified degree). Their construction also yields locally testable codes with similar parameters.

More recently, Dinur introduced a novel gap amplification technique to yield a fully combinatorial proof of the PCP theorem [Din06]. In addition, by applying the gap-amplification technique to the PCP constructed by Ben-Sasson and Sudan [BS05], she obtained PCPs for NP of length $n \cdot \text{poly}(\log n)$ and verifiable with a constant number of probes into the proof. Thus while Ben-Sasson and Sudan [BS05] reduce the proof length while increasing the number of queries, Dinur shows how to reduce the query size back to a constant, thereby improving both our results and those of [BS05].

The PCP verifiers in the constructions of both our paper and [BS05] require time at least polynomial in the length of the proof, though the verifiers probe at most $\text{poly}(\log n)$ locations in the proof. In a subsequent paper [BGH$^+$05], we demonstrate that both these constructions can in fact be accompanied with *superefficient* verifiers, i.e., verifiers that run in time at most polylogarithmic in the length of the proof.

Finally, the question of whether there exist nearly linear-sized PCPs for NP that achieve strong query-soundness trade-offs (e.g., achieving the parameters of Håstad [Hås01]) remains open. The recent work of Moshkovitz and Raz [MR06] may be useful towards answering this.

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 (sections 2–4). 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 important, 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 (sections 5–8). 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 section 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 augment of the Hadamard-based PCP of [ALM+98], which we present in Appendix A.

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

1.7. Relation to previous versions of this work. The current version includes a discussion of Szegedy's work [Sze99], of which we were unaware when writing the first version [BGH+04b]. The relation of his work to ours is now discussed in section 1.3.

Section 4 has been extensively revised, adding formal definitions and providing more precise descriptions of the main constructions and proofs. In addition, we identified a weaker form of the definition of a relaxed locally decodable code, proved that it essentially implies the original form, and restructured our presentation accordingly (see section 4.2).

The parameters of Theorem 1.2 in this version are stronger (to a limited extent) than that of earlier versions of this paper [BGH+04b, BGH+04a]. More specifically, we show that satisfiability of circuits of size n can be verified by probing $o(\log\log n)$ bit locations in an NP witness of length $\exp(o(\log\log n)^2)\cdot n$ as opposed to an NP witness of length $\exp(\widetilde{O}(\log\log n)^2)\cdot n$, as was claimed in earlier versions. We achieve this strengthening by improving the robustness parameter of the ALMSS-type robust PCP of proximity (Theorem 7.2) constructed in Part II of this paper, taking advantage of the greater slackness allowed in the randomness complexity of this PCP. (ALMSS-type robust PCP of proximity is one of the PCPs of proximity constructed in Part II. It is so called as it has parameters similar to the PCP constructed in [ALM+98].)

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 $\textsc{CktVal} = \{(C,w) : C(w) = 1\}$; the NP-complete Circuit Satisfiability, defined as $\textsc{CktSAT} = \{C : \exists w C(w) = 1\}$; and the NP-complete Nondeterministic CircuitValue, defined as $\textsc{NCktVal} = \{(C,w) : \exists z C(w,z) = 1\}$. (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 from 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 $\overline{\Delta}(u,v) \triangleq |\{i : u_i \neq v_i\}|$. We will also use the t-repetition x^t of a string x to denote the string formed by concatenating t copies of x (i.e., $x^t = x \ldots t$ *times* $\ldots x$).

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

section 2.2 and 2.3, respectively). We then 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, not only for simplicity but also 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 $I = (i_1, \ldots, i_q)$ and a circuit $D : \{0,1\}^q \to \{0,1\}$ 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 $\pi|_I \triangleq (\pi_{i_1}, \ldots, \pi_{i_q})$, and accepts iff $D(\pi|_I) = 1$.
- *We write $(I, D) \xleftarrow{R} V(x)$ to denote the queries and circuit generated by V on input x and random coin tosses and write $(I, D) = V(x; R)$ 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 $s : \mathbb{Z}^+ \to [0,1]$, 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,*

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

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

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

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 CKTSAT $= \{C : \exists w \; C(w) = 1\}$. Recall that every language in $\mathbf{NTIME}(t(n))$ reduces to CKTSAT in time $O(t(n) \log t(n))$ (cf. [HS66, PF79, Coo88]), and so a nearly linear-sized PCP for CKTSAT implies PCPs for $\mathbf{NTIME}(t(n))$ of length nearly linear in $t(n)$ for every polynomial $t(n)$.

2.2. PCPs of proximity. We now present a relaxation of PCPs that verify only that the input is *close* to an element of the language. The advantage of this relaxation is that it allows for 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 count only the verifier's queries to the latter. Thus we consider languages consisting of pairs of strings, which we refer to as *pair languages*. One pair language to keep in mind is the CircuitValue problem CktVal $= \{(C, w) : C(w) = 1\}$. For a pair language L, we define $L(x) = \{y : (x, y) \in L\}$. For example, 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 $s, \delta : \mathbb{Z}^+ \to [0, 1]$, 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,*

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

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,

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

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 charged only 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 CktVal has a trivial PCP (since it is in P), a PCPP for it implies a PCP for CktSAT as follows.

Proposition 2.4. *If* CktVal *has a PCPP, then* CktSAT *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 for "$C \in$ CktSAT" can be taken to be $w \circ \pi'$, where w is a satisfying assignment to C and π' is a PCPP for $(C,w) \in$ CktVal. This proof π can be verified using the PCPP verifier. The key observation is that if $C \notin$ CktSAT, then there exists no w that is 1-close to CktVal(C), because the latter set is empty. ∎

Note that we obtain only a standard PCP for CktSAT rather than a PCP of proximity. Indeed, CktSAT 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 it is the language 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* CktVal *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 for $(C', w^t \circ z) \in$ CktVal satisfies 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 of 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 required only 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* [RS96, GGR98].[7] 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

[7] This notion of approximation (of decision problems) should not be confused with the approximation of (search) optimization problems, which is also closely related to PCPs [FGL+96, ALM+98].

$L = \{(n, G) : \text{the } n\text{-vertex graph } G \text{ is bipartite}\}$, where the first (i.e., explicit) input is used only to specify the length of the second (i.e., nonexplicit) input G, to which the tester has oracle access (measured in its query complexity). We comment that the formulation of pair languages allows us 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 PCPPs, 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 PCPPs. (Some of the following examples were pointed out in [EKR04].) In the adjacency matrix model (cf. [GGR98]), bipartiteness has a PCPP 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 to any generalized graph partition property (cf. [GGR98]) with an efficient one-sided tester.) In contrast, Bogdanov and Trevisan [BT04] 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 [GR02]): testing bipartiteness (resp., 3-colorability) of n-vertex graphs has query complexity $\Omega(\sqrt{n})$ [GR02] (resp., $\Omega(n)$ [BOT02]), but again a PCPP will use only $O(1/\delta)$ queries.

Another example comes from the domain of codes. For any good code (or "even" any code of linear distance), there exists a PCPP with constant queries for checking whether a given word is a codeword.[8] 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, Harsha, and Raskhodnikova [BHR05].

Needless to say, there may be interesting cases in which PCPPs do not outperform property testers.

Queries versus 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 previously in the literature [BFLS91, EKR04]. 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 $L' \in \mathrm{NP}$), 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 [DR04]; 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 [ALM+98]) 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 which the verifier reads from the oracle be merely rejected with high probability, we ask that the bits which

[8] Indeed, this is a special case of our extension of the result of Babai et al. [BFLS91], 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 PCPP is not part of the codeword.

the verifier reads be *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 $s, \rho : \mathbb{Z}^+ \to [0,1]$, 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,*

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

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 of 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, Rompel, and Sipser [FRS94] performs less poorly than usual (i.e., the 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" PCPP. 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 $r_{\text{out}}, r_{\text{in}}, d_{\text{out}}, d_{\text{in}}, q_{\text{in}} : \mathbb{N} \to \mathbb{N}$, and $\varepsilon_{\text{out}}, \varepsilon_{\text{in}}, \rho_{\text{out}}, \delta_{\text{in}} : \mathbb{N} \to [0,1]$, 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}.*
- CKTVAL *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, there exists a universal algorithm with black-box access to V_{out} and V_{in} that can perform the actions of V_{comp} (i.e., evaluating $(I, D) \leftarrow V_{\text{comp}}(x; R)$). On inputs of length n, this algorithm runs in time n^c for a 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 ith position in the part of the oracle that represents the outer verifier's proof oracle, and (R, j) denotes the jth position in the Rth auxiliary block (which represents the Rth possible proof oracle (for the inner verifiers), 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} are with respect to length $d_{\text{out}}(n)$. With these conventions, the following is the description of the composed verifier $V_{\text{comp}}(x)$:

1. Choose $R \overset{\text{R}}{\leftarrow} \{0,1\}^{r_{\text{out}}}$.
2. Run $V_{\text{out}}(x; R)$ to obtain $I_{\text{out}} = (i_1, \ldots, i_{q_{\text{out}}})$ and D_{out}.
3. Run $V_{\text{in}}(D_{\text{out}})$ (on random coin tosses) to obtain $I_{\text{in}} = ((b_1, j_1), \ldots, (b_{q_{\text{in}}}, j_{q_{\text{in}}}))$ 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 $\ell = 1, \ldots, q_{\text{in}}$, determine the queries of the composed verifier as follows:
   (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 jth bit in V_{in}'s input oracle is the jth bit in the input to D_{out}, which in turn is the i_jth 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 Rth possible proof oracle are directed to the corresponding locations in the auxiliary proof. Recall that the jth bit in the proof oracle that V_{in} is using to verify the claim referring to the outer verifier's coins R is the jth bit in the Rth block of the auxiliary proof.
5. Output $I_{\text{comp}} = (k_1, \ldots, k_{q_{\text{in}}})$ 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 oracle π_{out} with 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 it is also ρ_{in}-far from being accepted by V_{in} (and hence also by V_{comp}). $\square$

The above theorem can serve as a substitute for the original composition theorem in the derivation of the original PCP theorem [ALM+98]. Specifically, one simply needs to modify the (precomposition) verifiers of [ALM+98] 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 [ALM+98]). In addition, the PCPs [ALM+98] can easily be made PCPPs 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 PCPPs (and thus PCPs) 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 PCPPs.

Queries versus 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 versus 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 $L =$ 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 CIRCUIT VALUE, 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 CIRCUIT VALUE, let $C : \{0,1\}^m \to \{0,1\}$ be a circuit of size n that accepts only the all-zeroes string 0^m, for $m = \Omega(n)$. Then we have $(C, 0^m) \in$ 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 $\rho : \mathbb{Z}^+ \to [0,1]$, a PCP has* expected robustness *ρ if for every $x \notin L$, we have*

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

Expected robustness for PCPPs is defined analogously.

We now present several generic transformations regarding robustness and soundness. Although we state them only 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 versus 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 $\varepsilon \le \rho$, 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 (i.e., oblivious) sampler (cf. [Gol97]). Combined with Proposition 2.10, we get a (soundness) error reduction for robust PCPs. For example, using the expander-neighborhood sampler of [GW97], we have the following.

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 $s, \gamma : \mathbb{Z}^+ \to [0,1]$ then L has a PCP with*

- *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 as follows.

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 $\rho \cdot 2^r \ge 2$, then L has a PCP with*

- *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 in which 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 multiprover 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-sized circuits for encoding (cf. [Spi96]). Using more classical codes would only give decision complexity $d + \widetilde{O}(a \cdot q)$, 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 $\mathrm{ECC} : \{0,1\}^a \to \{0,1\}^b$ 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 $\mathrm{ECC}(\pi_i)$. (We note that including the original oracle simplifies the argument as well as frees us from assuming a noiseless decoding algorithm.)

Our new verifier $V'(\mathrm{x})$, 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 $I = (i_1, \ldots, i_q)$, V' will query positions $I' = ((0, i_1), \ldots, (0, i_q), (1, i_1), \ldots, (1, i_q))$. If V outputs a decision circuit $D : (\{0,1\}^a)^q \to \{0,1\}$, V' will output the circuit $D' : (\{0,1\}^a)^q \times (\{0,1\}^b)^q \to \{0,1\}$ defined by

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

where

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

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 be an accepting oracle for V' by setting $\tau_i = \mathrm{ECC}(\pi_i)$ 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 $\mathrm{ECC}(x)$ 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 $I = (i_1, \ldots, i_q)$. Suppose that $(\pi_{i_1}, \ldots, \pi_{i_q}, \tau_{i_1}, \ldots, \tau_{i_q})$ is $\alpha\beta$-close to some $(\pi'_{i_1}, \ldots, \pi'_{i_q},$

$\tau'_{i_1}, \ldots, \tau'_{i_q}$) 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 $(\pi'_{i_j}, \tau'_{i_j}) = (\pi'_{i_j}, \mathrm{ECC}(\pi'_{i_j}))$. For such j, the choice of β implies that $\mathrm{ECC}(\pi'_{i_j})$ 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 the $\hat{\pi}$'s are α-close to satisfying D, as desired. ∎

The usual "parallelization" paradigm of PCPs [LS97, ALM+98] 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 as follows.

PROPOSITION 2.14 (parallelization versus 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 parameters, 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 item 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 item 2, let V be a robust PCP verifier for L with the stated parameters. We use the usual query-reduction technique for PCPs [FRS94] 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.[9] Specifically, the 2-query non-Boolean PCP verifier V' is defined as follows:

- 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 $j \overset{\mathrm{R}}{\leftarrow} [q]$ and computes $(I, D) = V(x; R)$, where $I = (i_1, \ldots, i_q)$. It sets $I' = (R, i_j)$ (which means 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 $[D(\tau_R) = 1] \wedge [(\tau_R)_j = \pi_{i_j}]$.

[9] It may be more instructive (but 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 $s' = 1 - \rho \cdot d$.

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 $\mathrm{E}_R[\Delta(\tau_R, \pi|_I)]$). 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 versus proximity. Finally, for PCPPs, we prove that the robustness parameter is upper-bounded by the proximity parameter.

PROPOSITION 2.15 (robustness versus 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 $L = \textsc{CktVal}$, 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 taken 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 taken 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 for 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 Theorems 1.2 and 1.3. Our starting point is the following robust PCPP, which is constructed in Part II, sections 5–8.

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 $\delta \leq \gamma/c$, $\textsc{CktVal}$ has a robust PCPP (for circuits of size n) with the following parameters:*

- *randomness $\left(1 - \frac{1}{m}\right)\log n + O(m \log m) + O(\log\log n) + O(\log(1/\delta))$;*
- *decision complexity $n^{1/m} \cdot \mathrm{poly}(\log n, 1/\delta)$; which also upper-bounds the query complexity;*[10]
- *perfect completeness; and*
- *for proximity parameter δ, a verifier having robust-soundness error γ with robustness parameter $(1-\gamma)\delta$.*

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

We comment that the condition $\delta < \gamma/c$ merely means that we present robust PCPPs 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 values of the query and randomness complexities, which yield a proof length that is upper-bounded by $\mathrm{poly}(m^m \log n) \cdot n$ (for δ and γ as small as $1/\mathrm{poly}(m^m, \log n)$), as well as to the small values of the soundness error and the small deterioration of robustness with respect to proximity.

Note that the main construct (of Theorem 3.1) works only when n, the size of the input circuit, is not too small (more precisely, when $n^{1/m} \geq m^{cm}/\delta^3$). While constructing our short PCPs (via proof composition), we need robust PCPPs that work for even smaller values of n. For this purpose, we also construct the following robust PCPP (of Theorem 3.2) that has parameters similar to a PCP constructed in [ALM+98]. In comparison to the main construct (of Theorem 3.1), this PCPP is not as efficient in randomness (i.e., it has randomness complexity $O(\log n)$ rather than $(1 - o(1)) \log_2 n$). However, since we plan to use the latter (robust) PCPP only towards the final stages of composition, we can afford to pay this cost in randomness. Theorem 3.2 will be proved in the second part of this work by modifying the proof of Theorem 3.1. An alternate construction of this robust PCPP can be obtained by adding a suitable proximity test to the "parallelized PCPs" of [ALM+98].

THEOREM 3.2 (ALMSS-type robust PCPP). *For all $n \in \mathbb{Z}^+$ and $\delta \in (0,1)$,* CKTVAL *has a robust PCPP (for circuits of size n) with the following parameters:*

- *randomness $O(\log n)$;*
- *decision complexity* $\mathrm{poly} \log n$*, which also upper-bounds the query complexity;*
- *perfect completeness; and*
- *for proximity parameter δ, a verifier having robust-soundness error $1 - \Omega(\delta)$ with robustness parameter $\Omega(1)$.*

Theorems 3.1 and 3.2 differ also in their robustness parameters. Theorem 3.2 provides a better bound on the robustness parameter (i.e., $\Omega(1)$ rather than $(1-\gamma)\delta$ provided by Theorem 3.1), while guaranteeing only a much weaker robust-soundness error (i.e., $1-\Omega(\delta)$ rather than γ), where $\gamma > \delta > 0$ is typically small. It is instructive to compare the expected robustness provided by the two results: The expected robustness in Theorem 3.1 is at least $(1-\gamma)^2\delta$, while that in Theorem 3.2 is $\Omega(\delta)\cdot\Omega(1) = \Omega(\delta)$. Thus, for $\gamma \ll 1$, the expected robustness in Theorem 3.1 can be very close to the proximity parameter δ (which is close to optimal; see Proposition 2.15), whereas in Theorem 3.2 the expected robustness is always a constant factor smaller than the proximity parameter. Hence, the robust PCPP of Theorem 3.1 is suitable for a large number of proof composition operations, whereas the one in Theorem 3.2 is useful when the query complexity of the outer verifier is already very small (and Theorem 3.1 can no longer be applied). Indeed, this is exactly how these two theorems are used in the construction of our short PCPs. Using Theorems 3.1 and 3.2, we derive a general trade-off between the length of PCPs and their query complexity as follows.

THEOREM 3.3 (randomness versus query complexity trade-off for PCPPs). *For every parameter $n, t \in \mathbb{N}$ such that $3 \leq t \leq \frac{2 \log\log n}{\log\log\log n}$ there exists a PCPP for* CKTVAL *(for circuits of size n) with the following parameters:*

- *randomness complexity $\log_2 n + A_t(n)$, where*

$$A_t(n) \triangleq O(t + (\log n)^{1/t}) \log\log n + O((\log n)^{2/t}); \tag{3.1}$$

- *query complexity $O(1)$;*
- *perfect completeness; and*

- *soundness error* $1 - \Omega(1/t)$ *with respect to proximity parameter* $\Theta(1/t)$.

Alternatively, we can have query complexity $O(t)$ *and soundness error* $1/2$ *maintaining all other parameters.*

For $t \in [3, \ldots, \frac{0.99 \log\log n}{\log\log\log n}]$, we have $(\log n)^{1/t} > (\log\log n)^{1/0.99}$, and thus $A_t(n) = O((\log n)^{2/t})$. On the other hand, for $t \geq \frac{1.01 \log\log n}{\log\log\log n}$, we have $(\log n)^{1/t} \leq (\log\log n)^{1/1.01}$, and thus $A_t(n) = O\big(\frac{(\log\log n)^2}{\log\log\log n}\big) = o(\log\log n)^2$.

Theorem 3.3 actually asserts a PCPP (for CktVal), but a PCP for CktSAT and a PCPP for Nondeterministic Circuit Value (of the same complexity) follow; see Propositions 2.4 and 2.5. Theorems 1.2 and 1.3 follow by suitable settings of the parameter t. Further details as well as a corollary appear in section 3.2.

3.1. Proof of Theorem 3.3. Theorem 3.3 is proved by using the robust PCPP described in Theorem 3.1. Specifically, this robust PCPP 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 PCPP with itself $O(t(n))$ times, we compose the resulting robust PCP with the ALMSS-type robust PCPP thrice to reduce the query complexity to $\operatorname{poly}\log\log\log n$. Finally, we compose this resultant robust PCPP with a PCPP parameter roughly $\Omega(1/t)$ that has query complexity $O(1)$ and exponential length. The latter PCPP can be obtained by a suitable modification of the Hadamard-based PCP of [ALM+98], as shown in Appendix A. We now turn to the actual proof.

Proof. We construct the PCPP of Theorem 3.3 by composing the robust PCPP 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 $m = (\log n)^{1/t}$ 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 last 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 $n^{1/m^{t-2}} = (n^{1/m^t})^{m^2} = 2^{m^2} = \exp(\log^{2/t} n) \ll n$, while the soundness error is bounded away from 1 and robustness is $\Omega(1/t)$. 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 $\operatorname{poly}\log\log\log n$ (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 the 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 the parameters of this verifier: For circuits of size ℓ and any proximity parameter $\delta_0 \in (\gamma/3c, \gamma/c)$, 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 $d_0(\ell) \triangleq \ell^{\frac{1}{m}} \cdot \mathrm{poly}(\log \ell, t)$, 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.[11] While doing so, we will use the largest possible proximity parameter for the inner verifier (V_0) in each step; that is, in the ith 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 $V_1, \ldots, V_{t-2}$ 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 $i < t-2$).[12] We first claim that the decision complexity, proximity, soundness-error, robustness, and proof length parameters satisfy the following conditions:

(1) Decision complexity: $d_i(n) \leq a(n,m)^2 \cdot n^{1/m^i}$, where $a(\ell, m) \triangleq d_0(\ell)/\ell^{1/m} = \mathrm{poly}(\log \ell, t)$. 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 $b(\ell, m) \triangleq 2^{r_0(\ell)} \cdot d_0(\ell)/\ell = \mathrm{poly}(m^m, \log \ell, t)$.

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, $d_1(n) \leq \mathrm{poly}(\log n, t) \cdot n^{1/m}$ and $2^{r_1(n)} d_1(n) \leq \mathrm{poly}(m^m, \log n, t) \cdot n$. Turning to the induction step, assuming that

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

[12] We 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 $\delta_0 = \rho_i \leq \rho_0 \leq \gamma/c$. Thus, the only condition that needs to be verified is $d_i^{1/m} \geq m^{cm}/(\gamma\rho_i)^3$ for $i < t-2$.

these claims holds for V_i, we prove that they hold also for V_{i+1}. For (1), note that

$$\begin{aligned} d_{i+1}(n) &= d_0(d_i(n)) && \text{[by the composition theorem]} \\ &= a(d_i(n), m) \cdot d_i(n)^{1/m} && \text{[by definition of } a(\cdot,\cdot)\text{]} \\ &\le a(n, m) \cdot d_i(n)^{1/m} && \text{[by monotonicity of } a(\cdot,\cdot) \text{ and } d_i(n) \le n\text{]} \\ &\le a(n, m) \cdot \left(a(n, m)^2 \cdot n^{1/m^i}\right)^{1/m} && \text{[by induction]} \\ &\le a(n, m)^2 \cdot n^{1/m^{i+1}} && \text{[using } m \ge 2\text{]} \end{aligned}$$

and $d_{i+1}(n) \ge d_i(n)^{1/m} \ge n^{1/m^{i+1}}$ 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, we see that 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, $\rho_i = (1-\gamma)^i\delta \ge (1-\frac{1}{t})^t\delta \ge e^{-1}\delta = \gamma/O(1)$. We now move to the last condition (essentially bounding the randomness). Notice first that $r_{i+1}(n) = r_i(n) + r_0(d_i(n))$, and thus

$$\begin{aligned} 2^{r_{i+1}(n)} \cdot d_{i+1}(n) &= 2^{r_i(n)} \cdot 2^{r_0(d_i(n))} \cdot d_0(d_i(n)) && \text{[by the composition theorem]} \\ &= 2^{r_i(n)} \cdot d_i(n) \cdot b(d_i(n), m) && \text{[by definition of } b(\cdot,\cdot)\text{]} \\ &\le b(n, m)^i \cdot n \cdot b(n, m) && \text{[by induction and monotonicity of } b(\cdot,\cdot)\text{]} \\ &= n \cdot b(n, m)^{i+1}. \end{aligned}$$

Thus, Part (5) is verified. Recall that we have to verify that $d_i^{1/m} \ge m^{cm}/(\gamma\rho_i)^3$ for $i < t-2$ as promised. We have $d_i^{1/m} \ge (n^{1/m^i})^{1/m} = n^{1/m^{i+1}} \ge n^{1/m^{t-2}}$ (since $i < t-2$). Since $m = (\log n)^{1/t}$, we have $n^{1/m^t} = 2$. Hence, $d_i^{1/m} \ge (n^{1/m^t})^{m^2} = 2^{m^2}$. On the other hand, $m^{cm}/(\gamma\rho_i)^3 \le m^{cm}/(e^{-1}\gamma\delta)^3 = m^{cm} \cdot \mathrm{poly}(t)$ because $\delta = \gamma/c$ and $\gamma = 1/t$. Thus it suffices to verify that $2^{m^2}/m^{cm} \ge \mathrm{poly}(t)$ for $2 \le t \le 2\log\log n/\log\log\log n$, which is straightforward.[13]

Lastly, we consider the running time of V_i, denoted T_i, which ought to be polynomial. A careful use of the composition theorem (Theorem 2.7) indicates that $T_i(n) = \mathrm{poly}(n) + T_{i-1}(n)$ for every $i = 2, \ldots, t-2$, where $T_1(n) = \mathrm{poly}(n)$ (since $V_1 = V_0$). Alternatively, unraveling the inductive composition, we note that V_i consists of invoking V_0 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 $i \le t-2 < n$ 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 $d_{t-2} \ge 1$), we have $r_c(n) = r_{t-2}(n) \le \log n + (t-2)\cdot\log b(n,m) \le \log n + t\log b(n,m)$. 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 $r_c(n) \le \log_2 n + O(t\cdot\log\log n) + t\cdot O(m\log m) = \log_2 n + O(t + (\log n)^{\frac{1}{t}})\cdot\log\log n$. The decision complexity of V_c is $d_c(n) = d_{t-2}(n) \le a(n,m)^2 \cdot n^{1/m^{t-2}} = a(n,m)^2 \cdot 2^{m^2}$, because $n^{1/m^t} = 2$. Using

[13]Note 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 $\mathrm{poly}(t)$ is $\mathrm{poly}(\log\log n/\log\log\log n) = \mathrm{poly}(\log\log n)$. 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 $2^{\sqrt{\log\log n}^2/2} = \sqrt{\log n} \gg \mathrm{poly}(\log\log n)$.

