

Toward Probabilistic Checking against Non-Signaling Strategies with Constant Locality

Mohammad Mahdi Jahanara
mjahanar@sfu.ca
Simon Fraser University

Sajin Koroth
sajin_koroth@sfu.ca
Simon Fraser University

Igor Shinkar
ishinkar@sfu.ca
Simon Fraser University

September 7, 2020

Abstract

Non-signaling strategies are a generalization of quantum strategies that have been studied in physics over the past three decades. Recently, they have found applications in theoretical computer science, including to proving inapproximability results for linear programming and to constructing protocols for delegating computation. A central tool for these applications is probabilistically checkable proof (PCPs) systems that are *sound against non-signaling strategies*.

In this paper we show, assuming a certain geometrical hypothesis about noise robustness of non-signaling proofs (or, equivalently, about robustness to noise of solutions to the Sherali-Adams linear program), that a slight variant of the parallel repetition of the exponential-length constant-query PCP construction due to Arora et al. (JACM 1998) is sound against non-signaling strategies with *constant locality*.

Our proof relies on the analysis of the *linearity test* and *agreement test* (also known as the *direct product test*) in the non-signaling setting.

Keywords: direct product testing; linearity testing; non-signaling strategies; parallel repetition; probabilistically checkable proofs

Contents

1	Introduction	3
1.1	Informal statement of the result	5
1.2	Roadmap	6
2	Preliminaries	6
2.1	Probabilistically Checkable Proofs	6
2.2	Parallel repetition	7
2.3	Non-signaling functions	7
2.4	Permutation folded repeated non-signaling functions	8
2.5	Linear non-signaling functions	9
2.6	Consistent repeated non-signaling functions	10
2.7	Flattening of a t -nsPCP	11
3	Main result	12
4	The PCP construction	13
4.1	The linear ALMSS verifier	14
4.2	Parallel repetition of the linear ALMSS verifier	15
4.3	From linear PCPs to standard PCPs using linearity and consistency testing	15
5	Proof overview: Soundness	17
6	Soundness of the linear PCP verifier against structured proofs	18
7	Testing and self-correcting repeated non-signaling functions	19
7.1	Definitions of tests and the self-correction	19
7.2	Self-correction of a t -repeated k -non-signaling function	20
7.3	Self-correction is almost linear and almost consistent	21
7.3.1	$\widehat{\mathcal{F}}^{(t)}$ is permutation folded	21
7.3.2	$\widehat{\mathcal{F}}^{(t)}$ is almost linear	21
7.3.3	$\mathcal{F}^{(t)}$ is almost consistent	23
8	Proof of Theorem 2	24
9	Conclusions and open problems	26
	References	26
A	Discussion on Hypothesis 2	29

1 Introduction

Probabilistically Checkable Proofs (PCPs) [BFLS91; FGLSS96; AS98; ALMSS98] are proofs that can be verified by a probabilistic verifier that queries only a few locations of the proof. PCPs have been a powerful tool in the theory of computing, with applications in diverse areas such as hardness of approximation [FGLSS96] and delegation of computation [Kil92; Mic00]. A seminal result of [AS98; ALMSS98], known as the PCP theorem, says that every language decidable by a non-deterministic Turing machine in time $T(n)$ has a PCP system which allows to check if a given input of length n is in the language by using $O(\log(T(n)))$ random bits and making only $O(1)$ queries to the given proof.

Recall that in the classical setting of PCPs the two standard requirements are completeness and soundness. *Completeness* requires that if a given input is in the language, then there is some proof that convinces the prover with probability 1. *Soundness* requirement states that if the input is not in the language, the prover rejects *any* proof with some significant probability. In this paper we study PCP systems that are sound against *non-signalling proofs* or *non-signalling strategies*, i.e., we require the prover to reject *any* non-signalling proof with some significant probability.

Non-signalling strategies are a certain restricted class of probabilistic oracles. When such oracle is given a set of queries, the response to the queries is sampled from a distribution such that the answer to each query may depend on all queries. More precisely, a non signalling strategy with locality k is a collection $\mathcal{F} = \{\mathcal{F}_S\}_{S \subseteq D, |S| \leq k}$, where each \mathcal{F}_S is a distribution over Σ^S (i.e., over functions $f : S \rightarrow \Sigma$), and for any two subsets $S, T \subseteq D$ of size at most k , the restrictions of \mathcal{F}_S and \mathcal{F}_T to $S \cap T$ are equal as distributions. This setting stands in contrast to the standard notion of a classical proof, where the answer to each query is deterministic. Note that if the locality is the maximum possible, i.e., $k = |D|$, then \mathcal{F} is a distribution over functions, which is (essentially) equivalent to the classical notion of a proof.

We note that one may think about k -non-signalling functions, equivalently, as the class of all feasible solutions to the linear program arising from the k 'th level relaxation of the Sherali-Adams hierarchy [SA90]. This implies that computing the maximum acceptance probability of an nsPCP verifier that uses r random bits, where the maximum is taken over all k -non-signaling proofs, reduces to a linear program with $2^{O(r \cdot k^2)} \cdot \Sigma^{O(k)}$ variables and constraints. In particular, if a language L has a PCP verifier that on an input of length n uses $r = O(\log(n))$ random bits, and is sound against $O(1)$ -non-signaling proofs over an alphabet of constant size, then L is decidable in time $\text{poly}(n)$.

Non-signaling strategies have been studied in physics since 1980's [Ras85; KT85; PR94] in order to better understand quantum entanglement. Indeed, these strategies strictly generalize quantum strategies and capture minimal requirements on “non-local” correlations that rule out instantaneous communication.

PCP systems that are sound against *non-signalling proofs* have recently found numerous applications in theoretical computer science, including schemes for 1-round delegation of computation from cryptographic assumptions [KRR13; KRR14], and hardness of approximation for linear programming [KRR16]. However, as opposed to the well studied setting of the classical PCP theorem, where there are many constructions achieving best parameters possible, in the non-signalling setting many parameters of the known PCP constructions appear to be far from optimal.

One of the most important parameters associated with a non-signalling proof is the locality parameter, denoted by k . Indeed, [KRR13; KRR14] have studied the related notion of multi-prover interactive proofs that are sound against non-signaling strategies (nsMIPs). They have shown that nsPCPs are essentially equivalent to nsMIPs where k , the locality of the proof in the nsPCP setting,

exactly corresponds to the number of provers in the nsMIP setting.

Despite the importance of the locality parameter, the exact complexity of languages admitting nsPCPs that are sound against k -non-signaling proofs is still open for most k 's. Note that as the locality of the proof decreases, there are fewer constraints imposed on the proof, and hence the task of the verifier becomes more challenging. The seminal result of Kalai, Raz, and Rothblum [KRR13; KRR14] showed that every language in $\text{DTIME}(T)$ has an nsPCP verifier that uses $\text{polylog}(T)$ random bits, makes $\text{polylog}(T)$ queries to a proof of length $\text{poly}(T)$, and is sound against $\text{polylog}(T)$ -non-signaling proofs. In particular, every language in EXP is captured by a nsMIP with a polynomial time randomized verifier who communicates with $\text{poly}(n)$ non-signaling provers. For the limitations of nsPCPs, Ito [Ito10] proved that for $k = 2$ the corresponding linear program is solvable in PSPACE , which is tight by the result of [IKM09], and hence the class PSPACE is captured by PCPs that are sound against 2-non-signaling proofs. Much less is known about the power of PCP systems that are sound against k -non-signaling proofs for $k > 2$. Recently, Holden and Kalai [HK20] proved that $o(\sqrt{\log(n)})$ -prover non-signalling proofs with *negligible soundness* is contained in PSPACE .

All these results give rise to the following question, raised in [CMS19], asking for the non-signaling analogue of the PCP theorem.

Question 1.1. *Is it true that every language in $\text{DTIME}(T)$ has an nsPCP verifier that uses $O(\log(T))$ random bits, makes $O(1)$ queries to the proof, and is sound against $O(1)$ -non-signalling functions?*

Motivated by this problem, Chiesa et. al [CMS19; CMS20] started a systematic study of non-signalling PCPs. They proposed studying the classical (algebraic) PCP constructions and their building blocks (which are very well understood in the classical setting), and adapting each of the building blocks to the non-signaling setting. In particular, focusing on the PCP construction of [ALMSS98] they made an appropriate definition of *linear* non-signalling functions and analyzed the linearity test of [BLR93] against non-signalling strategies [CMS20]. Then, building on the linearity test, they proved in [CMS19] that the classical exponential length $O(1)$ -query PCP of [ALMSS98] is sound against $O(\log^2(N))$ -non-signalling proofs. We emphasize, that even for exponential length nsPCPs (corresponding to nsPCPs with $r = \text{poly}(n)$ randomness), there are no known constructions that are sound against $O(1)$ -non-signaling proofs. Given this state of affairs, it is natural to ask the following question, that is simpler than Question 1.1

Question 1.2. *Is it true that every language in $\text{DTIME}(T(n))$ has an nsPCP verifier that uses $O(\text{poly}(T(n)))$ random bits, makes $O(1)$ queries to the proof, and is sound against $O(1)$ -non-signalling functions?*

One must be careful with the precise formulation of Question 1.2. Note that if the verifier uses more than $T(n)$ random bits, the runtime spent on reading the randomness is more than $T(n)$, which is the time complexity of the problem. To recover a nontrivial question, we require the verifier to be *input oblivious*. That is, in order to decide whether an instance x belongs to the given language $L \in \text{DTIME}(T(n))$, the verifier generates the queries based only on the length of the input x and its randomness (but not the input itself), and then rules according to an $o(T)$ -time decision predicate (where the predicate does depend on x). Indeed, the [ALMSS98] verifier studied in [CMS19] is input oblivious.

In this work we build on the work of [CMS19] and provide a positive answer to Question 1.2 *assuming a certain geometric hypothesis*. Specifically, we construct an input oblivious nsPCP verifier for any language $L \in \text{DTIME}(T(n))$ that uses $\text{poly}(T(n))$ random bits, makes $O(1)$ queries to a given proof, and is sound against $O(1)$ -non-signalling functions, with *two caveats*.

1. The first is that the alphabet of the nsPCP system is $\Sigma = \{0, 1\}^{\text{polylog}(T(n))}$, instead of the binary alphabet employed by [CMS19; KRR14; ALMSS98]. Still, this means that the verifier reads a total of $\text{polylog}(T(n))$ bits from the proof, which makes our result non-trivial. Also, recall that in the classical setting, we have the alphabet reduction technique using proof composition, and it is plausible that we can apply similar ideas also in the non-signaling setting. Indeed, proof composition is an important building block in the classical PCP literature, and we believe it will also be an important step toward resolving Question 1.1.
2. The second caveat is that our result depends on a certain quantitative geometric hypothesis about proximity between almost non-signaling proofs and exactly non-signaling proofs. Equivalently, the hypothesis says that every feasible solution for the noisy Sherali-Adams LP is close (in some precise, rather weak, sense) to a feasible solution for the (exact) Sherali-Adams LP. See Hypothesis 2 for details, and the discussion in Appendix A.

Our work follows the general philosophy of [CMS19; CMS20], who proposed building modular analogues of tools and techniques from the classical PCP literature. A classical tool used in the construction of PCPs is parallel repetition [Raz98; Hol09]. In the classical setting of 2-query PCP, parallel repetition is used to reduce the soundness error. In this work we use parallel repetition for non-signaling proofs to reduce the locality to $O(1)$, while the soundness stays in the “high-probability acceptance regime”. In addition to parallel repetition, we study additional tools from the PCP literature. Specifically, we use the modular approach that is typical for the classical setting. Specifically, we show first that the parallel repetition of the [ALMSS98] verifier is sound against “nicely structured” proofs. Then, we use *linearity test* and *direct product test*, and claim that proofs that satisfy both tests with high probability must be nicely structured, and hence we essentially reduce the analysis to the structured case.

Another interesting feature of our proof is the reduction from the parallel repetition of the [ALMSS98] verifier to the non-repeated [ALMSS98] verifier. Specifically, we show that if for some input x , the parallel repetition of the [ALMSS98] verifier accepts a proof with high probability, and the proof is “nicely structured”, then it is possible to “flatten” the repeated proof into a proof over the binary alphabet, that satisfies the (non-repeated) [ALMSS98] verifier with high probability. Therefore, by applying the result of [CMS19] about the soundness of the [ALMSS98] verifier, we conclude that the input x is in the language.

1.1 Informal statement of the result

Below we discuss the main result of the paper. Our result depends on an hypothesis about approximating *almost non-signaling* functions using *exactly non-signaling* functions.

Hypothesis 1 (Informal). *Any almost linear, almost non-signaling function $\mathcal{F}: \{0, 1\}^n \rightarrow \{0, 1\}$ can be well approximated by some non-signaling function $\mathcal{F}': \{0, 1\}^n \rightarrow \{0, 1\}$ of slightly lower locality.*

Equivalently, any solution to the noisy Sherali-Adams LP can be well approximated by a solution to the (exact) Sherali-Adams LP of slightly lower level in the hierarchy.

The exact formulation of the hypothesis relies on the precise definitions of *non-signaling* and *almost non-signaling* functions (or, equivalently, the related notions of *noisy Sherali-Adams LP*), as well as the appropriate definitions of distance. For the formal statement of the hypothesis see Hypothesis 2 following the required definitions in Section 2.

We are now ready to state our main theorem.

Theorem 1 (Main theorem - informal). *Assuming Hypothesis 1 every language $L \in \text{DTIME}(T)$ has an input oblivious nsPCP verifier that on an input of length n uses $\tilde{O}(T^2)$ random bits, makes $O(1)$ queries to proofs over the alphabet $\Sigma = \{0, 1\}^{\text{polylog}(T)}$, and is sound against $O(1)$ -non-signaling proofs. The query sampler runs in time $\tilde{O}(T^2)$, and the decision predicate runs in time $O(n \cdot \text{polylog}(T))$.*

To the best of our knowledge, this is the first result that constructs a PCP system that is sound against non-signaling proofs with constant locality.

1.2 Roadmap

The rest of the paper is organized as follows. In Section 2 we formally define the notions that we utilize throughout this work, and use them to formally state our hypothesis and the main theorem in Section 3. In Section 4 we recall the ALMSS verifier, and define our variant of its parallel repetition. In Section 5 we provide an overview of the soundness proof. In Section 6 we prove soundness of our verifier against structured proofs. In Section 7 we discuss our local testing and self-correction, which enables us to reduce soundness against general proofs to soundness against structured proofs. Finally, in Section 8 we prove the main result.

2 Preliminaries

2.1 Probabilistically Checkable Proofs

We start with the definition of Probabilistically Checkable Proofs (PCPs). Recall that a classical PCP verifier for a language L is given an input x , and an oracle access to a proof. The verifier reads the input, uses randomness, queries the proof in a small number of coordinates, and based on the answers to the queries decides whether to accept or reject. Completeness requires that if $x \in L$, then there exists a proof that makes the verifier always accept. Soundness requires that if $x \notin L$, then for any proof the verifier will reject with high probability.

