Electronic Colloquium on Computational Complexity, Report No. 11 (2015)

On Monotonicity Testing and Boolean Isoperimetric type Theorems

Subhash Khot* Dor Minzer† Muli Safra‡

January 22, 2015

Abstract

We show a directed and robust analogue of a boolean isoperimetric type theorem of Talagrand [15]. 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.

1 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 [15], generalizing many prior related theorems and (2) a monotonicity tester that is optimal in terms of its query complexity (see Section 1.5.2 for subtle issues regarding its optimality).

1.1 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 1.1.

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

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

*Department of Computer Science, Courant Institute of Mathematical Sciences, New York University. Research supported by NSF grants CCF 1422159, 1061938, 0832795 and Simons Collaboration on Algorithms and Geometry grant.

†School of Computer Science, Tel Aviv University.

‡School of Computer Science, Tel Aviv University.

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

ISSN 1433-8092

Theorem 1.2.

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

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 [15] shows that:

Theorem 1.3.

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

It is easily seen that Theorem 1.1 is implied by Theorem 1.2 which in turn is implied by Theorem 1.3. 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

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

1.2 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 [11, 3, 7, 2, 9, 1, 14, 4, 5, 12]. 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 [11] show that:

Theorem 1.4.

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

A hypercube can be thought of as a directed graph by orienting all its edges "monotonically upwards". In hindsight, Theorem 1.4 is viewed as a "directed" analogue of Theorem 1.1, 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 1.2:

Theorem 1.5.

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

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 1.3:

Theorem 1.6.

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

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 1.6 implies Theorem 1.5 (up to the poly-log factor), which in turn implies Theorem 1.4. 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, see Section 9.4). 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 1.4, 1.5, 1.6 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 1.6 is very different from that of Theorems 1.4 and 1.5 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 [12, 3] and instead relies on a new "split operator" on functions. The split operator and its properties are presented in Section 3 and the main proof appears in Section 4. 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 1.5. 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 1.7.

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

In this paper, we are faced with a similar issue. We do not know how to use Theorem 1.6 directly towards an application to the monotonicity testing problem. Also, we do not know how to deduce Theorem 1.7 from Theorem 1.6. However it turns out that a "robust" version holds both for Theorem 1.3 (i.e. the undirected case) and Theorem 1.6 (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 1.7 (up to the poly-log factor). Since the specific robust version wasn't considered before, we first describe it in an undirected setting.

1.3 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 1.3 is as follows:

Theorem 1.8. *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 1.3 by considering the coloring that colors all sensitive edges red. The theorem is proved by adapting Talagrand's proof appropriately, see Section 2. 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 1.6), stated next.

1.4 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 1.9. *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 1.6 along with a careful manipulation of underlying edge-coloring and the undirected robust version, i.e. Theorem 1.8. The theorem implies Theorem 1.6 by considering a coloring that colors all negatively sensitive edges red. It also implies Theorem 1.7 (up to the poly-log factor), see Section 9.3.

1.5 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 [11] 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 1.4. Chakrabarty and Seshadhri [3] present a pair tester that picks a number $\tau \in \{1, 2, \dots, \sqrt{n}\}$ 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 [11], 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 1.5. In this paper, equipped with our Theorem 1.9, we present and analyze a path tester[2] whose rejection probability is $\tilde{\Omega}(\frac{\varepsilon(f)^2}{\sqrt{n}})$:

Theorem 1.10. *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 1.9 leads to the said tester and then comment on the optimality of our tester.

1.5.1 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 1.5 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 1.9 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'

[2] Our path tester chooses τ uniformly from $\{1, 2, 4, 8, \dots, 2^{\lfloor \frac{\log n}{2} \rfloor}\}$, 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 $\{1, 2, 3, \dots, n^{1/8}\}$ 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].

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 1.11. **(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 1.9 to deduce that $G_f^-(V, W, E)$ has a $(\sigma \cdot 2^n, d)$-good subgraph with $\sigma\sqrt{d} \geqslant \tilde{\Omega}(\varepsilon(f))$ (see Section 6 for the necessary combinatorial argument). Combined with the informal statement of our tester above, we get a tester with rejection probability $\tilde{\Omega}(\frac{\varepsilon(f)^2}{\sqrt{n}})$ as claimed. Additional new ingredients used are bounds on the total influence of the function and on the fraction of "non-persistent" inputs (see Sections 9.1 and 9.2 respectively). 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].

1.5.2 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 [11] 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 1.12. *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 [11] 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 [9] 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 [9, 7, 6] for non-adaptive testers immediately imply a lower bound of $\Omega(\log n)$ for possibly adaptive testers.

2 Proof of Theorem 1.8

In this section, we present a proof of Theorem 1.8. The proof is an easy adaptation of Talagrand's proof [15] of Theorem 1.3. Our presentation is a bit different and more reader-friendly in our opinion.

We recall that for a function $f : \{0,1\}^n \mapsto \{0,1\}$, $\mathcal{S}_f$ denotes the set of sensitive edges. Let $\mathsf{col} : \mathcal{S}_f \mapsto \{\mathsf{red}, \mathsf{blue}\}$ be an arbitrary 2-coloring of the edges in $\mathcal{S}_f$. For $x \in \{0,1\}^n$, $I_{f,\mathsf{red}}(x)$ is equal to 0 if $f(x) = 0$ and equal to the number of red sensitive edges incident on x if $f(x) = 1$. For $y \in \{0,1\}^n$, $I_{f,\mathsf{blue}}(y)$ is equal to 0 if $f(y) = 1$ and equal to the number of blue sensitive edges incident on y if $f(y) = 0$. We intend to show that

$$\mathbb{E}_x\left[\sqrt{I_{f,\mathsf{red}}(x)}\right] + \mathbb{E}_y\left[\sqrt{I_{f,\mathsf{blue}}(y)}\right] \geqslant \Omega(\mathsf{var}(f)). \tag{1}$$

The proof is by induction on n. We show that the L.H.S. of the inequality above is lower bounded by $\frac{1}{8} \cdot \mathsf{var}(f)$. The claim is correct when $n = 1$. For any $n \geqslant 2$, let $f : \{0,1\}^n \mapsto \{0,1\}$ be the given function and $f_0 = f(0, x_2, \ldots, x_n)$ and $f_1 = f(1, x_2, \ldots, x_n)$ be the two sub-functions on $n-1$ co-ordinates. Let $\mathcal{H}_0$ and $\mathcal{H}_1$ denote the "bottom" and the "top" hypercubes on which the functions f_0 and f_1 are defined so that $\{0,1\}^n = \mathcal{H}_0 \cup \mathcal{H}_1$. Let $\mathcal{I}$ be the set of sensitive edges along the first co-ordinate so that:

$$\mathcal{S}_f = \mathcal{S}_{f_0} \cup \mathcal{S}_{f_1} \cup \mathcal{I}.$$

Also, a given coloring of edges in $\mathcal{S}_f$ induces a coloring of edges in $\mathcal{S}_{f_0}$, $\mathcal{S}_{f_1}$ and $\mathcal{I}$. Let $s = \frac{|\mathcal{I}|}{2^{n-1}}$ be the influence of the first co-ordinate. By Lemma 2.1,

$$\mathsf{var}(f) \;\leqslant\; \frac{1}{2} \cdot \mathsf{var}(f_0) + \frac{1}{2} \cdot \mathsf{var}(f_1) + \frac{1}{4} \cdot s^2.$$

The main idea is to express the L.H.S. of Equation (1), denoted T, as a sum of four terms A, B, C, D denoting similar quantities from functions f_0 and f_1 and in addition, a term Φ that corresponds to the "incremental contribution" from the edges in $\mathcal{I}$ (along the first co-ordinate). Towards this end, let

$$A = \mathbb{E}_x\left[\sqrt{I_{f_0,\mathsf{red}}(x)}\right], \quad B = \mathbb{E}_y\left[\sqrt{I_{f_0,\mathsf{blue}}(y)}\right],$$

$$C = \mathbb{E}_z\left[\sqrt{I_{f_1,\mathsf{red}}(z)}\right], \quad D = \mathbb{E}_w\left[\sqrt{I_{f_1,\mathsf{blue}}(w)}\right].$$

In the following x, y will always denote vertices in the bottom hypercube $\mathcal{H}_0$ and z, w will always denote vertices in the top hypercube $\mathcal{H}_1$. Also, $f(x) = f(z) = 1$ and we are concerned about the number of red edges incident on them. Similarly, $f(y) = f(w) = 0$ and we are concerned about the number of blue edges incident on them. We may assume, inductively, that

$$A + B \;\geqslant\; \frac{1}{8} \cdot \mathsf{var}(f_0),$$

$$C + D \;\geqslant\; \frac{1}{8} \cdot \mathsf{var}(f_1).$$

It is clear that T, i.e. the L.H.S. of Equation (1), equals

$$T = \frac{1}{2}(A+B) + \frac{1}{2}(C+D) + \Phi,$$

where Φ denotes the incremental contribution from edges in $\mathcal{I}$ as follows. The edges in $\mathcal{I}$ have the effect that for vertices $x, y \in \mathcal{H}_0$ and $z, w \in \mathcal{H}_1$, the relevant degrees $I_{f_0,\mathsf{red}}(x)$, $I_{f_0,\mathsf{blue}}(y)$, $I_{f_1,\mathsf{red}}(z)$, $I_{f_1,\mathsf{blue}}(w)$ may increase by one. Each (colored) edge in $\mathcal{I}$ leads to an increase in the degree of exactly one vertex x, y, z or w. Partition the edges in $\mathcal{I}$ as $\mathcal{I} = \mathcal{I}_A \cup \mathcal{I}_B \cup \mathcal{I}_C \cup \mathcal{I}_D$ depending on whether an edge leads to an increase in the degree of x, y, z or w "type" of vertex. Defining $s_A = \frac{|\mathcal{I}_A|}{2^{n-1}}$ and similarly s_B, s_C, s_D, we have $s = s_A + s_B + s_C + s_D$.

