2015 IEEE 56th Annual Symposium on Foundations of Computer Science

On Monotonicity Testing and Boolean Isoperimetric type Theorems

Subhash Khot
Department of Computer Science
New York University
New York, U.S.A.
Email: khot@cs.nyu.edu

Dor Minzer, Muli Safra
Department of Computer Science
Tel Aviv University
Tel Aviv, Israel
Email: dorminzer@mail.tau.ac.il, muli.safra@gmail.com

Abstract

We show a directed and robust analogue of a boolean isoperimetric type theorem of Talagrand [13]. As an application, we give a monotonicity testing algorithm that makes $\tilde{O}(\sqrt{n}/\varepsilon^2)$ non-adaptive queries to a function $f : \{0,1\}^n \mapsto \{0,1\}$, always accepts a monotone function and rejects a function that is ε-far from being monotone with constant probability.

Keywords

Boolean functions; mononotne functions; property testing; isoperimetry.

I. INTRODUCTION

In this paper, we study the problem of testing whether a given boolean function $f : \{0,1\}^n \mapsto \{0,1\}$ is monotone. We also study certain isoperimetric type theorems on the boolean hypercube that are closely related. Our main results are: (1) a directed and robust analogue of a theorem of Talagrand [13], generalizing many prior related theorems and (2) a monotonicity tester that is optimal in terms of its query complexity (see Section I-E2 for subtle issues regarding its optimality).

A. Boolean Isoperimetric Type Theorems

Given a function $f : \{0,1\}^n \mapsto \{0,1\}$, define the variance of the function as $\mathsf{var}(f) = p(1-p)$ where $p = \mathsf{Pr}_x[f(x) = 1]$. Let $\mathcal{S}_f$ denote the set of sensitive edges, i.e. the set of pairs (x, y) such that $x, y \in \{0,1\}^n$ differ in exactly one co-ordinate, $f(x) = 1$ and $f(y) = 0$. Let $I_f = \frac{|\mathcal{S}_f|}{2^n}$ denote the "total influence" of the function. A folk-lore theorem states:[1]

Theorem I.1.

$$I_f \geqslant \Omega(\mathsf{var}(f)).$$

The parameter I_f reflects the size of the edge boundary of the function f (or more precisely of the subset $\{x|f(x) = 1\}$ of the hypercube). The size of the vertex boundary Γ_f is defined as

$$\Gamma_f = \frac{1}{2^n} \cdot |\{x|f(x) = 1, \exists (x, y) \in \mathcal{S}_f\}|.$$

Margulis [11] shows that the size of the edge boundary and that of the vertex boundary cannot both be small. Specifically,

Research supported by NSF grants CCF 1422159, 1061938, 0832795 and Simons Collaboration on Algorithms and Geometry grant.

[1] A Fourier analytic proof: $\mathsf{var}(f) = \sum_{S \subseteq \{1,\ldots,n\}, S \neq \phi} \hat{f}(S)^2$ whereas $I_f = 2 \cdot \sum_{S \subseteq \{1,\ldots,n\}, S \neq \phi} \hat{f}(S)^2 \cdot |S|$.

0272-5428/15 $31.00 © 2015 IEEE
DOI 10.1109/FOCS.2015.13

Theorem I.2.

$$I_f \cdot \Gamma_f \;\geqslant\; \Omega(\mathsf{var}(f)^2).$$

It is instructive to note that the inequality above is tight up to a constant factor, as shown by a dictatorship function as well as the majority function. Both functions have a constant variance. For the dictatorship function, both I_f and Γ_f are $\Theta(1)$. For the majority function, $I_f = \Theta(\sqrt{n})$ and $\Gamma_f = \Theta(\frac{1}{\sqrt{n}})$.