$a(n,m) = \text{poly}(\log n, t)$, 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 $s_c = s_{t-2} = 1-(1-\gamma)^{t-2} = 1-(1-1/t)^{t-2} < 1-\Omega(1)$, 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; that is, V' equals V_c composed with V_a, whereas V'' equals V' composed with V_a, and V''' equals V'' composed with V_a. We apply composition as before, setting the proximity parameter of the inner verifier to equal the robustness parameter of the outer verifier. Recall from Theorem 3.2 that the ALMSS-type verifier V_a has the following parameters: randomness $r_a(\ell, \delta_a) = O(\log \ell)$, decision complexity $d_a(\ell, \delta_a) = \text{poly}\log \ell$, soundness error $s_a(\ell, \delta_a) = 1 - \Omega(\delta_a)$, and robustness $\rho_a(\ell, \delta_a) = \Omega(1)$ for input size ℓ and proximity parameter δ_a. Recall that when composing V_c with V_a we set $\delta_a = \rho_c = \Omega(1/t)$, whereas when composing V' (resp., V'') with V_a we set $\delta_a = \rho' = \Omega(1)$ (resp., $\delta_a = \rho'' = \Omega(1)$). Each composition with the inner verifier V_a adds $O(\log d)$ to the randomness, while reducing the query complexity to $\text{poly}\log d$, where d is the decision complexity of the outer verifier. Furthermore, when composing any of these outer verifiers (i.e., either V_c, V', or V'') with V_a, the resulting verifier has robustness parameter $\Omega(1)$ while its robust-soundness error is $1 - \Omega((1-s)\rho)$, where ρ and s are the robustness parameter and soundness error of the outer verifier. 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 $\rho_c = \Omega(\delta)$):

$r' = r_c + O(\log d_c(n)) = r_c + O(m^2 + \log\log n)$,
$d' = \text{poly}(\log d_c(n)) = \text{poly}(m, \log\log n)$,
$s' = 1 - \Omega((1-s_c)\rho_c) = 1 - \Omega(\delta)$,
and $\rho' = \Omega(1)$.

Parameters of V'':

$r'' = r' + O(\log d') = r' + O(\log m + \log\log\log n)$,
$d'' = \text{poly}(\log d') = \text{poly}(\log m, \log\log\log n)$,
$s'' = 1 - \Omega((1-s')\rho') = 1 - \Omega(\delta)$,
and $\rho'' = \Omega(1)$.

Parameters of V''':

$r''' = r'' + O(\log d'') = r'' + O(\log\log m + \log\log\log\log n)$,
$d''' = \text{poly}(\log d'') = \text{poly}(\log\log m, \log\log\log\log n)$,
$s''' = 1 - \Omega((1-s'')\rho'') = 1 - \Omega(\delta)$,
and $\rho''' = \Omega(1)$.

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

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

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

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 $r_f(n) \triangleq r'''(n) + r_h(q'''(n)) = r'''(n) + \text{poly}(\log\log\log n)$, because $r_h(\ell) = O(\ell^2)$. Thus, $r_f(n) = \log_2 n + O(t + (\log n)^{1/t}) \cdot \log\log n + O((\log n)^{2/t})$. On proximity parameter δ_h, the soundness error of V_h is $s_h = 1 - \Omega(\delta_h)$. Setting $\delta_h = \rho''' = \Omega(1)$, we conclude that

the soundness error of V_f on proximity parameter δ is $1 - \Omega(\delta) = 1 - \Omega(1/t)$, since the soundness error of V''' is $1 - \Omega(\delta)$.

To obtain soundness error $1/2$, we perform $O(t)$ repetitions of V_f, 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 [Gol97, Apdx. C.4]). ∎

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 PCPPs, the proximity parameter of the second must be upper-bounded by the robustness parameter of the first. Thus, when we compose many robust PCPPs, 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 PCPP 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 to each proof oracle encountered in the composition an 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.

3.2. Corollary to Theorem 3.3. Recall that Theorem 3.3 asserts a PCPP (for CktVal) with randomness complexity $\log_2 n + A_t(n)$, where $A_t(n) \triangleq O(t + (\log n)^{1/t}) \log\log n + O((\log n)^{2/t})$ and query complexity $O(t)$ (for soundness error $1/2$). For constant $t \geq 3$, we have $A_t(n) = O((\log n)^{2/t})$. On the other hand, for $t \geq \frac{1.01 \log\log n}{\log\log\log n}$, we have $A_t(n) = o(\log\log n)^2$. Using Proposition 2.4, these PCPPs yield corresponding PCPs for CktSAT.

Deriving Theorems 1.2 *and* 1.3. 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 $O(\frac{1}{\varepsilon})$ 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 $O\left(\frac{\log\log n}{\log\log\log n}\right)$ and randomness $\log_2 n + O\left(\frac{(\log\log n)^2}{\log\log\log n}\right)$. Thus, both Theorems 1.2 and 1.3 follow from Theorem 3.3.

Deriving a PCPP for Nondeterministic Circuit Value. By Proposition 2.5, we conclude that for every $3 \leq t(n) \leq \frac{2\log\log n}{\log\log\log n}$, there exists a PCPP 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 $\delta = \Omega(1/t(n))$).

A more flexible notion of a PCPP. Our definition of a PCPP (see Definition 2.3) specifies for each system a unique proximity parameter. In many settings (see, e.g., section 4.1), it is better that the proximity parameter is given as an input to the verifier and that the latter behaves accordingly (i.e., makes an adequate number of

queries). We refrain from presenting either a formal definition of such relaxed PCPPs or a general transformation of PCPPs into their more relaxed form. Instead, we state the following corollary to Theorem 3.3.

COROLLARY 3.4. *For all parameters* $n, t_1, t_2 \in \mathbb{N}$ *such that* $3 \leq t_1 \leq t_2 \leq \frac{2 \log\log n}{\log\log\log n}$ *there exists a PCPP for* CKTVAL (*for circuits of size* n) *with proof length* $2^{A_{t_1}(n)} \cdot n$, *where* $A_t(n)$ *is as in* (3.1), *query complexity* $O(t_2)$, *perfect completeness, and soundness error* $1/2$ *with respect to proximity parameter* $1/t_2$. *Furthermore, when given* (as auxiliary input) *a proximity parameter* $\delta \geq 1/t_2$, *the verifier makes only* $O(\max\{1/\delta, t_1\})$ *queries and rejects any input oracle that is* δ*-far from satisfying the circuit with probability at least* $1/2$.

Underlying the following proof is a general transformation of PCPPs to the more relaxed form as stated in Corollary 3.4.

Proof. The proof oracle consists of a sequence of proofs for the system of Theorem 3.3, when instantiated with proximity parameter 2^{-i} for $i = \lfloor \log_2 t_1 \rfloor, \ldots, \lceil \log_2 t_2 \rceil$. When the new verifier is invoked with proximity parameter δ, it invokes the original verifier with proximity parameter 2^{-i}, where $i = \lceil \log_2 1/\delta \rceil$, and emulates the answers using the ith portion of its proof oracle. ∎

4. Applications to coding problems. In this section we show that, combined with any good code, any PCPP yields a locally testable code (LTC). Using our PCPPs, we obtain an improvement in the rate of LTCs (improving over the results of [GS02, BSVW03]). We also introduce a relaxed notion of locally decodable codes (LDCs) and show how to construct such codes using any PCPP (and ours in particular).

Preliminaries For a string $w \in \{0,1\}^n$ and $i \in [n] \triangleq \{1, 2, \ldots, n\}$, unless stated differently, w_i denotes the ith bit of w.

We consider codes mapping sequences of k (input) bits into sequences of $n \geq k$ (output) bits. Such a generic code is denoted by $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$, and the elements of $\{\mathrm{C}(x) : x \in \{0,1\}^k\} \subseteq \{0,1\}^n$ 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 $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$ is the minimum (Hamming) distance between its codewords; that is, $\min_{x \neq y}\{\overline{\Delta}(\mathrm{C}(x), \mathrm{C}(y))\}$, 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 $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$ of distance $\Omega(n)$. The *distance of* $w \in \{0,1\}^n$ *from a code* $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$, denoted $\overline{\Delta}_{\mathrm{C}}(w)$, is the minimum distance between w and the codewords, that is, $\overline{\Delta}_{\mathrm{C}}(w) \triangleq \min_x\{\overline{\Delta}(w, \mathrm{C}(x))\}$. For $\delta \in [0,1]$, the n-bit long strings u and v are said to be δ*-far* (resp., δ*-close*) if $\overline{\Delta}(u, v) > \delta \cdot n$ (resp., $\overline{\Delta}(u, v) \leq \delta \cdot n$). Similarly, w is δ*-far from* C (resp., δ*-close to* C) if $\overline{\Delta}_{\mathrm{C}}(w) > \delta \cdot n$ (resp., $\overline{\Delta}_{\mathrm{C}}(w) \leq \delta \cdot n$).

As in the case of PCPs, all oracle machines considered below are *nonadaptive*. Here these oracle machines will model highly efficient testing and decoding procedures, which probe their input $w \in \{0,1\}^n$ in relatively few places. Thus, these procedures are modeled as oracle machines having oracle access to w (which is viewed as a function $w : \{1, \ldots, n\} \to \{0,1\}$).

4.1. LTCs. Loosely speaking, by a *codeword test* (for the code $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$) we mean a randomized (nonadaptive) oracle machine, also called a *tester*,

that is given oracle access to $w \in \{0,1\}^n$. 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.) The first definition below provides a general template (in terms of several parameters) for the rejection condition. Later we will discuss the kinds of asymptotic parameters we would like to achieve.

DEFINITION 4.1 (codeword tests). *A randomized* (nonadaptive) *oracle machine M is called a (δ, s)-codeword test for* $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$ *if it satisfies the following two conditions:*

1. Accepting codewords (i.e., completeness)*: For every $x \in \{0,1\}^k$, given oracle access to $w = \mathrm{C}(x)$, machine M accepts with probability* 1*. That is,* $\Pr[M^{\mathrm{C}(x)} = 1] = 1$ *for every $x \in \{0,1\}^k$.*

2. Rejection of noncodeword (i.e., soundness)*: Given oracle access to any $w \in \{0,1\}^n$ that is δ-far from* C*, machine M accepts with probability at most s. That is,* $\Pr[M^w = 1] \leq s$ *for every $w \in \{0,1\}^n$ that is δ-far from* C*.*

The parameter δ is called the proximity parameter *and s is called the* soundness error*. The* query complexity *q of M is the maximum number of queries it makes (taken over all sequences of coin tosses).*

Note that *this definition requires nothing with respect to noncodewords that are relatively close to the code* (i.e., are δ-close to C). In addition to the usual goals in constructing error-correcting codes (e.g., maximizing minimum distance and minimizing the blocklength $n = n(k)$), here we are also interested in simultaneously minimizing the query complexity q, the proximity parameter δ, and the soundness error s. More generally, we are interested in the trade-off between q, δ, and s. (As usual, the soundness error can be reduced to s^k by increasing the query complexity to $k \cdot q$.) A minimalistic goal is to have a family of codes with q, δ, and s all fixed constants. However, note that this would be interesting only if δ were sufficiently small with respect to the distance parameters of the code, e.g., smaller than half the relative minimum distance. (For example, if δ is larger than the "covering radius" of the code, then there does not exist any string that is δ-far from the code, and the soundness condition becomes vacuous.) A stronger definition requires the tester to work for any *given* proximity parameter $\delta > o(1)$, but allows its query complexity to depend on δ as follows.

DEFINITION 4.2 (LTCs). *A family of codes $\{\mathrm{C}_k : \{0,1\}^k \to \{0,1\}^n\}_{k \in \mathbb{N}}$ is* locally testable *if it satisfies the following.*

1. Linear distance*: There is a constant $\rho > 0$, such that for every k, C_k has minimum distance at least $\rho \cdot n$.*

2. Local testability*: There is a randomized, nonadaptive oracle machine M such that for every constant $\delta > 0$, there is a constant $q = q(\delta)$ such that for all sufficiently large k, $M^w(1^k, \delta)$ is a $(\delta, 1/2)$-codeword test for C_k with query complexity q.*

The family is called explicit *if both C_k and $M^w(1^k, \delta)$ can be evaluated with computation time polynomial in k.*

We comment that Definition 4.2 is somewhat weaker than the definitions used in [GS02].[14]

[14] In the stronger of the definitions in [GS02], the tester is not given δ as input (and thus has query complexity that is a fixed constant independent of δ) but is required to be a $(\delta, 1 - \Omega(\delta))$-codeword test for every constant $\delta > 0$ and sufficiently large k. That is, strings that are δ-far from the code

Using an adequate PCPP, we can transform any code to a related code that has a codeword tester. This is done by appending each codeword with a PCPP proving the codeword is indeed the encoding of a message. One technical problem that arises is that the PCPP constitutes most of the length of the new encoding. Furthermore, we cannot assume much about the Hamming distance between different proofs of the same statement, and thus the distance of the new code may deteriorate. But this is easily fixed by repeating the codeword many times so that the PCPP constitutes only a small fraction of the total length.[15] Specifically, given a code $\mathrm{C}_0 : \{0,1\}^k \to \{0,1\}^m$, we consider the code $\mathrm{C}(x) \triangleq (\mathrm{C}_0(x)^t, \pi(x))$, where $t = (d(k)-1) \cdot |\pi(x)|/|\mathrm{C}_0(x)|$ such that (say) $d(k) = \log k$, and $\pi(x)$ is a PCPP that asserts that an m-bit string (given as an input oracle) is a codeword (of C_0).

CONSTRUCTION 4.3. *Let d be a free parameter to be determined later, let $\mathrm{C}_0 : \{0,1\}^k \to \{0,1\}^m$ be a code, and let V be a PCPP verifier for membership in $S_0 = \{\mathrm{C}_0(x) : x \in \{0,1\}^k\}$. Let $\pi(x)$ be the proof oracle corresponding to the claim that the input oracle equals $\mathrm{C}_0(x)$; that is, $\pi(x)$ is the canonical proof obtained by using x as an NP witness for membership of $\mathrm{C}_0(x)$ in S_0. Consider the code $\mathrm{C}(x) \triangleq (\mathrm{C}_0(x)^t, \pi(x))$, where $t = (d-1) \cdot |\pi(x)|/|\mathrm{C}_0(x)|$.*

The codeword test emulates the PCP verifier in the natural way. Specifically, given oracle access to $w = (w_1, \ldots, w_t, \pi) \in \{0,1\}^{t \cdot m + \ell}$, the codeword tester selects uniformly $i \in [t]$ and emulates the PCP verifier, providing it with 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 randomly selected positions in some q_{rep} randomly selected pairs of m-bit long blocks, where q_{rep} is a free parameter to be optimized later). Let us denote this tester by T. That is, T^w

1. uniformly selects $i \in [t]$ and invokes $V^{w_i, \pi}$.
2. repeats the following q_{rep} times: Uniformly selects $i_1, i_2 \in [t]$ and $j \in [m]$ and checks whether $(w_{i_1})_j = (w_{i_2})_j$.

PROPOSITION 4.4. *Let d and q_{rep} be the free parameters in the above construction of the code C and tester T. Suppose that the code $\mathrm{C}_0 : \{0,1\}^k \to \{0,1\}^m$ has a relative minimum distance of ρ_0, and that the PCPP has a proof length of $\ell > m$, soundness error $1/4$ for proximity parameter δ_{pcpp}, and query complexity q_{pcpp}. Then, the code C and tester T have the following properties:*

1. *The blocklength of C is $n \triangleq d \cdot \ell$ and its relative minimum distance is at least $\rho_0 - 1/d$.*
2. *The oracle machine T is a $(\delta, \frac{1}{2})$-codeword tester for the C, where $\delta = \delta_{\mathrm{pcpp}} + \frac{4}{q_{\mathrm{rep}}} + \frac{1}{d}$.*
3. *The query complexity of T is $q = q_{\mathrm{pcpp}} + 2q_{\mathrm{rep}}$.*

Proof. The parameters of the code C are obvious from the construction. In particular, C has blocklength $t \cdot m + \ell = d \cdot \ell = n$, and the PCPP $\pi(x)$ constitutes only an $\ell/n = 1/d$ fraction of the length of the codeword $\mathrm{C}(x)$. Since the remainder consists of replicated versions of $\mathrm{C}_0(x)$, it follows that the relative minimum distance of C is at least $(n-\ell)\rho_0/n > \rho_0 - 1/d$.

The query complexity of T is obvious from its construction, and so we need only show that it is a good codeword tester. Completeness follows immediately from

are rejected with probability $\Omega(\delta)$. Such a tester implies a tester as in Definition 4.2, with query complexity $q(\delta) = O(1/\delta)$.

[15] Throughout 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.

the completeness of the PCPP, and so we focus on the soundness condition. We consider an arbitrary $w = (w_1, \ldots, w_t, \pi) \in \{0,1\}^{t \cdot m + \ell}$ that is δ-far from C, and observe that $w' = (w_1, \ldots, w_t)$ must be δ'-far from $\mathrm{C}' = \{\mathrm{C}_0(x)^t : x \in \{0,1\}^k\}$, where $\delta' \geq (\delta n - \ell)/n = \delta - (1/d)$. Let $u \in \{0,1\}^m$ be a string that minimizes $\Delta(w', u^t) = \sum_{i=1}^t \Delta(w_i, u)$; that is, u^t is the "repetition sequence" closest to w'. We consider two cases:

Case 1. $\Delta(w', u^t) \geq 1/q_{\mathrm{rep}}$. In this case, a single execution of the basic repetition test (comparing two locations) rejects with probability:

$$\begin{aligned} \mathrm{E}_{r,s \in [t]} \left[\Delta(w_r, w_s)\right] &\geq \mathrm{E}_{r \in [t]} \left[\Delta(w_r, u)\right] \\ &= \Delta(w', u^t) \\ &\geq 1/q_{\mathrm{rep}}, \end{aligned}$$

where the last inequality is due to the case hypothesis. It follows that q_{rep} executions of the repetition test would accept with probability at most $(1 - 1/q_{\mathrm{rep}})^{q_{\mathrm{rep}}} < 1/e < 1/2$.

Case 2. $\Delta(w', u^t) \leq 1/q_{\mathrm{rep}}$. In this case

$$\Delta_{\mathrm{C}_0}(u) = \Delta_{\mathrm{C}'}(u^t) \geq \Delta_{\mathrm{C}'}(w') - \Delta(w', u^t) \geq \delta' - \frac{1}{q_{\mathrm{rep}}},$$

where the last inequality is due to the case hypothesis. Also, recalling that on the average (i.e., average i) w_i is $1/q_{\mathrm{rep}}$-close to u, it holds that at least two-thirds of the w_i's are $3/q_{\mathrm{rep}}$-close to u. Recalling that u is $(\delta' - (1/q_{\mathrm{rep}}))$-far from C_0 and using $\delta_{\mathrm{pcpp}} = \delta' - (4/q_{\mathrm{rep}})$, it follows that at least two-thirds of the w_i's are δ_{pcpp}-far from C_0. Thus, by the soundness condition of the PCP of proximity, these w_i will be accepted with probability at most $1/4$. Thus, in the current case, the tester accepts with probability at most $\frac{1}{3} + \frac{2}{3} \cdot \frac{1}{4} = \frac{1}{2}$. The soundness condition follows. ∎

To prove Theorem 1.4, we instantiate the above construction as follows. We let $\mathrm{C}_0 : \{0,1\}^k \to \{0,1\}^m$ come from a family of codes with a constant relative minimum distance of $\rho_0 > 0$ and nearly linear blocklength $m = \widetilde{O}(k)$, where encoding can be done by circuits of nearly linear size $s_0 = s_0(k) = \widetilde{O}(k)$. We take the PCPP from Corollary 3.4, setting $t_1 = O(1/\varepsilon)$ (for an arbitrarily small constant $\varepsilon > 0$) and $t_2 = 2 \log \log s_0 / \log \log \log s_0 = \omega(1)$. Thus, we obtain proof length $\ell = s_0 \cdot \exp(\log^{\varepsilon/2} s_0)$ and query complexity $q_{\mathrm{pcpp}} = O(\max\{1/\delta_{\mathrm{pcpp}}, t_1\}) = O(1/\delta_{\mathrm{pcpp}})$ for any proximity parameter $\delta_{\mathrm{pcpp}} \geq 1/t_2 = o(1)$. We actually invoke the verifier twice to reduce its soundness error to $1/4$. Setting $d = \log k = \omega(1)$, we obtain a final blocklength of $n = d \cdot \ell < k \cdot \exp(\log^{\varepsilon} k)$ and relative distance $\rho_0 - o(1)$. We further specify the test T as follows. Given a proximity parameter $\delta \geq 6/t_2 = o(1)$, the tester T invokes the aforementioned PCPP with $\delta_{\mathrm{pcpp}} = \delta/6$ and performs the repetition test $q_{\mathrm{rep}} = 6/\delta$ times. Observing that $\delta_{\mathrm{pcpp}} + (4/q_{\mathrm{rep}}) + (1/d) < \delta$, we conclude that the resulting test (i.e., $T = T(1^k, \delta_{\mathrm{pcpp}})$) is a $(\delta, 1/2)$-codeword tester of query complexity $O(1/\delta_{\mathrm{pcpp}}) + 2q_{\mathrm{rep}} = O(1/\delta)$. Thus we conclude as follows.

Conclusion (restating Theorem 1.4). *For every constant $\varepsilon > 0$, there exists a family of LTCs $\mathrm{C}_k : \{0,1\}^k \to \{0,1\}^n$, where $n = \exp(\log^{\varepsilon} k) \cdot k$, with query complexity $q(\delta) = O(1/\delta)$ (i.e., for every $\delta > 0$, the tester rejects words that are δ-far from C with probability $1/2$ by querying at most $q(\delta) = O(1/\delta)$ queries).*

4.2. Relaxed LDCs. We first recall the definition of LDCs, as formally stated by Katz and Trevisan [KT00]. A code $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$ 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 $\Delta(w, \mathrm{C}(x)) \leq \delta$, recovers the ith bit of x with probability at least $2/3$ while making a constant number of queries to w. That is, whenever relatively few locations 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 [KT00].[16] This lower-bound is quite far from the best known upper-bound, due to Beimal et al. [BIKR02], that asserts $n = O(\exp(k^{\varepsilon(q)}))$, where $\varepsilon(q) = O((\log\log q)/(q \log q)) = o(1/q)$, which improves (already for $q = 4$) a previous upper-bound where $\varepsilon(q) = 1/(2q+1)$. 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 LDCs, with the hope of obtaining more efficient constructions (e.g., $n = \mathrm{poly}(k)$).

We relax the definition of LDCs 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 remaining locations the decoder outputs either the right message bit or 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 amplification to reduce the error from $1/3$ to any desired constant).[17] Furthermore, it is desirable to recover all bits of the information whenever the codeword is not corrupted.

Definition 4.5 (relaxed LDC). *A code* $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$ *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 three conditions with respect to any* $w \in \{0,1\}^n$ *and* $x \in \{0,1\}^k$ *such that* $\Delta(w, \mathrm{C}(x)) \leq \delta$*:*

1. *If* $w = \mathrm{C}(x)$ *is a codeword, then the decoder correctly recovers every bit of* x *with probability at least* $\frac{2}{3}$*. That is, for every* $x \in \{0,1\}^k$ *and* $i \in [k]$*, it holds that* $\Pr[M^{\mathrm{C}(x)}(i) = x_i] \geq \frac{2}{3}$.
2. *On input, any index* $i \in [k]$ *and given access to the oracle* w*, with probability at least* $\frac{2}{3}$ *machine* M *outputs either the* i*th bit of* x *or a special failure symbol, denoted* $\perp$*. That is, for every* i*, it holds that* $\Pr[M^w(i) \in \{x_i, \perp\}] \geq \frac{2}{3}$.
3. *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* $\frac{2}{3}$*, machine* M *outputs the* i*th bit of* x*. That is, there exists a set* $I_w \subseteq [k]$ *of size at least* ρk *such that for every* $i \in I_w$ *it holds that* $\Pr[M^w(i) = x_i] \geq \frac{2}{3}$.

We call δ *the* proximity parameter.

One may strengthen the definition by requiring that ρ be greater than $1/2$ or any other favorite constant smaller than 1 (but one should probably refrain from setting $\rho > 1 - \delta$, for example). A different strengthening is for condition 1 to hold with

[16] Their lower-bound refers to nonadaptive 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 [DJK+02], and a lower-bound of $n = \Omega(k^{1+1/(q/2-1)})$ for nonadaptive decoders was presented in [KdW04]. (We note that we use a nonadaptive (relaxed) decoder.)

[17] Here error reduction may be performed by estimating the probability that the machine outputs each of the possible bits, and outputting the more frequent bit only if it has sufficient statistical support (e.g., 50% support, which the wrong bit cannot have). Otherwise, one outputs the `don't know` symbol.

probability 1 (i.e., $\Pr[M^{\mathrm{C}(x)}(i) = x_i] = 1$). In fact, we achieve both the stronger forms.

Remark 4.6. The above definition refers only to strings w that are δ-close to the code. However, using Construction 4.3, 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). This can be achieved with only a nearly linear increase in the length of the code (from length n to length $n \cdot \exp(\log^{\varepsilon} n)$).

Remark 4.7. We stress that condition 2 does not mean that, for every i, and for every w that is δ-close to $\mathrm{C}(x)$, either $\Pr[M^w(i) = x_i] \geq \frac{2}{3}$ or $\Pr[M^w(i) = \perp] \geq \frac{2}{3}$ holds. We refer to the latter condition as condition X and conjecture that the seemingly minor difference between conditions 2 and X is actually substantial. This conjecture is enforced by a recent work of Buhrman and de Wolf [BdW04] who showed that codes that satisfy condition X are actually locally decodable in the standard, nonrelaxed sense (i.e., according to the definition of [KT00]).

4.2.1. Definitional issues and transformations. Note that it is very easy to come up with constructions that satisfy each one of the three conditions of Definition 4.5. For example, condition 2 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 conditions 1 and 3.[18] Our aim, however, is to obtain a construction that satisfies all conditions and beats the performance of the known LDCs.

It turns out that codes that satisfy conditions 1 and 2 can be converted into "equally good" codes that satisfy all three conditions. Let us start with a key definition, which refers to the distribution of the decoder's queries *when asked to recover a random bit position.*

Definition 4.8 (average smoothness). *Let M be a randomized nonadaptive oracle machine having access to an oracle $w \in \{0,1\}^n$ and getting input $i \in [k]$. Further suppose that M always makes q queries. Let $M(i,j,r)$ denote the jth query of M on input i and coin tosses r. We say that M satisfies the* average smoothness *condition if, for every $v \in [n]$,*

$$\frac{1}{2n} < \Pr_{i,j,r}[M(i,j,r) = v] < \frac{2}{n},$$

where the probability is taken uniformly over all possible choices of $i \in [k]$, $j \in [q]$, and coin tosses r.

By having M randomly permute its queries, average smoothness implies that for every $j \in [q]$ and $v \in [n]$, it holds that $\frac{1}{2n} < \Pr_{i,r}[M(i,j,r) = v] < \frac{2}{n}$, where now the probability is taken uniformly over all possible choices of $i \in [k]$ and the coin tosses r. We stress that *average smoothness is different from the notion of smoothness as defined by Katz and Trevisan* [KT00]: They require that *for every $i \in [k]$* (and for every $j \in [q]$ and $v \in [n]$), it holds that $\frac{1}{2n} < \Pr_r[M(i,j,r) = v] < \frac{2}{n}$. Indeed, average smoothness is a weaker requirement, and (as we will shortly see) any code and decoder pair can be easily modified to satisfy it, while preserving the decoding properties (of Definition 4.5). (In contrast, Katz and Trevisan [KT00] present a

[18] In case one wishes the code to have a linear distance this can be achieved too, consider $\mathrm{C}(x) = (x, \mathrm{C}'(x))$, where C' is any code of linear length and linear distance, and a decoder that merely retrieves the desired bit from the first part.

modification that achieves smoothness while preserving strict local decodability, but their transformation does not preserve Definition 4.5.)