Let us focus on the edges in $\mathcal{I}_A$. For $d = 1, 2, \ldots, n$, consider the vertices $x \in \mathcal{H}_0$ such that the edges in $\mathcal{I}_A$ increase their degree from $d-1$ to d and let the number of such vertices be $s_d \cdot 2^{n-1}$. Specifically, for such a vertex $I_{f_0,\mathsf{red}}(x) = d-1$ and $I_{f,\mathsf{red}}(x) = d$ (when x is viewed as a vertex from $\{0,1\}^n = \mathcal{H}_0 \cup \mathcal{H}_1$). The incremental contribution of x is:

$$\sqrt{I_{f,\mathsf{red}}(x)} - \sqrt{I_{f_0,\mathsf{red}}(x)} = \sqrt{d} - \sqrt{d-1} \;\geqslant\; \frac{1}{4\sqrt{d}}.$$

We note that $s_A = \sum_{d=1}^{n} s_d$. Denoting by Φ_A the incremental contribution averaged over all $x \in \mathcal{H}_0$, we get

$$\Phi_A \;\geqslant\; \sum_{d=1}^{n} s_d \cdot \frac{1}{4\sqrt{d}}.$$

We also note that (by considering the contribution to T from only the x "type" vertices)

$$T \;\geqslant\; \frac{1}{2} \cdot \sum_{d=1}^{n} s_d \sqrt{d}.$$

By Cauchy-Schwartz, we have

$$\Phi_A \cdot T \geqslant \frac{1}{2} \cdot \left(\sum_{d=1}^{n} s_d \cdot \frac{1}{4\sqrt{d}} \right) \left(\sum_{d=1}^{n} s_d \sqrt{d} \right) \geqslant \frac{1}{8} \cdot \left(\sum_{d=1}^{n} s_d \right)^2 = \frac{1}{8} \cdot s_A^2.$$

If $T \geqslant \frac{1}{8}$ we are done, since we only intend to prove a lower bound of $T \geqslant \frac{1}{8} \cdot \mathsf{var}(f)$. Thus we may assume that $\Phi_A \geqslant s_A^2$. Similarly we may assume that $\Phi_B \geqslant s_B^2$, $\Phi_C \geqslant s_C^2$, $\Phi_D \geqslant s_D^2$ and

$$\Phi = \frac{1}{2} \cdot (\Phi_A + \Phi_B + \Phi_C + \Phi_D) \;\geqslant\; \frac{1}{2} \cdot (s_A^2 + s_B^2 + s_C^2 + s_D^2) \;\geqslant\; \frac{1}{8} \cdot (s_A + s_B + s_C + s_D)^2 = \frac{1}{8} s^2.$$

Finally, we get inductively that

$$\begin{aligned}
T &= \frac{1}{2}(A+B) + \frac{1}{2}(C+D) + \Phi \\
&\geqslant \frac{1}{2}\left(\frac{1}{8} \cdot \mathsf{var}(f_0)\right) + \frac{1}{2}\left(\frac{1}{8} \cdot \mathsf{var}(f_1)\right) + \frac{1}{8} \cdot s^2 \\
&\geqslant \frac{1}{8} \cdot \left(\frac{1}{2} \cdot \mathsf{var}(f_0) + \frac{1}{2} \cdot \mathsf{var}(f_0) + \frac{1}{4} \cdot s^2\right) \\
&\geqslant \frac{1}{8} \cdot \mathsf{var}(f).
\end{aligned}$$

This completes the inductive proof.

Lemma 2.1. *For a function $f : \{0,1\}^n \mapsto \{0,1\}$, let s denote the influence of the first co-ordinate and f_0 and f_1 denote the sub-functions on the remaining $n-1$ co-ordinates. Then*

$$\mathsf{var}(f) \leqslant \frac{1}{2} \cdot \mathsf{var}(f_0) + \frac{1}{2} \cdot \mathsf{var}(f_1) + \frac{1}{4} \cdot s^2.$$

Proof. Let $p_0 = \mathbb{E}_x[f_0(x)]$ and $p_1 = \mathbb{E}_x[f_1(x)]$ so that $\mathbb{E}_z[f(z)] = \frac{p_0+p_1}{2}$. We have, by definition, $\mathsf{var}(f_0) = p_0(1-p_0)$ and $\mathsf{var}(f_1) = p_1(1-p_1)$. Thus we have the identity:

$$\mathsf{var}(f) = \frac{p_0+p_1}{2}\left(1 - \frac{p_0+p_1}{2}\right) = \frac{1}{2} \cdot \mathsf{var}(f_0) + \frac{1}{2} \cdot \mathsf{var}(f_1) + \frac{1}{4}(p_0 - p_1)^2.$$

The proof is complete by observing that

$$|p_0 - p_1| = |\mathbb{E}_x[f_0(x)] - \mathbb{E}_x[f_1(x)]| \leqslant \mathbb{E}_x[|f_0(x) - f_1(x)|] = s.$$

□

3 The Switch and the Split Operators

Goldreich et al [11] define a "switch operator" S_i for a co-ordinate $i \in \{1, \ldots, n\}$ that constructs a function $S_i(f) : \{0,1\}^n \mapsto \{0,1\}$ from a given function $f : \{0,1\}^n \mapsto \{0,1\}$. In this paper, we define a new operator that we call a "split operator" ∇_i for a co-ordinate $i \in \{1, \ldots, n\}$ that constructs a function $\nabla_i(f) : \{0,1\}^{n+1} \mapsto \{0,1\}$ from a given function $f : \{0,1\}^n \mapsto \{0,1\}$. Note that $\nabla_i(f)$ is a function of $n+1$ co-ordinates. Both the operators S_i and ∇_i are "applied" on co-ordinate i and can be sequentially applied on co-ordinates 1 through n in any desired order. The operators are non-commutative in the sense that the resulting function, in general, depends on the order in which the operators are applied on multiple co-ordinates. In this section, we define both the operators and prove several important properties of the latter.

3.1 The Switch Operator

Definition 3.1. *For a function $f\colon \{0,1\}^n \to \{0,1\}$ and a co-ordinate $i \in \{1, \ldots, n\}$, the function $S_i(f)\colon \{0,1\}^n \to \{0,1\}$ is defined as follows. For each $\alpha \in \{0,1\}^{i-1}$, $\beta \in \{0,1\}^{n-i}$, and a co-ordinate variable $x_i \in \{0,1\}$,*

$$S_i(f)(\alpha, x_i, \beta) \stackrel{def}{=} \begin{cases} \min\{f(\alpha,0,\beta), f(\alpha,1,\beta)\} & \text{if } x_i = 0 \\ \max\{f(\alpha,0,\beta), f(\alpha,1,\beta)\} & \text{if } x_i = 1. \end{cases}$$

In the following, we will abuse notation for the sake of conciseness and convenience: with the understanding that for $x_{-i} \in \{0,1\}^{n-1}$, written as $x_{-i} = \alpha \circ \beta$, $\alpha \in \{0,1\}^{i-1}, \beta \in \{0,1\}^{n-i}$, and that the co-ordinate x_i is inserted in the i^{th} position as $\alpha \circ x_i \circ \beta$, we will write the above definition of the switch operator as:

$$S_i(f)(x_i, x_{-i}) \stackrel{def}{=} \begin{cases} \min\{f(0,x_{-i}), f(1,x_{-i})\} & \text{if } x_i = 0 \\ \max\{f(0,x_{-i}), f(1,x_{-i})\} & \text{if } x_i = 1. \end{cases}$$

The switch operator considers the edges of the hypercube along the i^{th} dimension and for every edge that violates the monotonicity of f in that dimension, switches the values of the function at the two endpoints of that edge. The "corrected" function $S_i(f)$ is now monotone along the i^{th} dimension. A remarkable property of the switch operator, as shown by Goldreich et al [11], is that if another switch is now applied on co-ordinate $j \neq i$, the resulting function (i.e. $S_j(S_i(f))$) stays monotone along the i^{th} dimension. In particular, if the switch operator is applied on co-ordinates 1 through n, one after another in some order, the final function is monotone. Another remarkable property is that the Hamming distance between a pair of functions can only decrease after applying the switch operator to both. We state and prove these properties below.