For $x \in \{0,1\}^n$, the sensitive edges incident on x are precisely the edges in $\mathcal{S}_f$ that are incident on x. Let $I_f(x)$ be equal to 0 if $f(x) = 0$ and equal to the number of sensitive edges incident on x if $f(x) = 1$. Talagrand [13] shows that:

Theorem I.3.

$$\mathbb{E}_x \left[\sqrt{I_f(x)} \right] \geqslant \Omega(\mathsf{var}(f)).$$

It is easily seen that Theorem I.1 is implied by Theorem I.2 which in turn is implied by Theorem I.3. For the former implication, one observes that (here $\mathbf{1}_{(\cdot)}$ denotes indicator of an event)

$$I_f = \mathbb{E}_x\left[I_f(x)\right] \geqslant \mathbb{E}_x\left[\mathbf{1}_{I_f(x)>0}\right] = \Gamma_f.$$

For the latter, one observes using Cauchy-Schwartz that

$$I_f \cdot \Gamma_f = \mathbb{E}_x\left[I_f(x)\right] \cdot \mathbb{E}_x\left[\mathbf{1}_{I_f(x)>0}\right] \geqslant \mathbb{E}_x\left[\sqrt{I_f(x)}\right]^2 \geqslant \Omega(\mathsf{var}(f)^2).$$

B. Directed Analogues of Boolean Isoperimetric Type Theorems

A function $h : \{0,1\}^n \mapsto \{0,1\}$ is called monotone if for any two inputs x and y where y is obtained by changing a co-ordinate of x from 0 to 1 it holds that $h(x) = 1 \Longrightarrow h(y) = 1$. Equivalently, writing $x \leqslant y$ to mean that $x_i \leqslant y_i$ for every co-ordinate $i \in \{1, \ldots, n\}$, f is monotone if and only if

$$\forall\; x, y \in \{0,1\}^n, \quad x \leqslant y \Longrightarrow f(x) \leqslant f(y).$$

For a function $f : \{0,1\}^n \mapsto \{0,1\}$, let $\varepsilon(f)$ denote the distance of f from the class of monotone functions, i.e. minimum fraction of its values that need to be changed to turn f into a monotone function. The monotonicity testing problem asks for an algorithm that queries a given function f at a "few" places and distinguishes whether the function is monotone or is far from being monotone (more on this later). The problem has been very well-studied since late 1990s, on the boolean hypercube as well as over more general posets, for functions that take non-negative integer values instead of boolean values, and also in the context of related problems such as estimating distance to monotonicity, approximating total influence and shortest path routing on the hypercube [9], [3], [7], [2], [8], [1], [12], [4], [5], [10]. Still, designing an optimal tester for boolean functions on the boolean hypercube (the most basic and interesting case in our opinion) remained open.

Let $\mathcal{S}_f^-$ denote the set of negatively sensitive edges, i.e. the set of pairs (x, y) such that y is obtained by changing a single co-ordinate of x from 0 to 1 and $f(x) = 1, f(y) = 0$. These are precisely the edges that violate the monotonicity property. Let $I_f^- = \frac{|\mathcal{S}_f^-|}{2^n}$ be the "total negative influence". Motivated by an application to the monotonicity testing problem, Goldreich et al [9] show that:

Theorem I.4.

$$I_f^- \;\geqslant\; \Omega(\varepsilon(f)).$$

A hypercube can be thought of as a directed graph by orienting all its edges "monotonically upwards". In hindsight, Theorem I.4 is viewed as a "directed" analogue of Theorem I.1, where I_f is replaced by its analogue I_f^- and $\mathsf{var}(f)$ is replaced by its analogue $\varepsilon(f)$. As far as we know, Chakrabarty and Seshadhri [3] are the first to suggest this analogy. Also motivated by an application to the monotonicity testing problem, they show the following directed analogue of Margulis' Theorem I.2:

Theorem I.5.

$$I_f^- \cdot \Gamma_f^- \;\geqslant\; \Omega(\varepsilon(f)^2).$$

We note that again I_f is replaced by its analogue I_f^- (which sounds intuitive) and $\mathsf{var}(f)$ is replaced by its analogue $\varepsilon(f)$ (which is not so intuitive, and hence quite remarkable, in our opinion). Lastly, Γ_f is replaced by its analogue Γ_f^-, size of the negative vertex boundary, defined as:

$$\Gamma_f^- = \frac{1}{2^n} \cdot \left| \left\{ x | f(x) = 1, \exists (x,y) \in \mathcal{S}_f^- \right\} \right|.$$