LEMMA 4.9. *Let* $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$ *be a code and* M *a machine that satisfies conditions* 1 *and* 2 *of Definition* 4.5 *with respect to proximity parameter* δ. *Then, for some* $n' \in [3n, 4n]$, *there exists a code* $\mathrm{C}' : \{0,1\}^k \to \{0,1\}^{n'}$ *and a machine* M' *that satisfies average smoothness as well as conditions* 1 *and* 2 *of Definition* 4.5 *with respect to proximity parameter* $\delta' = \delta/20$. *Furthermore, the query complexity of* M' *is twice the one of* M, *and if* M *satisfies also condition* 3, *with respect to a constant* ρ, *then so does* M'.

Jumping ahead, we mention that, for a decoder that satisfies average smoothness, conditions 1 and 2 essentially imply condition 3. Hence our interest in average smoothness and in Lemma 4.9.

Proof. As noted above, we may assume without loss of generality that each of M's queries is distributed identically. Throughout the analysis, we refer to the distribution of queries for a uniformly distributed index $i \in [k]$. Let q denote the query complexity of M.

We first modify M such that for a random $i \in [k]$, each query probes each possible location with probability $\Omega(1/n)$. This is done by adding q dummy queries, each being uniformly distributed. Thus, each location gets probed by each query with probability at least $1/2n$.

Next we modify the code and the decoder such that each location is probed with almost uniform distribution. The idea is to repeat heavily probed locations for an adequate number of times and have the decoder probe a random copy. Specifically, let p_v be the probability that location v is probed (i.e., $p_v \triangleq \Pr_{i\in[k],r}[M(i,1,r) = v]$ or equivalently $p_v = \sum_{i\in[k],j\in[2q]} \Pr_r[M(i,j,r) = v]/2kq$). By the above modification, we have $p_v \geq 1/2n$. Now, we repeat location v $r_v = \lfloor 4np_v \rfloor$ times. Note that $r_v \leq 4np_v$ and $r_v > 4np_v - 1 \geq 2 - 1$ (and so $r_v \geq 2$). We obtain a new code C' of length $n' = \sum_v r_v \leq 4n$. (Note that $n' > 3n$.) The relative distance of C' is at least one-fourth that of C, and the rate changes in the same way. The new decoder, M', when seeking to probe location v, will select and probe at random one of the r_v copies of that location. (Interestingly, there is no need to augment this decoder by a testing of the consistency of the copies of an original location.)

Each new location is probed with probability $p'_v \triangleq p_v \cdot \frac{1}{r_v}$ (by each of these queries). Recalling that $\frac{p_v}{r_v} = \frac{p_v}{\lfloor 4np_v \rfloor}$, it follows that $p'_v \geq 1/4n$ and $p'_v \leq \frac{p_v}{4np_v - 1} \leq 1/2n$ (using $p_v \geq 1/2n$). Recalling that $n' \in [3n, 4n]$, each p'_v is in $[(3/4)\cdot(1/n'), 2\cdot(1/n')]$, i.e., within a factor of 2 from uniform.

Clearly, M' satisfies condition 1 (of Definition 4.5), and we show that it (essentially) satisfies condition 2 as well. Let $w = (w_1, \ldots, w_n) \in \{0,1\}^{n'}$ be δ'-close to $\mathrm{C}'(x)$, where $|w_v| = r_v$. Let Y_v be a 0-1 random variable that represents the value of a random bit in w_v; that is, $\Pr[Y_v = 1]$ equals the fraction of 1's in w_v. Then, $\Pr[Y_v \neq \mathrm{C}(x)_v] > 0$ implies that $\overline{\Delta}(w_v, c_v) \geq 1$, where $\mathrm{C}'(x) = (c_1, \ldots, c_n)$ and $|c_v| = r_v$. For $Y = Y_1 \cdots Y_n$, it follows that $\mathrm{E}(\overline{\Delta}(Y, \mathrm{C}(x))) \leq \overline{\Delta}(w, \mathrm{C}'(x))$, and so $\mathrm{E}(\overline{\Delta}(Y, \mathrm{C}(x))) \leq \delta' n' \leq (\delta/5)\cdot n$ (since $\delta' = \delta/20$ and $n' \leq 4n$). Thus, with probability at least 4/5, the random string Y is δ-close to $\mathrm{C}(x)$, in which case the M must succeed with probability at least 2/3. Noting that $M'^{w}(i)$ merely invokes $M^Y(i)$, we conclude that

$$\begin{aligned}
\Pr[M'^{w}(i) \in \{x_i, \perp\}] &= \Pr[M^Y(i) \in \{x_i, \perp\}] \\
&\geq \Pr[\overline{\Delta}(Y, \mathrm{C}(x)) \leq \delta n] \cdot \Pr[M^Y(i) \in \{x_i, \perp\} \mid \overline{\Delta}(Y, \mathrm{C}(x)) \leq \delta n] \\
&\geq \frac{4}{5} \cdot \frac{2}{3} = \frac{8}{15}.
\end{aligned}$$

An analogous argument can be applied in the case M satisfies condition 3. In both cases, additional error reduction is needed in order to satisfy the actual conditions, which require success with probability at least 2/3. (For details see footnote 17.) ∎

LEMMA 4.10. *Let* $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$ *be a code and* M *be a machine that satisfies conditions* 1 *and* 2 *of Definition* 4.5 *with respect to a constant* δ. *Suppose that* M *satisfies the average smoothness condition and has query complexity* q. *Then, invoking* M *for a constant number of times (and ruling as in footnote* 17*) yields a decoder that satisfies all three conditions of Definition* 4.5. *Specifically, condition* 3 *holds with respect to a constant* $\rho = 1 - 18q\delta$. *Furthermore, for any* w *and* x, *for a* $1 - 18q\overline{\Delta}(w, \mathrm{C}(x))$ *fraction of the* i*'s, it holds that* $\Pr[M^w(i) = x_i] \geq 5/9$.

Our usage of the average smoothness condition actually amounts to using the hypothesis that, for a uniformly distributed $i \in [k]$, each query hits any fixed position with probability at most $2/n$.

Proof. By condition 1, for any $x \in \{0,1\}^k$ and every $i \in [k]$, it holds that $\Pr[M^{\mathrm{C}(x)}(i) = x_i] \geq 2/3$. Considering any w that is δ-close to $\mathrm{C}(x)$, the probability that on input a *uniformly distributed* $i \in [k]$ machine M queries a location on which w and $\mathrm{C}(x)$ disagree is at most $q \cdot (2/n) \cdot \delta n = 2q\delta$. This is due to the fact that, for a uniformly distributed i, the queries are almost uniformly distributed; specifically, no position is queried with probability greater than $2/n$ (by a single query).

Let p_i^w denote the probability that on input i machine M queries a location on which w and $\mathrm{C}(x)$ disagree. We have just established that $(1/k) \cdot \sum_{i=1}^k p_i^w \leq 2q\delta$. For $I_w \triangleq \{i \in [k] : p_i^w \leq 1/9\}$, it holds that $|I_w| \geq (1 - 18q\delta) \cdot k$. Observe that for any $i \in I_w$, it holds that $\Pr[M^w(i) = x_i] \geq (2/3) - (1/9) = 5/9$. Note that, by replacing δ with $\overline{\Delta}(w, \mathrm{C}(x))/n$, the above argument actually establishes that for a $1 - 18q \cdot \overline{\Delta}(w, \mathrm{C}(x))$ fraction of the i's, it holds that $\Pr[M^w(i) = x_i] \geq 5/9$.

Additional error reduction is needed in order to satisfy the actual definition (of condition 3), which requires success with probability at least 2/3. The error reduction should be done in a manner that preserves conditions 1 and 2 of Definition 4.5. For details see footnote 17. ∎

In view of the last sentence of Lemma 4.10, it makes sense to state a stronger definition of relaxed LDCs.

DEFINITION 4.11 (relaxed LDC, revisited). *A code* $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$ *is* relaxed locally decodable *if for some constants* $\delta > 0$ *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* $w \in \{0,1\}^n$ *and* $x \in \{0,1\}^k$ *such that* $\overline{\Delta}(w, \mathrm{C}(x)) \leq \delta n$:

1. *For every* $i \in [k]$ *it holds that* $\Pr[M^w(i) \in \{x_i, \perp\}] \geq \frac{2}{3}$.
2. *There exists a set* $I_w \subseteq [k]$ *of density at least* $1 - O(\overline{\Delta}(w, \mathrm{C}(x))/n)$ *such that for every* $i \in I_w$ *it holds that* $\Pr[M^w(i) = x_i] \geq \frac{2}{3}$.

Note that the "everywhere good" decoding of codewords (i.e., condition 1 of Definition 4.5) is implied by condition 2 of Definition 4.11. By combining Lemmas 4.9 and 4.10, we get the following theorem.

THEOREM 4.12. *Let* $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$ *be a code and* M *be a machine that makes a constant* q *number of queries and satisfies conditions* 1 *and* 2 *of Definition* 4.5 *with respect to a constant* δ. *Then, for some* $n' \in [3n, 4n]$, *there exists a code* $\mathrm{C}' : \{0,1\}^k \to \{0,1\}^{n'}$ *that is relaxed locally decodable with respect to proximity parameter* $\delta' = \delta/20$. *Furthermore, this code satisfies Definition* 4.11.

4.2.2. Constructions. In view of Lemma 4.10, we focus on presenting codes with decoders that satisfy conditions 1 and 2 of Definition 4.5 as well as the average

smoothness property. (The latter property will save us the need to invoke Lemma 4.9.) We will start with a code that has nearly quadratic length (i.e., $n = k^{2+o(1)}$), which serves as a good warm-up towards our final construction in which $n = k^{1+\varepsilon}$ for any desired constant $\varepsilon > 0$.

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 allowing for "local decodability" (e.g., using the identity map) and the second providing the distance property (e.g., using any code of linear distance). It is obvious that a third part that guarantees the consistency of the first two parts should be added, and it is natural to try to use a PCPP 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 first part. Indeed, the first part may not be "robust to corruption" but the second 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 PCPP is unlikely to detect a small change in the first part. Thus, if we use the identity map in the first part, then the decoder may output the wrong value of some (although few) bits. In other words, the "proximity relaxation" in PCPPs makes sense for the second part of the codewords but not for the first part. Our solution is to provide, *for each bit* (position) in the first part, a proof of the consistency of this bit (value) with the entire second part. The second problem is that the PCPPs (let alone all of them combined) are much longer than the first two parts, whereas the corruption rate is measured in terms of the entire codeword. This problem 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 (as per condition 3 of Definition 4.5) even at low corruption rates (measured in terms of the entire codeword). In other words, if the third part is too short, then we have no chance to satisfy the average smoothness condition.

The actual construction. Let $\mathrm{C}_0 : \{0,1\}^k \to \{0,1\}^m$ be a good code of relative distance δ_0; then we encode $x \in \{0,1\}^k$ by $\mathrm{C}(x) \triangleq (x^t, \mathrm{C}_0(x)^{t'}, \pi_1(x), \ldots, \pi_k(x))$, where $t = |\pi_1(x), \ldots, \pi_k(x)|/|x|$ (resp., $t' = |\pi_1(x), \ldots, \pi_k(x)|/|\mathrm{C}_0(x)|$), and $\pi_i(x)$ is a PCPP to be further discussed. We first note that the replicated versions of x (resp., $\mathrm{C}_0(x)$) take a third of the total length of $\mathrm{C}(x)$. As for $\pi_i(x)$, it is a PCPP 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 $z_2 = \mathrm{C}_0(x)$.[19] We use our PCPP from Theorem 3.3, while setting its parameters such that the proximity parameter is small enough but the query complexity is a constant. Specifically, let $\delta_{\mathrm{pcpp}} > 0$ be the proximity parameter of the PCPP, which will be set to be sufficiently small, and let $q = O(1/\delta_{\mathrm{pcpp}})$ denote the number of queries the verifier makes in order to support a soundness error of $1/6$ (rather than the standard $1/2$). A key observation regarding this verifier is that its queries to its input-oracle are uniformly distributed. The queries to the proof oracle can be made almost uniform by a modification analogous to the one used in the proof of Lemma 4.9.

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 PCPP proof and $s_0(m)$ denotes the size of the circuit for encoding relative to C_0. Using a good code $\mathrm{C}_0 : \{0,1\}^k \to \{0,1\}^m$ (i.e., of constant relative distance δ_0, linear length $m = O(k)$, and $s_0(m) = \widetilde{O}(m)$), we obtain $n = k^{2+o(1)}$. The relative distance of C is at least $\delta_0/3$.

[19] Indeed, 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.

We now turn to the description of the decoder D. Recall that a valid codeword has the form $(x^t, \mathrm{C}_0(x)^{t'}, \pi_1(x), \ldots, \pi_k(x))$. The decoding of the ith information bit (i.e., x_i) will depend on a random (possibly wrong) copy of x_i located in the first part (which supposedly equals x^t), a random (possibly corrupted) copy of $\mathrm{C}_0(x)$ located in the second part, and the relevant (i.e., ith) proof located in the third part (which is also possibly corrupted). On input $i \in [k]$ and oracle access to $w = (w_1, w_2, w_3) \in \{0,1\}^n$, where $|w_1| = |w_2| = |w_3|$, the decoder invokes the PCPP verifier while 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 ith block (ℓ-bit long) of w_3. That is, when the verifier asks to access the jth 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.

The above construction can be performed for any sufficiently small constant proximity parameter $\delta \in (0, \delta_0/18)$. All that this entails is setting the proximity parameter of the PCPP to be sufficiently small but positive (e.g., $\delta_{\mathrm{pcpp}} = (\delta_0 - 18\delta)/2$). We actually need to augment the decoder such that it makes an equal number of queries to each of the three (equal length) parts of the codeword, which is easy to do by adding (a constant number of) dummy queries. Let us denote the resulting decoder by D.

PROPOSITION 4.13. *The above code and decoder satisfy conditions* 1 *and* 2 *of Definition* 4.5 *with respect to proximity parameter* $\delta \in (0, \delta_0/18)$. *Furthermore, this decoder satisfies the average smoothness property.*

Proof. Condition 1 (of Definition 4.5) is obvious from the construction (and the completeness property of the PCPP). In fact, the perfect completeness of the PCPP implies that bits of an uncorrupted codeword are recovered with probability 1 (rather than with probability at least 2/3). The average smoothness property of the decoder is obvious from the construction and the smoothness property of the PCPP. We thus turn to establish condition 2 (of Definition 4.5).

Fixing any $x \in \{0,1\}^k$, we consider an arbitrary oracle $w = (w_1, w_2, w_3)$ that is δ-close to $\mathrm{C}(x)$, where w_1 (resp., w_2) denotes the alleged replication of x (resp., $\mathrm{C}_0(x)$) and $w_3 = (u_1, \ldots, u_k)$ denotes the part of the PCPPs. Note that w_2 is 3δ-close to $\mathrm{C}_0(x)^{t'}$. 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 ith block of w_3 (i.e., $\pi = u_i$). We bound the probability that the decoder outputs $\neg x_i$ by considering three cases as follows.

Case 1. $\sigma = x_i$. Recall that σ is the bit read by D from w_1, and that by construction D always outputs either σ or $\perp$. Thus, in this case, Condition 2 is satisfied (because, regardless of whether D outputs σ or $\perp$, the output is always in $\{x_i, \perp\}$).

Case 2. Z_2 is 18δ-far from $\mathrm{C}_0(x)$. Recall that w_2 is 3δ-close to $\mathrm{C}_0(x)^{t'}$, which means that the expected relative distance of Z_2 and $\mathrm{C}_0(x)$ is at most 3δ. Thus, the current case occurs with probability at most 1/6.

Case 3. Z_2 is 18δ-close to $\mathrm{C}_0(x)$ and $\sigma \neq x_i$. Then, on one hand, (Z_1, Z_2) is 1/2-

far from $(x_i^m, \mathrm{C}_0(x))$, because $Z_2 = \sigma^t$. On the other hand, Z_2 is $(\delta_0 - 18\delta)$-far from any other codeword of C_0, because Z_2 is 18δ-close to $\mathrm{C}_0(x)$ and the codewords of C_0 are δ_0-far from one another. Thus, (Z_1, Z_2) is $(\delta_0 - 18\delta)/2$-far from any string of the form $(y_i^m, \mathrm{C}_0(y))$. Using $\delta_{\mathrm{pcpp}} \leq (\delta_0 - 18\delta)/2$, we conclude that the PCPP verifier accepts (Z_1, Z_2) with probability at most $1/6$. It follows that, in the current case, 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 $1/6 + 1/6 = 1/3$. ∎

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 $\Omega(k)$. Now, consider the following improvement: Partition the message into $\sqrt{k}$ blocks, each of length $\sqrt{k}$. Encode the original message, as well as each of the smaller blocks, via a good error-correcting code. Let w be the encoding of the entire message, and let w_i $(i = 1, \ldots, \sqrt{k})$ be the encodings of the blocks. For every $i = 1, \ldots, \sqrt{k}$, append a PCPP for the claim "w_i is the encoding of the ith block of a message encoded by w." In addition, for each message bit $x_{(i-1)\sqrt{k}+j}$ residing in block i, append a PCPP of the statement "$x_{(i-1)\sqrt{k}+j}$ is the jth 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 ith partition uses $k^{i/\ell}$ blocks, each of length $k^{1-(i/\ell)}$. Encoding is done by providing, for each $i = 0, 1, \ldots, \ell$, encodings of each of the blocks in the ith partition by 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 PCPPs that assert the consistency of "directly related" blocks (i.e., blocks of consecutive levels that contain one another). That is, for every $i = 1, \ldots, \ell$ and $j \in [k^{i/\ell}]$, we place a proof that the encoding of the jth block in the ith level is consistent with the encoding of the $\lceil j/k^{1/\ell} \rceil$th block in the $(i-1)$st level. The ith 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 proof length 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. We stress that, as before, the proofs in use are PCPPs. In the current context these proofs refer to two input oracles of vastly different lengths, 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.[20]

CONSTRUCTION 4.14. *Let C_0 be a code of minimal relative distance δ_0, constant rate, and nearly linear-sized encoding circuits. For simplicity, assume that a single bit is encoded by repetitions; that is, $\mathrm{C}_0(\sigma) = \sigma^{O(1)}$ for $\sigma \in \{0, 1\}$. Let*

[20]Indeed, this was also done in the simpler code analyzed in Proposition 4.13.

V be a PCPP of membership in $S_0 = \{\mathrm{C}_0(x) : x \in \{0,1\}^*\}$ having almost linear proof length, query complexity $O(1/\delta_{\mathrm{pcpp}})$, and soundness error $1/9$, for proximity parameter δ_{pcpp}. Furthermore, V's queries to both its input oracle and proof oracle are distributed almost uniformly.[21] *For a fixed parameter $\ell \in \mathbb{N}$, let $b \triangleq k^{1/\ell}$. For $x \in \{0,1\}^k$, we consider ℓ different partitions of x, such that the jth partition denoted $(x_{j,1}, \ldots, x_{j,b^j})$, where $x_{j,j'} = x_{(j'-1)\cdot b^{\ell-j}+1} \cdots x_{j'\cdot b^{\ell-j}}$. We define $\mathrm{C}_j(x) \triangleq (\mathrm{C}_0(x_{j,1}), \mathrm{C}_0(x_{j,2}), \ldots, \mathrm{C}_0(x_{j,b^j}))$, and $\pi_j(x) = (\pi_{j,1}(x), \ldots, \pi_{j,b^j}(x))$, where $p_{j,j'}(x)$ is a PCPP proof oracle that asserts the consistency of j'th block of $\mathrm{C}_j(x)$ and the $\lceil j'/b \rceil$th block of $\mathrm{C}_{j-1}(x)$. That is, $p_{j,j'}(x)$ refers to an input oracle of the form (z_1, z_2), where $|z_1| = |z_2| = O(b^{\ell-j+1})$, and asserts the existence of x such that $z_1 = \mathrm{C}_0(x_{j,j'})^b$ and $z_2 = \mathrm{C}_0(x_{j-1,\lceil j'/b \rceil})$. We consider the code*

$$\mathrm{C}(x) \triangleq (\mathrm{C}_0(x)^{t_0}, \mathrm{C}_1(x)^{t_0}, \ldots, \mathrm{C}_\ell(x)^{t_0}, \pi_1^{t_1}, \ldots, \pi_\ell^{t_\ell}),$$

where the t_j's are selected such that each of the $2\ell+1$ parts of $\mathrm{C}(x)$ has the same length. The decoder, denoted D, operates as follows on input $i \in [k]$ and oracle access to $w = (w_0, w_1, \ldots, w_\ell, v_1, \ldots, v_\ell)$, where $|w_0| = |w_j| = |v_j|$ for all j:

- *D selects uniformly $r_0, r_1, \ldots, r_\ell \in [t_0]$, and $(r'_1, r'_2, \ldots, r'_\ell) \in [t_1] \times [t_2] \times \cdots \times [t_\ell]$.*
- *For $j = 1, \ldots, \ell$, D invokes the PCPP verifier, providing it with access to an input oracle (z_1, z_2) and a proof oracle π that are defined as follows:*
  - *$z_1 = u^b$, where u is the $((r_j - 1) \cdot b^j + \lceil i/b^{\ell-j} \rceil)$th block of w_j.*
  - *z_2 is the $((r_{j-1} - 1) \cdot b^{j-1} + \lceil i/b^{\ell-j+1} \rceil)$th block of w_{j-1}.*
  - *π is the $((r'_j - 1) \cdot b^j + \lceil i/b^{\ell-j} \rceil)$th block of v_j.*

  The PCPP verifier is invoked with proximity parameter $\delta_{\mathrm{pcpp}} = 13\ell\delta > 0$, where $\delta \le \delta_0/81\ell$ is the proximity parameter sought for the decoder.
- *If the PCPP verifier rejects, in any of the aforementioned ℓ invocations, then the decoder outputs a special* (failure) *symbol. Otherwise, the decoder outputs a random value in the $((r_\ell - 1) \cdot k + i)$th block of w_ℓ* (which is supposedly a repetition code of x_i).
- *D satisfies the average smoothness property when we issue some dummy queries that are uniformly distributed in adequate parts of w that are queried less by the above. (In other words, suppose that V makes q_1 (resp., q_2) queries to the first (resp., second) part of its input oracle and q' queries to its proof oracle. Then, w_0 is accessed q_2 times, w_ℓ is accessed q_1 times, each other w_j is accessed $q_1 + q_2$ times, and each v_j is accessed q' times. Thus, we may add dummy queries to make each part accessed $\max(q_1 + q_2, q')$ times, which means increasing the number of queries by a factor of at most $(2\ell+1)/(\ell-1)$ assuming $\ell \ge 2$.)*

Using an adequate PCPP, it holds that $|\mathrm{C}(x)| = \ell \cdot (|x|^{1+(1/\ell)})^{1+o(1)} < |x|^{1+\varepsilon}$ for $\varepsilon = 2/\ell$. The query complexity of D is $O(\ell) \cdot O(1/\delta_{\mathrm{pcpp}}) = O(\ell^2)$. The proof of Proposition 4.13 can be extended, obtaining the following.

PROPOSITION 4.15. *The code and decoder of Construction 4.14 satisfy conditions 1 and 2 of Definition 4.5 with respect to proximity parameter $\delta \le \delta_0/81\ell$. Furthermore, this decoder satisfies the average smoothness property.*

Using Lemma 4.10, Theorem 1.5 follows.

[21] Recall that all these conditions hold for the PCPP of Theorem 3.3, where almost uniformly distributed queries to the proof oracle are obtained by a modification analogous to the proof of Lemma 4.9.

Proof. Again, condition 1 as well as the average smoothness property are obvious from the construction, and thus we focus on establishing condition 2. Thus, we fix an arbitrary $i \in [k]$ and follow the outline of the proof of Proposition 4.13.

We consider an oracle $(w_0, w_1, \ldots, w_\ell, \pi_1, \ldots, \pi_\ell)$ that is δ-close to an encoding of $x \in \{0,1\}^k$, where each w_j is assumed to consist of encodings of the $k^{j/\ell}$ (nonoverlapping) $k^{1-(j/\ell)}$-bit long blocks of x, and π_i consists of the corresponding proofs of consistency. It follows that each w_j is $(2\ell+1)\cdot\delta$-close to $\mathrm{C}_j(x)^{t_0}$. Let Z_j denote the block of w_j that was selected and accessed by D. Thus, the expected relative distance of Z_0 from $\mathrm{C}_0(x)$ is at most $(2\ell+1)\cdot\delta$, but we do not know the same about the other Z_j's because their choice depends on i (or rather on $\lceil i/b^{\ell-j}\rceil$). Assuming, without loss of generality, that $\delta_0 < 1/3$ (and $\ell \geq 1$), we consider three cases as follows.

Case 1. Z_ℓ is 1/9-close to $\mathrm{C}_0(x_i)$. In this case, D outputs either $\perp$ or a uniformly selected bit in Z_ℓ, and so D outputs $\neg x_i$ with probability at most 1/9.

Using $\delta \leq \delta_0/81\ell$ and $\delta_0 < 1/3$, it follows that $27\ell\delta < 1/9$. Thus, if Case 1 does not hold, then Z_ℓ is $27\ell\delta$-far from $\mathrm{C}_0(x_i)$.

Case 2. Z_0 is $27\ell\delta$-far from $\mathrm{C}_0(x)$. This case may occur with probability at most 1/9, because $\mathrm{E}[\overline{\Delta}(Z_0, \mathrm{C}_0(x))] \leq 3\ell\delta \cdot |\mathrm{C}_0(x)|$.

Note that if both Cases 1 and 2 do not hold, then Z_0 is $27\ell\delta$-close to $\mathrm{C}_0(x)$ but Z_ℓ is $27\ell\delta$-far from $\mathrm{C}_0(x_i)$. Also note that $x = x_{0,1}$ and $x_i = x_{\ell,i}$.

Case 3. For some $j \in [\ell]$, it holds that Z_{j-1} is $27\ell\delta$-close to $\mathrm{C}_{j-1}(x_{j-1,\lceil i/b^{\ell-j+1}\rceil})$ but Z_j is $27\ell\delta$-far from $\mathrm{C}_j(x_{j,\lceil i/b^{\ell-j}\rceil})$. In this case, the pair (Z_j^b, Z_{j-1}) is $27\ell\delta/2$-far from the consistent pair $(\mathrm{C}_j(x_{j,\lceil i/b^{\ell-j}\rceil}), \mathrm{C}_{j-1}(x_{j-1,\lceil i/b^{\ell-j+1}\rceil}))$ and is $(\delta_0 - 27\ell\delta)/2$-far from any other consistent pair. Using $\delta_{\mathrm{pcpp}} = 13\ell\delta < \min(27\ell\delta/2, (\delta_0 - 27\ell\delta)/2)$, which holds because $\delta \leq \delta_0/81\ell$, it follows that in the current case the PCPP verifier accepts (and the decoder does not output $\perp$) with probability at most 1/9.