Lemma 3.2. *If $f : \{0,1\}^n \mapsto \{0,1\}$ is monotone in the j^{th} co-ordinate, then so is $S_i(f)$ for a co-ordinate $i \in \{1,\ldots,n\}, i \neq j$.*

Proof. It is enough to fix a setting $x_{-i-j} \in \{0,1\}^{n-2}$ of co-ordinates except i and j and consider the behavior of four values

$$a = f(x_i = 0, x_j = 0, x_{-i-j}), \quad b = f(x_i = 1, x_j = 0, x_{-i-j}),$$
$$c = f(x_i = 0, x_j = 1, x_{-i-j}), \quad d = f(x_i = 1, x_j = 1, x_{-i-j}).$$

Monotonicity of f in the j^{th} co-ordinate implies that $a \leqslant c,\ b \leqslant d$. After the switch on the i^{th} co-ordinate, the four values change to

$$S_i(f)(0,0,x_{-i-j}) = \min\{a,b\}, \quad S_i(f)(1,0,x_{-i-j}) = \max\{a,b\},$$
$$S_i(f)(0,1,x_{-i-j}) = \min\{c,d\}, \quad S_i(f)(1,1,x_{-i-j}) = \max\{c,d\}.$$

Now the monotonicity of $S_i(f)$ in the j^{th} co-ordinate amounts to saying that

$$\min\{a,b\} \leqslant \min\{c,d\}, \quad \max\{a,b\} \leqslant \max\{c,d\},$$

which follows since $\min\{\cdot,\cdot\}, \max\{\cdot,\cdot\}$ are monotone functions and $(a,b) \leqslant (c,d)$ by hypothesis. $\square$

Definition 3.3. *For functions $f_1, f_2 : \{0,1\}^n \mapsto \{0,1\}$, their Hamming distance is*

$$\Delta(f_1, f_2) \stackrel{def}{=} \mathbb{E}_x \left[\mathbf{1}_{f_1(x) \neq f_2(x)}\right].$$

Lemma 3.4. *For functions $f, h : \{0,1\}^n \mapsto \{0,1\}$ and a co-ordinate $i \in \{1,\ldots,n\}$,*

$$\Delta(S_i(f), S_i(h)) \leqslant \Delta(f,h).$$

Proof. For any fixed setting of $x_{-i} \in \{0,1\}^{n-1}$, we show that the inequality holds as far as contribution from that setting of x_{-i} is concerned. Indeed, denoting

$$a = f(0, x_{-i}), \quad b = f(1, x_{-i}),$$
$$c = h(0, x_{-i}), \quad d = h(1, x_{-i}),$$

we see that the contribution to the R.H.S. is

$$\mathbf{1}_{a \neq c} + \mathbf{1}_{b \neq d},$$

whereas the contribution to the L.H.S. is

$$\mathbf{1}_{\min\{a,b\} \neq \min\{c,d\}} + \mathbf{1}_{\max\{a,b\} \neq \max\{c,d\}}.$$

The inequality can be checked by case analysis. For example, if $a = b$ then L.H.S = R.H.S. The same holds if $c = d$ or if $(a,b) = (c,d)$. The only remaining case is when $(a,b) = (0,1) \wedge (c,d) = (1,0)$ (and the other way round) in which case the Hamming distance actually decreases. $\square$

We recall that for a function f, the parameter $\varepsilon(f)$ denotes the distance of f from the class of monotone functions, i.e. the minimum fraction of values of f that need to be changed to turn f into a monotone function. The parameter $\varepsilon(f)$ seems difficult to characterize in a "constructive" manner. However it turns out that it can be well-approximated in a constructive manner. Goldreich et al [11] show, thanks to Lemma 3.2, that if the switch operator is applied to a function f, on co-ordinates 1 through n, say in that order, then the resulting function is monotone. We observe that this transformation is efficient in the sense that the fraction of values of f changed is at most $2\,\varepsilon(f)$. This observation also appears in a paper of Fattal and Ron [8, Lemma 4.3], but several researchers (including us) seemed unaware of this (thanks to Andrej Bogdanov for pointing out).

Lemma 3.5. *For any function* $f : \{0,1\}^n \mapsto \{0,1\}$,

$$\varepsilon(f) \leqslant \Delta(f, S_n(S_{n-1}(\ldots S_2(S_1(f))\ldots))) \leqslant 2\,\varepsilon(f).$$

Proof. By a repeated application of Lemma 3.2, the function $S_n(S_{n-1}(\ldots S_2(S_1(f))\ldots))$ is monotone and hence the left inequality holds by definition of $\varepsilon(f)$. Towards the right inequality, let h be a monotone function such that $\Delta(f,h) = \varepsilon(f)$. By a repeated application of Lemma 3.4,

$$\Delta(S_n(S_{n-1}(\ldots S_2(S_1(f))\ldots)), S_n(S_{n-1}(\ldots S_2(S_1(h))\ldots))) \leqslant \Delta(f,h).$$

We note however that since h is already monotone, applying switch operators keeps it unaffected. Thus, the above inequality is same as

$$\Delta(S_n(S_{n-1}(\ldots S_2(S_1(f))\ldots)), h) \leqslant \Delta(f,h) = \varepsilon(f).$$

Now by triangle inequality,

$$\Delta(f, S_n(S_{n-1}(\ldots S_2(S_1(f))\ldots))) \leqslant \Delta(f,h) + \Delta(S_n(S_{n-1}(\ldots S_2(S_1(f))\ldots)), h) \leqslant 2\,\varepsilon(f).$$

□

In this paper, it will be useful to consider the scenario when the switch operator is applied to co-ordinates 1 through n in a random order. Let $\rho \in \mathbb{S}_n$ denote a permutation of $\{1, 2, \ldots, n\}$ and $\gamma(f)$ denote the expected fraction of values of f changed by selecting $\rho \in \mathbb{S}_n$ at random and then applying switches according to the order ρ. Since in the proof of Lemma 3.5, we didn't use the fact that the switches were in a specific order, it follows that $\gamma(f)$ is sandwiched between $\varepsilon(f)$ and $2\,\varepsilon(f)$.

Definition 3.6. *For a permutation* $\rho \in \mathbb{S}_n$, $\rho \circ f$ *denotes the function* $S_{\rho(n)}(S_{\rho(n-1)}(\ldots S_{\rho(2)}(S_{\rho(1)}(f))\ldots))$.

$$\gamma(f) \stackrel{def}{=} \mathbb{E}_{\rho \in \mathbb{S}_n}\left[\Delta(f, \rho \circ f)\right].$$

We have

$$\varepsilon(f) \leqslant \gamma(f) \leqslant 2\,\varepsilon(f).$$

3.2 The Split Operator

Now we define the split operator ∇_i applied on a co-ordinate $i \in \{1, \ldots, n\}$. For a function $f\colon \{0,1\}^n \mapsto \{0,1\}$, the function $\nabla_i(f)$ is a function of $n+1$ co-ordinates. It is best to think that the i^{th} co-ordinate is now split into two co-ordinates indexed as $(i,+)$ and $(i,-)$.

Definition 3.7. *For a function $f\colon \{0,1\}^n \to \{0,1\}$ and a co-ordinate $i \in \{1,\ldots,n\}$, the function $\nabla_i(f)$ is defined as follows. For each $x_{-i} \in \{0,1\}^{n-1}$*

$$\nabla_i(f)(x_{i,+},x_{i,-},x_{-i}) \stackrel{def}{=} \begin{cases} f(0,x_{-i}) & \text{if} \quad x_{i,+}=0, x_{i,-}=0 \\ \min\{f(0,x_{-i}),f(1,x_{-i})\} & \text{if} \quad x_{i,+}=0, x_{i,-}=1 \\ \max\{f(0,x_{-i}),f(1,x_{-i})\} & \text{if} \quad x_{i,+}=1, x_{i,-}=0 \\ f(1,x_{-i}) & \text{if} \quad x_{i,+}=1, x_{i,-}=1. \end{cases}$$

If the function f is written as $f(x_1,\ldots,x_{i-1},x_i,x_{i+1},\ldots,x_n)$, then the function $\nabla_i(f)$ is written as $(\nabla_i(f))(x_1,\ldots,x_{i-1},x_{i,+},x_{i,-},x_{i+1},\ldots,x_n)$. In particular, the co-ordinates of $\nabla_i(f)$ have the same indices as that of f, except that the co-ordinate i is split into two co-ordinates indexed as $(i,+)$ and $(i,-)$. We start with a preliminary observation about the split operator.