For $x \in \{0,1\}^n$, the negatively sensitive edges incident on x are precisely the edges in $\mathcal{S}_f^-$ that are incident on x. Let $I_f^-(x)$ be equal to 0 if $f(x) = 0$ and equal to the number of negatively sensitive edges incident on x if $f(x) = 1$. Carrying the analogy between the undirected and directed case further and still motivated by an application to the monotonicity testing problem, we show a directed analogue of Talagrand's Theorem I.3:

Theorem I.6.

$$\mathbb{E}_x \left[\sqrt{I_f^-(x)} \right] \geqslant \tilde{\Omega}(\varepsilon(f)).$$

Here the notation $\tilde{\Omega}(\varepsilon(f))$ hides factors that are poly-logarithmic in n and $\frac{1}{\varepsilon}$. The precise lower bound we obtain is $\Omega\left(\frac{\varepsilon(f)}{\log n + \log(1/\varepsilon(f))}\right)$. Note that unlike previous theorems, our lower bound has dependence on the dimension n, which might just be an artifact of our proof method and not inherent.

Just like the undirected case, it is easily observed that Theorem I.6 implies Theorem I.5 (up to the poly-log factor), which in turn implies Theorem I.4. We note that even though an informal analogy holds between the theorems in the undirected and directed settings, the proofs in the directed setting are completely different and much more involved (as an aside, we do show that the theorems in the directed setting imply the corresponding theorems in the undirected setting and hence are more general). One difficulty is that the parameter $\varepsilon(f)$ is not too friendly to work with (as opposed to its analogue $\mathsf{var}(f)$). In particular, there is no straightforward way to characterize or estimate $\varepsilon(f)$. Proofs of Theorems I.4, I.5, I.6 proceed in reverse: assuming an upper bound on the L.H.S. of the respective inequality, one gives a sequence of transformations that turns the given function f into a monotone function and hence upper bounding $\varepsilon(f)$.

We also remark that our proof of Theorem I.6 is very different from that of Theorems I.4 and I.5 and involves several new technical ingredients that might be useful towards further research. In particular, our proof does not use routing schemes on the hypercube as in [10], [3] and instead relies on a new "split operator" on functions. The proof involves applying the split operator on random restrictions of f.

Towards an application to the monotonicity testing problem, Chakrabarty and Seshadhri [3] actually need and prove a stronger form of Theorem I.5. Let $\Gamma_{f,\mathsf{matching}}^-$ denote the size of the maximum matching among the edges in $\mathcal{S}_f^-$ (divided by a normalizing factor of 2^n), which is clearly at most Γ_f^- since the endpoints x of the matching with $f(x) = 1$ are also points on the negative vertex boundary. Chakrabarty and Seshadhri [3] show that:

Theorem I.7.

$$I_f^- \cdot \Gamma_{f,\mathsf{matching}}^- \;\geqslant\; \Omega(\varepsilon(f)^2).$$

In this paper, we are faced with a similar issue. We do not know how to use Theorem I.6 directly towards an application to the monotonicity testing problem. Also, we do not know how to deduce Theorem I.7 from Theorem I.6. However it turns out that a "robust" version holds both for Theorem I.3 (i.e. the undirected case) and Theorem I.6 (i.e. the directed case). The latter is now enough for our application to the monotonicity testing problem and if one wishes, to deduce Theorem I.7 (up to the poly-log factor). Since the specific robust version wasn't considered before, we first describe it in an undirected setting.