In the non-signaling setting, a *non-signaling PCP verifier* is a verifier, whose soundness is further required to hold against any non-signaling proof of prescribed locality. More precisely, an nsPCP verifier V for a language L gets an input x and an oracle access to a non-signaling function $\mathcal{F}: D \rightarrow \Sigma$. The verifier reads the input x , uses random bits to decide on a subset $S \subseteq D$ on which \mathcal{F} is queried. Then, based on the answer $\mathcal{F}(S) \in \Sigma^S$ it decides to accept or reject.

Definition 2.1. *A nsPCP verifier for a language $L \subseteq \{0, 1\}^*$ is a randomized algorithm V that gets an input $x \in \{0, 1\}^n$ and oracle access to a k -non-signaling proof $\mathcal{F}: D \rightarrow \Sigma$. The verifier uses randomness to decide on a subset $S \subseteq D$ of size $|S| \leq k$, and queries \mathcal{F} on S . Then, based on the answer $\mathcal{F}(S) \in \Sigma^S$ it decides to accept or reject. We say that V has perfect completeness and soundness error γ against k -non-signaling proofs if the following holds.*

Completeness: *For all $x \in L$ there exists a (classical) proof π such that $\Pr[V_\pi(x) = 1] = 1$.*

Soundness: *If $x \notin L$, then for all k -non-signaling proofs \mathcal{F} it holds that $\Pr[V_\mathcal{F}(x) = 1] \leq \gamma$.*

*We say that verifier V is **input oblivious** if the choice of the query set S depends only on the input length n , the randomness of the verifier, but is independent of x .*

Remark 2.2. Note that in the non-signaling setting the locality parameter k upper bounds the number of queries made by the verifier, and it is possible that the actual predicate used by the verifier depends on significantly less than k coordinates of the proof. For example, [CMS19] proved that the 11-queries verifier of [ALMSS98] is sound against $O(\log^2(n))$ -non-signaling proofs, and it is not known whether the verifier is sound against $O(1)$ -non-signaling proofs, or even $o(\log^2(n))$ -non-signaling proofs.

2.2 Parallel repetition

In the classical setting a proof is assumed to be a string, or equivalently, a static function $\pi : D \rightarrow \Sigma$ committed by the prover. A t -parallel repetition of a proof π is a mapping $\pi^t : D^t \rightarrow \Sigma^t$ that allows accessing t locations of the (supposed) proof by making only 1 query to a (longer) proof over a larger alphabet. That is, the intended proof $\pi^{(t)}$ corresponds to some “base” proof $\pi : D \rightarrow \Sigma$ defined as $\pi^{(t)}((x_1, \dots, x_t)) = (\pi(x_1), \dots, \pi(x_t))$. Analogously, given a verifier V , a t -repeated verifier which is denoted by $V^{(t)}$, runs t parallel *independent* instances of V and accepts if and only if all instances accept.

The original motivation for using parallel repetition was to reduce the soundness error of a proof system, while keeping the number of queries fixed. In the classical setting, if the repeated proof is indeed a parallel repetition of some base proof π , then it is not hard to see that the soundness error of $V_{\pi^t}^{(t)}$ is exponentially smaller than the soundness error of V_π . The soundness analysis of the repeated proof need not be based on this comparison to the soundness error of the *base* proof, and analyzing such proofs in both classical and non-signalling settings has been a subject of a long line of research [Ver96; Raz98; Hol09; DS14a; BG15; LW16; HY19].

In this work, we use parallel repetition to improve the *minimum locality parameter* of non-signaling proofs required for the soundness of the verifier, rather than its soundness error. Next, we formally define non-signaling proofs, and some properties of such proofs that we will need in the paper.

2.3 Non-signaling functions

In this work we consider PCP verifiers that are sound against non-signaling proofs. Below, we formally define the notion of non-signaling functions, and introduce some notation we will use in the paper. Throughout the paper we will use terms *non-signaling function*, *non-signaling proof*, and *non-signaling strategy* interchangeably.

Definition 2.3. Fix a domain D , an alphabet Σ , and a parameter $k \in \mathbb{N}$. A ***k*-non-signaling function** $\mathcal{F} : D \rightarrow \Sigma$ is a collection $\mathcal{F} = \{\mathcal{F}_S\}_{S \subseteq D, |S| \leq k}$, where each \mathcal{F}_S is a distribution over assignments $f_S : S \rightarrow \Sigma$, such that for every two subsets $S, T \subseteq D$ each of size at most k , the marginal distributions of \mathcal{F}_S and \mathcal{F}_T restricted to $S \cap T$ are equal.

Unlike a classic function, we can use a k -non-signaling function only once in the sense that one has to present the set of at most k queries all at once. In other words, it is not possible to use the non-signaling function adaptively.

Remark 2.4. Throughout the paper we will consider non-signaling functions of two types:

- functions over the domain $D = \{0, 1\}^N$ for some $N \in \mathbb{N}$ and alphabet $\Sigma = \{0, 1\}$;

- functions over the domain $D = (\{0, 1\}^N)^t$ and alphabet $\Sigma = \{0, 1\}^t$ for some parameters $N, t \in \mathbb{N}$.

Next, we define a relaxed notion of non-signaling functions, that allows the marginal distributions induced by different query sets to be only *statistically close* rather equal on the intersection. This relaxation arises in our analysis. It has also appeared naturally in other works in this area, especially in cryptographic applications [ABOR00; DLNRR04; KRR13; KRR14].

Definition 2.5. Fix a domain D , an alphabet Σ , and parameters $k \in \mathbb{N}$ and $\varepsilon \in [0, 1]$. A (ε, k) -**non-signaling function** over a domain D and an alphabet Σ , is a collection $\mathcal{F} = \{\mathcal{F}_S\}_{S \subseteq D, |S| \leq k}$, where each \mathcal{F}_S is a distribution over assignments $f_S: S \rightarrow \Sigma$, such that for every two subsets $S, T \subseteq D$ each of size at most k , the marginal distributions of \mathcal{F}_S and \mathcal{F}_T restricted to $S \cap T$ are ε -close with respect to total variation distance, i.e.,

$$\max_{E \subseteq \Sigma^{S \cap T}} \left| \Pr_{\mathcal{F}_S}[\mathcal{F}_S|_{S \cap T} \in E] - \Pr_{\mathcal{F}_T}[\mathcal{F}_T|_{S \cap T} \in E] \right| \leq \varepsilon .$$

In particular, a $(\varepsilon = 0, k)$ -non-signaling-function coincides with the definition of k -non-signaling function from Definition 2.3.

Next we define non-signaling and almost non-signaling counterpart of parallel repeated functions.

Definition 2.6. Fix a domain D , an alphabet Σ , and parameters $k, t \in \mathbb{N}$. A t -**repeated** (δ, k) -**non-signaling function** is an (δ, k) -non-signaling function $\mathcal{F}^{(t)}: D^t \rightarrow \Sigma^t$. Namely, a t -repeated (δ, k) -non-signaling function $\mathcal{F}^{(t)}: D^t \rightarrow \Sigma^t$ is a collection $\mathcal{F}^{(t)} = \{\mathcal{F}_S^{(t)}\}_{S \subseteq D^t, |S| \leq k}$, where each $\mathcal{F}_S^{(t)}$ is a distribution over assignments $f_S^{(t)}: S \rightarrow \Sigma$, such that for every two subsets $S, T \subseteq D$ each of size at most k , the marginal distributions of $\mathcal{F}_S^{(t)}$ and $\mathcal{F}_T^{(t)}$ restricted to $S \cap T$ are δ -close with respect to total variation distance.

We will also need the definition of distance between non-signaling or almost non-signaling functions.

Definition 2.7 (Statistical distance). Let $\mathcal{F}, \mathcal{F}': D \rightarrow \Sigma$ be two non-signaling or almost non-signaling functions with locality k . For $\ell \leq k$ the Δ_ℓ -distance between \mathcal{F} and \mathcal{F}' is defined as

$$\Delta_\ell(\mathcal{F}, \mathcal{F}') = \max_{S \subseteq D, |S| \leq \ell} \Delta(\mathcal{F}_S, \mathcal{F}'_S) ,$$

where $\Delta(\mathcal{F}_S, \mathcal{F}'_S) = \max_{E \subseteq \Sigma^S} |\Pr[\mathcal{F}_S \in E] - \Pr[\mathcal{F}'_S \in E]|$ is the total variation distance between \mathcal{F}_S and \mathcal{F}'_S .

We say that \mathcal{F} and \mathcal{F}' are ε -close in the Δ_ℓ -distance if $\Delta_\ell(\mathcal{F}, \mathcal{F}') \leq \varepsilon$, and say that they are ε -far otherwise.

2.4 Permutation folded repeated non-signaling functions

Folding is a technique used to impose some structure on the given proof without really making extra queries. The idea of using folded proofs was first introduced by [BGS98]. We formally define the *permutation folding* property, and then explain why we can impose this property without making extra queries.

Definition 2.8. Let $Q = (q_1, \dots, q_t) \in D^t$ be a D -values vector, and let $\pi \in S_t$ be a permutation of the indices $[t]$. Define $\pi(Q) = (q_{\pi(1)}, \dots, q_{\pi(t)})$ to be the vector obtained from Q by permuting the coordinates according to π .

Let $\mathcal{F}^{(t)}: (D^n)^t \rightarrow \Sigma^t$ be a t -repeated k -non-signaling function. $\mathcal{F}^{(t)}$ is said to be **permutation folded** or **permutation invariant** if for any $S = \{Q_1, \dots, Q_\ell\} \subseteq (D^n)^t$ with $1 \leq \ell \leq k$, for any $T = \{\pi_1(Q_1), \dots, \pi_\ell(Q_\ell)\}$ for some permutations $\pi_1, \dots, \pi_\ell \in S_t$, and for any $b_1, \dots, b_\ell \in \Sigma^t$ it holds that

$$\Pr \left[\forall i \in [\ell] \quad \mathcal{F}_S^{(t)}(Q_i) = b_i \right] = \Pr \left[\forall i \in [\ell] \quad \mathcal{F}_T^{(t)}(\pi_i(Q_i)) = \pi_i(b_i) \right] .$$

Observation 2.9. It is important to note that we can fold any given t -repeated k -non-signaling function $\mathcal{F}^{(t)}: D^t \rightarrow \Sigma^t$ by partitioning D^t into equivalence classes, where Q and Q' belong to the same class if $Q' = \pi(Q)$ for some permutation π .

We defined the folding of $\mathcal{F}^{(t)}$, denoted by $\overline{\mathcal{F}^{(t)}}$ as follows. For any query Q to $\overline{\mathcal{F}^{(t)}}$, let $\pi \in S_t$ be a uniformly random permutation, and define the distribution of $\overline{\mathcal{F}^{(t)}}(Q)$ as the distribution of $\pi^{-1}(\mathcal{F}^{(t)}(\pi(Q)))$.

It is easy to see that $\overline{\mathcal{F}^{(t)}}$ is indeed k -non-signaling and permutation folded. Furthermore, note that if $\mathcal{F}^{(t)}$ is permutation folded, then $\overline{\mathcal{F}^{(t)}} = \mathcal{F}^{(t)}$.

2.5 Linear non-signaling functions

In this part, we define *linear* t -repeated non-signaling functions. Linear non-signaling *boolean* functions have been studied in [CMS20; CMS19], and played a key role in the proving that the PCP verifier of [ALMSS98] is sound against non-signaling proofs. We also use such structured non-signaling proofs in this paper. See Section 4 for details.

Definition 2.10 (Linear t -repeated functions). Let $\mathcal{L}^{(t)}: (\{0, 1\}^n)^t \rightarrow \{0, 1\}^t$ be a t -repeated (ε, k) -non-signaling function. We say that $\mathcal{L}^{(t)}$ is *linear* if for all $X, Y \in (\{0, 1\}^n)^t$, and $X + Y \in (\{0, 1\}^n)^t$ defined by the coordinate-wise addition modulo 2, and for all $S \subseteq (\{0, 1\}^n)^t$ containing $X, Y, X + Y$ of size at most $|S| \leq k$, it holds that

$$\Pr_{\mathcal{L}_S^{(t)}} \left[\mathcal{L}^{(t)}(X) + \mathcal{L}^{(t)}(Y) = \mathcal{L}^{(t)}(X + Y) \right] = 1 .$$

Remark 2.11. Note that in the degenerate case of $t = 1$ if a (non-repeated) k -non-signaling function \mathcal{F} satisfies the linearity condition in Definition 2.10 then $\Pr[\mathcal{F}(x) + \mathcal{F}(y) = \mathcal{F}(x + y)] = 1$ for all $x, y \in \{0, 1\}^n$, i.e., \mathcal{F} satisfies the linearity test of [BLR93] with probability 1. Non-signaling functions satisfying this property have been the subject of work on linearity testing in the non-signaling setting [CMS20].

Next we extend Definition 2.10 by introducing the notion of an *almost* linear t -repeated non-signalling function.

Definition 2.12 (Almost linear t -repeated functions). Let $\mathcal{L}^{(t)}: (\{0, 1\}^n)^t \rightarrow \{0, 1\}^t$ be a t -repeated (δ, k) -non-signaling function. We say that $\mathcal{L}^{(t)}$ is $(1 - \varepsilon)$ -linear if for all $X, Y \in (\{0, 1\}^n)^t$, and $X + Y \in (\{0, 1\}^n)^t$ defined by the coordinate-wise addition modulo 2, and for all $S \subseteq (\{0, 1\}^n)^t$ containing $X, Y, X + Y$ of size at most $|S| \leq k$, it holds that

$$\Pr_{\mathcal{L}_S^{(t)}} \left[\mathcal{L}^{(t)}(X) + \mathcal{L}^{(t)}(Y) = \mathcal{L}^{(t)}(X + Y) \right] \geq 1 - \varepsilon .$$

We will allow ourselves to use the informal term *almost linear*, when referring to a non-signaling function $\mathcal{L}^{(t)}$ that is $(1 - \varepsilon)$ -linear for some small ε .

2.6 Consistent repeated non-signaling functions

In this part, we define the notion of *consistency* for t -repeated k -non-signaling function.

Definition 2.13 (Consistent t -repeated functions). *Let $\mathcal{C}^{(t)}: D^t \rightarrow \Sigma^t$ be a t -repeated k -non-signaling function. We say that $\mathcal{C}^{(t)}$ is consistent, if for any $Q, Q' \in D^t$ it holds that*

$$\Pr_{\mathcal{C}^{(t)}} \left[\mathcal{C}^{(t)}(Q)_j = \mathcal{C}^{(t)}(Q')_j \quad \forall j \in [t] \text{ such that } Q_j = Q'_j \right] = 1 .$$

Similarly to the almost linear property, we define the relaxed notion of almost consistent non-signalling function.