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

Conclusion (restatement of Theorem 1.5). *For every constant $\varepsilon > 0$, there exists a code $\mathrm{C} : \{0,1\}^k \to \{0,1\}^n$, where $n = k^{1+\varepsilon}$, that is relaxed locally decodable under Definition* 4.11. *The query complexity of the corresponding decoder is $O(1/\varepsilon^2)$ and the proximity parameter is $\varepsilon/O(1)$.*

Open problem. We wonder whether one can obtain a relaxed LDC that can be decoded using q queries while having length $n = o(k^{1+/(q-1)})$. The existence of such a 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. [KT00]). Alternatively, it may be possible to improve the lower-bound for the (q-query) LDC to $n > k^{1+\sqrt{c/q}}$ 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 $\mathbb{F}_2$-linear codes whenever the base code C_0 is also $\mathbb{F}_2$-linear. Proving this assertion reduces to proving that the PCPPs used (in the aforementioned constructions) have proof oracles in which each bit is a linear function 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 main construct (i.e., the PCPPs stated in Theorems 3.1 and 3.2) when applied to a linear circuit yields an $\mathbb{F}_2$-linear transformation of assignments (satisfying the circuit) to proof oracles (accepted by the verifier). In addition, we also need to show that the construction underlying the proof of Theorem 3.3 satisfies this property. This is done next, and consequently we get the following.

PROPOSITION 4.16. *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.

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 3.1, 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 (i.e., bottom) verifier.

In each composition step, we append the proof oracle with new (inner) PCPPs 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 (Appendix A). The encoding defined there is not $\mathbb{F}_2$-linear (rather it is quadratic), but this was necessary for dealing with nonlinear gates. It can be verified that for a linear circuit, one can perform all necessary tests of Appendix A with the Hadamard encoding of the input. Thus, we conclude that this final phase of the encoding is also linear, and this completes the proof of Proposition 4.16. ∎

Part II. The main construct: A short, robust PCPP.

5. Overview of our main construct. Throughout this section, n denotes the length of the explicit input given to the PCPP verifier, which in the case of CKTVAL 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 PCPP.

THEOREM 3.1 (main construct restated). *There exists a universal constant c such that for all $n, m \in \mathbb{Z}^+$, $0 < \delta, \gamma < 1/2$ satisfying $n^{1/m} \geq m^{cm}/(\gamma\delta)^3$ and $\delta \leq \gamma/c$,* CKTVAL *has a robust PCPP (for circuits of size n) with the following parameters:*

- *randomness $\left(1 - \frac{1}{m}\right)\log n + O(m \log m) + O(\log\log n) + O(\log(1/\delta))$;*
- *decision complexity $n^{1/m} \cdot \mathrm{poly}(\log n, 1/\delta)$, which also upper-bounds the query complexity;*[22]
- *perfect completeness; and*
- *for proximity parameter δ, the verifier has robust-soundness error $\Omega(\gamma)$ with robustness parameter $(1-\gamma)\delta$.*

A (simplified) variant of Theorem 3.1 also yields the ALMSS-type robust PCPP (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 [HS00] and modify their construction. (An alternative approach would be to start from [PS94], but that construction does not seem amenable to achieving robust soundness.) The construction of [HS00] may be abstracted as follows: To verify the satisfiability of a circuit of size n, a verifier expects oracles $\mathrm{A}_i : \mathbb{F}^m \to \mathbb{F}$, $i \in \{1, \ldots, t = \mathrm{poly}\log n\}$, where $\mathbb{F}$ is a field and m is a parameter such that $|\mathbb{F}^m| \approx m^m \cdot n$. The verifier then needs to test that (1) each of the A_i's is close to an m-variate polynomial of *low degree* and (2) the polynomials satisfy some *consistency*

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

properties which verify that A_i is locally consistent with A_{i-1}.[23] (These consistency checks include tests which depend on the input circuit and verify that A_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 the A_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 [BSVW03], while using only the "canonical" representations of these lines (and avoiding any complication that was introduced towards proof composition). In particular, unlike in [HS00, GS02, BSVW03], we cannot afford to pack the polynomials $\mathrm{A}_1, \ldots, \mathrm{A}_t$ 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 PCPP by comparing the input oracle (i.e., the 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 by 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 [BFLS91].

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 $\mathrm{A}_1, \ldots, \mathrm{A}_t$. Selecting a random line $\ell : \mathbb{F} \to \mathbb{F}^m$ (from the aforementioned sample space), we can check that (for each i) the restriction of A_i to the line ℓ (i.e., the function $f_i(j) \triangleq \mathrm{A}_i(\ell(j))$) is a low-degree (univariate) polynomial. Each of these tests is *individually robust*; that is, if A_i is far from being a low-degree polynomial, then with high probability the restriction of A_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 A_i's is δ-far from being a low-degree polynomial, then it is guaranteed only 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 $\overline{\mathrm{A}} : \mathbb{F}^m \to \mathbb{F}^t$ such that if one of the A_i's is far from being a low-degree polynomial, then the restriction of $\overline{\mathrm{A}}$ to a random line will be far from being a bundling of t low-degree univariate polynomials. Specifically, for every $x \in \mathbb{F}^m$, define $\overline{\mathrm{A}}(x) \triangleq (\mathrm{A}_1(x), \ldots, \mathrm{A}_t(x))$. To test that $\overline{\mathrm{A}}$ is a bundling of low-degree polynomials, select a

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

random line ℓ (as above), and check that $\overline{f}_\ell(j) = \overline{\mathrm{A}}(\ell(j))$ is a bundling of low-degree univariate polynomials. Thus, we establish *robustness at the bundle level*; that is, if one of the A_i's is far from being low degree then, with high probability, one must modify $\overline{f}_\ell$ on a constant fraction of values in order to make the test accept. The point is that this robustness refers to the Hamming distance over the alphabet $\mathbb{F}^t$, rather than alphabet $\mathbb{F}$ as before. We can afford this increase in alphabet size, as we later encode the values of $\overline{\mathrm{A}}$ 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 $\overline{\mathrm{A}}$ restricted to a random line, we get the values of the individual A_i's restricted to a random line, which is exactly what a low-degree test of each A_i needs. This fact is not very surprising, given that we are subjecting all A_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 A_i and A_{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 A_i and A_{i+1} on (random) ith axis-parallel lines. That is, for every i, and a random $x' = (x_1, \ldots, x_{i-1}) \in \mathbb{F}^{i-1}$ and $x'' = (x_{i+1}, \ldots, x_m) \in \mathbb{F}^{m-i}$, we need to check some relation between $\mathrm{A}_i(x', \cdot, x'')$ and $\mathrm{A}_{i+1}(x', \cdot, x'')$.[24] Clearly, querying $\overline{\mathrm{A}}$ as above on the ith axis-parallel line, we can obtain the relevant values from $\overline{\mathrm{A}}(x', \cdot, x'')$, 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 is that we gain nothing from the bundling (i.e., from $\overline{\mathrm{A}}$) over using the individual A_i's, which yields a factor of t loss in the robustness.[25] Fortunately, a different bundling works in this case.

Consider $\overline{\mathrm{A}}'$ such that $\overline{\mathrm{A}}'(x) \triangleq (\mathrm{A}_1(x), \mathrm{A}_2(S(x)), \ldots, \mathrm{A}_t(S^{t-1}(x)))$, for every $x \in \mathbb{F}^m$, where S denotes a (right) cyclic-shift (i.e., $S(x_1, \ldots, x_m) = (x_m, x_1, \ldots, x_{m-1})$ and $S^i(x_1, \ldots, x_m) = (x_{m-(i-1)}, \ldots, x_m, x_1, x_2, \ldots, x_{m-i})$). Now, if we ask for the value of $\overline{\mathrm{A}}'$ on the first and last axis-parallel lines (i.e., on $(\cdot, x_2, \ldots, x_m)$ and $(x_2, \ldots, x_m, \cdot) = S^{-1}(\cdot, x_2, \ldots, x_m)$), then we get all we need for all the m tests. Specifically, for every i, the ith component in the bundled function $\overline{\mathrm{A}}'(\cdot, x_2, \ldots, x_m)$ is $\mathrm{A}_i(S^{i-1}(\cdot, x_2, \ldots, x_m)) = \mathrm{A}_i(x_{m-i+2}, \ldots, x_m, \cdot, x_2, \ldots, x_{m-i+1})$, whereas the $(i+1)$st component in $\overline{\mathrm{A}}'(S^{-1}(\cdot, x_2, \ldots, x_m))$ is $\mathrm{A}_{i+1}(S^i(S^{-1}(\cdot, x_2 \ldots, x_m))) = \mathrm{A}_{i+1}(x_{m-i+2}, \ldots, x_m, \cdot, x_2, \ldots, x_{m-i+1})$. Thus, we need only query two bundles (rather than t), and robustness drops only 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 $\overline{\mathrm{A}}$)? Note that we cannot use both $\overline{\mathrm{A}}$ and $\overline{\mathrm{A}}'$, because this will requires 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., $\overline{\mathrm{A}}'$), designed to serve the axis-parallel line comparisons, can also serve the low-degree tests. Indeed, the various A_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).

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

[25] It 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 $\log\log n$) and not for proving Theorem 1.3, which requires setting and $m = \log^\varepsilon n$ for constant $\varepsilon > 0$.

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 the comment at the end of section 3.1. To obtain the stronger bound on the robustness parameter, we take a closer look at the conjunction of the standard PCP test and 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 CktSAT (i.e., a robust PCPP without either the robustness or proximity properties). We favor this construction 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* $m \le \log n/\log\log n$. *Then there exists a PCP for* CktSAT *(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* $\widetilde{O}(q)$,
- *perfect completeness,*
- *and soundness error* $1-\varepsilon$.

The construction of the PCP for CktSAT proceeds in three steps. First, we transform the input circuit φ into a well-structured circuit φ' along the lines of Polishchuk and Spielman [PS94, Spi95] (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 into 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 [HS00] 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. [BSVW03]); 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 CktSAT 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 we are lucky to 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, Appendix A) incurs a quadratic blow-up in length. To get over this problem, Polishchuk and Spielman [PS94, Spi95] introduced a different, more efficient restructuring process that embeds the input circuit into well-structured graphs known as de Bruijn graphs. Indeed, the blow-up 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 [PS94, Spi95], 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 [PS94, Spi95]. However, it can be easily 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 [PS94, Spi95] 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 an 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 $0, 1, \ldots, l-1$. The node represented by $v = (b_0, \ldots, b_{i^*}, \ldots, b_{N-1})$ in layer i has edges pointing to the nodes represented by $\Gamma_{i,0}(v) = (b_0, \ldots, b_{i^*}, \ldots, b_{N-1})$ and $\Gamma_{i,1}(v) = (b_0, \ldots, b_{i^*} \oplus 1, \ldots, b_{N-1})$ 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 [Lei92]) 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 $1, \ldots, l-1$ 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).

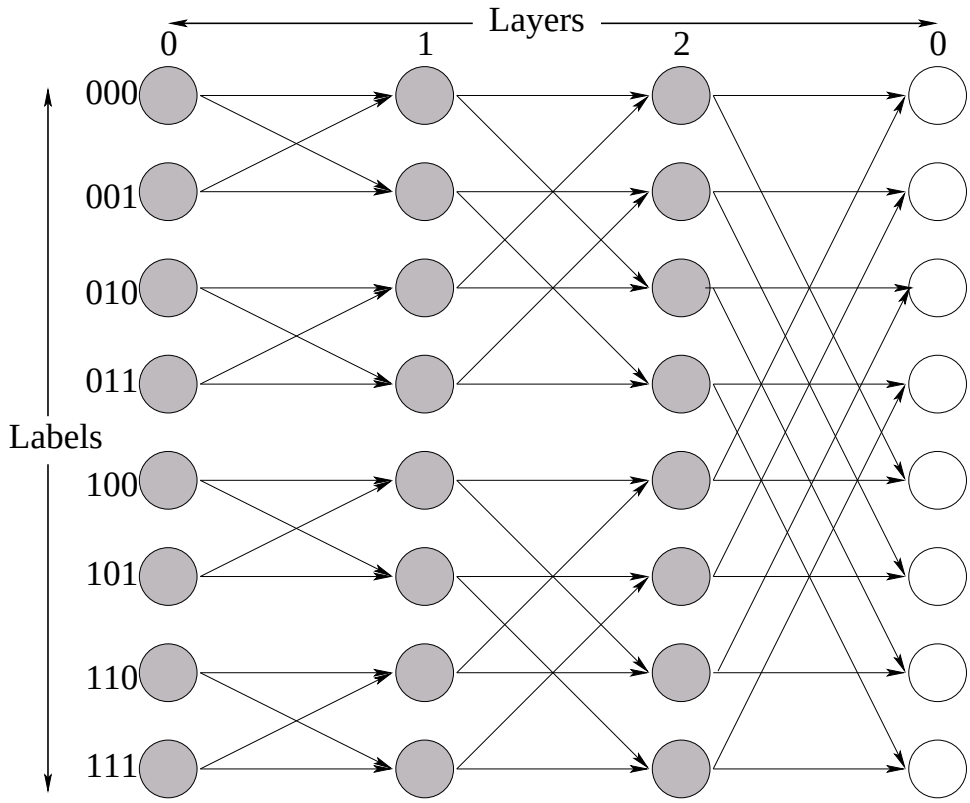

FIG. 1. *The wrapped de Bruijn graph $\mathcal{G}_{3,3}$. Notice the first and last layers are the same.*

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.

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 $T_0, T_1, \ldots, T_{l-1} : \{0,1\}^N \to \mathcal{C}$, 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

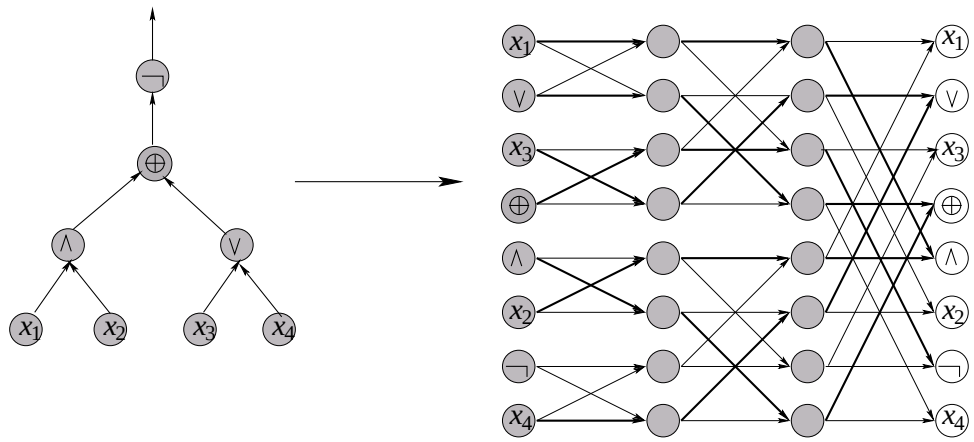

FIG. 2. *Embedding of a circuit into $\mathcal{G}_{3,3}$. In this example all paths between nodes at the* 0 *layer are vertex disjoint. For general circuits we merely need edge disjoint paths.*

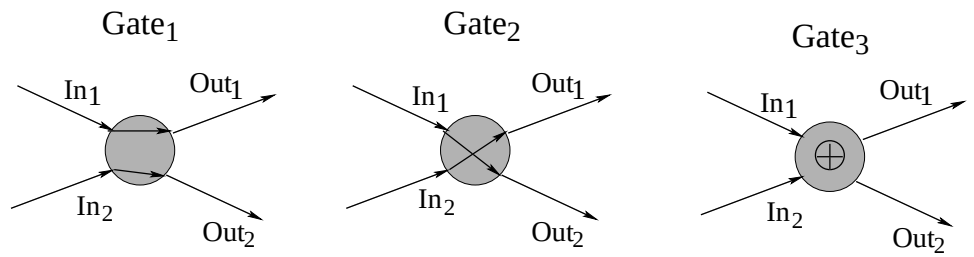

FIG. 3. *Some gates of a well-structured circuit. Gates* 1–2 *are switching gates, and gate* 3 *sits in layer* 0 *and computes the parity (xor) function.*

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 $A_0, A_1, \ldots, A_{l-1}$, where each function $A_i : \{0,1\}^N \to \{0,1\}^4$ specifies the values carried by the four edges incident at that vertex.

We now list the constraints on the embedded circuit that assure us that the only legal assignments are the 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 that*

- *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*

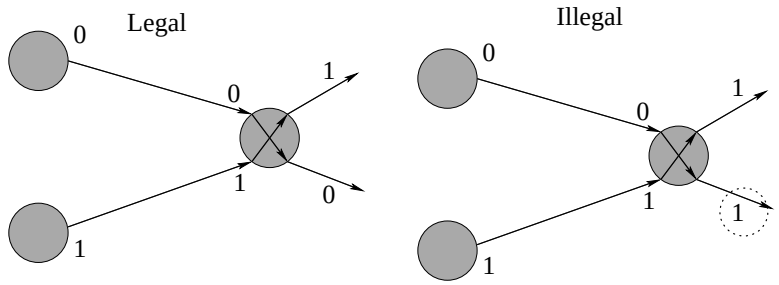

FIG. 4. *Example of legal and illegal assignments. The two vertices on the left are the inputs (at layer $i-1$) to a gate at layer i. Recall that assignments evaluate each incoming and outgoing edge of a gate.*

value equals 1.

Let $\psi : \mathcal{C} \times (\{0,1\}^4)^3 \to \{0,1\}$ 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. See Figure 4 *for an example of legal and illegal assignments.*

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

$$\psi\Big(T_i(v), A_i(v), A_{i+1}\big(\Gamma_{i,0}(v)\big), A_{i+1}\big(\Gamma_{i,1}(v)\big)\Big) = 0.$$

We are now ready to formally define the well-structured circuit satisfiability problem (STRUCTURED-CKTSAT).

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

$\langle \mathcal{G}_{N,l}, \{T_0, \ldots, T_{l-1}\}\rangle \in$ STRUCTURED-CKTSAT *if there exists a set of assignments $A_0, A_1, \ldots, A_{l-1}$, where $A_i : \{0,1\}^N \to \{0,1\}^4$ 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,*

$$\psi\Big(T_i(v), A_i(v), A_{i+1}\big(\Gamma_{i,0}(v)\big), A_{i+1}\big(\Gamma_{i,1}(v)\big)\Big) = 0.$$

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 $C \in$* CKTSAT *iff $\mathcal{R}(C) \in$* STRUCTURED-CKTSAT. *Moreover, if C is a circuit of size n, then $\mathcal{R}(C) = \langle \mathcal{G}_{N,l}, \{T_0, \ldots, T_{l-1}\}\rangle$, 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 $A_0, \ldots, A_l$ proving that the reduced instance $\mathcal{R}(C) = \langle \mathcal{G}_{N,l}, \{T_0, \ldots, T_{l-1}\}\rangle$ 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 $A_0, \ldots, A_{l-1}$ 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 $\mathcal{G}_{N,l}$ according to the switching actions to obtain $A_1, \ldots, A_{l-1}$ and the incoming bits of A_0. This observation will be vital while constructing PCPPs (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 1 (see Definition 8.13). Then it can be verified that the transformation mapping satisfying assignments w of C to legal assignments $A_0, \ldots, A_{l-1}$ of $\mathcal{R}(C)$ is $\mathbb{F}_2$-linear. The reason is that each gate in the computation of $C(w)$ is an $\mathbb{F}_2$-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 [HS00]. The broad overview of the arithmetization is as follows: We embed the nodes in each layer of the wrapped de Bruijn graph $\mathcal{G}_{N,l}$ into a vector space and extend the gate specifications and assignments to low-degree polynomials over this space. Finally, we express the assignment constraints (see Definition 6.3) as a pair of polynomial identities satisfied by these polynomials.

First, we have 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 $\mathbb{F}$ of $\mathbb{F}_2$ 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 $\mathbb{F} = \mathbb{F}_2^{\mho} = \mathbb{F}_{2^{\mho}}$ for $\mho = h + 2\log m + \log c_F$. Let $\{e_0, e_1, \ldots, e_{\mho-1}\}$ be a basis of $\mathbb{F}$ over $\mathbb{F}_2$. Set H to be a subspace of $\mathbb{F}_2^{\mho}$ (and hence a subset of $\mathbb{F}$) spanned by $\{e_0, \ldots, e_{h-1}\}$. Note that H^m is a subset of the space $\mathbb{F}^m$. Furthermore, $|H^m| = 2^N$. Hence, we can embed each layer of the graph $\mathcal{G}_{N,l}$ in $\mathbb{F}^m$ by identifying the node $v = (b_0, \ldots, b_{N-1}) \in \{0,1\}^N$ with the element $(b_0 e_0 + \cdots + b_{h-1} e_{h-1}, b_h e_0 + \cdots + b_{2h-1} e_{h-1}, \ldots, b_{(m-1)h} e_0 + \cdots + b_{mh-1} e_{h-1})$ 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 $S : H^m \to \mathbb{F}$ can be interpolated to obtain a polynomial $\tilde{S} : \mathbb{F}^m \to \mathbb{F}$ 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 $\tilde{S} : \mathbb{F}^m \to \mathbb{F}$ can be interpreted as an assignment from H^m to $\mathbb{F}$ by considering the function restricted to the subdomain 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 a fixed subset of $\mathbb{F}$ and identify $\{0,1\}^4$ with the set of elements spanned by $\{e_0, e_1, e_2, e_3\}$ over $\mathbb{F}_2$. With this identification, we can view the assignments A_i and gates T_i as functions $A_i : H^m \to \mathbb{F}$ and $T_i : H^m \to \mathbb{F}$, respectively. Furthermore, we can interpolate these functions, as mentioned above, to obtain polynomials $\tilde{A}_i : \mathbb{F}^m \to \mathbb{F}$ and $\tilde{T}_i : \mathbb{F}^m \to \mathbb{F}$ of degree at most $m|H|$ over $\mathbb{F}$.

We now express the neighborhood functions of the graph in terms of affine functions over $\mathbb{F}^m$. This is where the nice structure of the wrapped de Bruijn graph will be

useful. For any positive integer i, define affine transformations $\tilde{\Gamma}_{i,0}, \tilde{\Gamma}_{i,1} : \mathbb{F}^m \to \mathbb{F}^m$ 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 $\tilde{\Gamma}_{i,1}(z_0, \ldots, z_{m-1}) = (z_0, \ldots, z_{t-1}, z_t + e_u, z_{t+1}, \ldots, z_{m-1})$.[26] 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 $\mathbb{F}^m$.

Finally, we now express the assignment constraints (see 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 $\psi_0 : \mathbb{F} \to \mathbb{F}$ be the univariate polynomial of degree 2^4 given by

$$\psi_0(z) = \prod_{\alpha \in \{0,1\}^4} (z - \alpha). \tag{6.1}$$

This polynomial satisfies $\psi_0(z) = 0$ iff $z \in \{0,1\}^4$ (recall we identified $\{0,1\}^4$ with a subset of $\mathbb{F}$ spanned by $e_0, \ldots, e_3$). 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 $\psi_1 : \mathbb{F}^4 \to \mathbb{F}$. In other words, $\psi_1 : \mathbb{F}^4 \to \mathbb{F}$ 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.[27] 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 $\psi' = (\psi_0, \psi_1) : \mathbb{F}^4 \to \mathbb{F}^2$ such that $\psi'(x_1, x_2, x_2, x_4) = (\psi_0(x_2), \psi_1(x_1, x_2, x_3, x_4))$.[28] 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 $\kappa m^2 2^h / |\mathbb{F}|$ is low.

We have thus reduced Structured-CktSAT to an algebraic consistency problem, which we shall call the AS-CktSAT(Algebraic-Structured-CktSAT) problem.[29]

Definition 6.9. *The promise problem* AS-CktSAT = (AS-CktSAT^{YES}, AS-CktSAT^{NO}) *has as its instances* $\langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l-1}\}\rangle$, *where* $\mathbb{F}$ *is a finite extension field of* $\mathbb{F}_2$ *(i.e.,* $\mathbb{F} = \mathbb{F}_{2^{\mathfrak{v}}}$ *for some* $\mathfrak{v}$*),* H *an* $\mathbb{F}_2$*-linear subspace of* $\mathbb{F}$*, and* $\tilde{T}_i : \mathbb{F}^m \to \mathbb{F}$ *for* $i = 0, \ldots, l-1$ *a sequence of polynomials of degree* d*, where* $|H| = n^{1/m}$*,*

[26] An alternate description of $\tilde{\Gamma}_{i,1}$ is as follows: Since $\mathbb{F} = \mathbb{F}_2^{\mathfrak{v}}$, we can view $\mathbb{F}^m$ as an $m\mathfrak{v}$-dimensional space over $\mathbb{F}_2$. Hence, any vector $(z_0, \ldots, z_{m-1})$ can be written as $(b_{0,0}, \ldots, b_{0,\mathfrak{v}-1}, b_{1,0}, \ldots, b_{1,\mathfrak{v}-1}, \ldots, b_{m-1,0}, \ldots, b_{m-1,\mathfrak{v}-1})$. Furthermore, we note that for any vector $(z_0, \ldots, z_{m-1})$ 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 $u = i \bmod h$.

[27] 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 nonstandard choice in section 8.4. Specifically, we argue that if C is a linear circuit, then ψ_1 can be picked to be an $\mathbb{F}_2$-linear transformation, and we point out that ψ_0 is an $\mathbb{F}_2$-linear transformation. For more details see section 8.4.

[28] An alternative approach (which we do not follow) to combine ψ_0 and ψ_1 into a single polynomial ψ was suggested to us by Sergey Yekhanin and Jaikumar Radhakrishnan. Let $p(x)$ be a monic, quadratic, irreducible polynomial over the field $\mathbb{F}$. Let $Q(x, y) = y^2 \cdot p(x, y)$. Now Q satisfies the property that $Q(a, b) = 0$ iff $a = b = 0$. (Proof: Q is homogeneous and so $Q(0, 0) = 0$; if $b = 0$, $Q(a, 0) = a^2$ (since p is monic) and so is zero only if $a = 0$. If $b \neq 0$, then $Q(a, b) = 0$ only if $p(a/b) = 0$, but p has no roots in $\mathbb{F}$.) Now define $\psi(x) = Q(\psi_0(x), \psi_1(x))$. For instance, over fields of odd characteristic, $\psi(x) = \psi_0^2(x) - \alpha \cdot \psi_1^2(x)$, where α is a nonsquare in $\mathbb{F}$ will work. ψ can be thought of as an "algebraic AND" of ψ_0 and ψ_1, since ψ' satisfies the property that $\psi(x) = 0$ iff $\psi_0(x) = 0$ and $\psi_1(x) = 0$.

[29] AS-CktSAT is actually a promise problem.