Lemma 3.8. *Let $f\colon \{0,1\}^n \to \{0,1\}$, and $i,j \in \{1,\ldots,n\}$ be co-ordinates such that $i \neq j$. Then:*

1. *$\nabla_i(f)$ is monotone in co-ordinate $(i,+)$.*
2. *$\nabla_i(f)$ is anti-monotone in co-ordinate $(i,-)$.*
3. *If f is monotone in co-ordinate j, then so is $\nabla_i(f)$.*
4. *If f is anti-monotone in co-ordinate j, then so is $\nabla_i(f)$.*

Proof. We verify that $\nabla_i(f)$ is monotone in co-ordinate $(i,+)$ by verifying its monotonicity for every fixing of $x_{-i} \in \{0,1\}^{n-1}$ and $x_{i,-} \in \{0,1\}$. Fix some $x_{-i} \in \{0,1\}^{n-1}$. When $x_{i,-}=0$, by definition,

$$\begin{aligned} \nabla_i(f)(x_{i,+}=0,x_{i,-}=0,x_{-i}) &= f(0,x_{-i}) \\ \nabla_i(f)(x_{i,+}=1,x_{i,-}=0,x_{-i}) &= \max\{f(0,x_{-i}),f(1,x_{-i})\}, \end{aligned}$$

and the first value is less or equal the second. When $x_{i,-}=1$, by definition,

$$\begin{aligned} \nabla_i(f)(x_{i,+}=0,x_{i,-}=1,x_{-i}) &= \min\{f(0,x_{-i}),f(1,x_{-i})\} \\ \nabla_i(f)(x_{i,+}=1,x_{i,-}=1,x_{-i}) &= f(1,x_{-i}), \end{aligned}$$

and again, the first value is less or equal the second. This confirms that $\nabla_i(f)$ is monotone in co-ordinate $(i,+)$. Similarly we can confirm that $\nabla_i(f)$ is anti-monotone in co-ordinate $(i,-)$.

Now suppose that f is monotone in co-ordinate j. We will show that $\nabla_i(f)$ is also monotone in co-ordinate j. We write $\nabla_i(f)$ as $\nabla_i(f)(x_{i,+},x_{i,-},x_j,x_{-i-j})$ where x_{-i-j} denotes all co-ordinates except i and j. We verify monotonicity of $\nabla_i(f)$ in co-ordinate j by going over all four fixings of $x_{i,+}$ and $x_{i,-}$. When $x_{i,+}=x_{i,-}=0$, we have

$$\nabla_i(f)(x_{i,+}=0,x_{i,-}=0,x_j,x_{-i-j}) = f(x_i=0,x_j,x_{-i-j}),$$

and monotonicity of $\nabla_i(f)$ in co-ordinate j follows from that of f. Similarly we can verify the case $x_{i,+}=x_{i,-}=1$. Now consider the case $x_{i,+}=1$ and $x_{i,-}=0$. In this case,

$$\begin{aligned} \nabla_i(f)(x_{i,+}=1,x_{i,-}=0,x_j=0,x_{-i-j}) &= \max\{f(x_i=0,x_j=0,x_{-i-j}),f(x_i=1,x_j=0,x_{-i-j})\}, \\ \nabla_i(f)(x_{i,+}=1,x_{i,-}=0,x_j=1,x_{-i-j}) &= \max\{f(x_i=0,x_j=1,x_{-i-j}),f(x_i=1,x_j=1,x_{-i-j})\}, \end{aligned}$$

and using monotonicity of f in co-ordinate j, the first value is less or equal the second (as written, there appears to be a 2×2 array of values such that each column is increasing downwards; this implies that the maximum of the two values in the upper row is less or equal the maximum of the two values in the lower row). Lastly, the case when $x_{i,+} = 0$ and $x_{i,-} = 1$ is handled similarly, confirming the monotonicity of $\nabla_i(f)$ in co-ordinate j.

Finally, a similar argument shows that if f is anti-monotone in co-ordinate j, then so is $\nabla_i(f)$, completing the proof of the lemma. $\square$

Though we do not need the following fact, we do note that applying the split operator (just like the switch operator, see Lemma 3.4) can only decrease the distance to monotonicity.

Lemma 3.9.

$$\varepsilon(\nabla_i(f)) \leqslant \varepsilon(f).$$

Proof. Let h be a monotone function nearest to f, i.e. $\varepsilon(f) = \Delta(f, h)$. Since h is monotone, so is $\nabla_i(h)$ and it is enough to show that

$$\Delta(\nabla_i(f), \nabla_i(h)) \leqslant \Delta(f, h).$$

We observe that the inequality above holds for every fixed setting of $x_{-i} \in \{0,1\}^n$. Indeed, letting

$$a_0 = f(x_i = 0, x_{-i}), \quad a_1 = f(x_i = 1, x_{-i}),$$

$$b_0 = h(x_i = 0, x_{-i}), \quad b_1 = h(x_i = 1, x_{-i}),$$

the contribution to $\Delta(f, h)$ is $\Delta((a_0, a_1), (b_0, b_1))$, whereas the contribution to $\Delta(\nabla_i(f), \nabla_i(h))$ is

$$\Delta\left((a_0, \min\{a_0, a_1\}, \max\{a_0, a_1\}, a_1),\ (b_0, \min\{b_0, b_1\}, \max\{b_0, b_1\}, b_1)\right).$$

It can be checked by easy case analysis that the former always dominates the latter. $\square$

3.3 Pure Functions

We will need to consider functions $g_{f,S,\rho}$ that are obtained from function $f : \{0,1\}^n \mapsto \{0,1\}$ by applying the split operator on co-ordinates in some index set $S \subseteq \{1, \ldots, n\}$ in an order specified by a permutation ρ on the set S. In particular if $S = \{1, \ldots, n\}$ and ρ is arbitrary, then applying the split operator n times successively in the order ρ yields a function

$$g(x_{1,+}, x_{1,-}, x_{2,+}, x_{2,-}, \ldots, x_{n,+}, x_{n,-}),$$

which, by re-arranging the co-ordinates can be written as

$$g(x_{1,+}, x_{2,+}, \ldots, x_{n,+}, x_{1,-}, x_{2,-}, x_{n,-}).$$

We may write the function as $g(x, y)$ where $x, y \in \{0,1\}^n$, x denotes the '$+$' co-ordinates and y denotes the '$-$' co-ordinates. By Lemma 3.8, $g(x, y)$ is monotone in the x-co-ordinates and anti-monotone in the y-co-ordinates. Functions with this property will be important for us and we call them "pure".

Definition 3.10. *A function $g : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}$, written as $g(x, y)$, is called pure if it is monotone in every x-co-ordinate and anti-monotone in every y-co-ordinate.*

As we said, pure functions will arise as intermediate functions during our proof and the main point is that pure functions are also "simple" in the sense that several of the parameters and theorems are easily characterized and proved respectively for pure functions. We start with the observation that $\varepsilon(g)$ has a easy characterization for a pure function $g(x, y)$ (although this is not the case for a general function). For a fixed $x \in \{0,1\}^n$, let $g(x, \cdot)$ denote the restriction of g when the first argument is fixed to x.

Lemma 3.11. *For a pure function $g(x, y)$,*

$$\varepsilon(g) = \Theta\left(\mathop{\mathbb{E}}_{x}\left[\mathsf{var}(g(x, \cdot))\right]\right).$$

Proof. Define $p_x = \mathbb{E}_y\left[g(x, y)\right]$, and define a function $h(x, y)$ as

$$h(x, y) = \begin{cases} 1 & \text{if } p_x > \frac{1}{2} \\ 0 & \text{if } p_x \leqslant \frac{1}{2}. \end{cases}$$

We note that $h(x, y)$ depends only on x. Since g is monotone in x, whenever $x \leqslant x'$, it holds that $g(x, y) \leqslant g(x', y)$ and then taking expectation over y, one gets $p_x \leqslant p_{x'}$. Thus h is monotone and

$$\varepsilon(g) \leqslant \Delta(h, g) = \mathop{\mathbb{E}}_{x}\left[\Delta(h(x, \cdot), g(x, \cdot))\right] \leqslant 2 \cdot \mathop{\mathbb{E}}_{x}\left[\mathsf{var}(g(x, \cdot))\right].$$

On the other hand if h' is any monotone function, then since $g(x, \cdot)$ is anti-monotone, by Lemma 3.12,

$$\Delta(h', g) = \mathop{\mathbb{E}}_{x}\left[\Delta(h'(x, \cdot), g(x, \cdot))\right] \geqslant \mathop{\mathbb{E}}_{x}\left[\mathsf{var}(g(x, \cdot))\right].$$

□

Lemma 3.12. *If $\psi, \psi' : \{0,1\}^n \mapsto \{0,1\}$ are such that ψ is monotone and ψ' is anti-monotone, then*

$$\Delta(\psi, \psi') \geqslant \max\{\mathsf{var}(\psi), \mathsf{var}(\psi')\}.$$

Proof. Let $p = \mathbb{E}[\psi]$ and $p' = \mathbb{E}[\psi']$. We note that $(\psi, 1 - \psi')$ is a pair of monotone functions. The lemma follows using FKG inequality [10],

$$\begin{aligned} \Delta(\psi, \psi') &= \mathsf{Pr}[\psi = 1, \psi' = 0] + \mathsf{Pr}[\psi = 0, \psi' = 1] \\ &= \mathsf{Pr}[\psi = 1, (1 - \psi') = 1] + \mathsf{Pr}[\psi = 0, (1 - \psi') = 0] \\ &\geqslant p(1 - p') + (1 - p)p' \\ &\geqslant p(1 - p) + p'(1 - p') = \mathsf{var}(\psi) + \mathsf{var}(\psi'). \end{aligned}$$

□

We now show that Theorem 1.6 holds for a pure function $g(x, y)$, which has a simple enough structure that it follows immediately from the undirected version of the theorem, i.e. Talagrand's Theorem 1.3. We note that for a pure function, we do not lose any poly-log factor.

