Electronic Colloquium on Computational Complexity, Report No. 182 (2014)

Direct Product Testing With Nearly Identical Sets

Dana Moshkovitz *

November 25, 2014

Abstract

In this work we analyze a direct product test in which each of two provers receives a subset of size n of a ground set U, $|U| = \Theta(n)$, and the two subsets intersect in about $(1-\delta)n$ elements. We show that if each of the provers provides labels to the n elements it received, and the labels of the two provers agree in the intersection between the subsets with non-negligible probability, then the answers of the provers must correspond to a certain global assignment to the elements of U.

While previous results only worked for intersection of size at most $n/2$, in our model the questions and expected answers of the two provers are nearly identical. This is related to a recent construction of a unique games instance (ECCC TR14-142) where this setup arises at the "outer verifier" level.

Our main tool is a hypercontractive bound on the Bernoulli-Laplace model (aka a slice of the Boolean hypercube), from which we can deduce a "small set expansion"-type lemma. We then use ideas from a recent work of the author about "fortification" to reduce the case of large intersection to the already studied case of smaller intersection.

1 Direct Product Testing

Let U and R be finite sets, let $n \geq 1$ be a natural number, and let $0 \leq \alpha \leq 1$. We define the following two prover game. The provers are asked to agree on an assignment $f : U \to R$ ahead of time. Then a verifier picks at random two subsets $S, S' \subseteq U$ that intersect in αn elements on average. The verifier sends one subset to each prover. The prover is supposed to report what $f(x)$ is for every x in the subset it got, but is not necessarily truthful. The verifier checks that the assignments reported by the provers agree on the intersection between the two sets. Let $\mathcal{S}$ be the family of sized-n subsets of U.

Direct Product Test(α):

1. The verifier picks uniformly at random a set $S \in \mathcal{S}$.
2. The verifier picks a correlated set $S' \in \mathcal{S}$ as follows: for[1] $k \sim Poisson((1-\alpha)n)$ times, switch an element from S with a uniform element outside of S.

*`dmoshkov@csail.mit.edu`. Department of Electrical Engineering and Computer Science, MIT. This material is based upon work supported by the National Science Foundation under Grant Number 1218547.

[1]Several other choices of a correlated S' have been used in other works. The one we use follows a convention from probability literature.

ISSN 1433-8092

3. The verifier sends S to one prover, and sends S' to the other prover. Each prover responds with n elements in R that are supposed to be assignments to the n elements in the set the prover got.

4. The verifier checks that the two provers agree on all the assignments of the elements in $S \cap S'$.

If the provers indeed agree on $f : U \to R$ and respond accordingly, then the verifier always accepts. A direct product theorem shows that if a strategy of the provers makes the verifier accept with large probability, then the strategy necessarily gives rise to some global assignment $f : U \to R$ (the exact meaning of giving rise to a global assignment will be discussed shortly).

Many works prove direct product testing theorems [6, 3, 1, 8, 5]. Their main motivation is hardness amplification for PCP (see also [14, 2, 12] for applications). The idea is that the sets S, S' can span roughly n tests conducted on the global assignment f rather than one, thereby amplifying the soundness of the PCP test.

The different works proving direct product testing theorems differ in the size of the intersection they consider (determined by α), and in the success probability ε of the provers from which they can deduce a global assignment. The theorems become more challenging to prove as α increases (more intersection between the provers) and ε decreases (the test succeeds less often). The frontier of research today is $\alpha = \frac{1}{2}$ and $\varepsilon = \exp(-\Omega(n))$ [5].

The focus of the current work is the "nearly unique" case of $\alpha = 1-\delta$ for a possibly very small $\delta > 0$. In this case the sets that the two provers receive are nearly identical. This is a case that was not addressed in any previous work. The nearly identical questions to the provers make it easier for the provers to coordinate answers on the intersection. Indeed, if the provers had exactly identical subsets, they could simply decide on a different assignment for every subset ahead of time, and we could not have possibly expected a direct product testing theorem. The nearly unique setup is related to a recent construction of a unique games instance [9], where the setup arises at the "outer verifier" level.

Next we discuss the sense in which a successful strategy gives rise to a global assignment. To motivate our definition we consider two examples:

Example 1.1 (list decoding example). *For the provers to succeed with probability $p \leq \frac{1}{2}$, the provers can pick several global assignments $f_1, \ldots, f_{1/p} : U \to R$; and answer according to f_i in fraction p of the subsets.*

This example shows that for $p \leq \frac{1}{2}$ we are in the "list decoding" regime; there might not be a single global assignment, but rather a set of about $1/p$ global assignments.