C. Robust version of Talagrand's Theorem

The robust version concerns the scenario when the sensitive edges are colored with two colors, red or blue. Let $\mathsf{col} : \mathcal{S}_f \mapsto \{\mathsf{red}, \mathsf{blue}\}$ be an arbitrary 2-coloring of the edges in $\mathcal{S}_f$. For $x \in \{0,1\}^n$, let $I_{f,\mathsf{red}}(x)$ be equal to 0 if $f(x) = 0$ and equal to the number of red sensitive edges incident on x if $f(x) = 1$. For $y \in \{0,1\}^n$, let $I_{f,\mathsf{blue}}(y)$ be equal to 0 if $f(y) = 1$ and equal to the number of blue sensitive edges incident on y if $f(y) = 0$. The robust version of Talagrand's Theorem I.3 is as follows:

Theorem I.8. *For a function* $f : \{0,1\}^n \mapsto \{0,1\}$ *and an arbitrary coloring* $\mathsf{col} : \mathcal{S}_f \mapsto \{\mathsf{red}, \mathsf{blue}\}$*,*

$$\mathbb{E}_x\left[\sqrt{I_{f,\mathsf{red}}(x)}\right] + \mathbb{E}_y\left[\sqrt{I_{f,\mathsf{blue}}(y)}\right] \geqslant \Omega(\mathsf{var}(f)).$$

We note that this theorem implies Theorem I.3 by considering the coloring that colors all sensitive edges red. The theorem is proved by adapting Talagrand's proof appropriately. Our presentation is a bit different (in addition to being a proof of the more general robust version) and more reader-friendly in our opinion. Also, the theorem is needed in the proof of the robust version of the directed analogue of Talagrand's Theorem (i.e. of Theorem I.6), stated next.

D. A Robust and Directed Analogue of Talagrand's Theorem

We finally state the robust and directed analogue of Talagrand's Theorem, which is what we really need towards an application to the monotonicity testing problem.

As before, let $\mathcal{S}_f^-$ denote the set of negatively sensitive edges. The robust version concerns the scenario when the negatively sensitive edges are colored with two colors, red and blue. Let $\mathsf{col} : \mathcal{S}_f^- \mapsto \{\mathsf{red}, \mathsf{blue}\}$ be an arbitrary 2-coloring of the edges in $\mathcal{S}_f^-$. For $x \in \{0,1\}^n$, let $I^-_{f,\mathsf{red}}(x)$ be equal to 0 if $f(x) = 0$ and equal to the number of red negatively sensitive edges incident on x if $f(x) = 1$. For $y \in \{0,1\}^n$, let $I^-_{f,\mathsf{blue}}(y)$ be equal to 0 if $f(y) = 1$ and equal to the number of blue negatively sensitive edges incident on y if $f(y) = 0$. The robust and directed analogue of Talagrand's Theorem is as follows:

Theorem I.9. *For a function* $f : \{0,1\}^n \mapsto \{0,1\}$ *and an arbitrary coloring* $\mathsf{col} : \mathcal{S}_f^- \mapsto \{\mathsf{red}, \mathsf{blue}\}$*,*

$$\mathbb{E}_x\left[\sqrt{I^-_{f,\mathsf{red}}(x)}\right] + \mathbb{E}_y\left[\sqrt{I^-_{f,\mathsf{blue}}(y)}\right] \geqslant \tilde{\Omega}(\varepsilon(f)).$$

Again the precise bound is $\Omega\left(\frac{\varepsilon(f)}{\log n + \log(1/\varepsilon(f))}\right)$. This theorem is proved by combining (part of) proof of Theorem I.6 along with a careful manipulation of underlying edge-coloring and the undirected robust version, i.e. Theorem I.8. The theorem implies Theorem I.6 by considering a coloring that colors all negatively sensitive edges red. It also implies Theorem I.7 (up to the poly-log factor).