$d = m \cdot |H|$, and $|\mathbb{F}| = c_F \cdot md$. The field $\mathbb{F}$ is specified by an irreducible polynomial $p(x)$ of degree f over $\mathbb{F}_2$, 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 as follows:

- $\langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l-1}\}\rangle \in$ AS-CktSAT^{YES} *if there exists a sequence of degree d polynomials $\tilde{A}_i : \mathbb{F}^m \to \mathbb{F}, i = 0, \ldots, l-1$, such that for all $i = 0, \ldots, l-1$ and all $x \in H^m$,*

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

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

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

where the $\tilde{\Gamma}_{i,j}$'s and ψ' are as defined earlier. (Recall that the $\tilde{\Gamma}$'s are linear functions 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 $\mathcal{I} = \langle \mathcal{G}_{N,l}, \{T_0, T_1, \ldots, T_{l-1}\}\rangle$ of* Structured-CktSAT *and parameter $m \le \log n / \log\log n$ (where $n = |\mathcal{I}|$) to an instance $\mathcal{R}(\mathcal{I}, 1^m)$ of* AS-CktSAT *such that*

$$\mathcal{I} \in \text{Structured-CktSAT} \implies \mathcal{R}(\mathcal{I}, 1^m) \in \text{AS-CktSAT}^{\text{YES}},$$
$$\mathcal{I} \notin \text{Structured-CktSAT} \implies \mathcal{R}(\mathcal{I}, 1^m) \in \text{AS-CktSAT}^{\text{NO}}.$$

Moreover, if $\mathcal{R}(\mathcal{I}, 1^m) = \langle 1^{n'}, 1^{m'}, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l'-1}\}\rangle$, 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 $m \le \log n / \log\log n$ (where $n = |C|$) to an instance $\mathcal{R}(C, 1^m)$ of* AS-CktSAT *such that $C \in$ CktSAT $\iff \mathcal{R}(C, 1^m) \in$ AS-CktSAT.*

Moreover, if C is a circuit of size n, then $\mathcal{R}(C, 1^m) = \langle 1^{n'}, 1^{m'}, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l'-1}\}\rangle$ where $n' = \Theta(n)$, $m' = m$, and $l' \le 5 \log n'$. 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 $\tilde{A}_0, \ldots, \tilde{A}_{l-1}$ proving that the reduced instance $\mathcal{R}(C, 1^m) = \langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l-1}\}\rangle$ is a YES instance of AS-CktSAT contains 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 $\tilde{A}_0 : \mathbb{F}^m \to \mathbb{F}$ 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 $\tilde{A}_0, \ldots, \tilde{A}_{l-1}$ such that $\tilde{A}_0|_I$ contains w. This is done by taking low-degree extensions of the functions $A_0, \ldots, A_{l-1}$ 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 $\tilde{A}_0, \ldots, \tilde{A}_{l-1}$ satisfying $\mathcal{R}(C)$ is $\mathbb{F}_2$-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 $\mathbb{F}$, and from the fact that the interpolation from A_i to $\tilde{A}_i$ is $\mathbb{F}$-linear and hence $\mathbb{F}_2$-linear. For more information see section 8.4.

Comment. Recall that the arithmetization was obtained by considering low-degree extensions over $\mathbb{F}^m$ of functions from H^m to H. If H were a subfield of the field $\mathbb{F}$, this step would have caused a quadratic blow-up, and we avoid this problem by not insisting that H be a field. In [PS94, Spi95], H is a field and $\mathbb{F} = H^2$ is an extension of it, but the PCP system refers only to an $O(|H|)$-sized subset of $\mathbb{F}$. We cannot take this approach because we will be using a total low-degree test, which needs to refer to the entire vector space $\mathbb{F}^m$. In contrast, in [PS94, Spi95] an individual low-degree test is used, which can work with a subset of $\mathbb{F}^m$.

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. [BSVW03]. 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 $\tilde{A}_i : \mathbb{F}^m \to \mathbb{F}$ for $i = 0, \ldots, l-1$ and an auxiliary sequence of oracles $P_{i,j} : \mathbb{F}^m \to \mathbb{F}^2$ for $i = 0, \ldots, l-1$ and $j = 0, \ldots, m$. For each i and j, we view the auxiliary oracle $P_{i,j} : \mathbb{F}^m \to \mathbb{F}^2$ as a pair of functions $P^{(0)}_{i,j} : \mathbb{F}^m \to \mathbb{F}$ and $P^{(1)}_{i,j} : \mathbb{F}^m \to \mathbb{F}$ (i.e., $P_{i,j}(x) = (P^{(0)}_{i,j}(x), P^{(1)}_{i,j}(x))$). 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 $i = 0, \ldots, l-1$, to satisfy the following relation:

$$P_{i,0}(x) = \psi'\Big(\tilde{T}_i(x), \tilde{A}_i(x), \tilde{A}_{i+1}\big(\tilde{\Gamma}_{i,0}(x)\big), \tilde{A}_{i+1}\big(\tilde{\Gamma}_{i,1}(x)\big)\Big) \quad \forall x \in \mathbb{F}^m. \tag{6.2}$$

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

Lemma 6.14.

1. *If $\langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l-1}\}\rangle$ is a YES instance of* AS-CktSAT*, satisfied by polynomials $\tilde{A}_0, \ldots, \tilde{A}_{l-1}$, and $P_{0,0}, \ldots, P_{l-1,0}$ are defined according to (6.2), then $P_{i,0}(x) = (0,0)$ for all $x \in H^m$.*

2. *If $\langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l-1}\}\rangle$ is a NO instance of* AS-CktSAT*, then for any sequences of functions $\tilde{A}_0, \ldots, \tilde{A}_{l-1}$, $P_{0,0}, \ldots, P_{l-1,0}$, either (6.2) 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 item 1 above.

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. [LFKN92]. Specifically, the verifier expects the remaining set of auxiliary oracles $P_{i,j} = (P^{(0)}_{i,j}, P^{(1)}_{i,j})$ ($i = 0, \ldots, l-1$ and $j = 1, \ldots, m$) to satisfy the following relations: Let $H = \{h_0, \ldots, h_{|H|-1}\}$ be some fixed enumeration of the elements in H. For all $b \in \{0,1\}$,

$$P^{(b)}_{i,j}\Big(\underbrace{x_1, \ldots, x_{j-1}}, x_j, \underbrace{x_{j+1}, \ldots, x_m}\Big) = \sum_{k=0}^{|H|-1} P^{(b)}_{i,j-1}\Big(\underbrace{x_1, \ldots, x_{j-1}}, h_k, \underbrace{x_{j+1}, \ldots, x_m}\Big) x_j^k \quad \forall (x_1, \ldots, x_m) \in \mathbb{F}^m. \tag{6.3}$$

These relations ensure that for all i and $j \geq 1$, $P_{i,j}(\cdot)$ vanishes on $\mathbb{F}^j \times H^{m-j}$ iff the function $P_{i,j-1}(\cdot)$ vanishes on $\mathbb{F}^{j-1} \times H^{m-j+1}$. For future reference, we (re)state this fact as a lemma.

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

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

Thus, the verifier needs to make the following checks:

- LOW-DEGREE TEST.
  For $i = 0, \ldots, l-1$ and $j = 0, \ldots, m$, the functions $\tilde{A}_i$ are polynomials of degree at most $d = m \cdot |H|$ and the functions $P_{i,j}$ are pairs of polynomials of degree at most κd.
- EDGE-CONSISTENCY TEST.
  For $i = 0, \ldots, l-1$, the functions $P_{i,0}$ obey (6.2).
- ZERO-PROPAGATION TEST.
  For $i = 0, \ldots, l-1$ and $j = 1, \ldots, m$, the functions $P_{i,j}$ satisfy (6.3).
- IDENTITY TEST.
  For $i = 0, \ldots, l-1$, the functions $P_{i,m}$ are identically zero on the entire domain $\mathbb{F}^m$.

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 $x, y \in \mathbb{F}^m$ and setting $l = l_{x,y} = \{x + ty | t \in \mathbb{F}\}$. However, this requires $2m \log |\mathbb{F}|$ 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. [BSVW03] based on small-biased spaces (see Appendix B for more details). The low-degree test of [BSVW03] 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 $\mathbb{F}^m$, while the second point y is chosen from a λ-biased subset S_λ of $\mathbb{F}^m$. This needs only $\log |S_\lambda| + \log |\mathbb{F}|^m$ bits of randomness. We further save on randomness by the use of *canonical lines*.[30] Consider any pseudorandom line $l = l_{x,y}$, where $x \in \mathbb{F}^m$ and $y \in S_\lambda$. We observe that for every $x' \in l$, we have $l_{x',y} = l_{x,y}$. In other words, $|\mathbb{F}|$ different choices of random bits lead 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 y as specifying the direction (i.e., slope) of the line. This direction partitions the space $\mathbb{F}^m$ into $|\mathbb{F}|^{m-1}$ 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 $\log |S_\lambda| + \log |\mathbb{F}|^{m-1}$ 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_{ij}'s are close to low-degree polynomials), the verifier then performs the NODE-

[30] It 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.

Consistency Test, Zero-Propagation Test, and Identity Test 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 $|\mathbb{F}|$ points selected from a partition of $\mathbb{F}^m$ into $|\mathbb{F}|^{m-1}$ 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 $\mathbb{F}^m$ required for the low-degree tests of Ben-Sasson et al. [BSVW03]. We rely on the fact that λ-biased subsets of $\mathbb{F}^m$ of size at most $\mathrm{poly}(\log|\mathbb{F}|^m, 1/\lambda)$ can be constructed efficiently [NN93, AGHP92].

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

1. Use Proposition 6.11 to reduce the instance C of CktSAT, using parameter m, to an instance $\langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l-1}\}\rangle$ of AS-CktSAT, and set $d = m \cdot |H|$.
   Notation. We let $S_\lambda \subset \mathbb{F}^m$ be a λ-biased set of size at most $(\frac{\log|\mathbb{F}|^m}{\lambda})^2$ [AGHP92]. Let $\mathbb{F}^m = \biguplus_{\eta=1}^{|\mathbb{F}|^{m-1}} U_\eta$ and $\mathbb{F}^m = \biguplus_{\eta=1}^{|\mathbb{F}|^{m-1}} V_\eta$ be two arbitrary partitions of the space $\mathbb{F}^m$ into $|\mathbb{F}|$-sized sets each.
2. Choose a random string R of length $\log(|S_\lambda| \cdot |\mathbb{F}|^{m-1})$. (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 $\mathbb{F}^m$ using the λ-biased set S_λ.
   For $i = 0, \ldots, l-1$,
   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 $i = 0, \ldots, l-1$, $j = 0, \ldots, m$, 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 $\mathbb{F}^m = \biguplus_{\eta=1}^{|\mathbb{F}|^{m-1}} U_\eta$.
   For $i = 0, \ldots, l-1$,
   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 (6.2) is not satisfied.
5. Zero-Propagation Test.
   For $i = 0, \ldots, l-1$, $j = 1, \ldots, m$, and $b \in \{0,1\}$,
   use random string R to determine a random jth axis-parallel line in $\mathbb{F}^m$ of the form $\mathcal{L} = \{(a_1, \ldots, a_{j-1}, X, a_{j+1}, \ldots, a_m) : X \in \mathbb{F}\}$. 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 (6.3).
6. Identity Test.
   Use the random string R to determine a random set V_η of the partition

$\mathbb{F}^m = \biguplus_{\eta=1}^{|\mathbb{F}|^{m-1}} V_\eta$. For $i = 0, \ldots, l-1$,
for all $x \in V_\eta$, query $P_{i,m}(x)$. Reject if any of these $P_{i,m}(x)$ are not $(0,0)$.

Accept if none of the above tests reject.

Remark 6.16.

1. The Low-Degree Test requires $\log(|S_\lambda| \cdot |\mathbb{F}|^{m-1})$ random bits to generate a canonical line in $\mathbb{F}^m$ using the λ-biased set, while each of the other tests requires at most $\log(|\mathbb{F}|^{m-1})$ bits of randomness. Hence, the string R suffices for each of the tests. For the settings of parameters we use, $\log(|S_\lambda| \cdot |\mathbb{F}|^{m-1})$ is typically significantly smaller than $\log(|\mathbb{F}|^m)$, 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 $\mathbb{F}^m$ and checking whether the corresponding condition is satisfied. However, we state these tests in a "nonstandard" manner using partitions of the space $\mathbb{F}^m$ into $|\mathbb{F}|$-sized tests so that these tests can be easily adapted when we construct the robust PCP verifier (see section 8). The nonstandard 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 $\mathbb{F}^m$; 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 $m \le \log n / \log\log n$ and $\lambda \le 1/c \log n$ for a sufficiently large constant c. Theorem 6.1 can be deduced by setting $\lambda = 1/c \log n$.

Complexity. The PCP Verifier makes $O(lm|\mathbb{F}|) = O(m^3 n^{1/m} \log n)$ queries, each of which expects as an answer an element of $\mathbb{F}$ or $\mathbb{F}^2$ (i.e., a string of length $O(\log|\mathbb{F}|)$). Hence, the total (bit) query complexity is $O(lm|\mathbb{F}|\log|\mathbb{F}|) = O(lm \cdot c_F m^2 n^{1/m} \log(c_F m^2 n^{1/m}))$. Recalling that $l = 5\log n$, this quantity is at most $O(m^2 n^{1/m} \log^2 n)$ for $m \le \log n$. For the decision complexity, we note that the main computations required are (a) testing whether a function is a low-degree univariate polynomial over $\mathbb{F}$ (for Low-Degree Test and Zero-Propagation Test), (b) evaluating ψ' on $|\mathbb{F}|$ quadruples of points (for Edge-Consistency Test), and (c) univariate polynomial interpolation and evaluation (for testing (6.3) in Zero-Propagation Test). We now argue that each of these can be done with a nearly linear ($\widetilde{O}(|\mathbb{F}|)$) number of operations over $\mathbb{F}$, yielding a nearly linear ($\widetilde{O}(q)$) decision complexity overall. Each evaluation of ψ' can be done with a constant number of $\mathbb{F}$-operations because ψ' is of constant degree. Polynomial interpolation and evaluation can be done with a nearly linear number of $\mathbb{F}$-operations by [SS71, Sch77], and testing whether a function is of low degree reduces to polynomial interpolation (this is achieved by interpolating the function to represent it as a polynomial of degree $|\mathbb{F}|-1$ and checking that the high-degree coefficients are zero). Each $\mathbb{F}$-operation can be done with $\widetilde{O}(\log|\mathbb{F}|)$ bit operations, using the polynomial multiplication algorithm of [SS71, Sch77] (over $\mathbb{F}_2$).

The number of random bits used by the verifier is exactly $\log(|S_\lambda| \cdot |\mathbb{F}|^{m-1})$. Let $n' = |\mathbb{F}|^m$. Then $\log(|S_\lambda| \cdot |\mathbb{F}|^{m-1}) = \left(1 - \frac{1}{m}\right)\log n' + \log(\mathrm{poly}(\frac{\log n'}{\lambda})) = \left(1 - \frac{1}{m}\right)\log n' + O(\log\log n') + O\left(\log\left(\frac{1}{\lambda}\right)\right)$. 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 $\left(1 - \frac{1}{m}\right)\log n + O(m\log m) + O(\log\log n) + O\left(\log\left(\frac{1}{\lambda}\right)\right)$.

We summarize the above observations in the following proposition for future ref-

erence.

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

Completeness. If C is satisfiable, then the reduction reduces it to a YES instance of AS-CKTSAT. Then by definition there exist polynomials $\tilde{A}_i$ that satisfy constraint ψ'. Setting $P_{i,j}$ according to (6.2) and (6.3), 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 $\tilde{A}_0 : \mathbb{F}^m \to \mathbb{F}$ supposedly has the satisfying assignment embedded within it. Let $I \subset \mathbb{F}^m$ be the set of locations in $\mathbb{F}^m$ 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 $\varepsilon \le \varepsilon_0$, $m \le \log n/\log\log n$, and $\lambda \le 1/c\log n$, the following holds. If $\tilde{A}_0$ is 4ε-far from every polynomial $\widehat{A}_0$ of degree md such that $C(\widehat{A}_0|_I) = 1$, then for all proof oracles $\{\tilde{A}_i\}$ and $\{P_{i,j}\}$, the verifier accepts with probability at most $1 - \varepsilon$.*

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 $\kappa md/|\mathbb{F}| = \kappa/c_F \le \varepsilon_0$. 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 κmd. If this were not the case, then by Theorem B.4, 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 $i = 0, \ldots, l-1$, let $\widehat{A}_i : \mathbb{F}^m \to \mathbb{F}$ be the polynomial of degree at most md that is 4ε-close to $\tilde{A}_i$. Similarly, for each $i = 0, \ldots, l-1, j = 0, \ldots, m$ and $b \in \{0,1\}$, let $\widehat{P}_{i,j}^{(b)}$ 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 $1 - \frac{\kappa md}{|\mathbb{F}|} \ge 1 - \varepsilon_0 > 8\varepsilon$ fraction of points. As in the case of the $P_{i,j}$'s, let $\widehat{P}_{i,j} : \mathbb{F}^m \to \mathbb{F}^2$ be the function given by $\widehat{P}_{i,j}(x) = (\widehat{P}_{i,j}^{(0)}(x), \widehat{P}_{i,j}^{(1)}(x))$.

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

(a) *There exists $i \in \{0, \ldots, l-1\}$ and $b \in \{0,1\}$ such that $\widehat{P}_{i,m}^{(b)} \not\equiv 0$.*
Then for this i, 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 $P_{i,m}^{(b)}(x) \ne \widehat{P}_{i,m}^{(b)}(x)$ or $\widehat{P}_{i,m}^{(b)}(x) = 0$. Since the V_η's form a partition of $\mathbb{F}^m$, the probability of this occurring is upper-bounded by the probability that a random $x \in \mathbb{F}^m$ satisfies either $P_{i,m}^{(b)}(x) \ne \widehat{P}_{i,m}^{(b}(x)$ or $\widehat{P}_{i,m}^{(b)}(x) = 0$. This probability is at most $4\varepsilon + \frac{\kappa md}{|\mathbb{F}|} = 4\varepsilon + \frac{\kappa}{c_F} \le 5\varepsilon_0$, where we use the fact that $\widehat{P}_{i,m}^{(b)}$ is 4ε-close to $P_{i,m}^{(b)}$ and that a nonzero polynomial of degree κmd vanishes on at most a $\kappa md/|\mathbb{F}|$ fraction of points.

(b) *There exists $i \in \{0, \ldots, l-1\}$ such that $\widehat{P}_{i,0}$, $\widehat{A}_i$, and $\widehat{A}_{i+1}$ do not obey* (6.2).
In other words, $\widehat{P}_{i,0}(x) \not\equiv \psi'(\tilde{T}_i(x), \widehat{A}_i(x), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x)))$.
Then for this i, EDGE-CONSISTENCY TEST fails unless a random partition U_η is chosen such that for all $x \in U_\eta$, $P_{i,0}(x) = \psi'(\tilde{T}_i(x), \tilde{A}_i(x), \tilde{A}_{i+1}(\tilde{\Gamma}_{i,0}(x))$,

$\tilde{A}_{i+1}(\tilde{\Gamma}_{i,1}(x)))$. Hence, it must be the case that for every $x \in U_\eta$, one of the following (six) holds:

$$P^{(0)}_{i,0}(x) \neq \widehat{P}^{(0)}_{i,0}(x); \qquad P^{(1)}_{i,0}(x) \neq \widehat{P}^{(1)}_{i,0}(x); \qquad \tilde{A}_i(x) \neq \widehat{A}_i(x);$$
$$\tilde{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)) \neq \widehat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)); \qquad \tilde{A}_{i+1}(\tilde{\Gamma}_{i,1}(x)) \neq \widehat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x));$$
$$\widehat{P}_{i,0}(x) = \psi'(\tilde{T}_i(x), \widehat{A}_i(x), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))).$$

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

(c) *For some $i = 0, \dots, l-1$, $j = 1, \dots, m$, and $b \in \{0,1\}$, $\widehat{P}^{(b)}_{i,j}$ does not obey* (6.3).

In other words, $\widehat{P}^{(b)}_{i,j}(\dots, x_j, \dots) \not\equiv \sum_{k=0}^{|H|-1} \widehat{P}^{(b)}_{i,j-1}(\dots, h_j, \dots)x_j^k$. Then, for this i, j, Zero-Propagation Test rejects unless a random axis-parallel line $\mathcal{L}$ is chosen such that both $P^{(b)}_{i,j}|_{\mathcal{L}}$ and $P^{(b)}_{i,j-1}|_{\mathcal{L}}$ are polynomials of degree at most κd and for every $x \in \mathcal{L}$, $P^{(b)}_{i,j}(\dots, x, \dots) = \sum_{k=0}^{|H|-1} P^{(b)}_{i,j-1}(\dots, h_k, \dots)x^k$. Suppose we have that for all $x \in \mathcal{L}$, $P^{(b)}_{i,j}(x) = \widehat{P}^{(b)}_{i,j}(x)$ and $P^{(b)}_{i,j-1}(x) = \widehat{P}^{(b)}_{i,j-1}(x)$. Then, it must be the case that for all $x \in \mathcal{L}$, $\widehat{P}^{(b)}_{i,j}(\dots, x, \dots) = \sum_{k=0}^{|H|-1} \widehat{P}^{(b)}_{i,j-1}(\dots, h_k, \dots)x^k$. Since the axis-parallel lines cover $\mathbb{F}^m$ uniformly, the probability of this occurring is at most the probability of a random $x \in \mathbb{F}^m$ satisfying this condition, which is at most $\frac{\kappa m d}{c_F} \leq \varepsilon$. The probability that both $P^{(b)}_{i,j}|_{\mathcal{L}}$ and $P^{(b)}_{i,j-1}|_{\mathcal{L}}$ are polynomials of degree κd and either $P^{(b)}_{i,j}|_{\mathcal{L}} \neq \widehat{P}^{(b)}_{i,j}|_{\mathcal{L}}$ or $P^{(b)}_{i,j-1}|_{\mathcal{L}} \neq P^{(b)}_{i,j-1}|_{\mathcal{L}}$ is at most $2 \cdot 4\varepsilon/(1-\varepsilon_0) \leq 9\varepsilon_0$, since $P^{(b)}_{i,j}$ and $P^{(b)}_{i,j-1}$ are 4ε-close to $\widehat{P}^{(b)}_{i,j}$ and $\widehat{P}^{(b)}_{i,j-1}$, respectively, and any two distinct polynomials of degree $\kappa m d$ disagree on at least a $1 - \kappa m d/|\mathbb{F}| \geq 1 - \varepsilon_0$ fraction of points. Hence, 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 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 $\langle 1^{n'}, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\}\rangle$ 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 PCPP. In this section, we modify the PCP for CktSAT and construct a PCPP for CktVal while maintaining all the complexities. (Recall that, by Proposition 2.4, the latter implies the former.) We do so by adding a proximity test to 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 assumed 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 $0 < \delta < 1$ satisfying $n^{1/m} \geq m^{cm}/\delta^3$. There exists a PCPP for* CktVal *(for circuits of size n) with the following parameters:*

- *randomness* $\left(1 - \frac{1}{m}\right)\log n + O(m \log m) + O(\log\log n) + O(\log(1/\delta))$,

- *query complexity* $q = O(m^2 n^{1/m} \log^2 n)$ *and decision complexity* $d = \widetilde{O}(q)$,
- *perfect completeness,*
- *soundness error* $1 - \varepsilon$ *for proximity parameter* δ.

Note that the condition $n^{1/m} \geq m^{cm}/\delta^3$ (made in Theorem 7.1) implies the condition $m \leq \log n / \log\log n$ stated in Theorem 6.1. Thus, the PCPP of Theorem 7.1 works only when n, the size of the circuit, is not too small (more, precisely, when $n \geq m^{cm^2}/\delta^{3m}$). As explained in section 3, when applying multiple proof compositions, we need (at the last compositions) PCPPs that work for even smaller values of n. For this purpose, we construct the following PCPP that works for small values of n. This PCPP, however, performs relatively poorly with respect to randomness (i.e., it has randomness complexity $O(\log n)$ rather than $(1 - o(1)) \log_2 n$). This will not be a concern for us since this verifier (or rather the robust version of this verifier) is used only in the inner levels of composition.

THEOREM 7.2. *For all* $n \in \mathbb{Z}^+$ *and* $\delta \in (0,1)$, CKTVAL *has a PCPP (for circuits of size* n*) with the following parameters:*

- *randomness* $O(\log n)$;
- *decision complexity* $\operatorname{poly}\log n$, *which also upper-bounds the query complexity;*
- *perfect completeness; and*
- *soundness error* $1 - \Omega(\delta)$ *for proximity parameter* δ.

Preliminaries. Recall that a 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 $W : [k] \to \{0,1\}$. 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 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, and thus it better be the case that w is a nonnegligible fraction of the circuit. This assumption is not a major restriction, because if this is not the case then we work with the modified circuit C' and input w' that are defined as follows: For $t = \lceil n/k \rceil$, the circuit C' has $n' = n + 3tk$ gates and $k' = tk$ 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 consists of t copies of some satisfying assignment of C. (It can be verified that C' can indeed be implemented on a circuit of size $n + 3tk$.) 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 $W : [k] \to \{0,1\}$ in a straightforward manner; that is, it defines $W' : [tk] \to \{0,1\}$ such that $W'(i) \triangleq W(((i-1) \bmod k) + 1)$. Indeed, in view of the way W' is emulated based on W, testing that W' is a repetition of some k-bit string makes no sense. This test is incorporated into C' in order to maintain the distance features of C; that is, if w is δ-far from satisfying C, then $w' = w^t$ is δ-far from satisfying C' (without having C' explicitly run C on all t copies of w, because that would make its size larger than nt and defeat our goal of increasing the length of the input relative to the circuit size). We state this fact as a proposition for future reference.

PROPOSITION 7.3. *There exists a generic transformation from* CKTVAL *to* CKTVAL *that maps the instance* (C, w), *where* C *is a circuit with* n *gates and* k *input bits, to the instance* (C', w'), *where* C' *is a circuit on* $n' = n + 3tk$ *gates and* $k' = kt$ *input bits (where* $t = \lceil n/k \rceil$*) defined as follows:* $C'(w') = 1$ *iff* $w' = w^t$ *for some* w *such that* $C(w) = 1$ *and* $w' = w^t$. *This transformation increases the length of the input oracle*

compared to the proof oracle (here, the values of all gates in C). The transformation preserves the relative distance to the set of satisfying assignments; that is, if w is δ-far from the set of satisfying assignments of C, then $w' = w^t$ is δ-far from the satisfying assignments of C'.

We first describe PCPP Verifier which proves Theorem 7.1 and later describe ALMSS PCPP Verifier which proves Theorem 7.2.

7.1. The construction of PCPP Verifier (Theorem 7.1). As in the case of PCP Verifier described in section 6.3, PCPP Verifier is constructed by reducing the input circuit C, an instance of CktSAT, using parameter m, to an instance $\langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l-1}\}\rangle$ of AS-CktSAT. The proof oracle for PCPP Verifier is the same as that of the PCP Verifier (i.e., the proof oracle consists of a sequence of functions $\tilde{A}_i : \mathbb{F}^m \to \mathbb{F}, i = 0, \ldots, l-1$, and $P_{i,j} : \mathbb{F}^m \to \mathbb{F}^2, i = 0, \ldots, l-1, j = 0, \ldots, m$, where $l = 5 \log n$).

Recall that the function $\tilde{A}_0 : \mathbb{F}^m \to \mathbb{F}$ is assumed to contain within it an assignment (see Remarks 6.7 and 6.12). Let $I \subseteq H^m \subset \mathbb{F}^m$ be the set of locations in $\mathbb{F}^m$ that contain the assignment. PCPP Verifier, in addition to the tests of PCP Verifier, performs the following Proximity Test to check if the assignment given by $\tilde{A}_0|_I$ matches the input oracle W. Specifically, we have the following.