Example 1.2 (local correlation example). *The provers can use the intersection between the subsets to agree on assignments locally. They can agree on a different assignment $f_x : U \to R$ for every $x \in U$; then, given a subset of U, pick a random element x in the subset, and assign all n elements according to f_x. Since the intersection is of size αn on average, the provers have probability about α to agree, even though there are $|U|$ different global assignments, each agreeing only with a small fraction of the subsets.*

Because of this example, we aim for a global assignment that agrees with the strategy of the provers on a large fraction of the subsets *among those that pass through many common points.*

Fix a strategy for the provers in the direct product game. For a small set of points $C \subseteq X$, $|C| < n$, let $\mathcal{S}_C$ be those sets in $\mathcal{S}$ that contain C.

Definition 1.1 (Strategy list decoding). *We say that a strategy of a prover in response to questions $\mathcal{S}$ has an (l,ϵ)-list decoding $f_1,\ldots,f_l : X \to R$ if there exists a partition $\mathcal{S} = \mathcal{S}_1 \cup \ldots \cup \mathcal{S}_l$, where for $1 \le i \le l$, for every subset $S \in \mathcal{S}_i$, the strategy is consistent with f_i for at least $1-\epsilon$ fraction of the elements in S.*

The following was proved by Dinur and Steurer [5]:

Lemma 1.1 (Direct product test [5]). *Suppose $\alpha \le 1/2$, and one considers Direct Product Test(α) conditioned on $|S \cap S'| = \alpha n$. For any strategy for the provers that the verifier accepts with probability at least $(1-\eta)^{\alpha n/2}$, there exists $C \subseteq X$, $|C| \le (3/2)\alpha n$, where the strategy of each of the provers on $\mathcal{S}_C$ has an $(R^{|C|}, O(\eta))$-list decoding.*

In this work we prove the following theorem:

Theorem 1 (Direct product test with nearly identical subsets). *Let $0 < r \le 1/2$ such that $n \le r|U|/2$. Let $\delta' = \rho\delta/(1+\rho\delta)$. There exists err that is exponentially small in* $\mathrm{poly}(\epsilon)n$, *such that for any strategy for the provers that the verifier in Direct Product Test($1-\delta$) accepts with probability at least*

$$2\left((1-\eta)^{rn/2} + err\right)^{\delta'} + err,$$

there exists $C \subseteq X$, $|C| \le 3rn/2$, where $\mathcal{S}_C$ has an $(R^{|C|}, O(\eta))$-list decoding.

1.1 Our Technique

The proof of Theorem 1 uses a couple of ingredients: (i) A hypercontractive inequality for the Bernoulli-Laplace model and a "Small Sets Partitions" lemma. (ii) An idea from [12, 13] that is used to reduce the case of very large intersection to the case of lower intersection that Lemma 1.1 deals with. In this section we discuss those techniques and the contributions of this work.

1.1.1 Small Sets Partition

A "small set expansion" lemma for a graph $G = (V,E)$ says that any small set of vertices in the graph $A \subseteq V$, $|A| \le \delta|V|$, expands, that is, there can't be too many edges that stay inside A relative to the number of edges that touch A (See Ryan O'Donnell's book [15] for more details on such theorems and how they are proved for graphs of interest). For this work we prove a "small sets partition" lemma for the graph that corresponds to the direct product test, namely for the graph whose vertex set is the family of all sized-n subsets of U, and in which edges correspond to tests of the verifier. The lemma says that given two partitions of the sets of vertices into small sets, $A_1 \cup \ldots \cup A_k = V$ and $B_1 \cup \ldots \cup B_k = V$, $|A_i|, |B_i| \le \delta|V|$ for all $1 \le i \le k$, there can't be too many edges between matching A_i and B_i over all $1 \le i \le k$:

Lemma 1.2 (Small sets partitions). *Consider two partitions into small sets $A_1 \cup \ldots \cup A_k = \binom{U}{n}$ and $B_1 \cup \ldots \cup B_k = \binom{U}{n}$, $\alpha_i \doteq |A_i| / \binom{|U|}{n}$, $\beta_i \doteq |B_i| / \binom{|U|}{n}$. Then, for any $0 < \varepsilon < 1$, if we pick S uniformly at random from $\binom{U}{n}$ and S' by applying Poisson(δn) random transpositions on S, then*

$$\Pr\left[\exists i,\ S' \in B_i \wedge S \in A_i\right] \le \sqrt{(1/\varepsilon) \cdot \sum_{i=1}^{k} \alpha_i\beta_i + \varepsilon^{\rho\delta/2}}.$$

Note that ε should be very small in order for $\varepsilon^{\rho\delta/2}$ to be small, and, moreover, the sets in the partition should be small for the square-root term in the right hand side to be small as well. The small set partition lemma is proved in Section 2 using a hypercontractive inequality for slices of the hypercube (the Bernoulli-Laplace model) due to Lee-Yau [11] and Gross [7].