Lemma 3.13. *For a pure function $g(x, y)$,*

$$\mathop{\mathbb{E}}_{x,y}\left[\sqrt{I_g^-(x, y)}\right] \geqslant \Omega(\varepsilon(g)).$$

Proof. Since $g(x,y)$ is monotone in x, all its negatively sensitive edges are incident on the y-co-ordinates. Moreover, since $g(x,y)$ is anti-monotone in y, all sensitive edges incident on the y-co-ordinates are actually negatively sensitive. Thus,

$$\mathop{\mathbb{E}}_{x,y}\left[\sqrt{I_g^-(x,y)}\right] = \mathop{\mathbb{E}}_{x}\left[\mathop{\mathbb{E}}_{y}\left[\sqrt{I_{g(x,\cdot)}(y)}\right]\right].$$

Using Theorem 1.3 on function $g(x,\cdot)$ along with Lemma 3.11, we have

$$\mathop{\mathbb{E}}_{x}\left[\mathop{\mathbb{E}}_{y}\left[\sqrt{I_{g(x,\cdot)}(y)}\right]\right] \geqslant \mathop{\mathbb{E}}_{x}\left[\Omega(\mathsf{var}(g(x,\cdot)))\right] \geqslant \Omega(\varepsilon(g)).$$

□

Similarly, Theorem 1.9 holds for a pure function $g(x,y)$, which is noted as the lemma below. This follows in a similar manner from the undirected version of the theorem i.e. Theorem 1.8. We skip the proof. We note again that for a pure function, we do not lose any poly-log factor.

Lemma 3.14. *For a pure function $g(x,y)$ and for an arbitrary coloring $\mathsf{col} : \mathcal{S}_g^- \mapsto \{\mathsf{red}, \mathsf{blue}\}$,*

$$\mathbb{E}_{x,y}\left[\sqrt{I_{g,\mathsf{red}}^-(x,y)}\right] + \mathbb{E}_{x,y}\left[\sqrt{I_{g,\mathsf{blue}}^-(x,y)}\right] \geqslant \Omega(\varepsilon(g)).$$

Finally, we will need the notion of a strongly monotone function, which is defined only for functions of the type

$$h(x_{1,+}, x_{1,-}, x_{2,+}, x_{2,-}, \ldots, x_{n,+}, x_{n,-}),$$

i.e. for functions for which the co-ordinates are paired and in each pair, one co-ordinate is designated as '$+$' and the other as '$-$'.

Definition 3.15. *A function $h : \{0,1\}^{2n} \mapsto \{0,1\}$, written as,*

$$h(x_{1,+}, x_{1,-}, x_{2,+}, x_{2,-}, \ldots, x_{n,+}, x_{n,-}),$$

is strongly monotone if it is monotone and moreover, for any $i \in \{1,\ldots,n\}$, changing the pair of co-ordinates $((i,+),(i,-))$ from 01 to 10 can only change the function from 0 to 1.

For a function $g(x_{1,+}, x_{1,-}, x_{2,+}, x_{2,-}, \ldots, x_{n,+}, x_{n,-})$, let $\delta(g)$ denote the minimum distance of g from any strongly monotone function. Clearly $\delta(g) \geqslant \varepsilon(g)$, but it turns out that if $g = g(x,y)$ is a pure function, then $\delta(g)$ is same as $\varepsilon(g)$ up to a constant factor. Here, it is understood that we re-arrange the co-ordinates of g as

$$g(x_{1,+}, x_{2,+}, \ldots, x_{n,+}, x_{1,-}, x_{2,-}, \ldots, x_{n,-}),$$

and write $g = g(x,y)$ with x and y denoting the '$+$' inputs and '$-$' inputs respectively (and vice versa; from the representation $g(x,y)$ one can go back to the representation in terms of pairs of co-ordinates).

Lemma 3.16. *For a pure function $g(x,y)$,*

$$\varepsilon(g) \geqslant \Omega(\delta(g)).$$

Proof. The lemma can be equivalently stated as $\delta(g) \leqslant O(\varepsilon(g))$. By Lemma 3.11, and also examining its proof, we see that for a pure function $g(x,y)$,

$$\varepsilon(g) = \Theta\left(\mathbb{E}_x\left[\mathsf{var}(g(x,\cdot))\right]\right),$$

i.e. the most efficient way to turn g into a monotone function (up to a constant factor) is to change $g(x,\cdot)$ to identically 0 or identically 1 depending on whether $g(x,\cdot)$ is more likely to be 0 or 1. Note that, once this is done, the new function $h(x,y)$ does not depend on the y co-ordinates at all and hence is a strongly monotone function. □

3.4 Splitting only Decreases Talagrand Objective

In this section, we show that the "Talagrand objective", i.e. the L.H.S. of the inequality in Theorem 1.6 (and later the robust Theorem 1.9) can only decrease when a split operator is applied to f.

Lemma 3.17. *For a co-ordinate $i \in \{1, \ldots, n\}$,*

$$\mathop{\mathbb{E}}_{x}\left[\sqrt{I_f^-(x)}\right] \geqslant \mathop{\mathbb{E}}_{x'}\left[\sqrt{I_{\nabla_i(f)}^-(x')}\right].$$

Proof. First we examine how a typical contribution to the L.H.S. looks like. Fix some setting x_{-i} of co-ordinates except i and we look at the value of $I_f^-(\cdot)$ at the two inputs $(x_i = 0, x_{-i})$ and $(x_i = 1, x_{-i})$. Consider values of f at these two inputs:

$$a = f(x_i = 0, x_{-i}), \quad b = f(x_i = 1, x_{-i}).$$

Let $j \in \{1, \ldots, n\}, j \neq i$ be any co-ordinate such that the j^{th} co-ordinate of x_{-i} equals 0. Changing this co-ordinate to 1, let the values of f at the two new inputs be:

$$a' = f(x_i = 0, x_j = 1, x_{-i-j}), \quad b' = f(x_i = 1, x_j = 1, x_{-i-j}).$$

Consider the notation:

$$(a, b) \mapsto [(0,0), \#\alpha],\ [(1,0), \#\beta],\ [(0,1), \#\gamma],\ [(1,1), \#\delta],$$

meaning that the number of co-ordinates j for which $(a', b') = (0,0)$ is exactly α, the number of co-ordinates j for which $(a', b') = (1,0)$ is exactly β, and so on. Recall that the value of $I_f^-(\cdot)$ at an input equals zero if $f = 0$ at the input and the number of negatively sensitive edges on the input otherwise. The values of $I_f^-(\cdot)$ at the two inputs $(x_i = 0, x_{-i})$ and $(x_i = 1, x_{-i})$ depend on the values of f at these inputs, i.e. a and b, and the numbers $\alpha, \beta, \gamma, \delta$ in the following manner (the extra $+1$ at one place takes into account a negatively sensitive edge contributed by co-ordinate i itself):

$$I_f^-(x_i = 0, x_{-i}),\ I_f^-(x_i = 1, x_{-i}) \ = \ \begin{cases} 0, & 0 \qquad \text{if } a = 0, b = 0 \\ \alpha + \gamma + 1, & 0 \qquad \text{if } a = 1, b = 0 \\ 0, & \alpha + \beta \qquad \text{if } a = 0, b = 1 \\ \alpha + \gamma, & \alpha + \beta \qquad \text{if } a = 1, b = 1 \end{cases}$$

Now we consider the effect of applying the split operator on the co-ordinate i. Under the action of the operator, the two inputs $(x_i = 0, x_{-i})$ and $(x_i = 1, x_{-i})$ give rise to four inputs, namely

$$(x_{i,+} = 0, x_{i,-} = 0, x_{-i}), (x_{i,+} = 0, x_{i,-} = 1, x_{-i}), (x_{i,+} = 1, x_{i,-} = 0, x_{-i}), (x_{i,+} = 1, x_{i,-} = 1, x_{-i}),$$

with the values of $\nabla_i(f)$ at these inputs respectively as

$$a,\ \min\{a, b\},\ \max\{a, b\},\ b.$$

As before, we investigate the scenario when, for some $j \in \{1, \ldots, n\}, j \neq i$, the j^{th} co-ordinate of these four inputs is changed from 0 to 1. It is not difficult to see that the values of $\nabla_i(f)$ at these four new inputs are:

$$a',\ \min\{a', b'\},\ \max\{a', b'\},\ b',$$

if the values of f at the two new inputs are (a', b') as discussed above. Using a similar notation as before:

$$(a, \min\{a, b\}, \max\{a, b\}, b) \mapsto [(0,0,0,0), \#\alpha],\ [(1,0,1,0), \#\beta],\ [(0,0,1,1), \#\gamma],\ [(1,1,1,1), \#\delta],$$

meaning that the number of co-ordinates j for which $(a', \min\{a', b'\}, \max\{a', b'\}, b') = (0,0,0,0)$ is exactly α and so on. Now we are ready to complete the proof considering the four cases depending on values a and b. In each case, we compare the contribution of $I_f^-(\cdot)$ at the two relevant inputs against the contribution of $I_{\nabla_i(f)}^-(\cdot)$ at the four relevant inputs to the inequality in the statement of the lemma.