E. Monotonicity Testing

As mentioned before, the monotonicity testing problem asks for a randomized algorithm that queries a given function $f : \{0,1\}^n \mapsto \{0,1\}$ at a few places and distinguishes whether the function is monotone or is far from being monotone. Let us focus on the case when the tester is non-adaptive, has perfect completeness and is a "pair tester" (all testers studied, including one in this paper, have all the three properties). Here non-adaptive means that the queries of the tester do not depend on the answers to the previous queries. Perfect completeness means that a monotone function must be accepted with probability 1. A "pair tester" picks a pair of inputs (x, y) from a pre-determined distribution such that y is monotonically above x and rejects if a violation to monotonicity is detected, i.e. if $f(x) = 1$ and $f(y) = 0$. For a pair tester, a measure of its quality is its rejection probability $\mathsf{rej}(n, \varepsilon(f))$ expressed in terms of n and the distance of f from the class of monotone functions. If one desires, one can (non-adaptively) repeat a pair tester $\frac{1}{\mathsf{rej}(n,\varepsilon(f))}$ times and achieve a constant rejection probability. Thus, the number of queries is often expressed as $\frac{1}{\mathsf{rej}(n,\varepsilon(f))}$, with a constant rejection probability as the stated goal.

Goldreich et al [9] present a pair tester that picks a uniformly random edge (x, y) of the hypercube (i.e. x and y differ in one co-ordinate). This is referred to as an "edge tester". The rejection probability is exactly $\frac{I_f^-}{n}$ and hence $\Omega(\frac{\varepsilon(f)}{n})$ by their Theorem I.4. Chakrabarty and Seshadhri [3] present a pair tester that picks a number $\tau \in \{1, 2, \ldots, \sqrt{n}\}$ with a certain distribution and then a pair (x, y) is picked, roughly uniformly, so that y is monotonically above x by a distance τ. This is referred to as a "path tester" and its rejection probability is $\tilde{\Omega}(\frac{\varepsilon(f)^{3/2}}{n^{7/8}})$. As far as dependence on n is concerned, this is the first improvement over the work of Goldreich et al [9], further improved to $\tilde{\Omega}(\frac{\varepsilon(f)^4}{n^{5/6}})$ by Chen et al [7]. The analysis of the tester relies on their Theorem I.5. In this paper, equipped with our Theorem I.9, we present and analyze a path tester[2] whose rejection probability is $\tilde{\Omega}(\frac{\varepsilon(f)^2}{\sqrt{n}})$:

Theorem I.10. *Given a function $f : \{0,1\}^n \mapsto \{0,1\}$, there is a path tester that is non-adaptive, has perfect completeness and rejection probability $\tilde{\Omega}(\frac{\varepsilon(f)^2}{\sqrt{n}})$.*

In next sections, we elaborate a bit on how Theorem I.9 leads to the said tester and then comment on the optimality of our tester.

1) Monotonicity Testing from Good Subgraphs: Given a function $f : \{0,1\}^n \mapsto \{0,1\}$, let $G_f^-(V, W, E)$ denote the bipartite graph of negatively sensitive edges, i.e. $V, W \subseteq \{0,1\}^n$, $\forall x \in V \; f(x) = 1$, $\forall y \in W \; f(y) = 0$, E is precisely the set of negatively sensitive edges $\mathcal{S}_f^-$, and every vertex in $V \cup W$ has at least one negatively sensitive edge incident on it.

Roughly speaking, Chakrabarty and Seshadhri [3] use their Theorem I.5 to deduce that the graph $G_f^-(V, W, E)$ has a large matching and analyze their tester with reference to this matching. We, on the other hand, use our Theorem I.9 to deduce that the graph $G_f^-(V, W, E)$ has a "(K, d)-good subgraph" with appropriate parameters K and d (a matching corresponds to the case $d = 1$ and then K is the size of the matching). We analyze our tester with reference to this good subgraph. Here, a bipartite graph $G'(V', W', E')$ is called (K, d)-good if $|W'| = K$, every vertex in W' has degree d and every vertex in V' has degree at most $2d$ (or the symmetric case with the roles of the two sides of the bipartite graph reversed). Leaving out some important details and caveats, the analysis of our tester is informally stated as:

Theorem I.11. (Informal) *If for a function $f : \{0,1\}^n \mapsto \{0,1\}$, the graph $G_f^-(V, W, E)$ has a $(\sigma \cdot 2^n, d)$-good subgraph, then there is a pair tester with rejection probability $\tilde{\Omega}(\frac{\sigma^2 d}{\sqrt{n}})$.*

We use Theorem I.9 to deduce that $G_f^-(V, W, E)$ has a $(\sigma \cdot 2^n, d)$-good subgraph with $\sigma\sqrt{d} \geqslant \tilde{\Omega}(\varepsilon(f))$. Combined with the informal statement of our tester above, we get a tester with rejection probability

[2]Our path tester chooses τ uniformly from $\{1, 2, 4, 8, \ldots, 2^{\lfloor \frac{\log n}{2} \rfloor}\}$, attempting to guess the "correct" value for τ. A similar guess is made in [7] and their distribution of τ is, morally speaking, the same as ours. In [3], τ is chosen uniformly from $\{1, 2, 3, \ldots, n^{1/8}\}$ with probability $\frac{1}{2}$ and $\tau = 1$ with probability $\frac{1}{2}$. This apparent difference, however, is only because the authors did not try to guess τ, which is later fixed in [7].

$\tilde{\Omega}(\frac{\varepsilon(f)^2}{\sqrt{n}})$ as claimed. Additional new ingredients used are bounds on the total influence of the function and on the fraction of "non-persistent" inputs. We would like to emphasize that the analysis of our tester is qualitatively different and a bit simpler than that of Chakrabarty and Seshadhri [3]. We do not elaborate this point further, but as a demonstration, we note (omitting the proof) that just using the large matching as in [3] as a good subgraph, we already get a tester with rejection probability $\tilde{\Omega}(\frac{\varepsilon^{4/3}}{n^{5/6}})$, improving the bound in both [3], [7].

2) Lower Bounds for Monotonicity Testing: We now give an overview of lower bounds on the number of queries required by a monotonicity testing algorithm and compare our tester against these (and new) lower bounds. Towards a uniform comparison of known bounds, for a parameter ε, let us require that a tester rejects any function that is ε-far from being monotone with a constant probability. Seemingly, the dependence of the number of queries on the two parameters n and ε can be traded against each other, so the situation is a bit subtle.

Let us first consider the case of pair testers that are non-adaptive and have perfect completeness (the most interesting case in our opinion, especially since all known testers are of this kind). The tester of Goldreich et al [9] achieves $O(n/\varepsilon)$ queries and Briet et al [2] show that if a pair tester has $\frac{F(n)}{\varepsilon}$ query complexity, then the dependence on n must be $F(n) \geqslant \Omega(n)$. We show that if a pair tester makes $O(n^\alpha/\varepsilon^\beta)$ queries, then $\alpha + \frac{\beta}{2} \geqslant \frac{3}{2}$. This follows from:

Theorem I.12. *For $\varepsilon = \Theta(1/\sqrt{n})$, a pair tester that is non-adaptive, has perfect completeness and rejects a function that is ε-far from being monotone with constant probability must make $\Omega(n^{3/2})$ queries.*

We note that for any $\alpha, \beta \geqslant 0$ such that $\alpha + \frac{\beta}{2} \geqslant \frac{3}{2}$, we have $\frac{n^\alpha}{\varepsilon^\beta} \geqslant \min\left\{\frac{n}{\varepsilon}, \frac{\sqrt{n}}{\varepsilon^2}\right\}$. Hence, for any setting of α, β, either the $O(n/\varepsilon)$-tester of Goldreich et al [9] or our $\tilde{O}(\sqrt{n}/\varepsilon^2)$-tester performs as well as a potential $O(n^\alpha/\varepsilon^\beta)$-tester. Thus, our tester in conjunction with Goldreich et al's tester is optimal. Also, if only the dependence on n is concerned (which is more interesting in our opinion than the dependence on ε), our tester is optimal even if compared against testers that possibly have imperfect completeness and not necessarily pair testers (see below).