1.1.2 Fortification and Small Set Expanders

In [12, 13] the author shows how strong parallel repetition and direct product testing type theorems can be proved fairly easily for natural tests in which the underlying graph is a good extractor (see the definition of fortified games in [13]). In this work we use one of the ideas from these works to reduce the case of large intersection between sets to the already studied case of small intersection between sets. Interestingly, we show that an extractor is not needed for the reduction, only a graph that admits a "small sets partition" lemma as discussed in the previous section (see Section 3). This extends the possible applications of the general fortification paradigm.

1.2 Related Work

Direct product testing can be thought of as a generalization of parallel repetition of two prover games (in the sense that it can be used to simulate parallel repetition [8]). Part of our work is analogous to a recent work of Tulsiani, Wright and Zhou [17] on parallel repetition of projection games on small set expanders (the other part in our work is to argue the relevant small set expansion properties of the relevant graph). The work of Tulsiani et al does not apply to our setting, and we employ different techniques from theirs. More details follow.

Tulsiani et al used a linear algebra approach pioneered by Dinur and Steurer [4] to show that in the small set expander case one can obtain stronger parameters than possible for general projection games (this was already known for expanders [16]). The Dinur-Steurer analysis of parallel repetition relies on Cheeger's inequality, and the Tulsiani et al result relies on an extension of Cheeger's inequality to the k largest eigenvalues, which was proved by [10]. Dinur an Steurer were able to convert their parallel repetition [4] to the direct product testing setting [5]. However, there it does not require Cheeger's inequality and some of the other steps needed in the parallel repetition setting, but does require certain other steps in order to argue about the structure of the strategy.

2 Bernoulli-Laplace and Partitions Into Small Sets

The main technical tool we use in our analysis is a hypercontractive inequality in the Bernoulli-Laplace model, which is the model that corresponds to the random walk on the n-sized subsets of U performed by the verifier to pick S and S'. The hypercontractive inequality allows us to prove a small sets partition lemma (see Lemma 2.3) for partitions of $\mathcal{S}$. In this section we present the inequality and the small sets partition lemma.

The *Bernoulli-Laplace* model is as follows: Given a set U and a parameter n, consider the graph whose vertex set is $\binom{|U|}{n}$, the family of subsets of size n of U. A random walk is defined on the graph, where starting at a set $S \subseteq U$, one makes a random transposition, picking a random element $x \in S$ and switching it with a random element $x' \in U - S$. Let K be the transition matrix of this random walk, and let its Laplacian be $L = I - K$. The noise operator H_t induced by this random walk is defined by applying a number of random transpositions determined by a Poisson random variable whose mean is t:

Definition 2.1 (Noise operator)**.**

$$H_t \doteq e^{-tL} = e^{-t} \sum_{l=0}^{\infty} \frac{t^l}{l!} K^l.$$

Lee and Yau [11] analyzed the log Sobolev constant of this Markov chain and determined that it is:

$$\rho = \Theta\left(1/\left(\log \frac{|U|^2}{n(|U|-n)}\right)\right).$$

This implies via Gross's work [7] the following hypercontractive inequality (where $\|g\|_p = (\mathbf{E}_x [g(x)]^p)^{1/p}$):

Lemma 2.1 (Hypercontractivity for Bernoulli-Laplace [11, 7])**.** *For every $t \geq 0$, $1 \leq p \leq q \leq \infty$, satisfying $\sqrt{\frac{p-1}{q-1}} \geq e^{-\rho t/n}$, for any f defined over $\binom{U}{n}$, we have*

$$\|H_t f\|_q \leq \|f\|_p.$$

Using Hölder inequality, we can deduce a two-function version of Lemma 2.1:

Lemma 2.2 (Two function hypercontractivity for Bernoulli-Laplace)**.** *For every $t \geq 0$, $1 \leq p \leq q \leq \infty$, satisfying $\sqrt{\frac{p-1}{q-1}} \geq e^{-\rho t/n}$, for any f, g defined over $\binom{U}{n}$, we have*

$$\langle H_t f, g\rangle \leq \|f\|_p \|g\|_{q/(q-1)}.$$

We will further deduce a "small set expansion" type conclusion from Lemma 2.2. The lemma upper bounds the probability that correlated S, S' are in matching parts of the partitions in the case where the sets of the partitions are small.