When $a = 0, b = 0$, values of $I_{\nabla_i(f)}^-(\cdot)$ are:

$$0,\ 0,\ 0,\ 0,$$

and the inequality in the statement of the lemma amounts to saying that:

$$\frac{\sqrt{0} + \sqrt{0}}{2} \geqslant \frac{\sqrt{0} + \sqrt{0} + \sqrt{0} + \sqrt{0}}{4}.$$

When $a = 1, b = 0$, values of $I_{\nabla_i(f)}^-(\cdot)$ are:

$$\alpha + \gamma + 1,\ 0,\ \alpha + 1,\ 0,$$

and the inequality in the statement of the lemma amounts to saying that:

$$\frac{\sqrt{\alpha + \gamma + 1} + \sqrt{0}}{2} \geqslant \frac{\sqrt{\alpha + \gamma + 1} + \sqrt{0} + \sqrt{\alpha + 1} + \sqrt{0}}{4}.$$

When $a = 0, b = 1$, values of $I_{\nabla_i(f)}^-(\cdot)$ are:

$$0,\ 0,\ \alpha,\ \alpha + \beta,$$

and the inequality in the statement of the lemma amounts to saying that:

$$\frac{\sqrt{0} + \sqrt{\alpha + \beta}}{2} \geqslant \frac{\sqrt{0} + \sqrt{0} + \sqrt{\alpha} + \sqrt{\alpha + \beta}}{4}.$$

When $a = 1, b = 1$, values of $I_{\nabla_i(f)}^-(\cdot)$ are:

$$\alpha + \gamma,\ \alpha + \beta + \gamma,\ \alpha,\ \alpha + \beta,$$

and the inequality in the statement of the lemma amounts to saying that:

$$\frac{\sqrt{\alpha + \gamma} + \sqrt{\alpha + \beta}}{2} \geqslant \frac{\sqrt{\alpha + \gamma} + \sqrt{\alpha + \beta + \gamma} + \sqrt{\alpha} + \sqrt{\alpha + \beta}}{4}.$$

□

The lemma below shows that the L.H.S. of inequality in the robust Theorem 1.9 can only decrease when a split operator is applied to f. This is a bit subtle as one shows that for any coloring of the negatively sensitive edges of f, there is a coloring of the negatively sensitive edges of $\nabla_i(f)$ such that the objective can only decrease.

Lemma 3.18. *For any coloring* $\mathsf{col} : \mathcal{S}_f^- \mapsto \{\mathsf{red}, \mathsf{blue}\}$ *and co-ordinate* $i \in \{1, \ldots, n\}$*, there exists a coloring* $\mathsf{col} : \mathcal{S}_{\nabla_i(f)}^- \mapsto \{\mathsf{red}, \mathsf{blue}\}$ *such that*

$$\mathbb{E}_x \left[\sqrt{I_{f,\mathsf{red}}^-(x)} \right] + \mathbb{E}_x \left[\sqrt{I_{f,\mathsf{blue}}^-(x)} \right] \geqslant \mathbb{E}_{x'} \left[\sqrt{I_{\nabla_i(f),\mathsf{red}}^-(x')} \right] + \mathbb{E}_{x'} \left[\sqrt{I_{\nabla_i(f),\mathsf{blue}}^-(x')} \right].$$

Proof. After defining the appropriate coloring $\mathsf{col} : \mathcal{S}_{\nabla_i(f)}^- \mapsto \{\mathsf{red}, \mathsf{blue}\}$, we will in fact show two separate inequalities:

$$\mathbb{E}_x \left[\sqrt{I_{f,\mathsf{red}}^-(x)} \right] \geqslant \mathbb{E}_{x'} \left[\sqrt{I_{\nabla_i(f),\mathsf{red}}^-(x')} \right],$$
$$\mathbb{E}_x \left[\sqrt{I_{f,\mathsf{blue}}^-(x)} \right] \geqslant \mathbb{E}_{x'} \left[\sqrt{I_{\nabla_i(f),\mathsf{blue}}^-(x')} \right].$$

The coloring $\mathsf{col} : \mathcal{S}_{\nabla_i(f)}^- \mapsto \{\mathsf{red}, \mathsf{blue}\}$ is derived naturally from the coloring $\mathsf{col} : \mathcal{S}_f^- \mapsto \{\mathsf{red}, \mathsf{blue}\}$ in the following manner. Fix a setting of $x_{-i} \in \{0,1\}^{n-1}$ for the rest of the proof. We recall that there are two relevant inputs for f with values

$$a = f(x_i = 0, x_{-i}), \quad b = f(x_i = 1, x_{-i}),$$

and four relevant inputs for $\nabla_i(f)$ with values

$$\begin{array}{rlcl} a = (\nabla_i(f))(x_{i,+} = 0, x_{i,-} = 0, x_{-i}), & \min\{a,b\} & = (\nabla_i(f))(x_{i,+} = 0, x_{i,-} = 1, x_{-i}), \\ \max\{a,b\} = (\nabla_i(f))(x_{i,+} = 1, x_{i,-} = 0, x_{-i}), & b & = (\nabla_i(f))(x_{i,+} = 1, x_{i,-} = 1, x_{-i}). \end{array}$$

Towards the desired coloring, we first consider the potential, negatively sensitive edge between inputs $(x_i = 0, x_{-i})$ and $(x_i = 1, x_{-i})$. This edge, denoted e, is negatively sensitive if and only if $a = 1, b = 0$ and in that case, it induces two negatively sensitive edges of $\nabla_i(f)$, between two inputs in the same row, for each of the two rows above, and these two edges are colored with the same color as e.

Next, we consider potential, negatively sensitive edges along dimension $j \in \{1, \ldots, n\}$, $j \neq i$. We assume, w.l.o.g. that the j^{th} co-ordinate of x_{-i} equals 0. When this co-ordinate is changed to 1, let us denote the values of f at the two new inputs as (a', b') and the values of $\nabla_i(f)$ at the four new inputs as

$$\begin{array}{cc} a', & \min\{a',b'\}, \\ \max\{a',b'\}, & b'. \end{array}$$

Let e_0 denote the edge whose endpoints have f-values (a, a') and e_1 denote the edge whose endpoints have f-values (b, b'). Similarly, Let $e_{00}, e_{\min}, e_{\max}, e_{11}$ denote the edges whose endpoints have $(\nabla_i)(f))$-values (a, a'), $(\min\{a,b\}, \min\{a',b'\})$, $(\max\{a,b\}, \max\{a',b'\})$, (b, b') respectively. We now assign colors to $e_{00}, e_{\min}, e_{\max}, e_{11}$ depending on the colors of e_0 and e_1. We note that colors are assigned to only negatively sensitive edges.

- If e_0 and e_1 are uncolored, so are $e_{00}, e_{\min}, e_{\max}, e_{11}$.
- If exactly one of the two edges e_0 and e_1 is colored, say by color $\mathsf{c} \in \{\mathsf{red}, \mathsf{blue}\}$, then all negatively sensitive edges among $e_{00}, e_{\min}, e_{\max}, e_{11}$ are colored with the same color c (depending on the case, there are one or two such negatively sensitive edges).
- If both the edges e_0 and e_1 are colored (which happens if and only if $a = b = 1, a' = b' = 0$), then:

– If e_0 and e_1 are colored with the same color, say $\mathsf{c} \in \{\mathsf{red}, \mathsf{blue}\}$, then $e_{00}, e_{\min}, e_{\max}, e_{11}$ all get the same color c.
– If e_0 and e_1 are colored with different colors, then e_{00} gets the color of e_0, e_{11} gets the color of e_1, $e_{\min}$ is colored red and $e_{\max}$ is colored blue.

We now prove that (the other "blue inequality" being symmetric):

$$\mathbb{E}_x \left[\sqrt{I^{-}_{f,\mathsf{red}}(x)} \right] \;\geqslant\; \mathbb{E}_{x'} \left[\sqrt{I^{-}_{\nabla_i(f),\mathsf{red}}(x')} \right]. \tag{2}$$

We note that the above inequality only concerns the number of red, negatively sensitive edges that are incident on an input with f or $\nabla_i(f)$-value equal to 1. For the fixed setting of x_{-i} we examine the contribution to L.H.S. of the inputs $(x_i \in \{0,1\}, x_{-i})$ and to R.H.S. of the inputs $(x_{i,+} \in \{0,1\}, x_{i,-} \in \{0,1\}, x_{-i})$, depending on four cases according to values of a and b. Also, in the following, co-ordinate j always refers to a co-ordinate $j \in \{1, \ldots, n\}, j \neq i$ such that the j^{th} co-ordinate of x_{-i} equals 0.