Definition 2.14 (Almost consistent t -repeated functions). *Let $\mathcal{C}^{(t)}: D^t \rightarrow \Sigma^t$ be a t -repeated k -non-signaling function. We say that $\mathcal{C}^{(t)}$ is $(1 - \varepsilon)$ -consistent, if for any $Q, Q' \in D^t$*

$$\Pr_{\mathcal{C}^{(t)}} \left[\mathcal{C}^{(t)}(Q)_j = \mathcal{C}^{(t)}(Q')_j \quad \forall j \in [t] \text{ such that } Q_j = Q'_j \right] \geq 1 - \varepsilon .$$

We will allow ourselves to use the informal term *almost consistent*, when referring to a non-signaling function $\mathcal{C}^{(t)}$ that is $(1 - \varepsilon)$ -consistent for some small ε .

Claim 2.15. *Let $\mathcal{C}^{(t)}: D^t \rightarrow \Sigma^t$ be a t -repeated k -non-signaling function for $k \geq 3$, and suppose that $\mathcal{C}^{(t)}$ is $(1 - \varepsilon)$ -consistent. Fix $Q, Q' \in D^t$ and let $J = \{j \in [t] : Q_j = Q'_j\}$. Then, for any event $E \subseteq \Sigma^J$ it holds that*

$$\left| \Pr[\mathcal{C}^{(t)}(Q)_J \in E] - \Pr[\mathcal{C}^{(t)}(Q')_J \in E] \right| \leq \varepsilon .$$

Proof. Note that

$$\begin{aligned} \Pr[\mathcal{C}^{(t)}(Q)_{|J} \in E] &\geq \Pr[\mathcal{C}^{(t)}(Q)_{|J} \in E \wedge \mathcal{C}^{(t)}(Q)_{|J} = \mathcal{C}^{(t)}(Q')_{|J}] \\ &= \Pr[\mathcal{C}^{(t)}(Q')_{|J} \in E \wedge \mathcal{C}^{(t)}(Q)_{|J} = \mathcal{C}^{(t)}(Q')_{|J}] \\ &\geq \Pr[\mathcal{C}^{(t)}(Q')_{|J} \in E] - \varepsilon , \end{aligned}$$

where the last inequality is by the assumption that $\mathcal{C}^{(t)}$ is $(1 - \varepsilon)$ -consistent. By symmetry, we also get the inequality in the other direction, and the claim follows. \square

We observe that for $D = \{0, 1\}^n$ and $\Sigma = \{0, 1\}$ (almost) linearity implies (almost) consistency. Specifically, we prove the following claim.

Claim 2.16. *Let $\mathcal{L}^{(t)}: (\{0, 1\}^n)^t \rightarrow \{0, 1\}^t$ be a t -repeated k -non-signaling function, and suppose that (i) $\mathcal{L}^{(t)}$ is $(1 - \varepsilon)$ -linear, and (ii) $\Pr[\mathcal{L}^{(t)}(Q)_j = 0 \quad \forall j \in [t] \text{ such that } Q_j = 0^n] > 1 - \varepsilon$ for all $Q \in (\{0, 1\}^n)^t$. Then, $\mathcal{L}^{(t)}$ is $(1 - 2\varepsilon)$ -consistent when treated as a $(k - 1)$ -non-signaling function.*

Proof. Let $S \in (\{0, 1\}^n)^t$ be a set of queries of size $|S| \leq k - 1$. Let $Q, Q' \in S$, and let $J = \{j \in [t] : Q_j = Q'_j\}$. We show below that

$$\Pr_{\mathcal{C}_{Q, Q'}^{(t)}} \left[\mathcal{C}^{(t)}(Q)_j = \mathcal{C}^{(t)}(Q')_j \quad \forall j \in J \right] \geq 1 - \varepsilon .$$

Consider the set of queries $S' = S \cup \{Q''\}$, where $Q'' = Q + Q'$. In particular, $Q''_j = 0^n$ for all $j \in J$. By the assumption of the claim we get that $\Pr[\mathcal{L}^{(t)}(Q''_j) = 0 \quad \forall j \in J] \geq 1 - \varepsilon$. Therefore, using the assumption that $\mathcal{L}^{(t)}$ is $(1 - \varepsilon)$ -linear it follows that

$$\begin{aligned} \Pr[\mathcal{C}^{(t)}(Q)_j \neq \mathcal{C}^{(t)}(Q')_j \quad \forall j \in J] &\geq \Pr[\mathcal{L}^{(t)}(Q)_j + \mathcal{L}^{(t)}(Q')_j = \mathcal{L}^{(t)}(Q''_j) \wedge \mathcal{L}^{(t)}(Q''_j) = 0 \quad \forall j \in J] \\ &\geq 1 - 2\varepsilon . \end{aligned}$$

Therefore, $\mathcal{L}^{(t)}$ is $(1 - 2\varepsilon)$ -consistent, as required. \square

2.7 Flattening of a t -nsPCP

Below we define the *flattening* operation, which transforms a given t -repeated proof into a non-repeated proof in the natural way. Namely, given a query set S to the non-repeated proof, we create a vector Q^S containing all the elements of S , query the repeated proof on Q^S , and respond according to the received answer.

Definition 2.17. Let $\mathcal{F}^{(t)}: D^t \rightarrow \Sigma^t$ be a k -non-signaling t -repeated proof. Define the flattening of $\mathcal{F}^{(t)}$, denoted by $\tilde{\mathcal{F}} = \text{Flat}[\mathcal{F}^{(t)}]: D \rightarrow \Sigma$ as follows. For a query set $S = \{q_1, \dots, q_s\} \subseteq D$ of size $s \leq t$, define a vector Q^S whose first s entries are (q_1, \dots, q_s) and the rest are set arbitrarily, query $\mathcal{F}^{(t)}$ on the single query Q^S , and let the distribution of $\tilde{\mathcal{F}}(S)$ be

$$\tilde{\mathcal{F}}(S) = (\mathcal{F}^{(t)}(Q^S)_1, \dots, \mathcal{F}^{(t)}(Q^S)_s) .$$

Claim 2.18. Let $\mathcal{C}^{(t)}: D^t \rightarrow \Sigma^t$ be a k -non-signaling function that is permutation folded and $(1 - \varepsilon)$ -consistent for $k \geq 2$. Then $\tilde{\mathcal{F}} = \text{Flat}[\mathcal{C}^{(t)}]$ is a (ε, t) -non-signaling function.

Furthermore, fix a query $Q = (w_1, \dots, w_t) \in D^t$ for $\mathcal{C}^{(t)}$, a query set $S \subseteq D$ of size s for \mathcal{F} , also let $1 \leq \ell \leq t$ such that w_1, \dots, w_ℓ are distinct and $w_j \in S$ for all $j \in [\ell]$. Then, the distribution of $\mathcal{F}_S(\{w_1, \dots, w_\ell\})$ and $(\mathcal{C}^{(t)}(Q)_1, \dots, \mathcal{C}^{(t)}(Q)_\ell)$ are ε -close in total variation distance.

Proof. To prove that $\tilde{\mathcal{F}}$ is (ε, t) -non-signaling function let $S, T \in D$ be two sets of queries, and suppose $S \cap T = \{w_1, \dots, w_\ell\}$. We want to show that for any event $E \subseteq \Sigma^{S \cap T}$ it holds that

$$\left| \Pr_{\tilde{\mathcal{F}}_S}[\tilde{\mathcal{F}}_S|_{S \cap T} \in E] - \Pr_{\tilde{\mathcal{F}}_T}[\tilde{\mathcal{F}}_T|_{S \cap T} \in E] \right| \leq \varepsilon . \quad (1)$$

Define $Q^S, Q^T \in D^t$ as in Definition 2.17, let $\pi, \pi' \in S_t$ be permutations such that for all $j \in [\ell]$ it holds that $\pi(Q^S)_j = \pi'(Q^T)_j = w_j$. By non-signaling and permutation invariance of $\mathcal{C}^{(t)}$, if we query it on $\{\pi(Q^S), \pi'(Q^T)\}$ we have:

$$\begin{aligned} \Pr_{\tilde{\mathcal{F}}_S}[\tilde{\mathcal{F}}_S|_{S \cap T} \in E] &= \Pr \left[\left(\mathcal{C}^{(t)}(\pi_1(Q^S)), \dots, \mathcal{C}^{(t)}(\pi_\ell(Q^S)) \right) \in E \right] \\ \Pr_{\tilde{\mathcal{F}}_T}[\tilde{\mathcal{F}}_T|_{S \cap T} \in E] &= \Pr \left[\left(\mathcal{C}^{(t)}(\pi'_1(Q^T)), \dots, \mathcal{C}^{(t)}(\pi'_\ell(Q^T)) \right) \in E \right] . \end{aligned}$$

Then, by Claim 2.15 we get the following:

$$\left| \Pr \left[\left(\mathcal{C}^{(t)}(\pi_1(Q^S)), \dots, \mathcal{C}^{(t)}(\pi_\ell(Q^S)) \right) \in E \right] - \Pr \left[\left(\mathcal{C}^{(t)}(\pi'_1(Q^T)), \dots, \mathcal{C}^{(t)}(\pi'_\ell(Q^T)) \right) \in E \right] \right| \leq \varepsilon$$

which proves Eq. (1). Therefore, $\tilde{\mathcal{F}}$ is a (ε, t) -non-signaling function.

Next we prove the second part of the claim. Given S , define $Q^S \in D^t$ as in Definition 2.17, and consider the query set $\{Q^S, Q\}$ to $\mathcal{C}^{(t)}$. Since $\mathcal{C}^{(t)}$ is permutation folded, we may assume that $Q_j = Q_j^S = w_j$ for all $j \in [\ell]$. Therefore, for any $E \subseteq \Sigma^\ell$ we have:

$$\begin{aligned} & \left| \Pr \left[\left(\tilde{\mathcal{F}}_S(w_1), \dots, \tilde{\mathcal{F}}_S(w_\ell) \right) \in E \right] - \Pr \left[\left(\mathcal{C}^{(t)}(Q)_1, \dots, \mathcal{C}^{(t)}(Q)_\ell \right) \in E \right] \right| \\ &= \left| \Pr \left[\left(\mathcal{C}^{(t)}(Q^S)_1, \dots, \mathcal{C}^{(t)}(Q^S)_\ell \right) \in E \right] - \Pr \left[\left(\mathcal{C}^{(t)}(Q)_1, \dots, \mathcal{C}^{(t)}(Q)_\ell \right) \in E \right] \right| , \end{aligned}$$

which is upper bounded by ε by Claim 2.15. This completes the proof of Claim 2.18 \square

The following claim follows rather immediately from Claim 2.18 above.

Claim 2.19. *Let $k \geq 4$, and let $\mathcal{L}^{(t)}: (\{0, 1\}^n)^t \rightarrow \{0, 1\}^t$ be a k -non-signaling function that is permutation folded, $(1 - \varepsilon_1)$ -linear, and $(1 - \varepsilon_2)$ -consistent. Then $\tilde{\mathcal{L}} = \text{Flat}[\mathcal{L}^{(t)}]$ is a (non-repeated) (ε_2, t) -non-signaling $(1 - \varepsilon_1 - 3\varepsilon_2)$ -linear function.*

Proof. By applying Claim 2.18 on $\mathcal{L}^{(t)}$, we get that $\tilde{\mathcal{L}} = \text{Flat}[\mathcal{L}^{(t)}]$ is a (ε_2, t) -non-signaling function. Next we prove that $\tilde{\mathcal{L}}$ is $(1 - (\varepsilon_1 + 3\varepsilon_2))$ -linear. Fix $x, y \in \{0, 1\}^n$, and let $S \subseteq \{0, 1\}^n$ be a query set for $\tilde{\mathcal{L}}$ such that $\{x, y, x + y\} \subseteq S$. We want to prove that

$$\Pr[\tilde{\mathcal{L}}(x) + \tilde{\mathcal{L}}(y) = \tilde{\mathcal{L}}(x + y)] \geq 1 - \varepsilon_1 - 3\varepsilon_2 . \quad (2)$$

Let Q^S be as in Definition 2.17. By the permutation folding property of $\mathcal{L}^{(t)}$ we may assume that the first three coordinates of Q^S are $x, y, x + y$. That is $Q_1^S = x, Q_2^S = y$, and $Q_3^S = x + y$.

By definition of Q^S we have $\Pr[\tilde{\mathcal{L}}(x) + \tilde{\mathcal{L}}(y) = \tilde{\mathcal{L}}(x + y)] = \Pr[\mathcal{L}^{(t)}(Q^S)_1 + \mathcal{L}^{(t)}(Q^S)_2 = \mathcal{L}^{(t)}(Q^S)_3]$. Consider now the vectors $Q^x = (x, 0^n, 0^n, \dots, 0^n)$, $Q^y = (y, 0^n, 0^n, \dots, 0^n)$, and $Q^{x+y} = (x + y, 0^n, 0^n, \dots, 0^n)$. Since $\mathcal{L}^{(t)}$ is $(1 - \varepsilon_1)$ -linear, we get that $\Pr[\mathcal{L}^{(t)}(Q^x) + \mathcal{L}^{(t)}(Q^y) = \mathcal{L}^{(t)}(Q^{x+y})] \geq 1 - \varepsilon_1$. Since $\mathcal{L}^{(t)}$ is $(1 - \varepsilon_2)$ -consistent, it follows that

$$\begin{aligned} \Pr[\tilde{\mathcal{L}}(x) + \tilde{\mathcal{L}}(y) = \tilde{\mathcal{L}}(x + y)] &= \Pr[\mathcal{L}^{(t)}(Q^S)_1 + \mathcal{L}^{(t)}(Q^S)_2 = \mathcal{L}^{(t)}(Q^S)_3] \\ &\geq \Pr[\mathcal{L}^{(t)}(Q^x) + \mathcal{L}^{(t)}(Q^y) = \mathcal{L}^{(t)}(Q^{x+y})] - 3\varepsilon_2 \\ &\geq 1 - \varepsilon_1 - 3\varepsilon_2 , \end{aligned}$$

as required. \square

3 Main result

In this section we formally state the main result of the paper. In order to describe the result we need to first state the hypothesis conditioned on which our main theorem holds.

Hypothesis 2. *Fix integers n and $k \leq 2^n$, and let $\varepsilon \in (0, 1)$. For any (ε, k) -almost non-signaling function $\mathcal{F}: \{0, 1\}^n \rightarrow \{0, 1\}$ that is $(1 - \varepsilon)$ -linear there exists a k' -non-signaling function $\mathcal{F}': \{0, 1\}^n \rightarrow \{0, 1\}$ such that $\Delta_4(\mathcal{F}, \mathcal{F}') \leq \varepsilon'$, where $k' \geq k^{\text{hyp}}$ for some positive absolute constant $\varepsilon_{\text{hyp}} > 0$, and $\varepsilon' = \varepsilon'_{\text{hyp}}(\varepsilon)$ is some function that depends only on ε such that $\varepsilon'_{\text{hyp}}(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.*

Remark 3.1. We make two remarks regarding the hypothesis.

- A statement analogous to Hypothesis 2 has been proven in [CMS20], showing that there exist a k -non-signaling function $\mathcal{F}' : \{0, 1\}^n \rightarrow \{0, 1\}$ such that $\Delta_k(\mathcal{F}, \mathcal{F}') \leq O(4^k \cdot \varepsilon)$. The multiplicative factor of 4^k is too large, which makes it insufficient for our applications.
- For our applications, we need a much weaker version of Hypothesis 2. We elaborate more on the hypothesis in Appendix A.