Note that we work with partitions into small sets rather than with two small sets $A, B \subseteq \binom{U}{n}$, because we might have sets of varying sizes, and if B contains A as well as its boundary, we might get that the probability that $S \in A$ and $S' \in B$ is upper bounded by the probability that $S \in A$, or that the probability that $S' \in B$ given that $S \in A$ is upper bounded by 1, which is a trivial statement.

Lemma 2.3 (Small sets partitions)**.** *Consider two partitions into small sets $A_1 \cup \ldots \cup A_k = \binom{U}{n}$ and $B_1 \cup \ldots \cup B_k = \binom{U}{n}$, $\alpha_i \doteq |A_i| / \binom{|U|}{n}, \beta_i \doteq |B_i| / \binom{|U|}{n}$. Then, for any $0 < \varepsilon < 1$, if we pick S uniformly at random from $\binom{U}{n}$ and S' by applying Poisson(δn) random transpositions on S, then*

$$\Pr\left[\exists i,\ S' \in B_i \wedge S \in A_i\right] \leq \sqrt{(1/\varepsilon) \cdot \sum_{i=1}^{k} \alpha_i \beta_i + \varepsilon^{\rho\delta/2}}.$$

Proof. Let $f_i = 1_{B_i}$ be the indicator function of B_i, and let $g_i = 1_{A_i}$ be the indicator function of A_i. Take $t = \delta n$, $q = 2$ and p so that $\sqrt{\frac{p-1}{q-1}} = e^{-\rho t/n}$. Note that $e^{-\rho t/n} \leq 1 - \frac{\rho\delta}{1+\rho\delta}$, so

$p \le 2(1-\frac{\rho\delta}{1+\rho\delta})$. By Lemma 2.2 and since $2/p \ge 1/(1-\frac{\rho\delta}{1+\rho\delta}) = 1+\rho\delta$ and $1/p - 1/q \ge \rho\delta/2$,

$$\begin{aligned}\Pr\left[\exists i,\ S \in A_i \wedge S' \in B_i\right] &= \sum_{i=1}^{k} \langle H_t f_i, g_i\rangle \\ &\le \sum_{i=1}^{k} \|f_i\|_p \|g_i\|_{q/(q-1)} \\ &= \sum_{i=1}^{k} \beta_i^{1/p} \alpha_i^{1-1/q}\end{aligned}$$

First let us bound the contribution from i's where $\beta_i \le \varepsilon$:

$$\begin{aligned}\sum_{i:\beta_i \le \varepsilon} \beta_i^{1/p} \alpha_i^{1-1/q} &\le \varepsilon^{\rho\delta/2} \cdot \sum_{i=1}^{k} \sqrt{\alpha_i \beta_i} \\ \text{Cauchy-Schwarz} &\le \varepsilon^{\rho\delta/2} \sqrt{\sum_{i=1}^{k} \alpha_i} \cdot \sqrt{\sum_{i=1}^{k} \beta_i} \\ &\le \varepsilon^{\rho\delta/2}\end{aligned}$$

There are less than $1/\varepsilon$ many i's where $\beta_i > \varepsilon$. We bound the contribution from such i's using concavity:

$$\begin{aligned}\sum_{i:\beta_i > \varepsilon} \beta_i^{1/p} \alpha_i^{1-1/q} &\le \sum_{i:\beta_i > \varepsilon} \sqrt{\alpha_i \beta_i} \\ &\le \sqrt{(1/\varepsilon) \cdot \sum_{i:\beta_i > \varepsilon} \alpha_i \beta_i} \\ &\le \sqrt{(1/\varepsilon) \cdot \sum_{i=1}^{k} \alpha_i \beta_i}\end{aligned}$$

The claim follows. □

3 Analysis of Direct Product Test with Nearly Identical Sets

In this section we'll prove Theorem 1. For our analysis we will use a generalization of Lemma 1.1 that was also proved by Dinur-Steurer [5]. The generalization considers tests that compare their subsets only on a random part of their intersection.