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

1. Run PCP Verifier$^{W;\ \tilde{A}_i, P_{i,j}}_{m,\lambda}(C)$ and reject if it rejects.
   Let R be the random string generated during the execution of this step.
2. Proximity Test.
   Use random string R to determine a random canonical line $\mathcal{L}$ in $\mathbb{F}^m$ 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 $I \cap \mathcal{L}$.

By inspection, the proximity test increases the query and decision complexities by (even less than) a constant factor. For the randomness complexity, the randomness is used only to generate a random canonical line (as in PCP Verifier), so the randomness complexity is $\log(|\mathbb{F}|^{m-1} \cdot |S_\lambda|)$ as before. However, in order to prove soundness, we need to assume not only that $\lambda \le 1/c \log n$ for some constant c (as before), but also that $\lambda \le \delta^3/m^{cm}$.[31] 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. Summarizing the above observations for future reference, we have the following proposition.

Proposition 7.4. *The randomness, query, and decision complexities of* PCPP Verifier *are* $r = \left(1 - \frac{1}{m}\right) \log n + O(m \log m) + O(\log \log n) + O(\log(1/\delta))$, $q = O(m^2 n^{1/m} \log^2 n)$, *and* $d = \widetilde{O}(q)$, *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 the encoding of the input in the proof oracle. In the 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

[31] Actually, for the proximity test we need only $\lambda \le \delta/m^{cm}$; however, to prove robustness of the proximity test (see section 8.1) we require $\lambda \le \delta^3/m^{cm}$.

explained in detail in the following lemma which proves the soundness of the verifier.

LEMMA 7.5. *There exists a constant c and a constant $\varepsilon > 0$ such that for all m, λ, δ satisfying $n \geq 8000|\mathbb{F}|^{m-1}/\delta^3$, $\lambda \leq 1/c \log n$, and $\lambda \leq \delta/m^{cm}$, the following holds. If the input w given by the input oracle $W : [k] \to \{0,1\}$ 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 $\widehat{A}_0$ of degree md such that $C(\widehat{A}_0|_I) = 1$. 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 $\widehat{A}_0$ of degree md such that $C(\widehat{A}_0|_I) = 1$. Since w is δ-far from any satisfying assignment, the assignment given by $\widehat{A}_0|_I$ must be at least δ-far from w. Let $B \subset \mathbb{F}^m$ denote the set of locations in I where the assignment given by $\widehat{A}_0$ disagrees with w (i.e., $B = \{x \in I | \widehat{A}_0(x)$ disagrees with w at $x\}$). Hence, $|B|/|I| \geq \delta$. Since $|I| = k \geq n/5$, we have $|B| \geq \delta n/5$. Consider the following two events.

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

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

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

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

Suppose Event I does not occur. Then, if $\widehat{A}_0|_{\mathcal{L}} \neq \tilde{A}_0|_{\mathcal{L}}$, 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 $\widehat{A}_0$ and hence cannot be closer to any other polynomial (as $5\varepsilon \leq \frac{1}{2}(1 - \frac{d}{|\mathbb{F}|}) = \frac{1}{2}(1 - \frac{1}{c_F})$). Now if $\widehat{A}_0|_{\mathcal{L}} = \tilde{A}_0|_{\mathcal{L}}$ and Event II does not occur, then 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 $\max\left\{1 - \varepsilon, \frac{1}{2}\right\} = 1 - \varepsilon$. This completes the proof of the lemma. ∎

Proof of Theorem 7.1. Theorem 7.1 is proved using PCPP VERIFIER defined in this section, setting $\lambda = \min\{1/c \log n, \delta^3/m^{cm}\}$. The randomness and decision (resp., query) complexities follow from Proposition 7.4. The only fact to be verified is the soundness of the verifier. By the hypothesis of Theorem 7.1, $n^{1/m} \geq m^{cm}/\delta^3$ for a suitably large constant c. This implies that $n^{1/m} \geq 8000(c_F m^2)^{m-1}/\delta^3$ or, equivalently, $n \geq 8000|\mathbb{F}|^{m-1}/\delta^3$. Hence, Lemma 7.5 applies and we have that the verifier has soundness error $1 - \varepsilon$ for proximity parameter δ. This proves Theorem 7.1. ∎

7.2. The ALMSS-type PCPP (Theorem 7.2). We now turn to designing a PCPP that proves Theorem 7.2. We call this ALMSS PCPP VERIFIER as it has parameters similar to the "parallelized" PCPs of [ALM+98]. ALMSS PCPP VERIFIER is identical to PCPP VERIFIER of Theorem 7.1 except for the fact that it has a slightly different proximity test. All other details remain the same.

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

1. Set $m = \log n / \log\log n$ and $\lambda = 1/c \log n$.

2. Run PCP $\textsc{Verifier}_{m,\lambda}^{W;\ \tilde{A}_i, P_{i,j}}(C)$ and reject if it rejects.
3. ALMSS Proximity Test.
   Choose a random position $i \xleftarrow{\text{R}} \{1, \ldots, k\}$ in the input and a random direction $y \in \mathbb{F}^m$. Let $x \in I$ be the point corresponding to i in H^m. Let $\mathcal{L}$ be the line through x in the 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 $W[i] \neq \tilde{A}_0(x)$.

Unlike PCPP Verifier, we will not calculate the randomness used by this verifier exactly. An upper bound within a constant factor suffices for our purposes. The extra randomness used by ALMSS Proximity Test is $\log k + m \log |\mathbb{F}|$ (i.e., the randomness required to choose a random index in $\{1, \ldots, k\}$ and a random direction in $\mathbb{F}^m$). For our choice of m and δ, the randomness of PCP Verifier is at most $O(\log n)$ (see the analysis preceding Proposition 6.17). Hence, the total randomness of ALMSS PCPP Verifier is at most $O(\log n)$. The query and decision complexities are at most a constant times that of PCP Verifier which is turn is upper-bounded by $\operatorname{poly} \log n$. Summarizing the above observations for future reference, we have the following proposition.

Proposition 7.6. *The randomness and decision complexities of* ALMSS PCPP Verifier *are* $O(\log n)$ *and* $\operatorname{poly} \log n$*, respectively.*

The soundness of the verifier is given by the following lemma.

Lemma 7.7. *For all* $\delta \in (0, 1)$*, the following holds. If the input* w *given by the input oracle* $W : [k] \to \{0, 1\}$ *is* δ*-far from satisfying the circuit, then for any proof oracle the verifier rejects* W *with probability* $\Omega(\delta)$*.*

Proof. Let ε_0 be the constant that appears in Lemma 6.18.

Case (i). $\tilde{A}_0$ is not $4\varepsilon_0$-close to any polynomial $\widehat{A}_0$ of degree md such that $C(\widehat{A}_0|_I) = 1$. Then by Lemma 6.18, we conclude that the verifier rejects with probability at least ε_0.

Case (ii). $\tilde{A}_0$ is $4\varepsilon_0$-close to some polynomial $\widehat{A}_0$ of degree md such that $C(\widehat{A}_0|_I) = 1$. Since w is δ-far from any satisfying assignment, the assignment given by $\widehat{A}_0|_I$ must be δ-far from w. With probability greater than δ over the choice of $i \in \{1, \ldots, k\}$ (and the corresponding point $x \in I$ in H^m), we have $W[i] \neq \widehat{A}_0(x)$. If this occurs, the only way the verifier can accept is if $\tilde{A}_0|_{\mathcal{L}}$ is a degree md polynomial other than $\widehat{A}_0|_{\mathcal{L}}$. We will show that for any fixed point x in $\mathbb{F}^m$, with probability at least $1 - 16\varepsilon_0$ over the choice of random line $\mathcal{L}$ through x, $\tilde{A}_0|_{\mathcal{L}}$ cannot be a degree md polynomial different from $\widehat{A}_0|_{\mathcal{L}}$. We can then conclude that the verifier rejects with probability at least $\delta \cdot (1 - 16\varepsilon_0) = \Omega(\delta)$. Since $\mathcal{L}$ is a random line through x, every point on $\mathcal{L}$ other than x is a uniformly random point in $\mathbb{F}^m$. Recall that $\tilde{A}_0$ and $\widehat{A}_0$ are $4\varepsilon_0$-close. By a Markov argument it follows that for every fixed value of x and a random line $\mathcal{L}$ through x, with probability at least $1 - 16\varepsilon_0$, $\tilde{A}|_{\mathcal{L}\setminus\{x\}}$ and $\widehat{A}|_{\mathcal{L}\setminus\{x\}}$ are at least 1/4-close. This implies that $\tilde{A}|_{\mathcal{L}}$ cannot be a polynomial of degree md other than $\widehat{A}|_{\mathcal{L}}$ (since two distinct polynomials agree in at most md points, and $(md-1)/|\mathbb{F}| < 1/4$).

In either case, ALMSS PCPP Verifier rejects with probability at least $\min\{\varepsilon_0, \Omega(\delta)\} = \Omega(\delta)$. ∎

Conclusion. Theorem 7.2 follows from Proposition 7.6 and Lemma 7.7.

8. A randomness-efficient robust PCPP. In this section, we modify the PCPP for CktVal constructed in section 7 to design a robust PCPP, while essentially maintaining all complexities. Recall the definition of robustness: If the input oracle

W is δ-far from a satisfying assignment, then a "regular" PCPP verifier rejects the input for most choices of its random coins; that is, it observes an inconsistency in the (input and) proof. In contrast, for most choices of its random coins, a *robust* PCPP verifier not only notices an inconsistency in the (input and) proof but also observes that a considerable portion of the (input and) proof read by it has to be modified to remove this inconsistency.

We construct two robust PCPPs which are robust analogues of the two PCPPs presented in section 7. The first robust PCPP is the main construct (claimed in Theorem 3.1), which is the robust analogue of PCPP Verifier. The second robust PCPP is an ALMSS-type robust PCPP (claimed in Theorem 3.2), which is the robust analogue of ALMSS PCPP Verifier. Thus, we prove Theorems 3.1 and 3.2.

Overview of the proofs of Theorem 3.1 *and* 3.2. We "robustify" the PCPP verifier in three steps. Recall 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 (either Proximity Test or ALMSS Proximity Test, as the case may be)). 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. Indeed, bundling is analogous to "parallelization" except that it does *not* involve any increase in the randomness complexity (unlike parallelization, which introduces such an increase, which although small is too big for our purposes). We stress that bundling is tailored to the specific tests, in contrast to parallelization which is generic. The aforementioned 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, PCPP Verifier (resp., ALMSS 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 (resp., ALMSS Proximity Test). In this section, we prove that each of these tests is robust individually. In other words, we show that when one of these tests fails, 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, we introduce some notation. We view functions $g, g' : \mathbb{F}^m \to \mathbb{F}$ as strings of length $|\mathbb{F}|^m$ over the alphabet $\mathbb{F}$, so their relative Hamming distance $\Delta(g, g')$ is simply $\Pr_x[g(x) \neq g'(x)]$. As before, let $I \subseteq H^m \subset \mathbb{F}^m$ be the set of locations in $\mathbb{F}^m$ that contains the assignment.

Let $0 < \varepsilon < 1$ be a small constant to be specified later. As before, for $i = 0, \ldots, l-1$, $j = 0, \ldots, m$ and $b \in \{0,1\}$, let $\widehat{A}_i$ (resp., $\widehat{P}_{i,j}^{(b)}$) 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 verifiers, PCPP Verifier and ALMSS PCPP Verifier (see section 6 and 7), was along the following lines: If the input oracle $W : [k] \to \{0,1\}$ is δ-far from satisfying the circuit, then one of the following must happen (changing ε by a factor of 2):

1. There exists $i \in \{0, \ldots, l-1\}$ such that $\tilde{A}_i$ is 8ε-far from every degree md polynomial or there exists $i \in \{0, \ldots, l-1\}$, $j \in \{0, \ldots, m\}$, and $b \in \{0,1\}$ such that $P_{i,j}^{(b)}$ is 8ε-far from every degree κmd polynomial. In this case,

Low-Degree Test detects the error with probability at least 2ε.

2. There exists $i \in \{0, \ldots, l-1\}$ and $b \in \{0,1\}$, such that $\Delta(P_{i,m}^{(b)}, \widehat{P}_{i,m}^{(b)}) \leq 8\varepsilon$ and $\widehat{P}_{i,m} \not\equiv 0$. In this case, Identity Test detects the error with probability at least $1 - 10\varepsilon$.
3. There exists $i \in \{0, \ldots, l-1\}$, $j \in \{1, \ldots, m\}$, and $b \in \{0,1\}$ such that
$$\Delta(P_{i,j}^{(b)}, \widehat{P}_{i,j}^{(b)}) \leq 8\varepsilon, \Delta(P_{i,j-1}^{(b)}, \widehat{P}_{i,j-1}^{(b)}) \leq 8\varepsilon,$$
$$\text{and } \widehat{P}_{i,j}^{(b)}(\ldots, x_j, \ldots) \not\equiv \sum_{k=0}^{|H|-1} \widehat{P}_{i,j-1}^{(b)}(\ldots, h_k, \ldots) x_j^k.$$
In this case, Zero-Propagation Test detects the error with probability at least $1 - 20\varepsilon$.
4. There exists $i \in \{0, \ldots, l-1\}$ such that
$$\Delta(P_{i,0}^{(0)}, \widehat{P}_{i,0}^{(0)}) \leq 8\varepsilon, \Delta(P_{i,0}^{(1)}, \widehat{P}_{i,0}^{(1)}) \leq 8\varepsilon, \Delta(\tilde{A}_i, \widehat{A}_i) \leq 8\varepsilon, \Delta(\tilde{A}_{i+1}, \widehat{A}_{i+1}) \leq 8\varepsilon,$$
$$\text{and } \widehat{P}_{i,0}(x) \not\equiv \psi'(\tilde{T}_i(x), \widehat{A}_i(x), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))).$$
In this case, Edge-Consistency Test detects the error with probability at least $1 - 42\varepsilon$.
5. $\Delta(\tilde{A}_0, \widehat{A}_0) \leq 8\varepsilon$, but W and $\widehat{A}_0|_I$ disagree on at least δ fraction of the points. In this case, Proximity Test (or ALMSS Proximity Test, as the case may be) detects the error with probability at least 2ε (or $\Omega(\delta)$ in the case of ALMSS Proximity Test).

Claims 8.1–8.6 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 $(I, D) \xleftarrow{\text{R}} T$ equals ρ.

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

The robustness of Low-Degree Test can be easily inferred from the analysis of the λ-biased low-degree test due to Ben-Sasson et al. [BSVW03] as shown below.

Claim 8.1. *The following holds for all sufficiently small $\varepsilon > 0$. If $A : \mathbb{F}^m \to \mathbb{F}$ (resp., $P : \mathbb{F}^m \to \mathbb{F}$) is 8ε-far from every polynomial of degree md (resp., degree κmd), then the expected distance of A (resp., P) from satisfying* Low-Degree Test *with degree parameter d (resp., κd) is at least 2ε.*

Proof. Recall that 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 $\mathcal{L} \to \mathbb{F}$ 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. [BSVW03] works as follows (see Appendix B for more details): The test LDT has oracle access to a points-oracle $f : \mathbb{F}^m \to \mathbb{F}$ 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 queries 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 $\mathrm{LDT}^{A,A_{\mathrm{lines}}}$ rejects the points-oracle A and lines-oracle A_{lines} as defined above equals $\mathrm{E}_{\mathcal{L}}[\Delta(A|_{\mathcal{L}}, A_{\mathrm{lines}}(\mathcal{L}))]$. By Theorem B.4, if A is 8ε-far from every degree md polynomial, then $\mathrm{LDT}^{A,A_{\mathrm{lines}}}$ 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 8.2, 8.3, and 8.4, respectively.

Claim 8.2. *The following holds for all sufficiently small $\varepsilon > 0$. If for some $i = 0, \ldots, l-1$ and $b \in \{0,1\}$, $\Delta(P_{i,m}^{(b)}, \widehat{P}_{i,m}^{(b)}) \le 8\varepsilon$ and $\widehat{P}_{i,m}^{(b)}(\cdot) \not\equiv 0$, then $P_{i,m}$ has expected distance at least $1 - 9\varepsilon$ from satisfying* Identity Test.

Proof. The expected distance of $P_{i,m}$ from satisfying Identity Test equals

$$\begin{aligned}
\mathrm{E}_{V_\eta}[\Delta(P_{i,m}|_{V_\eta}, 0)] &= \Delta(P_{i,m}, 0) && \text{[since } \{V_\eta\} \text{ is a partition]}\\
&\ge \Delta(\widehat{P}_{i,m}, 0) - \Delta(P_{i,m}, \widehat{P}_{i,m}) \\
&\ge \left(1 - \tfrac{\kappa md}{|\mathbb{F}|}\right) - 8\varepsilon && \text{[by the Schwartz–Zippel lemma and hypothesis]}\\
&\ge 1 - 9\varepsilon. \quad ∎
\end{aligned}$$

Claim 8.3. *The following holds for all sufficiently small $\varepsilon > 0$. Suppose for some $i = 0, \ldots, l-1$, we have $\Delta(P_{i,0}^{(0)}, \widehat{P}_{i,0}^{(0)}) \le 8\varepsilon$, $\Delta(P_{i,0}^{(1)}, \widehat{P}_{i,0}^{(1)}) \le 8\varepsilon$, $\Delta(\tilde{A}_i, \widehat{A}_i) \le 8\varepsilon$, $\Delta(\tilde{A}_{i+1}, \widehat{A}_{i+1}) \le 8\varepsilon$, and $\widehat{P}_{i,0}(\cdot) \not\equiv \psi'(\tilde{T}_i(\cdot), \widehat{A}_i(\cdot), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,0}(\cdot)), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,1}(\cdot)))$. 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* 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 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 8.2, we have

$$\begin{aligned}
&\mathrm{E}_{U_\eta}[\Delta(P_{i,0}(\cdot)|_{U_\eta}, \psi'(\tilde{T}_i(\cdot), A_i(\cdot), A_{i+1}(\tilde{\Gamma}_{i,0}(\cdot)), A_{i+1}(\tilde{\Gamma}_{i,1}(\cdot)))|_{U_\eta})]\\
&\qquad \ge \left(1 - \frac{\kappa md}{|\mathbb{F}|}\right) - 5 \cdot 8\varepsilon \ge 1 - 41\varepsilon,
\end{aligned}$$

where the $(1-\kappa md/|\mathbb{F}|)$ 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 Edge-Consistency Test is at least $(1 - 41\varepsilon)/5$. ∎

Claim 8.4. *The following holds for all sufficiently small $\varepsilon > 0$. Suppose for some $i = 0, \ldots, l-1$, $j = 1, \ldots, m$, and $b \in \{0,1\}$, we have $\Delta(P_{i,j}^{(b)}, \widehat{P}_{i,j}^{(b)}) \le$*

8ε, $\Delta(P^{(b)}_{i,j-1}, \widehat{P}^{(b)}_{i,j-1}) \leq 8\varepsilon$, and $\widehat{P}^{(b)}_{i,j}(\dots, x_j, \dots) \not\equiv \sum_{k=0}^{|H|-1} \widehat{P}^{(b)}_{i,j-1}(\dots, h_k, \dots) x_j^k$. Then $(P^{(b)}_{i,j}, P^{(b)}_{i,j-1})$ has expected distance at least $(1-19\varepsilon)/2$ from satisfying Zero-Propagation Test.

Proof. Suppose that $\mathcal{L}$ is a jth axis-parallel line such that

$$\widehat{P}^{(b)}_{i,j}(\dots, x_j, \dots)|_{\mathcal{L}} \not\equiv \sum_{k=0}^{|H|-1} \widehat{P}^{(b)}_{i,j-1}(\dots, h_k, \dots) x_j^k|_{\mathcal{L}}.$$

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

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

fraction of points.

Thus the expected distance of $(P^{(b)}_{i,j}, P^{(b)}_{i,j-1})$ from satisfying Zero-Propagation Test is at least

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

We are now left with analyzing the robustness of the proximity tests (Proximity Test and ALMSS Proximity Test). Note that the input for either of these proximity tests comes in two parts: (a) the restriction of A_0 to the line $\mathcal{L}$ and (b) the input W restricted to the line $\mathcal{L}$ (or to a point on $\mathcal{L}$). Thus, the robustness of these tests refers to both parts (i.e., parts of each of the two oracles), and it is beneficial to decouple the corresponding robustness properties. We note that the robustness of Proximity Test is proved by repeated applications of the sampling lemma (Lemma B.3), while the robustness of ALMSS Proximity Test follows by a simple Markov argument.

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

Claim 8.5. *There exists a constant c and a constant $\varepsilon > 0$ such that for all $m, \lambda, \delta, \delta'$ satisfying $n \geq 8000|\mathbb{F}|^{m-1}/\delta^3$, $\lambda \leq 1/c \log n$, $\lambda \leq \delta^3/m^{cm}$, $\delta' > \delta$, the following holds. Suppose $\Delta(\tilde{A}_0, \widehat{A}_0) \leq 1/4$ and the input oracle W is δ'-far from*

$\widehat{A}_0|_I$ (i.e., $|B|/|L| \geq \delta'$); then with probability at least $1-\delta/4$ (over the choice of the canonical line $\mathcal{L}$) either at least an 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 PROXIMITY TEST *accept.*

This claim is the robust analogue of Lemma 7.5. 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 $((1/|\mathbb{F}|)+\lambda)\cdot(1/4)/(1/12)^2 \leq (\delta/12)$ since $|\mathbb{F}| \geq (8000|\mathbb{F}|^m)/(\delta^3 n) > (12^3/2)/\delta$ and $\lambda < 2\delta/12^3$.

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

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

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

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

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

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

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

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

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

Now for 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 $\widehat{A}_0|_{\mathcal{L}}$ or correct the input for all $x \in B\cap\mathcal{L}$. The former requires us to change at least a $(1-(d/|\mathbb{F}|)-1/3) \geq 1/2$-fraction of the points of $A_0|_{\mathcal{L}}$ while the latter requires us to change at least a $\delta'-\delta/4$-fraction of the input read (i.e., the input oracle W restricted to the line $\mathcal{L}$). This proves the claim. ∎

We now analyze the robustness of ALMSS PROXIMITY TEST.

CLAIM 8.6. *There exists a constant $\varepsilon_0 > 0$ such that for all $\delta \in (0,1)$, the following holds. Suppose $\Delta(\tilde{A}_0, \widehat{A}_0) \leq 4\varepsilon_0$ and the input oracle W is δ-far from $\widehat{A}_0|_I$* (i.e., $|B|/|L| \geq \delta$)*; then with probability at least $\Omega(\delta)$ (over the choice of index i and direction y), either at least a $1/2$-fraction of $A_0|_{\mathcal{L}}$ or $W[i]$ (i.e., the single symbol of the input oracle read by the verifier) needs to be changed to make* ALMSS PROXIMITY TEST *accept.*

This claim is the robust analogue of Lemma 7.7. 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 $\widehat{A}_0|_I$ must be δ-far from w. Thus, with probability greater than δ over the choice of

$i \in \{1, \ldots, k\}$ (and the corresponding point $x \in I$), we have $W[i] \neq \widehat{A}_0(x)$. 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 $\widehat{A}_0|_{\mathcal{L}}$. As in the proof of Lemma 7.7, for any fixed x, with probability at least $1 - 16\varepsilon_0$ (over the choice of the random direction y), $\tilde{A}_0|_{\mathcal{L}\setminus\{x\}}$ and $\widehat{A}_0|_{\mathcal{L}\setminus\{x\}}$ have distance at most $1/4$, and hence $\tilde{A}_0|_{\mathcal{L}}$ would have to be changed in at least $1 - ((md-1)/|\mathbb{F}|) - 1/4 \geq 1/2$ points to be a degree md polynomial other than $\widehat{A}_0|_{\mathcal{L}}$. Thus, with probability at least $\delta(1 - 16\varepsilon_0) = \Omega(\delta)$, 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 nonnegligible 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 guarantees only 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 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 $\Pi : \mathbb{F}^m \to \mathbb{F}^{l\cdot(2m+3)}$ is assumed to satisfy $\Pi(x) = (\tilde{A}_0(x), \ldots, \tilde{A}_{l-1}(x), P_{0,0}(x), \ldots, P_{l-1,m}(x))$ for all $x \in \mathbb{F}^m$. It can now be easily checked that over this proof oracle, the conjunction of all Low-Degree Tests is robust (over alphabet $\mathbb{F}^{l\cdot(2m+3)}$) with the same soundness and robustness parameters as a single Low-Degree Test (over alphabet $\mathbb{F}$). 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—because the l executions of these tests each have different query patterns (i.e., we cannot execute all of these tests by querying the same set of points of Π). Next, we provide an alternate bundling and massage our verifier slightly to work with this bundling.

To find a suitable bundling, we examine the query patterns of Zero-Propagation Test and Edge-Consistency Test more closely. The (i,j)th Zero-Propagation Test queries $P_{i,j-1}$ and $P_{i,j}$ on a random jth axis-parallel line. Also the ith Edge-Consistency Test queries $P_{i,0}$, $\tilde{A}_i$ for all points $x \in U_\eta$, and $\tilde{A}_{i+1}$ for all points $x \in \tilde{\Gamma}_{i,0}(U_\eta) \cup \tilde{\Gamma}_{i,1}(U_\eta)$ for U_η being a random subset of an arbitrary partition of F^m. The key observation is that the neighborhood functions $\tilde{\Gamma}_{i,0}$ and $\tilde{\Gamma}_{i,1}$ take any ith axis-parallel line to itself. Thus, if we choose the partition U_η to consist of ith axis-parallel lines, then the ith Edge-Consistency Test is also making queries entirely along axis-parallel lines. However, for the bundling to work, we need all the tests to

be making queries along the *same* axis-parallel line. We accomplish this by shifting our functions according to appropriate cyclic permutations of the coordinates so that the query patterns of the tests "line up" (at least into a constant number of groups).

To implement this idea, we first need 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 $\Sigma = \mathbb{F}^{l+2l\cdot(m+1)}$. Unlike before, the proof oracle for the robust PCPP verifier will consist of only one function $\Pi : \mathbb{F}^m \to \Sigma$. 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 to the queries of the original PCPP verifier. For convenience, we index the $l + 2l \cdot (m + 1)$ coordinates of $\Sigma = \mathbb{F}^{l+2l\cdot(m+1)}$ 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 $0, \ldots, l-1$, j ranges over $0, \ldots, m$, and $b \in \{0, 1\}$. Let $S : \mathbb{F}^m \to \mathbb{F}^m$ denote the (linear) transformation that performs one cyclic-shift to the right; that is, $S(x_0, \ldots, x_{m-1}) = (x_{m-1}, x_0, \ldots, x_{m-2})$. The bundling of the proof oracles $\tilde{A}_i$ and $P_{i,j}$ by the modified proof oracle Π is as follows:

(8.1)
$$\forall x \in \mathbb{F}^m, \begin{cases} \Pi(x)_i = \tilde{A}_i\left(S^{\lfloor \frac{i}{h} \rfloor}(x)\right) & i = 0, \ldots, l-1, \\ \Pi(x)_{(i,j,b)} = P_{i,j}^{(b)}\left(S^{j+\lfloor \frac{i}{h} \rfloor}(x)\right) & i = 0, \ldots, l-1;\ j = 0, \ldots, m;\ b \in \{0,1\}, \end{cases}$$

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 $\tilde{A}_i$ and $P_{i,j}$.