For a computable function $N : \mathbb{N} \rightarrow \mathbb{N}$ we denote by $\text{SIZE}(N)$ the complexity class of all languages L having a uniform family of boolean circuits $(C_n : \{0, 1\}^n \rightarrow \{0, 1\})_{n \in \mathbb{N}}$ of maximum fan-in 2 with AND, OR, and NOT gates, such that C_n has at most $N(n)$ wires for all $n \in \mathbb{N}$.¹

Theorem 2 (Main theorem). *Assuming Hypothesis 2 every language $L \in \text{SIZE}(N)$ has an input oblivious nsPCP verifier that on input of length n uses $\tilde{O}(N^2)$ random bits, makes $O(1)$ queries to proofs over the alphabet $\Sigma = \{0, 1\}^{\text{polylog}(N)}$, and is sound against $O(1)$ -non-signaling proofs. The query sampler runs in time $\tilde{O}(N^2)$, and the decision predicate runs in time $O(n \cdot \text{polylog}(N))$. That is,*

$$\text{SIZE}(N) \subseteq \text{nsPCP} \left[\begin{array}{l} \text{soundness error: } 1 - \Omega(1) \\ \text{randomness: } \tilde{O}(N^2) \\ \text{proof length: } 2^{\tilde{O}(N^2)} \\ \text{query complexity: } 4 \\ \text{locality: } O(1) \\ \text{query sampler time: } \tilde{O}(N^2) \\ \text{decision predicate time: } O(n \cdot \text{polylog}(N)) \end{array} \right].$$

It is clear that Theorem 1 follows from Theorem 2 since $\text{DTIME}(T) \subseteq \text{SIZE}(O(T \log(T)))$.

4 The PCP construction

In this section we formally describe our PCP construction. In one sentence, the PCP verifier gets a permutation invariant proof $\mathcal{F}^{(t)} : (\{0, 1\}^{N^2})^t \rightarrow \{0, 1\}^t$, runs on it linearity test, direct product test, and the parallel repetition of the ALMSS verifier, and accepts if and only if all tests accepts.

We start by recalling the setting of the PCP verifier of [ALMSS98] (the “linear ALMSS verifier”). Let $L \in \text{SIZE}(N)$ be a language, and let $\{C_n\}_{n \in \mathbb{N}}$ be a uniform family of boolean circuits with $N = N(n)$ wires that decides L . That is, for all inputs $x \in \{0, 1\}^n$ of length n it holds that $C_n(x) = 1$ if and only if $x \in L$.

For a given length n let $C := C_n$ be the circuit corresponding to the computation on inputs of length n . The computation of C on the input x is viewed as a system of $M := N + 1$ constraints $\{P_j(\mathbf{w}) = c_j\}_{j \in [M]}$ over N boolean variables $\mathbf{w} = (w_1, \dots, w_N) \in \{0, 1\}^N$, where $P_1, \dots, P_M : \{0, 1\}^N \rightarrow \{0, 1\}$ are quadratic polynomials (each involving at most three variables in \mathbf{w}) and c_1, \dots, c_M are boolean constants. Each variable represents the value of one of the wires of C during the computation on the input x . In particular, the first n variables, w_1, \dots, w_n , correspond to the n input wires, and the variable w_N corresponds to the output wire. The constraints are of three types:

Input consistency: For every $j \in \{1, \dots, n\}$, $P_j(\mathbf{w}) := w_i$ and $c_j := x_j$.

¹Note that our complexity measure for the size of a circuit is the number of wires, (and not the number of gates, which is more standard) as this measure directly affects the complexity of the PCP construction. However, for circuits with bounded fan-in, the two quantities are equal up to a multiplicative constant factor.

Gate consistency: For every $j \in \{n + 1, \dots, N\}$,

- If the wire represented by the variable w_j is an output of an AND gate g , where the inputs to g are given by w_{j_1}, w_{j_2} , then $P_j(\mathbf{w}) := w_j - w_{j_1} \cdot w_{j_2}$ and $c_j := 0$.
- If the wire represented by the variable w_j is an output of an OR gate g , where the inputs to g are given by w_{j_1}, w_{j_2} , then $P_j(\mathbf{w}) := w_j - w_{j_1} - w_{j_2} + w_{j_1} \cdot w_{j_2}$ and $c_j := 0$.
- If the wire represented by the variable w_j is an output of a NOT gate g , where the input to g is given by w_{j_1} , then $P_j(\mathbf{w}) := w_j - w_{j_1}$ and $c_j := 1$.

Accepting output: $P_M(\mathbf{w}) := w_N$ and $c_M := 1$.

We overload notation, and use P_j to also denote the upper triangular matrix in $\{0, 1\}^{N^2}$ with $P_j(\mathbf{w}) = \langle P_j, \mathbf{w} \otimes \mathbf{w} \rangle$. That is, if $P_j(\mathbf{w}) = \sum_{i=1}^N a_i w_i + \sum_{1 \leq i < i' \leq N} a_{i,i'} w_i w_{i'}$, then the corresponding matrix has a_i in the diagonal entry (i, i) and $a_{i,i'}$ in the entry (i, i') , for $1 \leq i < i' \leq N$. Also, for $a \in \{0, 1\}^N$, denote by D_a the diagonal matrix in $\{0, 1\}^{N^2}$ whose diagonal is a .

4.1 The linear ALMSS verifier

The *linear* ALMSS verifier of [ALMSS98] is defined as follows.

Algorithm 1: The linear ALMSS verifier

Explicit input: A circuit $C: \{0, 1\}^n \rightarrow \{0, 1\}$ with N wires, and input $x \in \{0, 1\}^n$ to C .

Oracle access: A k -non-signaling linear function $\mathcal{L}^{(t)}: \{0, 1\}^{N^2} \rightarrow \{0, 1\}$.

- 1 Use the circuit C and input x to construct the matrices $P_1, \dots, P_M \in \{0, 1\}^{N^2}$ and constants $c_1, \dots, c_M \in \{0, 1\}$ representing the computation of C on x .
 - 2 Sample $u, v \in \{0, 1\}^N$ and $s \in \{0, 1\}^M$ uniformly and independently at random.
 - 3 Query the oracle \mathcal{L} on the 4-element set $S = \{D_u, D_v, u \otimes v, \sum_{j=1}^M s_j P_j\}$.
 - 4 **return** ACCEPT if and only if $\mathcal{L}(D_u)\mathcal{L}(D_v) = \mathcal{L}(u \otimes v)$ and $\mathcal{L}(\sum_{j=1}^M s_j P_j) = \sum_{j=1}^M s_j c_j$.
-

That is, the verifier makes 4 queries to a linear proof $\mathcal{L}: \{0, 1\}^{N^2} \rightarrow \{0, 1\}$ (of exponential length).

Completeness. Completeness of the ALMSS verifier is the same as in the classical setting. Indeed, $C(x) = 1$, then the classical proof defined by the design is accepted with probability 1.

Soundness. For soundness, Chiesa et al. [CMS19] proved that this construction is indeed sound against linear $O(\log N)$ -non-signaling proofs with soundness error bounded below 1.

Theorem 4.1 (Theorem 6 in [CMS19]). *For any language $L \in \text{SIZE}(N)$ there is an input oblivious PCP system, where the verifier gets as an explicit input a circuit C of size $N = N(n)$ deciding L and an input $x \in \{0, 1\}^n$, and an oracle access to a linear proof $\pi: \{0, 1\}^{O(N^2)} \rightarrow \{0, 1\}$. The verifier uses $O(N)$ random coins, makes 4 queries to the proof that are independent of x . If $x \in L$, then there exists a (classical) proof that causes the verifier to accept with probability 1. If $x \notin L$, then for any $O(\log(N))$ -non-signaling linear proof the verifier to accept with probability at most $39/40$.*

That is, we have

$$\text{SIZE}(N) \subseteq \text{nsLPCP} \left[\begin{array}{ll} \text{soundness error:} & 39/40 \\ \text{randomness:} & O(N) \\ \text{proof length:} & 2^{O(N^2)} \\ \text{query complexity:} & 4 \\ \text{locality:} & O(\log N) \\ \text{query sampler time:} & O(N^2) \\ \text{decision predicate time:} & O(n) \end{array} \right].$$

4.2 Parallel repetition of the linear ALMSS verifier

Next, we consider the t -repeated parallel repetition of the linear ALMSS verifier. Specifically, the verifier samples t independent sets of queries, makes 4 queries to the PCP over the alphabet $\{0, 1\}^t$, and accepts if and only if all t sets of answers satisfy the basic linear ALMSS verifier. Formally, the t -repetition of the linear ALMSS verifier is defined as follows.

Algorithm 2: The t -repeated linear ALMSS verifier

Explicit input : A circuit $C: \{0, 1\}^n \rightarrow \{0, 1\}$ with N wires, and input $x \in \{0, 1\}^n$ to C .

Oracle access : A t -repeated k -non-signaling linear function $\mathcal{L}^{(t)}: (\{0, 1\}^{N^2})^t \rightarrow \{0, 1\}^t$.

- 1 Construct the matrices $P_1, \dots, P_M \in \{0, 1\}^{N^2}$ and constants $c_1, \dots, c_M \in \{0, 1\}$, representing the computation of C on x .
 - 2 Sample $u^{(1)}, \dots, u^{(t)}, v^{(1)}, \dots, v^{(t)} \in \{0, 1\}^N$ and $s^{(1)}, \dots, s^{(t)} \in \{0, 1\}^M$ independently and uniformly at random.
 - 3 Let $Q_1 = (D_{u^{(i)}})_{i \in [t]}$; $Q_2 = (D_{v^{(i)}})_{i \in [t]}$; $Q_3 = (u^{(i)} \otimes v^{(i)})_{i \in [t]}$; $Q_4 = (\sum_{j=1}^M s_j^{(i)} P_j)_{i \in [t]}$.
 - 4 Query the oracle $\mathcal{L}^{(t)}$ on the 4-element set $S = \{Q_1, Q_2, Q_3, Q_4\}$.
 - 5 Check that $\mathcal{L}^{(t)}(Q_1)_i \cdot \mathcal{L}^{(t)}(Q_2)_i = \mathcal{L}^{(t)}(Q_3)_i \quad \forall i \in [t]$.
 - 6 Check that and $\mathcal{L}^{(t)}(Q_4)_i = \sum_{j=1}^M s_j^{(i)} c_j \quad \forall i \in [t]$.
 - 7 **return** ACCEPT if and only if in the two previous steps all equalities hold.
-

Here, just as in the previous case, the verifier makes 4 queries to a linear proof. However, now the proof is over the alphabet $\{0, 1\}^t$.

Completeness. Completeness of the repeated linear ALMSS verifier is clear. Indeed, if $C(x) = 1$, then we can take the parallel repetition of the intended classical linear proof, and it will satisfy the repeated linear ALMSS verifier with probability 1.

Soundness. For soundness we prove in Section 6 that if $t \geq O(\log(N))$, then the verifier is sound against $O(1)$ -non-signaling proofs that are *linear* and *consistent*. The proof works by reducing to the soundness of the *non-repeated* linear ALMSS verifier. Specifically, we consider a circuit C and an input x to C , and consider a t -repeated k -non-signaling proof that is accepted with probability at least γ . We show that if the proof is linear and consistent, then its flattening is a t -non-signaling (non-repeated) linear proof that satisfies the non-repeated ALMSS verifier with the same probability. Therefore, if $\gamma \geq 39/40$, then by Theorem 4.1 we conclude that $C(x) = 1$.

4.3 From linear PCPs to standard PCPs using linearity and consistency testing

So far we have assumed that the given non-signaling proof is linear and consistent. Below we show how to discard this assumption, and prove Theorem 2 by constructing a PCP verifier that is sound

against arbitrary proofs. This is done by running (the parallel repetition of) the linearity test, the consistency test, and then feeding (the self-corrected version of) the proof to the linear repeated ALMSS verifier from Algorithm 2. We describe the verifier formally below.

Algorithm 3: The $2t$ -repeated ALMSS verifier + consistency test

Explicit input : A circuit $C: \{0, 1\}^n \rightarrow \{0, 1\}$ with N wires, and input $x \in \{0, 1\}^n$ to C

Oracle access : A $2t$ -repeated k -non-signaling linear function $\mathcal{F}^{(2t)}: (\{0, 1\}^{N^2})^{2t} \rightarrow \{0, 1\}^{2t}$

- 1 Sample uniformly random $X, Y \in (\{0, 1\}^{N^2})^{2t}$.
 - 2 Sample uniformly random $W, Z_1, Z_2 \in (\{0, 1\}^{N^2})^t$.
 - 3 Sample the four queries $Q_1, Q_2, Q_3, Q_4 \in (\{0, 1\}^{N^2})^t$ of the t -repeated linear ALMSS verifier from Algorithm 2. and let $D_{\text{LIN}}: (\{0, 1\}^t)^4 \rightarrow \{ACCEPT, REJECT\}$ be the corresponding predicate.
 - 4 Define $\widehat{\mathcal{F}}^{(t)}: (\{0, 1\}^{N^2})^t \rightarrow \{0, 1\}^t$ as in Definition 7.3, which makes two queries to $\mathcal{F}^{(2t)}$ for every query to $\widehat{\mathcal{F}}^{(t)}$.
 - 5 Sample an input $S \subseteq (\{0, 1\}^{N^2})^{2t}$ to $\mathcal{F}^{(2t)}$ corresponding to querying $\widehat{\mathcal{F}}^{(t)}$ on $\{Q_1, Q_2, Q_3, Q_4\}$.
 - 6 Query $\mathcal{F}^{(2t)}$ on the set $S \cup \{X, Y, Z + Y\} \cup \{[W; Z_1], [W; Z_2]\}$.
 - 7 *Linearity test:* Check that $\mathcal{F}^{(2t)}(X) + \mathcal{F}^{(2t)}(Y) = \mathcal{F}^{(2t)}(X + Y)$.
 - 8 *Consistency test:* Check that $\mathcal{F}^{(2t)}([W; Z_1])|_W = \mathcal{F}^{(2t)}([W; Z_2])|_W$.
 - 9 *Linear PCP verifier:* Interpret $\mathcal{F}^{(2t)}(S)$ as the answers of $\widehat{\mathcal{F}}^{(t)}$ on the query set $(\{Q_1, Q_2, Q_3, Q_4\})$, and check that $\widehat{\mathcal{F}}^{(t)}(\{Q_1, Q_2, Q_3, Q_4\})$ satisfies D_{LIN} .
 - 10 **return** ACCEPT if and only if all three steps above accept.
-

That is, the verifier is almost the parallel repetition of the classical ALMSS verifier. The only difference is that our verifier makes 2 additional queries for the consistency test.

Completeness. Completeness of the repeated ALMSS verifier is clear, as by design the expected proof is linear, and hence $\mathcal{F}^{(t)}$ satisfies the linearity constraint with probability 1. Furthermore, it follows that $\widehat{\mathcal{F}}^{(t)}$ is equal to $\mathcal{F}^{(t)}$, and thus the predicate D_{LIN} is also satisfied with probability 1.