Lemma 3.1 (Direct product test [5])**.** *Suppose $r \le 1/2$; $r/2 \le \alpha \le r$; $\beta \le r/2$. One considers a direct product test that picks at random two sized-n subsets of U that intersect in $(\alpha+\beta)n$ points, and compares their assignments on αn random points in the intersection.*

For any strategy for the provers that the verifier accepts with probability at least $(1-\eta)^{rn/2}$, there exists $C \subseteq X$, $|C| \le (\alpha+\beta)n$, where the strategy of each of the provers on $\mathcal{S}_C$ has an $(R^{|C|}, O(\eta))$-list decoding.

Theorem 1 follows from Lemma 3.1, as well as Lemma 3.2 below, which relates the acceptance probabilities in two tests:

- $\mathcal{T}_1 = \mathcal{T}_1(\delta)$ is Direct Product Test$(1-\delta)$ that compares subsets that intersect in $\approx (1-\delta)n$ elements.
- $\mathcal{T}_2 = \mathcal{T}_2(\alpha, \beta)$ compares subsets that intersect in $(\alpha+\beta)n$ elements on random αn elements.

Note that $\mathcal{T}_1(\delta)$ is the test we wish to analyze, while $\mathcal{T}_2(\alpha, \beta)$ is the test that Lemma 3.1 analyzes. We use the small sets partition lemma to relate the acceptance probabilities of the two tests:

Lemma 3.2 (Reduction lemma)**.** *Assume that $n = \gamma |U|$. Let $\gamma < \alpha < 1$, $0 < \delta < 1$ and $\epsilon > 0$ such that $\alpha + \gamma + \epsilon < 1 - \delta$. Let $\delta' = \rho\delta/(1 + \rho\delta)$. For any strategy for the provers, there exist α', β' such that $\alpha' = (\alpha - \gamma)/(1 - \gamma)$, $\alpha' + \beta' \leq \alpha + \epsilon$, and*

$$\Pr\left[\mathcal{T}_1(\delta) \text{ accepts}\right] - err \leq 2\left(\Pr\left[\mathcal{T}_2(\alpha', \beta') \text{ accepts}\right] + err\right)^{\delta'},$$

where err is exponentially small in $\mathrm{poly}(\epsilon)n$*.*

Proof. Fix a strategy for the provers. We consider randomized choices of sets S, S' in a direct product test such that S, S' contain a set $C \subseteq U$ of size $\alpha' n$. The first choice approximates the choice of sets in $\mathcal{T}_1$, while the second choice approximates the choice of sets in $\mathcal{T}_2$.

Choice 1:

1. Pick uniformly at random $C \subseteq U$ of size $\alpha' n$.
2. Pick uniformly at random $S \subseteq U$ of size n that contains C.
3. Pick $S' \subseteq U$ by repeatedly, for λ times where $\lambda \sim Poisson(\delta n)$, switching an element from $S - C$ with an element outside S.

Note that in $\mathcal{T}_1$ there is a low probability err that the intersection of S, S' is smaller than $\alpha' n$, while with choice 1 this cannot happen.

Choice 2:

1. Pick uniformly at random $C \subseteq U$ of size $\alpha' n$.
2. Pick uniformly and independently at random $S, S' \subseteq U$ of size n that contain C.

Since $n = \gamma |U|$, the set S' has expected intersection size $\alpha' n + (1 - \alpha')\gamma n = \alpha n$ with S. With choice 2, the probability that the intersection is of size larger than $(\alpha + \epsilon)n$ is at most err. Let $\mathcal{T}_1', \mathcal{T}_2'$ be the direct product testers that pick their sets S, S' according to choices 1 and 2, respectively, and compare S, S' on C. We have that

$$\Pr\left[\mathcal{T}_1' \text{ accepts}\right] \geq \Pr\left[\mathcal{T}_1(\delta) \text{ accepts}\right] - err,$$

and that there exists β' such that $\alpha' + \beta' \leq \alpha + \epsilon$ for which

$$\Pr\left[\mathcal{T}_2(\alpha', \beta') \text{ accepts}\right] \geq \Pr\left[\mathcal{T}_2' \text{ accepts}\right] - err.$$

Next we'll relate the probabilities that $\mathcal{T}_1', \mathcal{T}_2'$ accept.