Now we turn to more general testers, where there are still gaps between the upper and lower bounds. We already stated all the upper bounds before. We do not know a scenario where it helps to be adaptive, have imperfect completeness, or not be a pair tester. From the lower bound side, if a tester is non-adaptive and has perfect completeness (but is not necessarily a pair tester), a lower bound of $\Omega(\sqrt{n})$ is shown by Fischer et al [8] for a constant ε. For non-adaptive testers that possibly have imperfect completeness, a lower bound of $\tilde{\Omega}(n^{1/5})$ is shown by Chen et al [7] for a constant ε and further improved to $\Omega(n^{\frac{1}{2}-o(1)})$ by Chen et al [6]. The lower bounds in [8], [7], [6] for non-adaptive testers immediately imply a lower bound of $\Omega(\log n)$ for possibly adaptive testers.

II. Acknowledgment

We sincerely thank Andrej Bogdanov, Deeparnab Chakrabarty, Seshadhri Comandur, Oded Goldreich, Guy Kindler, Dana Ron, Ronitt Rubinfeld, Rocco Servedio, and Omri Weinstein for their valuable comments on an earlier draft of the paper and many discussions over the years.

Authors Khot and Safra acknowledge hospitality of the French Roast cafe (78 West 11th Street, New York) while working on the paper.

References

[1] Nir Ailon, Bernard Chazelle, Seshadhri Comandur, and Ding Liu. Estimating the distance to a monotone function. *Random Struct. Algorithms*, 31(3):371–383, 2007.

[2] Jop Briët, Sourav Chakraborty, David García-Soriano, and Arie Matsliah. Monotonicity testing and shortest-path routing on the cube. *Combinatorica*, 32(1):35–53, 2012.

[3] Deeparnab Chakrabarty and C. Seshadhri. A o(n) monotonicity tester for boolean functions over the hypercube. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 411–418, 2013.

[4] Deeparnab Chakrabarty and C. Seshadhri. Optimal bounds for monotonicity and lipschitz testing over hypercubes and hypergrids. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 419–428, 2013.

[5] Deeparnab Chakrabarty and C. Seshadhri. An optimal lower bound for monotonicity testing over hypergrids. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 16th International Workshop, APPROX 2013, and 17th International Workshop, RANDOM 2013, Berkeley, CA, USA, August 21-23, 2013. Proceedings*, pages 425–435, 2013.

[6] Xi Chen, Anindya De, Rocco Servedio, and Li-Yang Tan. Boolean function monotonicity testing requires (almost) $n^{1/2}$ non-adaptive queries. *arXiv, Report 1412.5657*, 2014.

[7] Xi Chen, Rocco Servedio, and Li-Yang Tan. New algorithms and lower bounds for monotonicity testing. In *Annual Symposium on Foundations of Computer Science, FOCS '14*, 2014.

[8] Eldar Fischer, Eric Lehman, Ilan Newman, Sofya Raskhodnikova, Ronitt Rubinfeld, and Alex Samorodnitsky. Monotonicity testing over general poset domains. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 474–483, 2002.

[9] Oded Goldreich, Shafi Goldwasser, Eric Lehman, Dana Ron, and Alex Samorodnitsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000.

[10] Eric Lehman and Dana Ron. On disjoint chains of subsets. *J. Comb. Theory, Ser. A*, 94(2):399–404, 2001.

[11] G. A. Margulis. Probabilistic characteristics of graphs with large connectivity. *Problemy Peredači Informacii*, 10(2):101–108, 1974.

[12] Dana Ron, Ronitt Rubinfeld, Muli Safra, Alex Samorodnitsky, and Omri Weinstein. Approximating the influence of monotone boolean functions in o($\sqrt{}$n) query complexity. *TOCT*, 4(4):11, 2012.

[13] M. Talagrand. Isoperimetry, logarithmic Sobolev inequalities on the discrete cube, and Margulis' graph connectivity theorem. *Geom. Funct. Anal.*, 3(3):295–314, 1993.