Case $a = b = 0$:

In this case, since the relevant inputs have f and $\nabla_i(f)$ values equal to 0, there is no contribution to either side of inequality (2).

Case $a = b = 1$:

We partition the set of co-ordinates j as in proof of Lemma 3.17, however we need a finer partition. Let α be the number of co-ordinates j such that $(a,b) \mapsto (0,1)$ and the edge e_0 is red. Let $\beta_1, \beta_2, \beta_3$ be the number of co-ordinates j such that $(a,b) \mapsto (0,0)$, and the coloring of (e_0, e_1) is (red,red), (blue,red), or (red,blue) respectively. Let γ be the number of co-ordinates j such that $(a,b) \mapsto (1,0)$ and the edge e_1 is red. Then,

$$I^{-}_{f,\mathsf{red}}(x_i = 0, x_{-i}) = \alpha + \beta_1 + \beta_3, \;\; I^{-}_{f,\mathsf{red}}(x_i = 1, x_{-i}) = \beta_1 + \beta_2 + \gamma$$

$$I^{-}_{\nabla_i(f),\mathsf{red}}(x_{i,+}, x_{i,-}, x) = \begin{cases} \alpha + \beta_1 + \beta_3 & \text{if } x_{i,+} = 0, x_{i,-} = 0 \\ \alpha + \beta_1 + \beta_2 + \beta_3 + \gamma & \text{if } x_{i,+} = 0, x_{i,-} = 1 \\ \beta_1 & \text{if } x_{i,+} = 1, x_{i,-} = 0 \\ \beta_1 + \beta_2 + \gamma & \text{if } x_{i,+} = 1, x_{i,-} = 1 \end{cases}$$

So the inequality (2) amounts to saying

$$\frac{\sqrt{\alpha + \beta_1 + \beta_3} + \sqrt{\beta_1 + \beta_2 + \gamma}}{2} \geqslant \frac{\sqrt{\alpha + \beta_1 + \beta_3} + \sqrt{\alpha + \beta_1 + \beta_2 + \beta_3 + \gamma} + \sqrt{\beta_1} + \sqrt{\beta_1 + \beta_2 + \gamma}}{4}.$$

Case $a = 0, b = 1$:

In this case, only $(x_i = 1, x_{-i})$ contributes to the L.H.S. of inequality (2) and only $(x_{i,+} = 1, x_{i,-} \in \{0,1\}, x_{-i})$ contributes to the R.H.S. Again, we look at a partition of the co-ordinates j. Let α be the number of co-ordinates j such that $(a,b) \mapsto (0,0)$ and the edge e_1 is red. Let β be the number of co-ordinates j such that $(a,b) \mapsto (1,0)$ and the edge e_1 is red. Then,

$$I^{-}_{f,\mathsf{red}}(x_i = 1, x_{-i}) = \alpha + \beta.$$

$$I^{-}_{\nabla_i(f),\mathsf{red}}(x_{i,+} = 1, x_{i,-}, x_{-i}) = \begin{cases} \alpha & \text{if } x_{i,-} = 0 \\ \alpha + \beta & \text{if } x_{i,-} = 1 \end{cases}$$

So the inequality (2) amounts to saying

$$\frac{\sqrt{0}+\sqrt{\alpha+\beta}}{2} \geqslant \frac{\sqrt{0}+\sqrt{0}+\sqrt{\alpha}+\sqrt{\alpha+\beta}}{4}.$$

Case $a=1, b=0$:

In this case, only $(x_i=0, x_{-i})$ contributes to the L.H.S. of inequality (2) and only $(x_{i,+}\in\{0,1\}, x_{i,-}=0, x_{-i})$ contributes to the R.H.S. Again, we look at a partition of the co-ordinates j. Let α be the number of co-ordinates j such that $(a,b)\mapsto(0,0)$ and the edge e_0 is red. Let β be the number of co-ordinates j such that $(a,b)\mapsto(0,1)$ and the edge e_0 is red. We note that in this case, we also have negatively sensitive edges along dimension i itself. Let $\chi=1$ if the negatively sensitive edge with endpoints $(x_i=0,x_{-i})$ and $(x_i=1,x_{-i})$ is red and $\chi=0$ otherwise. Then,

$$I^-_{f,\mathsf{red}}(x_i=0,x_{-i})=\alpha+\beta+\chi.$$

$$I^-_{\nabla_i(f),\mathsf{red}}(x_{i,+},x_{i,-}=0,x_{-i})=\begin{cases}\alpha+\beta+\chi & \text{if } x_{i,+}=0\\ \alpha+\chi & \text{if } x_{i,+}=1\end{cases}$$

So the inequality (2) amounts to saying

$$\frac{\sqrt{\alpha+\beta+\chi}+\sqrt{0}}{2} \geqslant \frac{\sqrt{\alpha+\beta+\chi}+\sqrt{0}+\sqrt{\alpha+\chi}+\sqrt{0}}{4}.$$

□

4 Proof of Theorem 1.6

In this section, we present a proof of Theorem 1.6, which we consider the most interesting part of the paper. The first subsection describes an informal overview of the proof and the formal proof appears in the next subsection. The overview omits several of the key ingredients and is meant for intuition only, and the reader should expect the formal proof to be quite different in terms of notation, additional ideas etc.

4.1 An Overview

We recall that for a function $f:\{0,1\}^n\mapsto\{0,1\}$ and input x, $I^-_f(x)$ is equal to 0 if $f(x)=0$ and equal to the number of negatively sensitive edges incident on x if $f(x)=1$. The minimum distance of f from any monotone function is denoted as $\varepsilon(f)$. We intend to show that (ignoring the difference between Ω and $\tilde{\Omega}$ notations):

$$\mathbb{E}_x\left[\sqrt{I^-_f(x)}\right] \geqslant \Omega(\varepsilon(f)). \tag{3}$$

We start by describing a preliminary attempt towards a proof, then point out why it doesn't quite work, and then describe how to extend this preliminary attempt to a correct proof. We attempt to "reduce" the inequality (3) concerning function f to the same inequality concerning the function $g:\{0,1\}^{2n}\mapsto\{0,1\}$ that is obtained by applying the split operator to f, successively on co-ordinates 1 through n, say in that order. Partitioning the set of co-ordinates of g into two blocks of size n each, we use the notation $g=g(x,y)$. As in Section 3.3, $g(x,y)$ is a pure function in the sense that g is monotone in x-co-ordinates and anti-monotone in y-co-ordinates. Let $\varepsilon(g)$ be the minimum distance of g from any monotone function. Towards proving inequality (3), we observe that:

- when one replaces the function f by the pure function g, the L.H.S. of the inequality (3) can only decrease (this is by Lemma 3.17; splits can only decrease the Talagrand objective).

- the inequality (3) holds for the pure function $g(x, y)$ (this is by Lemma 3.13).

Thus, we can conclude

$$\mathbb{E}_x\left[\sqrt{I_f^-(x)}\right] \geqslant \mathbb{E}_{x,y}\left[\sqrt{I_g^-(x,y)}\right] \geqslant \Omega(\varepsilon(g)). \tag{4}$$

Now, the inequality (3) would be proved if it were (always) the case that $\varepsilon(g) \geqslant \Omega(\varepsilon(f))$. Though we do not present a counter-example here, this turns out to be incorrect. Still, by a careful examination of the split operator and relating it to the switch operator, we are able to show a lower bound (up to a constant factor)

$$\varepsilon(g) \geqslant \Delta(f, \pi_n \circ f) \; - \; \mathbb{E}_{\pi_{n/2}}[\Delta(f, \pi_{n/2} \circ f)]. \tag{5}$$

We elaborate more on this lower bound. Here $\pi_n \circ f$ denotes the function obtained from f by applying the switch operator on all n co-ordinates (and π_n denotes the full set $\{1, \ldots n\}$). Also, $\pi_{n/2}$ denotes a (random) subset of $\frac{n}{2}$ co-ordinates and $\pi_{n/2} \circ f$ denotes the function obtained from f by applying the switch operator on precisely the co-ordinates in $\pi_{n/2}$. We certainly know that $\pi_n \circ f$ is a monotone function and we think of applying the switch operator on co-ordinates, one by one, as "progressing" towards the "target function" $\pi_n \circ f$. By definition, applying the switch operator on all n co-ordinates, attains the target. However, it is possible that applying the switch operator on only (random) $\frac{n}{2}$ co-ordinates gets us very close to the target. If so, $\pi_{n/2} \circ f \approx \pi_n \circ f$ for almost every choice of $\pi_{n/2}$ and one does not get a good lower bound in inequality (5). Nevertheless, combining inequalities (4), (5) and thinking of π_n itself as a random set of size n (though there is only one set of size n and no randomness is involved), we get

$$\mathbb{E}_x\left[\sqrt{I_f^-(x)}\right] \geqslant \mathbb{E}_{\pi_n}[\Delta(f, \pi_n \circ f)] \; - \; \mathbb{E}_{\pi_{n/2}}[\Delta(f, \pi_{n/2} \circ f)].$$