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

The first steps of PCPP Verifier (and ALMSS 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 $\langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \ldots, \tilde{T}_{l-1}\}\rangle$ of AS-CktSAT, sets $d = m \cdot |H|$, and generates a random string R of length $\log(|S_\lambda| \cdot |\mathbb{F}|^{m-1})$. 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 $\tilde{A}_0$. For this, we observe that $\Pi(x)_0$ is $\tilde{A}_0 \circ S^{\lfloor \frac{0}{h} \rfloor}(x) = \tilde{A}_0(x)$. The two different proximity tests (Robust Proximity Test and Robust ALMSS Proximity Test) can easily be describes as follows:

Robust Proximity Test$^{W;\,\Pi}(R)$.

Use random string R to determine a random canonical line $\mathcal{L}$ in $\mathbb{F}^m$ 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 $I \cap \mathcal{L}$.

Robust ALMSS Proximity Test$^{W;\,\Pi}$.

Choose a random position $i \xleftarrow{\text{R}} \{1, \ldots, k\}$ in the input and a direction $y \leftarrow \mathbb{F}^m$. Let $x \in I$ be the point corresponding to i in H^m, and let $\mathcal{L}$ be the 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 $\tilde{A}_i : \mathbb{F}^m \to \mathbb{F}$ to being degree k (for any $k \in \mathbb{Z}^+$) is exactly the same as that of $\tilde{A}_i \circ S^{\lfloor \frac{i}{h} \rfloor} : \mathbb{F}^m \to \mathbb{F}$ 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 $\mathbb{F}^m$ using the λ-biased set S_λ.

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

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

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 $i = 0, \ldots, l-1$, $j = 0, \ldots, m$ 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 Identity Test, we observe that $P_{i,m}^{(b)}$ vanishes on $\mathbb{F}^m$ iff $P_{i,m}^{(b)} \circ S^{m+\lfloor \frac{i}{h} \rfloor}$ vanishes on $\mathbb{F}^m$. Recall that we were allowed to use arbitrary partitions of the space $\mathbb{F}^m$. 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 $\mathbb{F}^m$ of the form $\mathcal{L} = (X, a_1, \ldots, a_{m-1})$. Query the oracle Π on all points along the line $\mathcal{L}$.

    For $i = 0, \ldots, l-1$ 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 are nonzero.

Edge-consistency test. For any $x \in \mathbb{F}^m$, we say that $P_{i,0}$ is *well formed* at x if (6.2) is satisfied for this x. Edge-Consistency Test verifies that $P_{i,0}$ is well formed for all $x \in U_\eta$ and $i = 0, \ldots, l-1$. 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 $x = (x_0, \ldots, x_{m-1}) \in \mathcal{L}$. 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 $x' = (x_0 + e_{(i \bmod h)}, x_1, \ldots, x_{m-1}) \in \mathcal{L}$ (recall that $\{e_0, \ldots, e_{f-1}\}$ are a basis for $\mathbb{F}$ over $\mathbb{F}_2$ and $\{e_0, \ldots, e_{h-1}\}$ span $H \subset \mathbb{F}$). For this purpose, we first recall the definition of $\tilde{\Gamma}_{i,1}$: $\tilde{\Gamma}_{i,1}(z_0, \ldots, z_{m-1}) =$

$(z_0, \ldots, z_{t-1}, z_t + e_u, z_{t+1}, \ldots, z_{m-1})$, 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:

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

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 $A_{i+1} \circ \tilde{\Gamma}_{i,1}(S^{\lfloor \frac{i}{h} \rfloor}(x)) = A_{i+1}(S^{\lfloor \frac{i}{h} \rfloor}(x')) = A_{i+1}(S^{\lfloor \frac{i+1}{h} \rfloor}(S^{-1}x'))$. 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 $\mathbb{F}^m$ of the form $\mathcal{L} = (X, a_2, \ldots, a_m)$. Query the oracle Π along all points in the lines $\mathcal{L}$ and $S^{-1}(\mathcal{L})$.

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

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 $i = 0, \ldots, l-1$ and $b \in \{0, 1\}$, ZERO-PROPAGATION TEST checks that $P_{i,0}^{(b)}$ vanishes on H^m by verifying that (6.3) is satisfied for all $j = 1, \ldots, m-1$ (we also need to check that $P_{i,m}^{(b)} \equiv 0$; however, this is taken care of by 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 $P_{i,0}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor}; i = 0, \ldots, l-1, b \in \{0, 1\}$, instead of the polynomials $P_{i,0}^{(b)}; i = 0, \ldots, l-1, b \in \{0, 1\}$. In other words, we need to verify the following equation instead (6.3):

(8.2)
$$\begin{aligned}
&P_{i,j}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor}\left(\underbrace{x_1, \ldots, x_{j-1}}, x_j, \underbrace{x_{j+1}, \ldots, x_m}\right) \\
&= \sum_{k=0}^{|H|-1} P_{i,j-1}^{(b)} \circ S^{\lfloor \frac{i}{h} \rfloor}\left(\underbrace{x_1, \ldots, x_{j-1}}, h_k, \underbrace{x_{j+1}, \ldots, x_m}\right) x_j^k \qquad \forall (x_1, \ldots, x_m) \in \mathbb{F}^m.
\end{aligned}$$

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

(8.3)
$$P_{i,j}^{(b)}\left(S^{\lfloor \frac{i}{h}\rfloor+j-1}(x_1,x_2,\ldots,x_m)\right) = \sum_{k=0}^{|H|-1} P_{i,j-1}^{(b)}\left(S^{\lfloor \frac{i}{h}\rfloor+j-1}(h_k,x_2,\ldots,x_m)\right)x_1^k \quad \forall (x_1,\ldots,x_m)\in\mathbb{F}^m.$$

This helps us to rewrite Zero-Propagation Test with bundling as follows:

Robust Zero-Propagation Test$^{\Pi(R)}$.

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

For $i=0,\ldots,l-1$, $j=1,\ldots,m$, 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 (8.3).

The integrated robust verifiers. Having presented the robust version of each of the tests, the integrated robust verifiers are as follows: Robust PCPP Verifier is the robust analogue of PCPP Verifier, while ALMSS Robust PCPP Verifier is that of ALMSS PCPP Verifier. Following are full descriptions of these verifiers as well as their analyses.

8.2.1. Robust PCPP Verifier. Using the robust tests presented above, we present a robust analogue of the PCPP of section 7.1.

Robust PCPP Verifier$^{W;\,\Pi}_{m,\lambda,\delta}(C)$.

1. Using Proposition 6.11, reduce the instance C of CktSAT, using parameter m, to an instance $\langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0,\ldots,\tilde{T}_{l-1}\}\rangle$ of AS-CktSAT, and set $d=m\cdot|H|$.
   We let $S_\lambda\subset\mathbb{F}^m$ be a λ-biased set of size at most $(\frac{\log|\mathbb{F}|^m}{\lambda})^2$ [AGHP92].
2. Choose a random string R of length $\log(|S_\lambda|\cdot|\mathbb{F}|^{m-1})$. (Note: We reuse R in all tests, but only 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$^{\Pi}(R)$.
7. Run Robust Proximity Test$^{W;\Pi}(R)$.

Reject if any of the above tests reject.

The randomness of Robust PCPP Verifier is exactly the same as before, whereas the query complexity and decision complexity increase by a constant factor.[32]

[32] Though 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 as the number of symbols (i.e., elements of $|\Sigma|$) to maintain consistency throughout calculation of the query complexity into the proof and input oracles.

PROPOSITION 8.7. *The randomness, query, and decision complexities of* ROBUST PCPP VERIFIER *are* $r = \left(1 - \frac{1}{m}\right)\log n + O(m \log m) + O(\log\log n) + O\left(\log\left(\frac{1}{\delta}\right)\right)$, $q = O(m^2 n^{1/m} \log^2 n)$ *and* $d = \widetilde{O}(q)$, *respectively.*

It is straightforward to check perfect completeness of this verifier.

Robustness analysis of the integrated verifier. For future use, it is beneficial (but, alas, more cumbersome) to state the robustness of the integrated verifier in a way that decouples the robustness with respect to the input oracle from the robustness with respect to the proof oracle. Let $W : [k] \to \{0,1\}$ 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 ROBUST PCPP VERIFIER accept on coin tosses R. The following lemma states the (expected) robustness property of our verifier.

LEMMA 8.8. *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* $m \le \log n / \log\log n$, $n^{1/m} \ge m^{cm}/\delta'^3$, $\lambda \le \min\{1/c\log n, \delta'^3/m^{cm}\}$, $\delta > \delta'$. *If* W *is* δ*-far from satisfying the circuit, then for any proof oracle* $\Pi : \mathbb{F}^m \to \Sigma$, *either* $\mathrm{E}_R[\Delta_{\text{pf}}^{W,\Pi}(R)] \ge \rho$ *or* $\mathrm{E}_R[\Delta_{\text{inp}}^{W,\Pi}(R)] \ge \delta - \delta'/2$.

That is, the expected robustness with respect to the input is $\delta - \delta'/2$ (which should be compared against the proximity parameter δ), whereas the expected robustness with respect to the proof is a universal constant. Note that combining the two bounds into a single expected robustness parameter depends on the relative number of queries made to the input and proof oracles. To obtain Theorem 3.1, we will later modify ROBUST PCPP VERIFIER such that the relative number of queries is optimized to yield the best result.

Proof. Unbundle the proof oracle Π to obtain the functions $\tilde{A}_i$ and $P_{i,j}$ using (8.2). Consider the action of PCPP VERIFIER (i.e., the nonrobust 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 8.1–8.5 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 PCPP VERIFIER rejects its input robustly (see Claims 8.1–8.5).

1. There exists a $i \in \{0, \dots, l-1\}$ such that $\tilde{A}_i$ is 8ε-far from every degree md polynomial or there exists $i \in \{0, \dots, l-1\}$, $j \in \{0, \dots, m\}$, 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 LOW-DEGREE TEST with degree parameter d (resp., κd) is at least 2ε (Claim 8.1).
2. There exists $i \in \{0, \dots, l-1\}$ and $b \in \{0,1\}$, such that $\Delta(P_{i,m}^{(b)}, \widehat{P}_{i,m}^{(b)}) \le 8\varepsilon$ and $\widehat{P}_{i,m} \not\equiv 0$. In this case, $P_{i,m}$ has expected distance at least $1 - 9\varepsilon$ from satisfying IDENTITY TEST (Claim 8.2).
3. There exists $i \in \{0, \dots, l-1\}$ such that

$$\Delta(P_{i,0}^{(0)}, \widehat{P}_{i,0}^{(0)}) \le 8\varepsilon, \Delta(P_{i,0}^{(1)}, \widehat{P}_{i,0}^{(1)}) \le 8\varepsilon, \Delta(\tilde{A}_i, \widehat{A}_i) \le 8\varepsilon, \Delta(\tilde{A}_{i+1}, \widehat{A}_{i+1}) \le 8\varepsilon,$$
$$\text{and } \widehat{P}_{i,0}(x) \not\equiv \psi'(\tilde{T}_i(x), \widehat{A}_i(x), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))).$$

In this case, $\left\{P_{i,0}(\cdot), A_i(\cdot), A_{i+1}(\tilde{\Gamma}_{i,0}(\cdot))A_{i+1}(\tilde{\Gamma}_{i,1}(\cdot))\right\}$ has expected distance at least $(1 - 41\varepsilon)/5$ from satisfying EDGE-CONSISTENCY TEST (Claim 8.3).

4. There exists $i \in \{0, \dots, l-1\}$, $j \in \{1, \dots, m\}$ and $b \in \{0, 1\}$ such that

$$\Delta(P_{i,j}, \widehat{P}_{i,j}) \leq 8\varepsilon, \Delta(P_{i,j-1}, \widehat{P}_{i,j-1}) \leq 8\varepsilon,$$
$$\text{and } \widehat{P}_{i,j}(\dots, x_j, \dots) \not\equiv \sum_{k=0}^{|H|-1} \widehat{P}_{i,j-1}(\dots, h_k, \dots) x_j^k.$$

   In this case, $(P_{i,j}^{(b)}, P_{i,j-1}^{(b)})$ has expected distance at least $(1 - 19\varepsilon)/2$ from satisfying Zero-Propagation Test (Claim 8.4).
5. $\Delta(\tilde{A}_0, \widehat{A}_0) \leq 8\varepsilon$, but W and $\widehat{A}_0|_I$ 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 an 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 Proximity Test accept (Claim 8.5).
   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 Proximity Test.

For instance, lets us assume $\tilde{A}_0$ is 8ε-far from being low degree so 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 Low-Degree Test is at least 2ε. Recall from (8.2) that $\tilde{A}_0(x)$ is one of the coordinates in the bundled $\Pi(x)$. Hence, if $\tilde{A}_0|_{\mathcal{L}}$ is ρ-far from satisfying Low-Degree Test, so is $\Pi_{\mathcal{L}}$ from satisfying Robust Low-Degree Test. Thus, Π has expected distance at least 2ε from satisfying Robust Low-Degree Test. Now, the oracle positions read by Robust Low-Degree Test constitute a constant fraction of the oracle positions read by Robust PCPP Verifier, so Π has expected distance $\Omega(\varepsilon)$ from satisfying 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 $\mathrm{E}_R[\Delta_{\mathrm{pf}}^{W,\Pi}(R)] \geq \Omega(\varepsilon)$ or $\mathrm{E}_R[\Delta_{\mathrm{inp}}^{W,\Pi}(R)] \geq \delta - \delta'/2$. The lemma then follows by setting $\rho = \Omega(\varepsilon)$. ∎

8.2.2. ALMSS Robust PCPP Verifier. We now describe ALMSS Robust PCPP Verifier(which is a robust analogue of the PCPP of section 7.2) and analyze its complexity. ALMSS Robust PCPP Verifier verifier is identical to Robust PCPP Verifier except that Robust Proximity Test is replaced by Robust ALMSS Proximity Test.

ALMSS Robust PCPP Verifier$_\delta^{W;\,\Pi}(C)$.

1. Set parameters $m = \log n / \log\log n$ and $\lambda = 1/c \log n$.
   Using Proposition 6.11, reduce the instance C of CktSAT, using parameter m, to an instance $\langle 1^n, 1^m, \mathbb{F}, H, \{\tilde{T}_0, \dots, \tilde{T}_{l-1}\}\rangle$ of AS-CktSAT, and set $d = m \cdot |H|$.
   We let $S_\lambda \subset \mathbb{F}^m$ be a λ-biased set of size at most $(\frac{\log |\mathbb{F}|^m}{\lambda})^2$ [AGHP92].
2. Choose a random string R of length $\log(|S_\lambda| \cdot |\mathbb{F}|^{m-1})$.
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$^\Pi(R)$.
7. Run Robust ALMSS Proximity Test$^{W;\Pi}$.

Reject if any of the above tests reject.

The randomness of the ALMSS Robust PCPP Verifier is exactly the same as before, whereas the query complexity and decision complexity increase by a constant

factor. Furthermore, it can easily be verified that ALMSS Robust PCPP Verifier has perfect completeness.

Proposition 8.9. *The randomness, and decision complexities of* ALMSS Robust PCPP Verifier *are* $O(\log n)$ *and* $\operatorname{poly}\log n$*, respectively.*

Robustness analysis of the integrated verifier. As in the case of Robust PCPP Verifier, it is beneficial to state the robustness of ALMSS Robust PCPP Verifier by decoupling the robustness with respect to the input oracle from the robustness with respect to the proof oracle. Here, however, we refer to the robustness and soundness parameters (rather than to expected robustness).

Lemma 8.10. *If W is δ-far from satisfying the circuit, then for any proof oracle $\Pi : \mathbb{F}^m \to \Sigma$, with probability at least $\Omega(\delta)$, either a constant fraction of the portion of the proof oracle Π read by the verifier or the single symbol of the input oracle W read by the verifier (i.e., $W[i]$) needs to be changed in order to make* ALMSS Robust PCPP Verifier *accept.*

Proof. This proof proceeds in the same way as the proof of Lemma 8.8. For the sake of completeness, we present the entire proof.

Let ε be a sufficiently small constant such that Claims 8.1–8.4 hold and $\varepsilon \le \varepsilon_0/8$, where ε_0 is the constant that appears in Claim 8.6. 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 ALMSS PCPP Verifier rejects its input robustly (see Claims 8.1–8.6).

1. There exists $i \in \{0, \dots, l-1\}$ such that $\tilde{A}_i$ is 8ε-far from every degree md polynomial or there exists $i \in \{0, \dots, l-1\}$, $j \in \{0, \dots, m\}$, 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 $\tilde{A}_i$ (resp., $P_{i,j}^{(b)}$) from satisfying Low-Degree Test with degree parameter d (resp., κd) is at least 2ε (Claim 8.1).
   Translating to the bundled alphabet, we have that the expected distance of Π from satisfying Robust Low-Degree Test is at least 2ε. Since the number of oracle positions read by ALMSS Robust PCPP Verifier is at least a constant fraction of the number of oracle positions read by ALMSS Robust PCPP Verifier, the expected distance of Π from satisfying ALMSS Robust PCPP Verifier in this case is at least $\Omega(\varepsilon)$. Hence, with probability at least $\Omega(\varepsilon)$ (= constant), at least $\Omega(\varepsilon)$ (= constant)-fraction of the proof oracle Π needs to be modified to make ALMSS Robust PCPP Verifier accept.
2. There exists $i \in \{0, \dots, l-1\}$ and $b \in \{0,1\}$ such that $\Delta(P_{i,m}^{(b)}, \widehat{P}_{i,m}^{(b)}) \le 8\varepsilon$ and $\widehat{P}_{i,m} \not\equiv 0$. In this case, $P_{i,m}$ has expected distance at least $1 - 9\varepsilon$ from satisfying Identity Test (Claim 8.2).
   Arguing as in the earlier case, we have that with probability at least $\Omega(1-9\varepsilon)$, at least $\Omega(1-9\varepsilon)$-fraction of the proof oracle Π needs to be modified in this case to make ALMSS Robust PCPP Verifier accept.
3. There exists $i \in \{0, \dots, l-1\}$ such that

$$\Delta(P_{i,0}^{(0)}, \widehat{P}_{i,0}^{(0)}) \le 8\varepsilon, \Delta(P_{i,0}^{(1)}, \widehat{P}_{i,0}^{(1)}) \le 8\varepsilon, \Delta(\tilde{A}_i, \widehat{A}_i) \le 8\varepsilon, \Delta(\tilde{A}_{i+1}, \widehat{A}_{i+1}) \le 8\varepsilon,$$
$$\text{and } \widehat{P}_{i,0}(x) \not\equiv \psi'(\tilde{T}_i(x), \widehat{A}_i(x), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,0}(x)), \widehat{A}_{i+1}(\tilde{\Gamma}_{i,1}(x))).$$

   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 Edge-Consistency Test (Claim 8.3).

Again, we have that with probability at least $\Omega(1-41\varepsilon)$, at least $\Omega(1-41\varepsilon)$-fraction of the proof oracle Π needs to be modified in this case to make ALMSS Robust PCPP Verifier accept.

4. There exists $i \in \{0,\dots,l-1\}$, $j \in \{1,\dots,m\}$, and $b \in \{0,1\}$ such that

$$\Delta(P_{i,j}, \widehat{P}_{i,j}) \le 8\varepsilon, \Delta(P_{i,j-1}, \widehat{P}_{i,j-1}) \le 8\varepsilon,$$
$$\text{and } \widehat{P}_{i,j}(\dots, x_j, \dots) \not\equiv \sum_{k=0}^{|H|-1} \widehat{P}_{i,j-1}(\dots, h_k, \dots) x_j^k.$$

   In this case, $(P_{i,j}^{(b)}, P_{i,j-1}^{(b)})$ has expected distance at least $(1-19\varepsilon)/2$ from satisfying Zero-Propagation Test (Claim 8.4).
   We have that with probability at least $\Omega(1-19\varepsilon)$, at least $\Omega(1-19\varepsilon)$-fraction of the proof oracle Π needs to be modified in this case to make ALMSS Robust PCPP Verifier accept.
5. $\Delta(\tilde{A}_0, \widehat{A}_0) \le 8\varepsilon \le \varepsilon_0$, but W and $\widehat{A}_0|_I$ disagree on at least δ-fraction of the points. In this case, with probability at least $\Omega(\delta)$ (over the choice of index i and direction y), 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 ALMSS Proximity Test accept.
   This implies that with probability at least δ, either a constant fraction of the proof oracle Π or $W[i]$ (i.e., the entire portion of the input oracle read by the verifier) needs to be modified to make the verifier accept.

Since we do not know which of the five cases occur, we can guarantee only the weakest of the five claims. Hence, with probability at least $\Omega(\delta)$, either a constant fraction of the portion of the proof oracle Π read by the verifier or $W[i]$ (i.e., the entire portion of the input oracle W read by the verifier) needs to be changed in order to make ALMSS Robust PCPP Verifier accept. ∎

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 concatenation." 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 because 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. Recall that applying the aforementioned transformation maintains the robustness of the proof oracle *up to a constant factor*, whereas the robustness of the input oracle remains unchanged (like the input oracle itself). Actually, in order to avoid decoding (by the modified decision circuit), we maintain the original proof oracle along with its encoded form. Thus, the complexity of this circuit will depend on the minimum between the complexity of encoding and decoding (rather than on the complexity of decoding). Details follow.

Let $\mathrm{ECC} : \{0,1\}^{\log|\Sigma|} \to \{0,1\}^b$ 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 $O(\log|\Sigma|)$ [Spi96]. 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 assumed to be $\mathrm{ECC}(\Pi(x))$.

The actual transformation. We describe the transformation to the binary alphabet in the case of Robust PCPP Verifier. ALMSS Robust PCPP Verifier

can be transformed similarly. Our new verifier V, on oracle access to the input W and proof $\Pi \circ \Upsilon$, will simulate ROBUST PCPP VERIFIER. The queries to the input oracle are performed just as before. However, for each query $x \in \mathbb{F}^m$ 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,[33] and the decision complexity will increase by at most the encoding time (which can even be linear). The randomness is exactly the same. The action of the new verifier V is as follows: Suppose ROBUST PCPP VERIFIER issues queries $x_1, \ldots, x_{q_1}$ to the proof oracle Π, and queries $i_1, \ldots, i_{q_2}$ to the input oracle; then V issues queries $x_1, \ldots, x_{q_1}$ to the proof oracle Π, a similar set of queries $x_1, \ldots, x_{q_1}$ to the proof oracle Υ, and $i_1, \ldots, i_{q_2}$ to the input oracle. V accepts $(\Pi(x_1), \ldots, \Pi(x_{q_1}), \Upsilon(x_1), \ldots, \Upsilon(x_{q_1}), W(i_1), \ldots, W(i_{q_2}))$ iff ROBUST PCPP VERIFIER accepts $(\Pi(x_1), \ldots, \Pi(x_{q_1}), W(i_1), \ldots, W(i_{q_2}))$ and $\Upsilon(x_i) = \mathrm{ECC}(\Pi(x_i))$ for all $i = 1, \ldots, q_1$. It is straightforward to check that V has perfect completeness if ROBUST PCPP VERIFIER has perfect completeness. For the robust soundness, we define $\Delta_{\mathrm{inp}}^{W,\Pi\circ\Upsilon}(R)$ and $\Delta_{\mathrm{pf}}^{W,\Pi\circ\Upsilon}(R)$ with respect to V analogously as in the statement immediately preceding Lemma 8.8, but referring to distance over $\{0,1\}$ (rather than Σ) for the proof oracle. The proof of the following claim regarding the robust soundness of V mimics the proof of Lemma 2.13.

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

A similar transformation for ALMSS ROBUST PCPP VERIFIER yields a verifier, the robustness of which is stated in Lemma 8.12 following. It is to be noted that the robustness of the proof oracle (i.e., ρ' in Lemma 8.11 and $\Omega(1)$ in Lemma 8.12) is a constant factor smaller than the corresponding parameter in the nonbinary verifier (i.e., the constant ρ in Lemma 8.8 and a different $\Omega(1)$ in Lemma 8.10). (Indeed, this constant factor appears also in Lemma 2.13.)

LEMMA 8.12. *If W is δ-far from satisfying the circuit, then for any proof oracles $\Pi : \mathbb{F}^m \to \{0,1\}^{\log|\Sigma|}, \Upsilon : \mathbb{F}^m \to \{0,1\}^b$, with probability at least $\Omega(\delta)$, either a constant (i.e., $\Omega(1)$) fraction of the portion of the proof oracle $\Pi\circ\Upsilon$ read by the verifier or $W[i]$ (i.e., the entire portion of the input oracle W read by the verifier) needs to be changed in order to make the transformed* ALMSS ROBUST PCPP VERIFIER *accept.*

We finally turn to deriving Theorem 3.1 (and Theorem 3.2).

Proof of Theorem 3.1. Theorem 3.1 is proved using 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 ROBUST PCPP VERIFIER (i.e., before the transformation to the binary alphabet) are as mentioned in Proposition 8.7. As mentioned earlier in this section, the transformation from the alphabet Σ to the binary alphabet maintains the randomness complexity, while the query (and decision) complexity increases by at most a constant factor.

The manner in which the robustness of the verifier is expressed in Lemma 8.11

[33] Recall that the query complexity of the old verifier was measured in terms of "bits of information" rather than in terms of queries. That is, each query, answered by an element of Σ, contributes $\log_2|\Sigma|$ to the query complexity.

differs from that in Theorem 3.1 in two aspects. First, Lemma 8.11 expresses the robustness for the proof and input oracles separately, while Theorem 3.1 expresses them together. Second, Lemma 8.11 expresses robustness in terms of expected robustness, while Theorem 3.1 does it in terms of standard robustness. We obtain robustness as claimed in Theorem 3.1 in two steps, first by combining the proof and input oracles and then by moving from expected robustness to standard robustness.