Soundness. We prove soundness of the PCP system in Algorithm 3 in Section 8. Specifically, we use Hypothesis 2, and prove that for $t \geq \text{polylog}(N)$ and $k \geq O(1)$, the verifier is sound against $O(1)$ -non-signaling *linear* proofs. The proof works, again, by reducing to the soundness of the *non-repeated* linear ALMSS verifier. Specifically, we consider a circuit C and an input x to C , and prove that if the PCP verifier accepts a t -repeated k -non-signaling proof $1 - \varepsilon$, then its flattening (or rather the flattening of its self-correction) is a (ε, t) -non-signaling (non-repeated) proof $\mathcal{F}: \{0, 1\}^{N^2} \rightarrow \{0, 1\}$ that is $(1 - \varepsilon)$ -linear, and it satisfies the *non-repeated* linear ALMSS verifier with high probability. By applying Hypothesis 2 we obtain a $t^{\Omega(1)}$ -non-signaling (non-repeated) proof that is $(1 - \varepsilon')$ -linear and satisfies the *non-repeated* linear ALMSS verifier from Algorithm 1 with high probability. By applying a result from [CMS20] we conclude that \mathcal{F} is close to a linear $t^{\Omega(1)}$ -non-signaling (non-repeated) proof $\widehat{\mathcal{F}}$ that also satisfies the linear ALMSS verifier with high probability, and thus, by Theorem 4.1 it follows that $C(x) = 1$. See Section 8 for details.

5 Proof overview: Soundness

In this section we give an overview of the soundness analysis of the parallel repetition of the PCP verifier from Algorithm 3. Before describing the actual proof, we first consider soundness against *structured* proofs. Indeed, this is a common approach in the analysis of PCP systems. Specifically, we show first that the PCP verifier from Algorithm 2 is sound against such structured proofs. Then we use local testing and self-correction to show the proofs that are accepted by the verifier with high probability satisfy the desired properties.

Soundness of the linear t -repeated ALMSS verifier. Fix a circuit $C: \{0, 1\}^N \rightarrow \{0, 1\}$ and let $x \in \{0, 1\}^N$ be an input to C . Consider a $2t$ -repeated linear ALMSS verifier for $C(x)$, and suppose that $\mathcal{L}^{(2t)}: (\{0, 1\}^{N^2})^{2t} \rightarrow \{0, 1\}^{2t}$ is a $2t$ -repeated linear and consistent proof such that $2t$ -repeated linear ALMSS verifier accepts $\mathcal{L}^{(2t)}$ with high probability. According to Claim 2.19 it follows that we can “flatten” $\mathcal{L}^{(2t)}$ into a t -non-signaling *linear* proof $\mathcal{L}: \{0, 1\}^{N^2} \rightarrow \{0, 1\}$. We show that since $\mathcal{L}^{(2t)}$ is accepted with high probability by the $2t$ -repeated linear verifier, it follows that the (non-repeated) linear ALMSS verifier accepts \mathcal{L} with high probability. Therefore, by applying the result of [CMS19] it follows that if $t > c \log(N)$ for some (sufficiently large) constant c , then $C(x) = 1$. See Section 6 for details.

General proofs - forcing consistency using extended linearity test. Next we prove soundness of the general (i.e., non-linear) parallel repetition PCP against arbitrary $O(1)$ -non-signaling proofs. The general approach is analogous to the approach used for analyzing PCPs, specifically, we first run a test that “forces” the proof to be (close to) linear and consistent, and then apply the analysis of the linear proof in the previous paragraph.

More concretely, we fix a circuit $C: \{0, 1\}^N \rightarrow \{0, 1\}$ and an input $x \in \{0, 1\}^N$ to C , and consider a $2t$ -repeated (non-linear) ALMSS verifier for $C(x)$. Suppose that $\mathcal{F}^{(2t)}: (\{0, 1\}^{N^2})^{2t} \rightarrow \{0, 1\}^{2t}$ is a $2t$ -repeated proof such that $2t$ -repeated ALMSS verifier accepts $\mathcal{F}^{(2t)}$ with high probability. Our goal is to prove that $C(x) = 1$, and our high level strategy to show it is the following:

1. First we assume that the proof is permutation invariant as in Definition 2.8.
2. Suppose the repeated ALMSS verifier accepts $\mathcal{F}^{(2t)}$ with high probability $1 - \varepsilon$. In particular, this implies that $\mathcal{F}^{(2t)}$ passes linearity test with at least same probability.
3. By applying the self-correction procedure, we obtain the self-correction of $\mathcal{F}^{(2t)}$. The self-correction of $\mathcal{F}^{(2t)}$, denoted by $\widehat{\mathcal{F}}^{(t)}$, is a t -repeated \widehat{k} -non-signaling proof for $\widehat{k} = \Omega(k)$ such that in order to make one query to $\widehat{\mathcal{F}}^{(t)}$ we make $O(1)$ queries to $\mathcal{F}^{(2t)}$. We prove that $\widehat{\mathcal{F}}^{(t)}$ satisfies the following two properties.

- (a) $\widehat{\mathcal{F}}^{(t)}$ is $(1 - O(\varepsilon))$ -linear, i.e., for all $X, Y \in (\{0, 1\}^{N^2})^t$ it holds that $\Pr[\widehat{\mathcal{F}}^{(t)}(X) + \widehat{\mathcal{F}}^{(t)}(Y) = \widehat{\mathcal{F}}^{(t)}(X + Y)]$. That is, the self-correction transforms an *average-case* guarantee about linearity testing $\mathcal{F}^{(2t)}$ into a guarantee that $\widehat{\mathcal{F}}^{(t)}$ satisfies the linearity constraints *for all* $X, Y, X + Y$.
- (b) $\widehat{\mathcal{F}}^{(t)}$ is $(1 - O(\varepsilon))$ -consistent, i.e., for all $Q, Q' \in (\{0, 1\}^{N^2})^t$ with high probability $\widehat{\mathcal{F}}^{(t)}(Q)_j = \widehat{\mathcal{F}}^{(t)}(Q')_j$ for all $j \in [t]$ such that $Q_j = Q'_j$. Here also, the *average-case* guarantee of the consistency test is converted into the *worst-case* guarantee holding for all $Q, Q' \in (\{0, 1\}^{N^2})^t$.

4. Next, we let $\tilde{\mathcal{F}} = \text{Flat}[\widehat{\mathcal{F}^{(t)}}]$ be the flattening of $\widehat{\mathcal{F}^{(t)}}$. By Claim 2.18, $\tilde{\mathcal{F}}: \{0, 1\}^{N^2} \rightarrow \{0, 1\}$ is an almost linear $(O(\varepsilon), t)$ -no-signaling function. Furthermore, using the fact that $\widehat{\mathcal{F}^{(t)}}$ is $(1 - O(\varepsilon))$ -linear and is accepted by the repeated ALMSS verifier from Algorithm 2 with high probability, we prove that $\tilde{\mathcal{F}}$ is accepted by the (non-repeated) ALMSS verifier from Algorithm 1 with high probability.

At this point we would like to apply Theorem 4.1 on $\widehat{\mathcal{F}^{(t)}}$, and say that since $\tilde{\mathcal{F}}$ is accepted by the ALMSS verifier with high probability, it follows that $C(x) = 1$. However, the difficulty in applying Theorem 4.1 is that $\tilde{\mathcal{F}}$ is not necessarily non-signaling, but only *almost non-signaling* (see Definition 2.5 for reference). In order to still apply this result we use Hypothesis 2 to “round” $\tilde{\mathcal{F}}$ into a non-signaling proof, and then apply Theorem 4.1 to conclude that $C(x) = 1$. Specifically, we do the following.

5. Assuming Hypothesis 2, there exist a t' -no-signaling function \mathcal{F} , which is close to $\tilde{\mathcal{F}}$. In particular, \mathcal{F} is an almost linear non-signaling function.
6. Using the result of [CMS19] on linearity testing we get that for some $\bar{k} = \Omega(\sqrt{t'})$ there exists a \bar{k} -non-signaling linear proof \mathcal{L} that is $O(q\varepsilon)$ -close to \mathcal{F}' on queries sets of size at most $q \leq \bar{k}$.
7. By our choice of parameters, the locality of \mathcal{L} is $\bar{k} = \Omega(\sqrt{t'}) > C \log(N)$, and by the previous item the linear ALMSS verifier accepts \mathcal{L} with high probability. Therefore, using Theorem 4.1 we conclude that $C(x) = 1$.

This completes the overview of the proof. Below we describe each step in detail.

6 Soundness of the linear PCP verifier against structured proofs

In this section we prove that the t -repeated linear PCP verifier from Algorithm 2 is sound against linear consistent proofs. Specifically, we prove the following theorem.

Theorem 3. *Fix a circuit $C: \{0, 1\}^n \rightarrow \{0, 1\}$ with N wires, and input $x \in \{0, 1\}^n$ to C . Let $k \geq 4$, and t be positive integers such that $t \geq K \log(N)$ for some sufficiently large constant $K > 0$. Let $\mathcal{L}^{(t)}: (\{0, 1\}^{N^2})^t \rightarrow \{0, 1\}^t$ be a k -non-signaling t -repeated linear consistent proof, and suppose that the t -repeated ALMSS verifier accepts $\mathcal{L}^{(t)}$ with probability $\geq 39/40$. Then $C(x) = 1$.*

Proof. Let $\mathcal{L}^{(t)}: (\{0, 1\}^{N^2})^t \rightarrow \{0, 1\}^t$ be a k -non-signaling t -repeated linear consistent proof, and suppose that the t -repeated ALMSS verifier accepts $\mathcal{L}^{(t)}$ with probability $\geq 39/40$.

Let $\tilde{\mathcal{L}} = \text{Flat}[\mathcal{L}^{(t)}]$ be the flattening of $\mathcal{L}^{(t)}$ as per Definition 2.17. Since $\mathcal{L}^{(t)}$ is linear and consistent, it follows by Claim 2.19 that $\tilde{\mathcal{L}}$ is a t -non-signaling linear function.

Next we show that the non-repeated linear ALMSS verifier accepts $\tilde{\mathcal{L}}$ with probability $> 39/40$, and hence, by Theorem 4.1 it follows that $C(x) = 1$.

Indeed, consider the random choices of $u, v \in \{0, 1\}^N$ and $s \in \{0, 1\}^M$ in Algorithm 1, and let $Q^* \in (\{0, 1\}^{N^2})^t$ be a query to $\mathcal{L}^{(t)}$ that contains the four queries $Q_{ALMSS} = \{D_u, D_v, u \otimes v, \sum_{j=1}^M s_j P_j\}$ in its first 4 coordinates. That is, $Q_1^* = D_u$, $Q_2^* = D_v$, $Q_3^* = u \otimes v$, and $Q_4^* = \sum_{j=1}^M s_j P_j$. Denoting by D_{LIN} the predicated in Algorithm 1, by definition of the flattening operation we have

$$\Pr[\text{D}_{\text{LIN}}(\tilde{\mathcal{L}}(Q_{ALMSS}))] = \Pr[\text{D}_{\text{LIN}}(\mathcal{L}^{(t)}(Q^*)_1, \mathcal{L}^{(t)}(Q^*)_2, \mathcal{L}^{(t)}(Q^*)_3, \mathcal{L}^{(t)}(Q^*)_4) = 1] . \quad (3)$$

Next, consider the random choices of $u^{(1)}, \dots, u^{(t)}, v^{(1)}, \dots, v^{(t)} \in \{0, 1\}^N$ and $s^{(1)}, \dots, s^{(t)} \in \{0, 1\}^M$ in Algorithm 2, and let $Q_1 = (D_{u^{(i)}})_{i \in [t]}$, $Q_2 = (D_{v^{(i)}})_{i \in [t]}$, $Q_3 = (u^{(i)} \otimes v^{(i)})_{i \in [t]}$, and $Q_4 = (\sum_{j=1}^M s_j^{(i)} P_j)_{i \in [t]}$ be the queries made by the repeated ALMSS verifier.

Since $u^{(1)}, v^{(1)}$ and $s^{(1)}$ are distributed identically to the random choices of $u, v \in \{0, 1\}^N$ and $s \in \{0, 1\}^M$ in Algorithm 1, it follows by consistency of $\mathcal{L}^{(t)}$ that $\Pr[\text{D}_{\text{LIN}}(\mathcal{L}^{(t)}(Q^*)_{\{1,2,3,4\}}) = 1]$ is equal to

$$\Pr[\text{D}_{\text{LIN}}(\mathcal{L}^{(t)}(Q_1)_1, \mathcal{L}^{(t)}(Q_2)_1, \mathcal{L}^{(t)}(Q_3)_1, \mathcal{L}^{(t)}(Q_4)_1) = 1] ,$$

i.e., to the probability that D_{LIN} accepts the responses of $\mathcal{L}^{(t)}$ in the first coordinate of the parallel repetition in Algorithm 2. However, since the verifier in Algorithm 2 accepts $\mathcal{L}^{(t)}$ with probability $\geq 39/40$, it follows in particular, that the first coordinate is accepted with probability $\geq 39/40$, and hence, by Eq. (3) we conclude that

$$\Pr[\text{D}_{\text{LIN}}(\tilde{\mathcal{L}}(Q_{\text{ALMSS}}))] \geq 39/40 ,$$

and hence, by Theorem 4.1 we have $C(x) = 1$. This completes the proof of Theorem 3. \square

7 Testing and self-correcting repeated non-signaling functions

As shown in Section 6, it is rather straightforward to construct a PCP system that is sound against repeated non-signaling proofs that are consistent and linear. Therefore, we would like to make sure that the given proof satisfies these properties. We “enforce” these properties in Algorithm 3 by first running linearity test and consistency test on a given t -repeated non-signaling proof, and then run the linear PCP on the self-correction of the given proof. Next, we show that if the tests accept a given proof with high probability, then its self-correction (almost) satisfies the desired properties, hence reducing the problem to the structured case. In this section we analyze the tests and prove guarantees about the self-correction of any non-signaling function that passes the test with high probability. Then, in Section 8 we use these results on testing and self-correction in order to analyze the PCP system from Algorithm 3.

7.1 Definitions of tests and the self-correction

Testing linearity. Linearity test is a randomized algorithm that given an input function f , queries it on 3 inputs and wishes decides whether f is linear or far from linear. The test was first analyzed in [BLR93]. Bellare et al. in [BCHKS96] simplified the analysis and proved for any boolean function f , the probability that it passes the test is at most $1 - \Delta(f)$, where $\Delta(f)$ is the normalized Hamming distance of f to the closest linear function. Extension of [BLR93] linearity test to general groups and many other closely related problems have been studied since then [AHRS01; SW04; BCLR08; BKSSZ10; DDGKS17]. More recently, [IV12; Vid14] and [CMS20] analyzed the linearity test against quantum strategies and non-signaling strategies. For our setting, when the functions are of the form $\mathcal{F}^{(t)}: (\{0, 1\}^n)^t \rightarrow \{0, 1\}^t$ the test is as follows.