For a fixed $C \subseteq U$ of size $(\alpha - \gamma)n$, let $\mathcal{S}_C$ be all the n-sized subsets of U that contain C. Partition $\mathcal{S}_C$ in two ways, according to the assignments the two provers give C. That is, for every $\vec{a} \in R^{|C|}$, let $A_{\vec{a}}$ denote all the sets in $\mathcal{S}_C$ such that the first prover assigns the elements in C the assignments $\vec{a}$. Similarly, let $B_{\vec{a}}$ all the sets in $\mathcal{S}_C$ such that the second prover assigns the elements in C the assignment $\vec{a}$. Let $\alpha_{\vec{a}} = |A_{\vec{a}}| / |\mathcal{S}_C|$, and let $\beta_{\vec{a}} = |B_{\vec{a}}| / |\mathcal{S}_C|$. Note that the acceptance probability of $\mathcal{T}_2'$ is precisely

$$\Pr\left[\mathcal{T}_2' \text{ accepts}\right] = \mathop{\mathbf{E}}_{C}\left[\sum_{\vec{a}} \alpha_{\vec{a}}\beta_{\vec{a}}\right].$$

Meanwhile, the acceptance probability of $\mathcal{T}_1'$ is

$$\Pr\left[\mathcal{T}_1' \text{ accepts}\right] = \mathop{\mathbf{E}}_{C}\left[\Pr_{S,S'\in\mathcal{S}_C}\left[\exists\vec{a},\ S \in A_{\vec{a}} \wedge S' \in B_{\vec{a}}\right]\right].$$

From Lemma 2.3 with $\varepsilon = \Pr\left[\mathcal{T}_2' \text{ accepts}\right]^{\frac{1}{1+\rho\delta}}$,

$$\Pr\left[\mathcal{T}_1' \text{ accepts}\right] \le 2\Pr\left[\mathcal{T}_2' \text{ accepts}\right]^{\delta'}.$$

The lemma follows. $\square$

Theorem 1 follows from Lemma 3.2 and Lemma 3.1 by taking $\alpha = r$.

References

[1] I. Dinur and E. Goldenberg. Locally testing direct product in the low error range. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, pages 613–622, 2008.

[2] I. Dinur and O. Meir. Derandomized parallel repetition via structured PCPs. *Computational Complexity*, 20(2):207–327, 2011.

[3] I. Dinur and O. Reingold. Assignment testers: Towards a combinatorial proof of the PCP theorem. *SIAM Journal on Computing*, 36(4):975–1024, 2006.

[4] I. Dinur and D. Steurer. Analytical approach to parallel repetition. In *Proc. 46th ACM Symp. on Theory of Computing*, 2014.

[5] I. Dinur and D. Steurer. Direct product testing. In *Computational Complexity Conference*, 2014.

[6] O. Goldreich and S. Safra. A combinatorial consistency lemma with application to proving the PCP theorem. *SIAM J. Comput.*, 29(4):1132–1154, 2000.

[7] L. Gross. Logarithmic sobolev inequalities. *Amer. J. Math.*, 97:1061–1083, 1975.

[8] R. Impagliazzo, V. Kabanets, and A. Wigderson. New direct-product testers and 2-query PCPs. *SIAM Journal on Computing*, 41(6):1722–1768, 2012.

[9] S. Khot and D. Moshkovitz. Candidate lasserre integrality gap for unique games. Technical report, ECCC TR14-142, 2014.

[10] T. C. Kwok, L. C. Lau, Y. T. Lee, S. Oveis Gharan, and L. Trevisan. Improved cheeger's inequality: Analysis of spectral partitioning algorithms through higher order spectral gap. In *Proc. 45th ACM Symp. on Theory of Computing*, STOC '13, pages 11–20, 2013.

[11] T.-Y. Lee and H.-T. Yau. Logarithmic sobolev inequality for some models of random walks. *Ann. Prob.*, 26(4):1855–1873, 1998.

[12] D. Moshkovitz. An approach to the sliding scale conjecture via parallel repetition for low degree testing. Technical Report 30, ECCC, 2014.

[13] D. Moshkovitz. Parallel repetition from fortification. In *Proc. 55th IEEE Symp. on Foundations of Computer Science*, 2014.

[14] D. Moshkovitz and R. Raz. Two query PCP with sub-constant error. *Journal of the ACM*, 57(5), 2010.

[15] R. O'Donnell. *Analysis of Boolean Functions.* Cambridge University Press, 2014.

[16] R. Raz and R. Rosen. A strong parallel repetition theorem for projection games on expanders. In *IEEE Conference on Computational Complexity*, pages 247–257, 2012.

[17] M. Tulsiani, J. Wright, and Y. Zhou. Optimal strong parallel repetition for projection games on low threshold rank graphs. In *ICALP (1)*, pages 1003–1014, 2014.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de