First, we combine the robustness of the proof and input oracles, which were expressed separately in Lemma 8.11. This is done by giving adequate weights to the two oracle portions in the decision circuits (i.e., repeating queries; see Proposition 7.3). Let n, m, δ, and γ be as specified in Theorem 3.1. We give weight $(1-\gamma/3)$ to the input oracle and $\gamma/3$ to the proof oracle. 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 $1-\gamma/3$. These repeated queries may increase the query (and decision) complexity by a factor of at most $O(1/\gamma)$. 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 $n \geq 8000|\mathbb{F}|^{m-1}/\delta^3$. Hence, Lemma 8.11 can be applied. Setting $\delta' = 2\delta\gamma/3$ in Lemma 8.11, we have that either $\mathrm{E}_R[\Delta_{\mathrm{pf}}^{W,\Pi\circ\Upsilon}(R)] \geq \rho'$ or $\mathrm{E}_R[\Delta_{\mathrm{inp}}^{W,\Pi\circ\Upsilon}(R)] \geq \delta - \delta'/2 = \delta(1-\gamma/3)$. 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 $\gamma/3$ to the input-part). Hence, the expected robustness with respect to the said weighting is

$$\frac{\gamma}{3} \cdot \mathrm{E}_R[\Delta_{\mathrm{pf}}^{W,\Pi\circ\Upsilon}(R)] + \left(1-\frac{\gamma}{3}\right) \cdot \mathrm{E}_R[\Delta_{\mathrm{inp}}^{W,\Pi\circ\Upsilon}(R)] \;\geq\; \min\left\{\frac{\gamma}{3}\cdot\rho', \left(1-\frac{\gamma}{3}\right)^2 \cdot \delta\right\}.$$

This quantity is lower-bounded by $\rho \triangleq (1-\gamma/3)^2\delta$ since $\delta \leq \gamma/c$ for a suitably large c (and $\rho' > 0$ is a constant). We have thus obtained a robust PCPP for CktVal with randomness and decision complexities as claimed in Proposition 8.7, perfect completeness, and $\rho = (1-\gamma/3)^2\delta$ expected robustness for proximity parameter δ.

We now move from expected robustness to standard robustness, by using Lemma 2.11. Applying Lemma 2.11 with a slackness parameter of $\gamma' \triangleq \gamma\rho/3$ and $s = \gamma/3$ yields a robust-soundness error of $\gamma/3 \leq \gamma$ with robustness parameter of $\rho - \gamma' = (1-\gamma/3)^3 \cdot \delta \geq (1-\gamma)\delta$ for proximity parameter δ. Using $\gamma \leq 1/2$, note that the randomness increases by an additive term of $O(\log(1/\gamma')) + O(\log(1/\gamma)) = O(\log(1/\delta))$, and the decision complexity increases by a multiplicative factor of $O\left(1/(\gamma\cdot(\gamma\rho)^2)\right) = \mathrm{poly}(1/\delta)$. Hence, the randomness, query, and decision complexities of the verifier are as claimed in Theorem 3.1 ∎

Proof of Theorem 3.2. For this purpose we use ALMSS Robust PCPP Verifier described in this section. This verifier is then transformed to one over the binary alphabet as indicated earlier in this section. We combine the robustness of the proof and input oracles by giving equal weights to both oracles. This weighting may increase the query (and decision) complexity by at most a factor of 2 and has no affect on any other parameter. Proposition 8.9 and Lemma 8.12 then imply Theorem 3.2. ∎

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 PCPP

is linear. Throughout this section, "linear" means $\mathbb{F}_2$-linear (yet, we will sometimes refer to $\mathbb{F}$-linearity, for an extension field $\mathbb{F}$ of $\mathbb{F}_2$). The main motivation of 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 PCPP is linear in the aforementioned sense.

We call a circuit *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, i.e., one for each linear constraint. See the following definition.

DEFINITION 8.13. *A multi-output circuit is* linear *if all its* internal *gates are parity gates and an input is accepted by it iff 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). A similar result is true for the PCPP verifier of Theorem* 3.2.

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).[34] 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 zeroth layer of the well-structured circuit (i.e., the last two bits of the mapping $A_0 : \{0,1\}^N \to \{0,1\}^4$ 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 a finite field $\mathbb{F}$ that is an extension field of $\mathbb{F}_2$ of degree $\mho$. 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 $e_0, \ldots, e_{\mho-1}$ is a basis for $\mathbb{F}$ viewed as a vector space over $\mathbb{F}_2$. We view A_i as a function $A_i : H^m \to \mathbb{F}$ and construct a low-degree extension $\tilde{A}_i : \mathbb{F}^m \to \mathbb{F}$ 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 $\mathbb{F}^m$. Notice that interpolation is $\mathbb{F}$-linear, and hence also $\mathbb{F}_2$-linear. We conclude that the values of $\tilde{A}_i$ on all points in $\mathbb{F}^m$ are 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 $\mathbb{F}^m$. The linearity we claim is not linearity in

[34] Recall that an input is accepted by the linear circuit iff all output gates evaluate to zero.

x (the free variables of the polynomial). Rather, we claim that the table of values $\{P(a) : a \in \mathbb{F}^m\}$ is linear in the initial assignment 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 that the table $\{\tilde{A}_i(a) : a \in \mathbb{F}^m\}$ is obtained by a linear transformation applied to the table $\{A_i(a') : a' \in H^m\}$ (but we certainly do not claim that $\tilde{A}_i(a)$ 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 (6.2) and show that 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 (6.1) 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 $\mathbb{F}_2^{\mho}$, 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 an $\mathbb{F}_2$-linear subspace of $\mathbb{F} = \mathbb{F}_2^{\mho}$ and $\psi_L(t) = \prod_{\alpha \in L}(t-\alpha)$. Then the mapping $\psi_L : \mathbb{F} \to \mathbb{F}$ 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 $L \subseteq \mathbb{F}_2^{\mho}$. Indeed, for $\dim(L) = 0$ (i.e., $L = \{0^{\mho}\}$), 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 a 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

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

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 $\mathbb{F}$ 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 (6.2) we did not go into the specific construction of the polynomial ψ' because only its functionality mattered, and we showed that there exists a constant-degree polynomial that does the job. But for our current purposes (of showing linearity) we need to present a specific polynomial ψ' that is linear (as an operator over $\mathbb{F}_2^{\mho}$) and has the desired

properties needed by the verification process. To do this, we recall that $\mathcal{C}$ is the set of allowable gates in the well-structured circuit, define $\delta_{s_0}(z)$ to be the (minimal degree) univariate polynomial of degree $|\mathcal{C}|$ that is 1 if $z = s_0$ and 0 if $z \in \mathcal{C} \setminus \{s_0\}$, and write ψ' as

$$\psi'(s, a, a_0, a_1) = \sum_{s_0 \in \mathcal{C}} \delta_{s_0}(s) \cdot \psi'_{s_0}(a, a_0, a_1). \tag{8.4}$$

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* (8.4) *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 assumed to represent whether or not the four least significant bits of the three inputs (denoted a', a'_0, and a'_1) satisfy some specified condition. By inspecting Definition 6.3, it can be verified that (in our case) this condition is linear. That is, $\psi'_{s_0}(a, a_0, a_1) = 0$ iff the triplet (a', a'_0, a'_1), viewed as a 12-bit vector over $\mathbb{F}_2$, belongs to a specific linear space $L_{s_0} \subseteq \mathbb{F}_2^{12}$.

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 $P^{(0)}_{i,0}$. Thus, we seek a polynomial (over $\mathbb{F}^3$) such that if each of its three inputs belongs to $Span(e_0, \ldots, e_3)$, 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 $\alpha \in \mathbb{F}$ 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 $\mathbb{F}$. Using this element we can write $\psi'_{s_0} : \mathbb{F}^3 \to \mathbb{F}$ 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 $Span(e_0, \ldots, e_3)$, 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 $\mathbb{F}$ (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 (8.4) 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 (6.2) is obtained by linear transformations applied to the input to the circuit.

Phase IV: *The sum-check polynomials.* In this phase (described by (6.3)) we apply a sequence of interpolations to previously constructed polynomials $P^{(b)}_{i,j}$. Each such interpolation is an $\mathbb{F}$-linear transformation and hence also an $\mathbb{F}_2$-linear one. Thus, the sequence of polynomials $P^{(b)}_{i,j}$ 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 $|\mathbb{F}|^{l+2l(m+1)}$. 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 $|\mathbb{F}|^{l+2l(m+1)}$) 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 3.1 (the verifier of Theorem 3.2 is dealt with in a

similar fashion). Notice that this transformation from $l+2l(m+1)$ polynomials (each evaluated in $\mathbb{F}$) 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 Theorems 3.1 and 3.2 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 $\mathbb{F}$-linear conditions, and applying the general result of Strassen [Str73], 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.

Part III. Appendices.

Appendix A. Hadamard code–based PCPP. In this section we note that the Hadamard code–based inner verifier from Arora et al. [ALM+98] can be converted into a PCPP. Recall that the inner verifier of [ALM+98] accesses $O(1)$ input oracles, where the ith oracle is assumed 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 PCPP for* CKTVAL *(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 $\delta \le \delta_0$. 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 PCPP. 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 PCPP 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 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 $W : [m] \to \{0,1\}$ (representing a string $w \in \{0,1\}^m$) and a proof oracle $\Pi = (A, B)$, with $A : \{0,1\}^n \to \{0,1\}$ and $B : \{0,1\}^{n\times n} \to \{0,1\}$.

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)). Without loss of generality, 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 $\mathbb{F}_2$), 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 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 $P_1(z), \ldots, P_n(z)$ as follows. Viewing the variables $\{z_i\}$ as representing the values at the individual gates of the circuit C (with $z_1, \ldots, z_m$ being the input gates), the polynomial $P_i(z)$ is the quadratic polynomial (over $\mathbb{F}_2$) expressing the constraint imposed by the ith 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. [BLR93]) 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 $\alpha_1, \ldots, \alpha_n \in \{0,1\}$ 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 jth coordinate and 0 everywhere else. Verify that $W[j] = A[x + e_j] - A[x]$.

The verifier *accepts* if all the tests above accept; otherwise 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 [ALM+98] that there exists a $\delta_0 > 0$ such that for every $\delta \leq \delta_0$, if *Codeword tests* and *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 *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. ∎

Appendix B. Randomness-efficient low-degree tests and the sampling lemma. Following [BSVW03], our construction makes heavy use of small-biased spaces [NN93] to save on randomness when choosing random lines. For a field $\mathbb{F}$ and parameters $m \in \mathbb{Z}^+$ and $\lambda > 0$, we require a set $S \subseteq \mathbb{F}^m$ that is λ-*biased* (with respect to the additive group of $\mathbb{F}^m$). Rather than define small-biased spaces here, we simply state the properties we need. (See, e.g., [BSVW03] for definitions and background on small-biased spaces.)

LEMMA B.1. *For every $\mathbb{F}$ of characteristic 2, $m \in \mathbb{Z}^+$, and $\lambda > 0$, there is an explicit construction of a λ-biased set $S \subseteq \mathbb{F}^m$ of size at most* $(\log |\mathbb{F}^m|)/\lambda^2$ [AGHP92].

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. See the following lemma.

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

Randomness-efficient line samplers. In [BSVW03], 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 $B \subseteq \mathbb{F}^m$ using lines in $\mathbb{F}^m$ as the sample sets, one does not need to pick a random line in $\mathbb{F}^m$ which costs $2 \log |\mathbb{F}^m|$ random bits. A pseudorandom line whose slope comes from a λ-biased set will do nearly as well, and the randomness is only $(1 + o(1)) \cdot \log |\mathbb{F}^m|$. In what follows, $l_{x,y}$ is the line passing through point x in direction y, formally; $l_{x,y} = \{x + ty : t \in \mathbb{F}\}$

LEMMA B.3 ([BSVW03, Sampling Lemma 4.3]). *Suppose $S \subseteq \mathbb{F}^m$ is λ-biased. Then, for any $B \subseteq \mathbb{F}^m$ of density $\mu = |B|/|\mathbb{F}^m|$, and any $\zeta > 0$,*

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

Randomness-efficient low-degree tests. Ben-Sasson et al. [BSVW03] use the randomness-efficient sampling lemma, Lemma B.3, to obtain randomness-efficient low-degree tests by performing a "line versus point" test only for pseudorandom lines with a direction y coming from a small λ-biased set. That is, for a set $S \subseteq \mathbb{F}^m$, we consider

lines of the form $l_{x,y}(t) = x + ty$ for $x \in \mathbb{F}^m$ and $y \in S$, and let $\mathbb{L}$ be the set of all such lines, where each line is parameterized in a canonical way.

Then for functions $f : \mathbb{F}^m \to \mathbb{F}$ and $g : \mathbb{L} \to P_d$, where P_d is the set of univariate polynomials of degree at most d over $\mathbb{F}$, we let $\mathrm{LDT}^{f,g}_{S,d}$ be the test that uniformly selects $l \xleftarrow{\mathrm{R}} \mathbb{L} \triangleq \{l_{x,y} : x \in \mathbb{F}^m, y \in S\}$ and $t \in \mathbb{F}$ 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 the main theorem of [BSVW03] and will use it in our constructions.

THEOREM B.4 ([BSVW03, Theorem 4.1]). *There exists a universal constant $\alpha > 0$ such that the following holds. Let $d \le |\mathbb{F}|/3, m \le \alpha|\mathbb{F}|/\log|\mathbb{F}|, S \subseteq \mathbb{F}^m$ be a λ-biased set for $\lambda \le \alpha/(m \log|\mathbb{F}|)$, and $\delta \le \alpha$. Then, for every $f : \mathbb{F}^m \to \mathbb{F}$ and $g : \mathbb{L} \to P_d$ such that f is at least 4δ-far from any polynomial of degree at most md, we have the following:*

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

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. We thank Sergey Yekhanin and Jaikumar Radhakrishnan for bringing to our attention the "algebraic AND" mentioned in footnote 28. We thank the two anonymous referees for their careful reading and many comments which improved the presentation.

REFERENCES

[AGHP92] N. ALON, O. GOLDREICH, J. HÅSTAD, AND R. PERALTA, *Simple constructions of almost k-wise independent random variables*, Random Structures Algorithms, 3 (1992), pp. 289–304.

[ALM+98] S. ARORA, C. LUND, R. MOTWANI, M. SUDAN, AND M. SZEGEDY, *Proof verification and the hardness of approximation problems*, J. ACM, 45 (1998), pp. 501–555. (Preliminary version in Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science, 1992, pp. 14–23.)

[AS98] S. ARORA AND S. SAFRA, *Probabilistic checking of proofs: A new characterization of NP*, J. ACM, 45 (1998), pp. 70–122. (Preliminary version in Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science, 1992, pp. 2–13.)

[BFLS91] L. BABAI, L. FORTNOW, L. A. LEVIN, AND M. SZEGEDY, *Checking computations in polylogarithmic time*, in Proceedings of the 23rd ACM Symposium on Theory of Computing, New Orleans, LA, 1991, pp. 21–31.

[Bar01] B. BARAK, *How to go beyond the black-box simulation barrier*, in Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, Las Vegas, NV, 2001, pp. 106–115.

[BIKR02] A. BEIMEL, Y. ISHAI, E. KUSHILEVITZ, AND J. F. RAYMOND, *Breaking the $O(n^{1/(2k-1)})$ barrier for information-theoretic private information retrieval*, in Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science, Vancouver, BC, Canada, 2002, pp. 261–270.

[BGS95] M. BELLARE, O. GOLDREICH, AND M. SUDAN, *Free bits, PCPs, and nonapproximability—towards tight results*, SIAM J. Comput., 27 (1998), pp. 804–915. (Preliminary version in Proceedings of the 36th IEEE Symposium on Foundations of Computer Science, 1995, pp. 422–431.)

[BGLR93] M. BELLARE, S. GOLDWASSER, C. LUND, AND A. RUSSELL, *Efficient probabilistically checkable proofs and applications to approximation*, in Proceedings of the 25th ACM Symposium on Theory of Computing, San Diego, CA, 1993, pp. 294–304.

[BGH+04a] E. BEN-SASSON, O. GOLDREICH, P. HARSHA, M. SUDAN, AND S. VADHAN, *Robust PCPs of proximity, shorter PCPs and applications to coding*, in Proceedings of the 36th ACM Symposium on Theory of Computing, Chicago, IL, 2004, pp. 1–10.

[BGH+04b] E. BEN-SASSON, O. GOLDREICH, P. HARSHA, M. SUDAN, AND S. VADHAN, *Robust PCPs of Proximity, Shorter PCPs and Applications to Coding*, Tech. report

TR04-021, Electronic Colloquium on Computational Complexity, 2004. Available online at http://eccc.hpi-web.de/eccc-reports/2004/TR04-021/index.html

[BGH+05] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. Vadhan, *Short PCPs verifiable in polylogarithmic time*, in Proceedings of the 20th IEEE Conference on Computational Complexity, San Jose, CA, 2005, pp. 120–134.

[BHR05] E. Ben-Sasson, P. Harsha, and S. Raskhodnikova, *Some 3CNF properties are hard to test*, SIAM J. Comput., 35 (2005), pp. 1–21. (Preliminary version in Proceedings of the 35th ACM Symposium on Theory of Computing, 2003, pp. 345–354.)

[BS05] E. Ben-Sasson and M. Sudan, *Simple PCPs with poly-log rate and query complexity*, in Proceedings of the 37th ACM Symposium on Theory of Computing, Baltimore, MD, 2005, pp. 266–275.

[BSVW03] E. Ben-Sasson, M. Sudan, S. Vadhan, and A. Wigderson, *Randomness-efficient low degree tests and short PCPs via epsilon-biased sets*, in Proceedings of the 35th ACM Symposium on Theory of Computing, San Diego, CA, 2003, pp. 612–621.

[BLR93] M. Blum, M. Luby, and R. Rubinfeld, *Self-testing/correcting with applications to numerical problems*, J. Comput. System Sci., 47 (1993), pp. 549–595. (Preliminary version in Proceedings of the 22nd ACM Symposium on Theory of Computing, 1990, pp. 73–83.)

[BOT02] A. Bogdanov, K. Obata, and L. Trevisan, *A lower bound for testing 3-colorability in bounded-degree graphs*, in Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science, Vancouver, Canada, 2002, pp. 93–102.

[BT04] A. Bogdanov and L. Trevisan, *Lower bounds for testing bipartiteness in dense graphs*, in Proceedings of the 19th IEEE Conference on Computational Complexity, Amherst, MA, 2004, pp. 75–81.

[BdW04] H. Buhrman and R. de Wolf, *On Relaxed Locally Decodable Codes*, manuscript, 2004.

[CGH04] R. Canetti, O. Goldreich, and S. Halevi, *The random oracle methodology, revisited*, J. ACM, 51 (2004), pp. 557–594. (Preliminary version in Proceedings of the 30th ACM Symposium on Theory of Computing, 1998, pp. 209–218.)

[Coo88] S. A. Cook, *Short propositional formulas represent nondeterministic computations*, Inform. Process. Lett., 26 (1988), pp. 269–270.

[DJK+02] A. Deshpande, R. Jain, T. Kavitha, S. V. Lokam, and J. Radhakrishnan, *Lower bounds for adaptive locally decodable codes*, Random Structures Algorithms, 27 (2005), pp. 358–378. (Preliminary version in Proceedings of the 17th IEEE Conference on Computational Complexity, 2002, pp. 152–161.)

[Din06] I. Dinur, *The PCP theorem by gap amplification*, in Proceedings of the 38th ACM Symposium on Theory of Computing, Seattle, WA, 2006, pp. 241–250.

[DR04] I. Dinur and O. Reingold, *Assignment-testers: Towards a combinatorial proof of the PCP-theorem*, in Proceedings of the 45th IEEE Symposium on Foundations of Computer Science, Rome, Italy, 2004, pp. 155–164.

[EKR04] F. Ergün, R. Kumar, and R. Rubinfeld, *Fast approximate probabilistically checkable proofs*, Inform. Comput., 189 (2004), pp. 135–159. (Preliminary version in Proceedings of the 31st ACM Symposium on Theory of Computing, 1999, pp. 41–50.)

[FGL+96] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy, *Interactive proofs and the hardness of approximating cliques*, J. ACM, 43 (1996), pp. 268–292. (Preliminary version in Proceedings of the 32nd IEEE Symposium on Foundations of Computer Science, 1991, pp. 2–12.)

[FRS94] L. Fortnow, J. Rompel, and M. Sipser, *On the power of multi-prover interactive protocols*, Theoret. Comput. Sci., 134 (1994), pp. 545–557. (Preliminary version in Proceedings of the 3rd IEEE Symposium on Structure in Complexity Theory, 1988, pp. 156–161)

[Gol97] O. Goldreich, *A Sample of Samplers—A Computational Perspective on Sampling*, Tech. report TR97-020, Electronic Colloquium on Computational Complexity, 1997. Available online at http://www.eccc.uni-trier.de/eccc-reports/1997/TR97-020.

[GGR98] O. Goldreich, S. Goldwasser, and D. Ron, *Property testing and its connection to learning and approximation*, J. ACM, 45 (1998), pp. 653–750. (Preliminary version in Proceedings of the 37th IEEE Symposium on Foundations of Computer Science, 1996, pp. 339–348.)

[GR02] O. Goldreich and D. Ron, *Property testing in bounded degree graphs*, Algorithmica, 32 (2002), pp. 302–343. (Preliminary version in Proceedings of the 29th ACM

Symposium on Theory of Computing, 1997, pp. 416–415.)

[GS00] O. Goldreich and S. Safra, *A combinatorial consistency lemma with application to proving the PCP theorem*, SIAM J. Comput., 29 (2000), pp. 1132–1154. (Preliminary version in RANDOM, Lecture Notes in Comput. Sci. 1269, Springer, Berlin, 1997, pp. 67–84.)

[GS02] O. Goldreich and M. Sudan, *Locally testable codes and PCPs of almost-linear length*, in Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science, Vancouver, Canada, 2002, pp. 13–22. Available online at http://www.wisdom.weizmann.ac.il/~oded/p_ltc.html.

[GW97] O. Goldreich and A. Wigderson, *Tiny families of functions with random properties: A quality-size trade-off for hashing*, Random Structures Algorithms, 11 (1997), pp. 315–343. (Preliminary version in Proceedings of the 26th ACM Symposium on Theory of Computing, 1994, pp. 574–584.)

[GLST98] V. Guruswami, D. Lewin, M. Sudan, and L. Trevisan, *A tight characterization of NP with 3-query PCPs*, in Proceedings of the 39th IEEE Symposium on Foundations of Computer Science, Palo Alto, CA, 1998, pp. 18–27.

[HS00] P. Harsha and M. Sudan, *Small PCPs with low query complexity*, Comput. Complex., 9 (2000), pp. 157–201. (Preliminary version in 18th STACS, Lecture Notes in Comput. Sci. 2010, Springer, Berlin, 2001, pp. 327–338.)

[Hås01] J. Håstad, *Some optimal inapproximability results*, J. ACM, 48 (2001), pp. 798–859. (Preliminary version in Proceedings of the 29th ACM Symposium on Theory of Computing, 1997, pp. 1–10.)

[HS66] F. C. Hennie and R. E. Stearns, *Two-tape simulation of multitape Turing machines*, J. ACM, 13 (1966), pp. 533–546.

[KT00] J. Katz and L. Trevisan, *On the efficiency of local decoding procedures for error-correcting codes*, in Proceedings of the 32nd ACM Symposium on Theory of Computing, Portland, OR, 2000, pp. 80–86.

[KdW04] I. Kerenidis and R. de Wolf, *Exponential lower bound for 2-query locally decodable codes via a quantum argument*, J. Comput. System Sci., 69 (2004), pp. 395–420. (Preliminary version in Proceedings of the 35th ACM Symposium on Theory of Computing, 2003, pp. 106–115.)

[Kil92] J. Kilian, *A note on efficient zero-knowledge proofs and arguments (extended abstract)*, in Proceedings of the 24th ACM Symposium on Theory of Computing, Victoria, BC, Canada, 1992, pp. 723–732.

[LS97] D. Lapidot and A. Shamir, *Fully parallelized multi prover protocols for NEXP-time*, J. Comput. System Sci., 54 (1997), pp. 215–220. (Preliminary version in Proceedings of the 32nd IEEE Symposium on Foundations of Computer Science, 1991, pp. 13–18.)

[Lei92] F. T. Leighton, *Introduction to Parallel Algorithms and Architectures*, Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1992.

[LFKN92] C. Lund, L. Fortnow, H. Karloff, and N. Nisan, *Algebraic methods for interactive proof systems*, J. ACM, 39 (1992), pp. 859–868. (Preliminary version in Proceedings of the 31st IEEE Symposium on Foundations of Computer Science, 1990, pp. 2–10.)

[Mic00] S. Micali, *Computationally sound proofs*, SIAM J. Comput., 30 (2000), pp. 1253–1298. (Preliminary version in Proceedings of the 35th IEEE Symposium on Foundations of Computer Science, 1994, pp. 436–453.)

[MR06] D. Moshkovitz and R. Raz, *Sub-constant error low degree test of almost linear size*, in Proceedings of the 38th ACM Symposium on Theory of Computing, Seattle, WA, 2006, pp. 21–30.

[NN93] J. Naor and M. Naor, *Small-bias probability spaces: Efficient constructions and applications*, SIAM J. Comput., 22 (1993), pp. 838–856. (Preliminary version in Proceedings of the 22nd ACM Symposium on Theory of Computing, 1999, pp. 213–233.)

[PF79] N. Pippenger and M. J. Fischer, *Relations among complexity measures*, J. ACM, 26 (1979), pp. 361–381.

[PS94] A. Polishchuk and D. A. Spielman, *Nearly-linear size holographic proofs*, in Proceedings of the 26th ACM Symposium on Theory of Computing, Montréal, Québec, Canada, 1994, pp. 194–203.

[Raz98] R. Raz, *A parallel repetition theorem*, SIAM J. Comput., 27 (1998), pp. 763–803. (Preliminary version in Proceedings of the 27th ACM Symposium on Theory of Computing, 1995, pp. 447–556.)

[RS96] R. Rubinfeld and M. Sudan, *Robust characterizations of polynomials with applications to program testing*, SIAM J. Comput., 25 (1996), pp. 252–271. (Preliminary versions in Proceedings of the 23rd ACM Symposium on Theory of Computing, 1991, pp. 33–42 and Proceedings of the 3rd ACM–SIAM Symposium on Discrete Algorithms, 1992, pp. 23–43.)

[ST00] A. Samorodnitsky and L. Trevisan, *A PCP characterization of NP with optimal amortized query complexity*, in Proceedings of the 32nd ACM Symposium on Theory of Computing, Portland, OR, 2000, pp. 191–199.

[Sch77] A. Schönhage, *Schnelle multiplikation von polynomen über Körpern der charakteristik* 2 *[Fast multiplication of polynomials over characteristic* 2*]*, Acta Inform., 4 (1977), pp. 395–398.

[SS71] A. Schönhage and V. Strassen, *Schnelle multiplikation großer zahlen [Fast multiplication of large numbers]*, Computing, 7 (1971), pp. 281–292.

[Spi96] D. A. Spielman, *Linear-time encodable and decodable error-correcting codes*, IEEE Trans. Inform. Theory, 42 (1996), pp. 1723–1732. (Preliminary version in Proceedings of the 27th ACM Symposium on Theory of Computing, 1995, pp. 388–397.)

[Spi95] D. A. Spielman, *Computationally Efficient Error-Correcting Codes and Holographic Proofs*, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, June 1995.

[Str73] V. Srassen, *Vermeidung von Divisionen [Avoiding divisions]*, J. Reine Angew. Math., 264 (1973), pp. 184–202.

[Sze99] M. Szegedy, *Many-valued logics and holographic proofs*, in Proceedings of the 26th International Colloquium of Automata, Languages, and Programming (ICALP '99), J. Wiedermann, P. van Emde Boas, and M. Nielsen, eds., Lecture Notes in Comput. Sci. 1644, Springer-Verlag, Berlin, pp. 676–686.