Definition 7.1 (Linearity test [BLR93]). *Let $\mathcal{F}^{(t)}: (\{0, 1\}^n)^t \rightarrow \{0, 1\}^t$ be a t -repeated k -non-signaling function. Linearity test works by uniformly sampling $X, Y \in (\{0, 1\}^n)^t$, querying $\mathcal{F}^{(t)}$ on set $\{X, Y, X + Y\}$, and checking that $\mathcal{F}^{(t)}(X) + \mathcal{F}^{(t)}(Y) = \mathcal{F}^{(t)}(X + Y)$, i.e., that for all $j \in [t]$ it holds that $\mathcal{F}^{(t)}(X)_j + \mathcal{F}^{(t)}(Y)_j = \mathcal{F}^{(t)}(X + Y)_j$.*

In the non-signaling setting, linearity test was analyzed by [CMS20] for *boolean* functions. They proved that any k -non-signaling boolean function \mathcal{F} that passes the linearity test with probability $1 - \varepsilon$ can be self-corrected to a $\lfloor k/2 \rfloor$ -non-signaling function $\widehat{\mathcal{F}}$ that is $2^{O(k)}\varepsilon$ -close to a linear $\lfloor k/2 \rfloor$ -non-signaling function \mathcal{L} . However, we cannot directly apply their result to our setting, as our functions are not boolean. Furthermore, adapting the approach of [CMS20] will give a linear non-signaling function with the guarantee that the distance between $\widehat{\mathcal{F}}$ and a truly linear function \mathcal{L} is at most $2^{O(tk)}\varepsilon$, which is too large for our application.

Testing consistency. Next we consider *consistency test*, whose goal is to check that a given t -repeated non-signaling function is (close to) consistent as per Definition 2.13. The test works as follows.

Definition 7.2 (Consistency test). *Let $\mathcal{F}^{(2t)}: (\{0, 1\}^n)^{2t} \rightarrow \{0, 1\}^{2t}$ be a $2t$ -repeated k -non-signaling function for an integer t . Consistency test chooses $W, Z_1, Z_2 \in (\{0, 1\}^n)^t$ uniformly at random, queries $\mathcal{F}^{(t)}$ on $\{[W; Z_1], [W; Z_2]\}$, and checks that $\mathcal{F}^{(t)}([W; Z_1])|_W = \mathcal{F}^{(t)}([W; Z_2])|_W$.*

Similar tests have been studied in the literature in the context of *Direct product testing* in a long series of work [DR04; IKW12; DS14b; DN17; GCS19].

We prove below that if a $2t$ -repeated k -non-signaling proof $\mathcal{F}^{(2t)}$ passes both the linearity test and the consistency test with probability $1 - \varepsilon$, then its *self-correction* $\widehat{\mathcal{F}}^{(t)}$ is $(1 - O(\varepsilon))$ -linear and $(1 - O(\varepsilon))$ -consistent. That is, $\widehat{\mathcal{F}}^{(t)}$ is close to having the properties we need in order to prove soundness against repeated non-signaling proofs. Next, we discuss the notion of self-correction, and prove if $\mathcal{F}^{(2t)}$ passes the tests with high probability, then $\widehat{\mathcal{F}}^{(t)}$ satisfies the desired properties.

7.2 Self-correction of a t -repeated k -non-signaling function

Below we define the self-correction of a given t -repeated k -non-signaling function $\mathcal{F}^{(t)}$. Observe that if $\mathcal{F}^{(t)}$ passes the linearity test with high probability $1 - \varepsilon$, it does not necessarily imply that it satisfies *all* linearity constraints with high probability, i.e., it does not imply that $\mathcal{F}^{(t)}$ is $(1 - \varepsilon')$ -linear. As a simple example, one may consider the case when $\mathcal{F}^{(t)}$ is a deterministic function that is obtained from a linear function by changing some small fraction of its outputs. The same applies to the consistency test, i.e., satisfying the consistency constraints on average as opposed to satisfying each consistency constraint with high probability.

A standard approach to transform the “average-case” guarantee of the tests into a “point-wise” guarantee is by employing the idea of *self-correction*. Next we define the notion of self-correction suitable for our tests.

Definition 7.3. *Let $\mathcal{F}^{(2t)}: (\{0, 1\}^n)^{2t} \rightarrow \{0, 1\}^{2t}$ be a $2t$ -repeated k -non-signaling function. The **self-correction** of $\mathcal{F}^{(2t)}$, is a t -repeated \widehat{k} -non-signaling function $\widehat{\mathcal{F}}^{(t)}: (\{0, 1\}^n)^t \rightarrow \{0, 1\}^t$, for $\widehat{k} \leq \frac{k}{2}$ defined as follows.*

Given a query $Q \in (\{0, 1\}^n)^t$, in order to sample $\widehat{\mathcal{F}}^{(t)}(Q)$ we uniformly choose $R, W \in (\{0, 1\}^n)^t$, query $\mathcal{F}^{(2t)}$ on the set $\{[R; W], [R + Q; W]\}$, and output the first half of $(\mathcal{F}^{(2t)}([R; W]) + \mathcal{F}^{(2t)}([R + Q; W]))$.

More generally, for a query set $\widehat{S} = \{Q_1, \dots, Q_s\}$ of size $s \leq \widehat{k}$ we sample $R_i, W_i \in (\{0, 1\}^n)^t$ independently, uniformly at random for each $i \in [s]$, query $\mathcal{F}^{(2t)}$ on the set

$$S = \bigcup_{i=1}^s \{[R_i; W_i], [Q_i + R_i; W_i]\} ,$$

and output

$$\widehat{\mathcal{F}}^{(t)}(Q_i)_j := (\mathcal{F}^{(2t)}([R_i; W_i]) + \mathcal{F}^{(2t)}([Q_i + R_i; W_i]))_j \quad \forall j \in [t] .$$

for all $i \in [s]$.

Observe that the self-correction of $\mathcal{F}^{(2t)}$ is indeed a non-signaling function with the appropriate locality parameter. Indeed, this follows immediately from the assumption that $\mathcal{F}^{(2t)}$ is k -non-signaling and the fact that the R_i, W_i 's are uniformly random and independent.

7.3 Self-correction is almost linear and almost consistent

Next we show that if $\mathcal{F}^{(2t)}$ passes both the linearity test and the agreement test with high probability then its self-correction $\widehat{\mathcal{F}}^{(t)}$ is almost linear and almost consistent. Indeed, this average-to-worst-case is a standard step in the analysis of non-signaling PCPs [KRR14; CMS19].

Theorem 4. *Let $\mathcal{F}^{(2t)} : (\{0, 1\}^n)^{2t} \rightarrow \{0, 1\}^{2t}$ be a $2t$ -repeated k -non-signaling function, and suppose that $\mathcal{F}^{(2t)}$ is permutation folded. If $\mathcal{F}^{(2t)}$ passes both the linearity and consistency tests with probability at least $1 - \varepsilon$, then $\widehat{\mathcal{F}}^{(t)}$ is \widehat{k} -non-signaling function that is permutation folded, $(1 - 4\varepsilon)$ -linear, and $(1 - 8\varepsilon)$ -consistent, for $\widehat{k} = k/2 - 5$.*

The rest of this section is devoted to the proof of Theorem 4

7.3.1 $\widehat{\mathcal{F}}^{(t)}$ is permutation folded

We first prove that if $\mathcal{F}^{(2t)}$ is permutation folded, then $\widehat{\mathcal{F}}^{(t)}$ is also permutation-folded. (Recall Definition 2.8 for the definition of the permutation folded property and the application of permutations on vectors.)

Lemma 7.4. *Assuming $\mathcal{F}^{(2t)}$ is permutation-folded, $\widehat{\mathcal{F}}^{(t)}$ is also permutation-folded.*

Proof. Fix $S = \{Q_1, \dots, Q_\ell\} \subseteq (\{0, 1\}^n)^t$ with $1 \leq \ell \leq k$, and let $T = \{\pi_1(Q_1), \dots, \pi_\ell(Q_\ell)\}$ for some permutations $\pi_1, \dots, \pi_\ell \in S_t$. By definition of $\widehat{\mathcal{F}}^{(t)}$ for any $b_1, \dots, b_\ell \in \{0, 1\}^t$ it holds that

$$\begin{aligned} \Pr \left[\forall i \in [\ell] \quad \widehat{\mathcal{F}}^{(t)}_S(Q_i) = b_i \right] &= \Pr_{R_i, W_i} \left[\forall i \in [\ell] \quad \mathcal{F}^{(2t)}([R_i; W_i]) + \mathcal{F}^{(2t)}([Q_i + R_i; W_i]) = b_i \right] \\ &= \Pr_{R_i, W_i} \left[\forall i \in [\ell] \quad \mathcal{F}^{(2t)}([\pi(R_i)_i; W_i]) + \mathcal{F}^{(2t)}([\pi_i(Q_i + R_i); W_i]) = b_i \right] \\ &= \Pr \left[\forall i \in [\ell] \quad \widehat{\mathcal{F}}^{(t)}_T(\pi_i(Q_i)) = \pi_i(b_i) \right] , \end{aligned}$$

as required. □

7.3.2 $\widehat{\mathcal{F}}^{(t)}$ is almost linear

Next, we show that if $\mathcal{F}^{(2t)}$ passes the linearity test with high probability, then its self-correction is almost linear as per Definition 2.12.

Lemma 7.5. *Let $\mathcal{F}^{(2t)}: (\{0, 1\}^n)^{2t} \rightarrow \{0, 1\}^{2t}$ be a $2t$ -repeated k -non-signaling function such that $k \geq 7$. If $\mathcal{F}^{(2t)}$ passes the linearity test with probability at least $1 - \varepsilon$, then $\widehat{\mathcal{F}^{(2t)}}$ is $(1 - 4\varepsilon)$ -linear. That is, for any query set $\widehat{S} = \{X, Y, X + Y\} \subseteq (\{0, 1\}^n)^t$ we have $\Pr[\widehat{\mathcal{F}^{(2t)}}(X) + \widehat{\mathcal{F}^{(2t)}}(Y) = \widehat{\mathcal{F}^{(2t)}}(X + Y)] \geq 1 - 4\varepsilon$.*

The proof is almost the same as in [CMS20] Theorem 12 ($1 \implies 2$). The idea is to define a constant number of intermediate events such that each of them holds with high probability by high acceptance probability of the linearity test. Then we put together these intermediate events and derive the desired statement.

Proof. For $X, Y \in (\{0, 1\}^n)^t$ define $Z = X + Y$, and sample $R_X, R_Y, R_Z, W_X, W_Y, W_Z \in (\{0, 1\}^n)^t$ uniformly at random independently of each other. By definition of $\widehat{\mathcal{F}^{(2t)}}$ we have

$$\begin{aligned} \Pr[\widehat{\mathcal{F}^{(2t)}}(X) + \widehat{\mathcal{F}^{(2t)}}(Y) = \widehat{\mathcal{F}^{(2t)}}(X + Y)] \\ \geq \Pr \left[\mathcal{F}^{(2t)}([R_X; W_X]) + \mathcal{F}^{(2t)}([X + R_X; W_X]) \right. \\ \quad \left. + \mathcal{F}^{(2t)}([R_Y; W_Y]) + \mathcal{F}^{(2t)}([Y + R_Y; W_Y]) \right. \\ \quad \left. = \mathcal{F}^{(2t)}([R_Z; W_Z]) + \mathcal{F}^{(2t)}([Z + R_Z; W_Z]) \right] . \end{aligned}$$

Define

$$\begin{aligned} S_1 &:= \{[R_X; W_X], [R_Y; W_Y], [R_Z; W_Z], [X + R_X; W_X], [Y + R_Y; W_Y], [X + Y + R_Z; W_Z]\} , \\ S_2 &:= \{[R_X; W_X], [R_Z; W_Z], [X + R_X + R_Y; W_X + W_Y], [Y + R_Y; W_Y], [X + Y + R_Z; W_Z]\} , \\ S_3 &:= \{[R_X; W_X], [X + R_X + R_Y; W_X + W_Y], [Y + R_Y + R_Z; W_Y + W_Z], [X + Y + R_Z; W_Z]\} , \\ S_4 &:= \{[X + R_X + R_Y; W_X + W_Y], [Y + R_Y + R_Z; W_Y + W_Z], [X + Y + R_X + R_Z; W_X + W_Z]\} . \end{aligned}$$

Note that $|S_i \cup S_{i+1}| \leq 7 \leq k$ for $i = 1, 2, 3$.

Let $\text{add}(\cdot)$ be the addition function, and consider the sets S_1 and S_2 . Then

$$\begin{aligned} \Pr[\text{add}(\mathcal{F}^{(2t)}(S_1)) = \text{add}(\mathcal{F}^{(2t)}(S_2))] \\ = \Pr[\mathcal{F}^{(2t)}([X + R_X + R_Y; W_X + W_Y]) + \mathcal{F}^{(2t)}([X + R_X; W_X]) = \mathcal{F}^{(2t)}([R_Y; W_Y])] . \end{aligned}$$

Observing that the distribution on the right hand side is exactly as in the linearity test, we get that

$$\Pr[\text{add}(\mathcal{F}^{(2t)}(S_1)) = \text{add}(\mathcal{F}^{(2t)}(S_2))] \geq 1 - \varepsilon .$$

Similarly, we have

$$\begin{aligned} \Pr[\text{add}(\mathcal{F}^{(2t)}(S_2)) = \text{add}(\mathcal{F}^{(2t)}(S_3))] \\ = \Pr[\mathcal{F}^{(2t)}([R_Z; W_Z]) + \mathcal{F}^{(2t)}([Y + R_Y; W_Y]) = \mathcal{F}^{(2t)}([Y + R_Y + R_Z; W_Y + W_Z])] \\ \geq 1 - \varepsilon , \end{aligned}$$

and

$$\Pr[\text{add}(\mathcal{F}^{(2t)}(S_3)) = \text{add}(\mathcal{F}^{(2t)}(S_4))]$$

$$\begin{aligned}
&= \Pr[\mathcal{F}^{(2t)}([R_X; W_X]) + \mathcal{F}^{(2t)}([X + R_Z; W_Z]) = \mathcal{F}^{(2t)}([X + Y + R_X + R_Z; W_X + W_Z])] \\
&\geq 1 - \varepsilon .
\end{aligned}$$

Therefore,

$$\begin{aligned}
\left| \Pr[\text{add}(\mathcal{F}^{(2t)}(S_1)) = 0] - \Pr[\text{add}(\mathcal{F}^{(2t)}(S_4)) = 0] \right| &\leq \sum_{i=1}^3 \left| \Pr[\text{add}(\mathcal{F}^{(2t)}(S_i)) = 0] - \Pr[\text{add}(\mathcal{F}^{(2t)}(S_{i+1})) = 0] \right| \\
&\leq 3\varepsilon .
\end{aligned}$$

Finally, note that

$$\Pr[\text{add}(\mathcal{F}^{(2t)}(S_4)) = 0] \geq 1 - \varepsilon ,$$

because the distribution of S_4 is equal to the distribution of a three tuple used for linearity testing. Therefore,

$$\Pr[\widehat{\mathcal{F}}^{(t)}(X) + \widehat{\mathcal{F}}^{(t)}(Y) = \widehat{\mathcal{F}}^{(t)}(X + Y)] \geq \Pr[\text{add}(\mathcal{F}^{(2t)}(S_1)) = 0] \geq 1 - 4\varepsilon ,$$

as required. \square

7.3.3 $\widehat{\mathcal{F}}^{(t)}$ is almost consistent

Finally, we prove in Lemma 7.7 that if $\mathcal{F}^{(t)}$ passes the consistency test with high probability, then its self-correction is almost consistent. Before proving it we need the following claim.

Claim 7.6. *Let $\mathcal{F}^{(2t)}: (\{0, 1\}^n)^{2t} \rightarrow \{0, 1\}^{2t}$ be a $2t$ -repeated k -non-signaling function such that $k \geq 6$. Suppose that $\mathcal{F}^{(2t)}$ passes both linearity and consistency tests with probability at least $1 - \varepsilon$. Then for any $Q \in (\{0, 1\}^n)^t$ it holds that*

$$\Pr \left[\widehat{\mathcal{F}}^{(t)}(Q)_j = 0 \quad \forall j \in [t] \text{ such that } Q_j = 0^n \right] > 1 - 4\varepsilon$$

Proof. The key observation here is that for a uniformly random $R, W \in (\{0, 1\}^n)^t$ it holds that

$$\Pr \left[\mathcal{F}^{(2t)}([Q + R; W])_j = \mathcal{F}^{(2t)}([R; W])_j \quad \forall j = t + 1, \dots, 2t \right] \geq 1 - 4\varepsilon . \quad (4)$$

(Note that Eq. (4) does not follow from consistency testing since R and $Q + R$ are not independent.) Indeed, let $R', R'', W' \in (\{0, 1\}^n)^t$ be sampled uniformly at random, independently of all other random variables. Then, since $\mathcal{F}^{(2t)}$ passes linearity test with probability at least $1 - \varepsilon$, it follows that with probability at least $1 - 2\varepsilon$ the following equalities hold:

$$\begin{aligned}
\mathcal{F}^{(2t)}([Q + R; W]) &= \mathcal{F}^{(2t)}([Q + R''; W']) + \mathcal{F}^{(2t)}([R + R''; W + W']) \\
\mathcal{F}^{(2t)}([R; W]) &= \mathcal{F}^{(2t)}([R'; W']) + \mathcal{F}^{(2t)}([R + R'; W + W'])
\end{aligned}$$

If these two equalities hold, then

$$\begin{aligned}
\mathcal{F}^{(2t)}([Q + R; W]) + \mathcal{F}^{(2t)}([R; W]) &= \mathcal{F}^{(2t)}([Q + R''; W']) + \mathcal{F}^{(2t)}([R'; W']) \\
&\quad + \mathcal{F}^{(2t)}([R + R''; W + W']) + \mathcal{F}^{(2t)}([R + R'; W + W']) .
\end{aligned}$$

Noting that the queries $\{[Q + R''; W'], [R'; W']\}$ are distributed as in the consistency test, it follows that

$$\Pr \left[\mathcal{F}^{(2t)}([Q + R''; W'])_j = \mathcal{F}^{(2t)}([R'; W'])_j \quad \forall j = t + 1, \dots, 2t \right] \geq 1 - \varepsilon . \quad (5)$$

By the same argument we have

$$\Pr \left[\mathcal{F}^{(2t)}([R + R''; W + W'])_j = \mathcal{F}^{(2t)}([R + R'; W + W'])_j \quad \forall j = t + 1, \dots, 2t \right] \geq 1 - \varepsilon . \quad (6)$$

These immediately imply Eq. (4).

In order to complete the proof let $\pi \in S_{2t}$ be an arbitrary permutation such that for all $j \in [t]$, $\pi(j) \in \{t + 1, \dots, 2t\}$. Then,

$$\begin{aligned} & \Pr \left[\widehat{\mathcal{F}}^{(t)}(Q)_j = 0 \quad \forall j \in [t] \text{ such that } Q_j = 0^n \right] \\ &= \Pr \left[\mathcal{F}^{(2t)}([Q + R; W])_j = \mathcal{F}^{(2t)}([R; W])_j \quad \forall j \in [t] \text{ such that } Q_j = 0^n \right] \\ &= \Pr \left[\mathcal{F}^{(2t)}(\pi([Q + R; W]))_{\pi(j)} = \mathcal{F}^{(2t)}(\pi([R; W]))_{\pi(j)} \quad \forall j \in [t] \text{ such that } Q_j = 0^n \right] \\ &\geq 1 - 4\varepsilon , \end{aligned}$$

where the last inequality follows from Eq. (4) together with the permutation invariance of $\mathcal{F}^{(2t)}$. \square

The following lemma, saying that $\widehat{\mathcal{F}}^{(t)}$ is $(1 - O(\varepsilon))$ -consistent, follows almost immediately from Claim 7.6.

Lemma 7.7. *Let $\mathcal{F}^{(t)}: (\{0, 1\}^n)^{2t} \rightarrow \{0, 1\}^{2t}$ be a $2t$ -repeated k -non-signaling function such that $k \geq 7$, and suppose that $\mathcal{F}^{(2t)}$ is permutation folded.*

If $\mathcal{F}^{(2t)}$ passes both linearity and consistency tests with probability $1 - \varepsilon$, then $\widehat{\mathcal{F}}^{(t)}$ is $(1 - 8\varepsilon)$ -consistent. That is, for any two queries $X, Y \in (\{0, 1\}^n)^t$ to $\widehat{\mathcal{F}}^{(t)}$ it holds that

$$\Pr \left[\widehat{\mathcal{F}}^{(t)}(X)_j = \widehat{\mathcal{F}}^{(t)}(Y)_j \quad \forall j \in [t] \text{ such that } X_j = Y_j \right] \geq 1 - 8\varepsilon .$$

Proof. Let $J = \{j \in [t] : X_j = Y_j\}$. Consider the query set $S = \{X, Y, Z = X + Y\}$, and note that $Z_j = 0^n$ for all $j \in J$. Therefore, by Claim 7.6 it follows that $\Pr \left[\widehat{\mathcal{F}}^{(t)}(Z)_j = 0 \quad \forall j \in J \right] > 1 - 4\varepsilon$. By applying Lemma 7.5 we have $\Pr \left[\widehat{\mathcal{F}}^{(t)}(X) + \widehat{\mathcal{F}}^{(t)}(Y) = \widehat{\mathcal{F}}^{(t)}(Z) \right] \geq 1 - 4\varepsilon$. Therefore, by the union bound we conclude that $\Pr \left[\widehat{\mathcal{F}}^{(t)}(X)_j = \widehat{\mathcal{F}}^{(t)}(Y)_j \quad \forall j \in J \right] \geq 1 - 8\varepsilon$, thus concluding the proof of Lemma 7.7. \square

Theorem 4 is an immediate conclusion from Lemma 7.4, Lemma 7.5, and Lemma 7.7.

8 Proof of Theorem 2

Below we prove Theorem 2. Specifically, we show that assuming Hypothesis 2 the PCP construction in Algorithm 3 is sound against $O(1)$ -non-signaling proofs. Theorem 2 follows immediately from the following statement.

Theorem 8.1. Fix a circuit $C: \{0, 1\}^n \rightarrow \{0, 1\}$ with N wires, and input $x \in \{0, 1\}^n$ to C . Let $k \geq 18$ be a sufficiently large positive constant, and t be a positive integer such that $t \geq K \log^{2/e_{\text{hyp}}}(N)$ for some sufficiently large constant $K > 0$. Let $\mathcal{F}^{(2t)}: (\{0, 1\}^{N^2})^{2t} \rightarrow \{0, 1\}^{2t}$ be a k -non-signaling $2t$ -repeated linear consistent proof, and suppose that $\mathcal{F}^{(2t)}$ is permutation invariant. If $2t$ -repeated ALMSS verifier from Algorithm 3 accepts $\mathcal{F}^{(2t)}$ with probability $\geq 1 - \varepsilon$ for some sufficiently small ε , then $C(x) = 1$.

The proof follows the steps outlined in Section 5.

Proof. Fix a $2t$ -repeated k -non-signaling proof $\mathcal{F}^{(2t)}$ that satisfies the repeated ALMSS verifier from Algorithm 3 with probability at least $1 - \varepsilon$. In particular, $\mathcal{F}^{(2t)}$ passes the linearity test and the consistency test with probability at least $1 - \varepsilon$.

By applying Theorem 4 we conclude that $\widehat{\mathcal{F}}^{(t)}$, the self-correction of $\mathcal{F}^{(2t)}$, is a 4-no-signaling function that is $(1 - 4\varepsilon)$ -linear and $(1 - 8\varepsilon)$ -consistent. Furthermore, $\widehat{\mathcal{F}}^{(t)}$ satisfies the linear t -repeated verifier from Algorithm 2 with probability at least $1 - \varepsilon$.

Define $\widetilde{\mathcal{F}} = \text{Flat}[\widehat{\mathcal{F}}^{(t)}]$ to be the flattening of $\widehat{\mathcal{F}}^{(t)}$, as per Definition 2.17. Then, by Claim 2.19 the function $\widetilde{\mathcal{F}}: \{0, 1\}^{N^2} \rightarrow \{0, 1\}$ is $(1 - (4 + 3 \cdot 8)\varepsilon)$ -linear $(8\varepsilon, t)$ -non-signaling. Furthermore, since $\widehat{\mathcal{F}}^{(t)}$ is $(1 - 8\varepsilon)$ -consistent, and satisfies the linear repeated verifier from Algorithm 2 with probability at least $1 - \varepsilon$, it follows that $\widetilde{\mathcal{F}}$ satisfies the (non-repeated) linear verifier from Algorithm 1 with probability at least $1 - 9\varepsilon$.

Next, we use Hypothesis 2 to round $\widetilde{\mathcal{F}}$ to an exactly non-signaling function \mathcal{F} close to it. Specifically, since $\widetilde{\mathcal{F}}$ is $(1 - 28\varepsilon)$ -linear $(8\varepsilon, t)$ -non-signaling, by Hypothesis 2 there exist t' -non-signaling function $\mathcal{F}: \{0, 1\}^{N^2} \rightarrow \{0, 1\}$ for $t' \geq t^{\text{hyp}} = K' \log^2(N)$, such that $\Delta_4(\widetilde{\mathcal{F}}, \mathcal{F}) \leq \varepsilon'$, where $\varepsilon' = \varepsilon'_{\text{hyp}}(28\varepsilon)$. In particular, since $\widetilde{\mathcal{F}}$ is $(1 - 28\varepsilon)$ -linear, it follows that \mathcal{F} is $(1 - 28\varepsilon - \varepsilon')$ -linear, and satisfies the PCP verifier from Algorithm 1 with probability at least $1 - 9\varepsilon - \varepsilon'$.

Next, we apply the following theorem on almost linear non-signaling functions from [CMS19]. The theorem says that any almost linear function \mathcal{F} can be “rounded” into an exactly non-signaling function \mathcal{L} , such that the two are close to each other on predicates that depend on a small number of coordinates.

Theorem 8.2 (Theorem 7 in [CMS19]). Let $t', \bar{k} \in \mathbb{N}$ and $\varepsilon \in (0, 1/400]$ be such that $t' = \Omega(\frac{\bar{k}}{\varepsilon} \cdot (\bar{k} + \log \frac{1}{\varepsilon}))$. Suppose that $\mathcal{F}: \{0, 1\}^n \rightarrow \{0, 1\}$ is a t' -non-signaling function such that for all $x, y \in \{0, 1\}^n$ it holds that $\Pr[\mathcal{F}(x) + \mathcal{F}(y) = \mathcal{F}(x + y)] \geq 1 - \varepsilon$. Then there exists a linear \bar{k} -non-signaling function $\mathcal{L}: \{0, 1\}^n \rightarrow \{0, 1\}$ such that for all query sets $Q \subseteq \{0, 1\}^n$ of size $|Q| \leq \bar{k}$ and for all events $E \subseteq \{0, 1\}^Q$ it holds that

$$|\Pr[\mathcal{F}(Q) \in E] - \Pr[\mathcal{L}(Q) \in E]| \leq (6|Q| + 3)\sqrt{\varepsilon} .$$

Remark 8.3. Actually, Theorem 7 in [CMS19] assumes that linearity test accepts \mathcal{F} with high probability, and the conclusion of the theorem holds for its self-correction $\widehat{\mathcal{F}}$. However, if we make the stronger assumption that $\Pr[\mathcal{F}(x) + \mathcal{F}(y) = \mathcal{F}(x + y)] \geq 1 - \varepsilon$ holds for all $x, y \in \{0, 1\}^n$, then by following the proof, it is easy to see that the conclusion holds for \mathcal{F} , without the self-correction.

By applying Theorem 8.2 on \mathcal{F} , and using it for all 4-ary predicates used by Algorithm 1, it follows that there exists a linear \bar{k} -non-signaling function $\mathcal{L}: \{0, 1\}^n \rightarrow \{0, 1\}$ that satisfies the PCP verifier from Algorithm 1 with probability at least $1 - \widehat{\varepsilon}$ for $\widehat{\varepsilon} = 1 - 9\varepsilon - \varepsilon' - (6 \cdot 4 + 3)\sqrt{9\varepsilon + \varepsilon'} = 1 - O(\sqrt{\varepsilon + \varepsilon'})$. In particular, if $t' > K' \log^2(N)$ for a sufficiently large constant K' , then $\bar{k} \geq \bar{C} \log(N)$. Therefore,

if $\varepsilon > 0$ is a sufficiently small constant, it follows that the PCP verifier from Algorithm 1 accepts \mathcal{L} with probability greater than $39/40$, and by Theorem 4.1 we conclude that $C(x) = 1$. This completes the proof of Theorem 8.1. \square

9 Conclusions and open problems

In this paper we establish a conditional result on the existence of a PCP system that is sound against non-signaling proofs with constant locality. There are several natural research directions left open for future work.

Resolving the hypothesis. The implications of Hypothesis 2 motivates the study of geometry of non-signaling proofs. In particular, as a natural intermediate step toward settling Hypothesis 2, one can study the validity of a weaker version of hypothesis, requiring that the rounded proof is close to the given almost non-signaling proof on all subsets of size at most 2 (instead of 4) assuming that \mathcal{F} is linear (instead of almost linear), i.e., requiring that $\Delta_2(\mathcal{F}, \mathcal{F}') \leq \varepsilon'$. We remark that although Hypothesis 2 requires that $\Delta_4(\mathcal{F}, \mathcal{F}')$ is small, in fact, it suffices to show that $\Delta_3(\mathcal{F}, \mathcal{F}')$ is small, i.e., prove the hypothesis for subsets of size at most 3.

Reducing the alphabet. While we answer Question 1.2 affirmatively up to Hypothesis 2, we may require the proof to be of smaller alphabet. In the classical PCPs literature, the standard technique for alphabet reduction is known as *proof composition*, where the given “outer” proof over large alphabet is composed with a collection of “inner proofs of proximity” over small alphabet [BGHSV06; DR04]. Indeed, this component plays an important role in the modular proof of the PCP theorem. It would be interesting to apply a similar approach to the non-signaling setting.

Extending our approach to polynomial size nsPCPs. Our PCP construction is based on *exponential-length* PCP construction of [ALMSS98], which encodes proofs using the Hadamard code of exponential length. The effective proof length or alternatively the number of random bits used by the verifier are very important parameters for downstream applications. In order to reduce the proof length, it is natural to replace the linear encoding with low-degree polynomial encoding [BFL91; BFLS91]. Indeed, [KRR13; KRR14] proved that such an approach gives a PCP system that is sound against non-signaling proofs, albeit with locality $\text{polylog}(T)$. It would be interesting to see if the parallel repetition of their verifier is sound against non-signaling proofs with constant locality.

References

- [ABOR00] William Aiello, Sandeep N. Bhatt, Rafail Ostrovsky, and Sivaramakrishnan Rajagopalan. “Fast Verification of Any Remote Procedure Call: Short Witness-Indistinguishable One-Round Proofs for NP”. In: *Proceedings of the 27th International Colloquium on Automata, Languages and Programming*. ICALP '00. 2000, pp. 463–474.
- [AHR01] Yonatan Aumann, Johan Håstad, Michael O. Rabin, and Madhu Sudan. “Linear-Consistency Testing”. In: *Journal of Computer and System Sciences* 62.4 (2001), pp. 589–607.
- [ALMSS98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. “Proof verification and the hardness of approximation problems”. In: *Journal of the ACM* 45.3 (1998). Preliminary version in FOCS '92., pp. 501–555.
- [AS98] Sanjeev Arora and Shmuel Safra. “Probabilistic checking of proofs: a new characterization of NP”. In: *Journal of the ACM* 45.1 (1998). Preliminary version in FOCS '92., pp. 70–122.

- [BCHKS96] Mihir Bellare, Don Coppersmith, Johan Håstad, Marcos A. Kiwi, and Madhu Sudan. “Linearity testing in characteristic two”. In: *IEEE Transactions on Information Theory* 42.6 (1996), pp. 1781–1795.
- [BCLR08] Michael Ben-Or, Don Coppersmith, Mike Luby, and Ronitt Rubinfeld. “Non-abelian homomorphism testing, and distributions close to their self-convolutions”. In: *Random Structures and Algorithms* 32.1 (2008), pp. 49–70.
- [BFL91] László Babai, Lance Fortnow, and Carsten Lund. “Non-Deterministic Exponential Time has Two-Prover Interactive Protocols”. In: *Computational Complexity* 1 (1991). Preliminary version appeared in FOCS ’90., pp. 3–40.
- [BFLS91] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. “Checking computations in polylogarithmic time”. In: *Proceedings of the 23rd ACM Symposium on Theory of Computing*. STOC ’91. 1991, pp. 21–32.
- [BG15] Mark Braverman and Ankit Garg. “Small Value Parallel Repetition for General Games”. In: *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*. STOC ’15. 2015, pp. 335–340.
- [BGHSV06] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil P. Vadhan. “Robust PCPs of Proximity, Shorter PCPs, and Applications to Coding”. In: *SIAM Journal on Computing* 36.4 (2006), pp. 889–974.
- [BGS98] Mihir Bellare, Oded Goldreich, and Madhu Sudan. “Free Bits, PCPs, and Nonapproximability—Towards Tight Results”. In: *SIAM J. Comput.* 27.3 (1998), pp. 804–915.
- [BKSSZ10] Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. “Optimal testing of Reed-Muller codes”. In: *Proceedings of the 51st IEEE Symposium on Foundations of Computer Science*. FOCS ’10. 2010, pp. 488–497.
- [BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. “Self-Testing/Correcting with Applications to Numerical Problems”. In: *Journal of Computer and System Sciences* 47.3 (1993), pp. 549–595.
- [CMS19] Alessandro Chiesa, Peter Manohar, and Igor Shinkar. “Probabilistic Checking against Non-Signaling Strategies from Linearity Testing”. In: *Proceedings of the 10th Innovations in Theoretical Computer Science Conference*. ITCS ’19. 2019.
- [CMS20] Alessandro Chiesa, Peter Manohar, and Igor Shinkar. “Testing Linearity against Non-Signaling Strategies”. In: *ACM Trans. Comput. Theory* 12.3 (2020).
- [DDGKS17] Roei David, Irit Dinur, Elazar Goldenberg, Guy Kindler, and Igor Shinkar. “Direct Sum Testing”. In: *SIAM Journal on Computing* 46 (4 2017), pp. 1336–1369.
- [DLNNR04] Cynthia Dwork, Michael Langberg, Moni Naor, Kobbi Nissim, and Omer Reingold. *Succinct NP Proofs and Spooky Interactions*. Available at www.openu.ac.il/home/mikel/papers/spooky.ps. 2004.
- [DN17] Irit Dinur and Inbal Livni Navon. “Exponentially Small Soundness for the Direct Product Z-Test”. In: *Proceedings of the 32nd Computational Complexity Conference*. CCC ’17. 2017.
- [DR04] Irit Dinur and Omer Reingold. “Assignment Testers: Towards a Combinatorial Proof of the PCP Theorem”. In: *Proceedings of the 45th IEEE Symposium on Foundations of Computer Science*. FOCS ’04. 2004, pp. 155–164.
- [DS14a] Irit Dinur and David Steurer. “Analytical Approach to Parallel Repetition”. In: *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. STOC ’14. 2014, pp. 624–633.
- [DS14b] Irit Dinur and David Steurer. “Direct Product Testing”. In: *2014 IEEE 29th Conference on Computational Complexity (CCC)*. 2014, pp. 188–196.

- [FGLSS96] Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. “Interactive proofs and the hardness of approximating cliques”. In: *Journal of the ACM* 43.2 (1996). Preliminary version in FOCS ’91., pp. 268–292.
- [GCS19] Elazar Goldenberg and Karthik C. S. “Toward a General Direct Product Testing Theorem”. In: *ACM Trans. Comput. Theory* 12.1 (2019).
- [HK20] Dhiraj Holden and Yael Tauman Kalai. “Non-Signaling Proofs with $o(\sqrt{\log(n)})$ Provers Are in PSPACE”. In: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2020. 2020, 1024–1037.
- [HY19] Justin Holmgren and Lisa Yang. “The parallel repetition of non-signaling games: counterexamples and dichotomy”. In: *Proceedings of the 51st ACM Symposium on Theory of Computing*. STOC ’19. 2019, pp. 185–192.
- [Hol09] Thomas Holenstein. “Parallel Repetition: Simplification and the No-Signaling Case”. In: *Theory of Computing* 5.1 (2009). Preliminary version appeared in STOC ’07., pp. 141–172.
- [IKM09] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. “Oracularization and Two-Prover One-Round Interactive Proofs against Nonlocal Strategies”. In: *Proceedings of the 24th IEEE Annual Conference on Computational Complexity*. CCC ’09. 2009, pp. 217–228.
- [IKW12] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. “New Direct-Product Testers and 2-Query PCPs”. In: *SIAM J. Comput.* (2012), pp. 1722–1768.
- [IV12] Tsuyoshi Ito and Thomas Vidick. “A Multi-prover Interactive Proof for NEXP Sound against Entangled Provers”. In: *Proceedings of the 53rd IEEE Symposium on Foundations of Computer Science*. FOCS ’12. 2012, pp. 243–252.
- [Ito10] Tsuyoshi Ito. “Polynomial-Space Approximation of No-Signaling Provers”. In: *Proceedings of the 37th International Colloquium on Automata, Languages and Programming*. ICALP ’10. 2010, pp. 140–151.
- [KRR13] Yael Kalai, Ran Raz, and Ron Rothblum. “Delegation for Bounded Space”. In: *Proceedings of the 45th ACM Symposium on the Theory of Computing*. STOC ’13. 2013, pp. 565–574.
- [KRR14] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. “How to delegate computations: the power of no-signaling proofs”. In: *Proceedings of the 46th ACM Symposium on Theory of Computing*. STOC ’14. Full version available at <https://ecc.weizmann.ac.il/report/2013/183/>. 2014, pp. 485–494.
- [KRR16] Yael Tauman Kalai, Ran Raz, and Oded Regev. “On the Space Complexity of Linear Programming with Preprocessing”. In: *Proceedings of the 7th Innovations in Theoretical Computer Science Conference*. ITCS ’16. 2016, pp. 293–300.
- [KT85] Leonid A. Khalfin and Boris S. Tsirelson. “Quantum and quasi-classical analogs of Bell inequalities”. In: *Symposium on the Foundations of Modern Physics* (1985), pp. 441–460.
- [Kil92] Joe Kilian. “A note on efficient zero-knowledge proofs and arguments”. In: *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*. STOC ’92. 1992, pp. 723–732.
- [LW16] Cécilia Lancien and Andreas Winter. “Parallel repetition and concentration for (sub-)no-signalling games via a flexible constrained de Finetti reduction”. In: *Chicago Journal of Theoretical Computer Science* (2016).
- [Mic00] Silvio Micali. “Computationally Sound Proofs”. In: *SIAM Journal on Computing* 30.4 (2000). Preliminary version appeared in FOCS ’94., pp. 1253–1298.
- [PR94] Sandu Popescu and Daniel Rohrlich. “Quantum nonlocality as an axiom”. In: *Foundations of Physics* 24.3 (1994), pp. 379–385.
- [Ras85] Peter Rastall. “Locality, Bell’s theorem, and quantum mechanics”. In: *Foundations of Physics* 15.9 (1985), pp. 963–972.

- [Raz98] Ran Raz. “A Parallel Repetition Theorem”. In: *SIAM Journal on Computing* 27.3 (1998), pp. 763–803.
- [SA90] Hanif D. Sherali and Warren P. Adams. “A Hierarchy of Relaxations between the Continuous and Convex Hull Representations for Zero-One Programming Problems”. In: *SIAM Journal on Discrete Mathematics* 3.3 (1990), pp. 411–430.
- [SW04] Amir Shpilka and Avi Wigderson. “Derandomizing Homomorphism Testing in General Groups”. In: *Proceedings of the 36th ACM Symposium on the Theory of Computing*. STOC ’04. 2004, pp. 427–435.
- [Sch86] Alexander Schrijver. *Theory of Linear and Integer Programming*. USA: John Wiley & Sons, Inc., 1986. ISBN: 0471908541.
- [Ver96] Oleg Verbitsky. “Towards the parallel repetition conjecture”. In: *Theoretical Computer Science* 157.2 (1996), pp. 277–282.
- [Vid14] Thomas Vidick. *Linearity testing with entangled provers*. http://users.cms.caltech.edu/~vidick/linearity_test.pdf. 2014.

A Discussion on Hypothesis 2

As mentioned in the introduction, we can think of k -non-signaling functions as points in the polytope $P_k \subseteq \mathbb{R}^d$, for $d = \sum_{i=0}^n \binom{n}{i} 2^i$, which corresponds the solutions of the k ’th level relaxation of the Sherali-Adams hierarchy. Analogously, we can think of (ε, k) -non-signaling functions as points in the polytope $P_k^\varepsilon \subseteq \mathbb{R}^d$, which corresponds the solutions of the noisy version of the k ’th level relaxation of Sherali-Adams hierarchy, where for any two sets $S, T \subseteq [n]$ the marginal distributions induced by P_S on $S \cap T$ is ε -close in total variation distance to the marginal distributions induced by P_T on $S \cap T$. Then, Hypothesis 2 can be rephrased as follows: for any $p \in P_k^\varepsilon$ there exists $p' \in P_{k'}$ such that $\Delta_4(p, p') \leq \varepsilon'$.

We remark that sensitivity analysis of linear programs has been studied in the past (see, e.g., [Sch86] Section 10). However, the parameters obtained by these results seem to be too weak for our application. Nonetheless, it is possible that this approach could still work for our setting, since we are looking for an approximate solution with respect to the Δ_4 distance, which is rather non-standard.

In [CMS20], the following lemma, in the same spirit as the hypothesis, was proved.

Lemma A.1 ((see [CMS20, Lemma C.3])). *For every (ε, k) -non-signalling function $\mathcal{F}: D \rightarrow \{0, 1\}$ there exists k -non-signalling function \mathcal{F}' such that $\Delta_k(\mathcal{F}, \mathcal{F}') \leq O(4^k \cdot \varepsilon)$*

While the guarantee of $O(4^k \cdot \varepsilon)$ on the distance in the lemma is too large for our applications, Hypothesis 2 is somewhat more specific, and it is plausible that proving it is easier than improving Lemma A.1. We discuss Hypothesis 2 below.

1. Note that unlike Lemma A.1, Hypothesis 2 assumes that \mathcal{F} is almost linear. We do not know whether this is essential, however, it is reasonable to believe that being almost linear adds constraints on the structure of \mathcal{F} , thus making it easier to prove Hypothesis 2.
2. In Hypothesis 2 the requirement on the distance between the given almost non-signaling function, and the rounded function is only on sets of size at most 4. In fact, it is not difficult to see that proving that $\Delta_3(\mathcal{F}, \mathcal{F}') \leq \varepsilon'$ also suffices for the applications. This seems to be a significant relaxation compared to Δ_k proved in Lemma A.1.

3. In fact, our proof of soundness would go through even with a weaker version of the hypothesis, where we replaced the “worst-case” notion of Δ_4 with the “average-case”. Specifically, given an (ε, k) -almost non-signaling proof \mathcal{F} that satisfies *every constraint* of the linear ALMSS verifier with high probability, we want the rounded proof to satisfy the linear ALMSS verifier with high probability with respect to the distribution induced by the verifier on the 4-query sets.

Furthermore, since our almost non-signaling proof \mathcal{F} is obtained by flattening the repeated proof $\widehat{\mathcal{F}}^{(t)}$, we may assume that \mathcal{F} satisfies every constraints of the $\Omega(k)$ -sequential repetition of the linearity test, i.e., for some $\ell = \Omega(k)$ it holds that

$$\forall x_1, y_1, \dots, x_\ell, y_\ell \in \{0, 1\}^n \quad \Pr [\mathcal{F}(x_i) + \mathcal{F}(y_i) = \mathcal{F}(x_i + y_i) \quad \forall i \in [\ell]] \geq 1 - \varepsilon ,$$

and, similarly, \mathcal{F} satisfies *every constraint* of the $\Omega(t)$ -sequential repetition of the linear ALMSS verifier with high probability, and the goal is to get a rounded proof to satisfy the linear ALMSS verifier with high probability with respect to the distribution induced by the verifier on the 4-query sets.

4. An alternative way to prove our main theorem is to prove that Theorem 4.1 holds for almost non-signaling proofs. This question seems to be well motivated by the application to *delegation of computation*. Indeed, Kalai et. al [KRR14] constructed PCP systems (of polynomial size) that are sound against $(\varepsilon, \text{polylog}(N))$ -non-signaling proofs, for some negligible $\varepsilon > 0$. However, their proof seems to break for constant $\varepsilon > 0$. Our work motivates studying the power of almost non-signaling proofs for constant $\varepsilon > 0$.