This inequality now suggests that we ought to get another inequality,

$$\mathbb{E}_x\left[\sqrt{I_f^-(x)}\right] \geqslant \mathbb{E}_{\pi_{n/2}}[\Delta(f, \pi_{n/2} \circ f)] \; - \; \mathbb{E}_{\pi_{n/4}}[\Delta(f, \pi_{n/4} \circ f)],$$

that reflects the difference between applying the switch operator on random $\frac{n}{2}$ co-ordinates versus applying it on random $\frac{n}{4}$ co-ordinates. Indeed, we do obtain such an inequality. Instead of working with the function $g = g_n$ that is obtained by applying the split operator to f on all n co-ordinates, we work with the function $g_{n/2}$ that is obtained by applying the split operator to f on (random) $n/2$ co-ordinates. More generally, we are able to obtain similar inequalities for $i = 0, 1, 2, \ldots, \lfloor \log n \rfloor$ as

$$\mathbb{E}_x\left[\sqrt{I_f^-(x)}\right] \geqslant \mathbb{E}_{\pi_{n/2^i}}[\Delta(f, \pi_{n/2^i} \circ f)] \; - \; \mathbb{E}_{\pi_{n/2^{i+1}}}[\Delta(f, \pi_{n/2^{i+1}} \circ f)].$$

Summing up these inequalities in a telescoping manner and ignoring the negative term for the last inequality numbered $i = \lfloor \log n \rfloor$ (corresponding to not applying switch operator at all and the term amounts to $\Delta(f, f) = 0$), we get

$$\lfloor \log n \rfloor \cdot \mathbb{E}_x\left[\sqrt{I_f^-(x)}\right] \geqslant \Delta(f, \pi_n \circ f) \geqslant \varepsilon(f),$$

since $\pi_n \circ f$ is a monotone function. This proves inequality (3), but with a loss of log-factor.

4.2 The Formal Proof

We now prove Theorem 1.6 formally. Let $f : \{0,1\}^n \mapsto \{0,1\}$ be the given function. Fix a non-empty set $S \subseteq \{1,\ldots,n\}$ and partition the set of co-ordinates $\{1,\ldots,n\}$ as $S \cup \overline{S}$. We write $f(x) = f(w,z)$ where the input x is partitioned into w, z denoting the inputs on co-ordinates in sets S and $\overline{S}$ respectively. For every fixed z, we will consider the function $f(\cdot, z)$, apply switch and split operators on it, obtain certain inequalities and finally take expectation over z. We start by observing that

$$\mathbb{E}_x\left[\sqrt{I_f^-(x)}\right] \geqslant \mathbb{E}_z\left[\mathbb{E}_w\left[\sqrt{I_{f(\cdot,z)}^-(w)}\right]\right].$$

We elaborate more on this inequality. The function $f(\cdot, z)$ is considered to be a function of only the co-ordinates in S. The inequality follows by observing that for any input $x = (w, z)$,

$$I_f^-(x) \geqslant I_{f(\cdot,z)}^-(w).$$

Indeed, if $f(x) = f(w,z) = 0$, both sides equal zero. Otherwise the L.H.S. equals the number of negatively sensitive edges of f incident on $x = (w,z)$, whereas the R.H.S. equals the number of negatively sensitive edges of $f(\cdot, z)$ incident on w, which is same as the number of negatively sensitive edges of f incident on $x = (w,z)$, but only considering the edges along co-ordinates in S.

Now we consider the function $g_{S,\rho}$ that is obtained from f by applying the split operator on co-ordinates in S in the order given by the permutation ρ of the set S. Since $f = f(w,z)$ and the split operator does not "touch" the co-ordinates in $\overline{S}$, we may write $g_{S,\rho} = g_{S,\rho}((u,v),z)$ where (u,v) denote the '$+$' and '$-$' co-ordinates obtained after splitting the co-ordinates of w. For every fixed z, the function $g_{S,\rho}((\cdot,\cdot),z)$ is pure in the sense that, regarded only as function of u and v, it is monotone in the u-co-ordinates and anti-monotone in the v-co-ordinates. Moreover, for every fixed z, we can consider $g_{S,\rho}((\cdot,\cdot),z)$ as the function obtained from $f(\cdot,z)$ by applying the split operator on co-ordinates in S in the order ρ. By Lemma 3.17, splitting can only decrease the Talagrand objective, and hence for every fixed z,

$$\mathbb{E}_w\left[\sqrt{I_{f(\cdot,z)}^-(w)}\right] \geqslant \mathbb{E}_{u,v}\left[\sqrt{I_{g_{S,\rho}((\cdot,\cdot),z)}^-(u,v)}\right].$$

Since $g_{S,\rho}((\cdot,\cdot),z)$ is a pure function, by Lemma 3.13, Definition 3.15 of strong monotonicity, and Lemma 3.16 regarding distance to strong monotonicity, we have

$$\mathbb{E}_{u,v}\left[\sqrt{I_{g_{S,\rho}((\cdot,\cdot),z)}^-(u,v)}\right] \geqslant \Omega\left(\delta(g_{S,\rho}((\cdot,\cdot),z))\right).$$

Combining the inequalities above, we have, for some absolute constant C,

$$C \cdot \mathbb{E}_x\left[\sqrt{I_f^-(x)}\right] \geqslant \mathbb{E}_z\left[\delta(g_{S,\rho}((\cdot,\cdot),z))\right]. \tag{6}$$

4.2.1 Relating Splits to Switches

Now we take a closer look at the function $g_{S,\rho}((\cdot,\cdot),z)$. This function is defined on input (u,v) where $|u| = |v| = |S|$. Re-arranging the input co-ordinates into $(+,-)$ pairs, and denoting by $\Sigma = \{00, 01, 10, 11\}$ as the four possible values that a pair of co-ordinates may take, we view the function as:

$$g_{S,\rho}(\cdot, z) : \Sigma^{|S|} \mapsto \{0,1\}.$$

We also write $g_{S,\rho}(\sigma, z)$ where $\sigma \in \Sigma^{|S|}$ and denote co-ordinates of σ as σ_i for $i \in S$. To avoid confusion, we emphasize that when we write $g_{S,\rho}(\sigma, z)$, the co-ordinates of σ are understood to be re-ordered according to the permutation ρ. For example, if $S = \{1, 2, 3, 4, 5, 6, 7, 8\}$ and $\rho = (5, 4, 8, 1, 7, 2, 3, 6)$, then $g_{S,\rho}(\sigma, z)$ is interpreted as:

$$g_{S,\rho}(\sigma_5, \sigma_4, \sigma_8, \sigma_1, \sigma_7, \sigma_2, \sigma_3, \sigma_6, z).$$

Let $\pi : \Sigma = \{00, 01, 10, 11\} \mapsto \{Y, N\}$ be defined as $\pi(00) = \pi(11) = N$ and $\pi(01) = \pi(10) = Y$. We extend π to $\pi : \Sigma^i \mapsto \{Y, N\}^i$ for integer i by its application on each co-ordinate. Let $\varphi : \Sigma = \{00, 01, 10, 11\} \mapsto \{0, 1\}$ be defined as $\varphi(00) = \varphi(01) = 0$ and $\varphi(10) = \varphi(11) = 1$. Thus φ simply selects the first co-ordinate of a pair of co-ordinates. We extend $\varphi : \Sigma^i \mapsto \{0, 1\}^i$ for integer i by its application on each co-ordinate.

The lemma below shows that $g_{S,\rho}(\cdot, z)$ is actually composed of copies of $f(\cdot, z)$ with suitable switch operators applied. Before stating the lemma, we need to explain further notation. For a permutation ρ of set S and a vector $\pi \in \{Y, N\}^{|S|}$, we denote by $\rho \star \pi$, the permutation ρ with the elements whose π-co-ordinate is 'N' "dropped". More explicitly, we think of the permutation ρ as an ordered list $(t_1, \ldots, t_{|S|})$ of elements of $S \subseteq \{1, \ldots, n\}$ and then $\rho \star \pi$ is this ordered list with element t_i dropped if $\pi_i = N$. Thus $\rho \star \pi$ is also a ordered list and then $(\rho \star \pi) \circ f(\cdot, z)$ denotes the function obtained from $f(\cdot, z)$ by applying the switch operator on co-ordinates in the ordered list $\rho \star \pi$. As an illustration, suppose

$$\begin{aligned} S &= \{1, 2, 3, 4, 5, 6, 7, 8\}, \\ \rho &= (5, 4, 8, 1, 7, 2, 3, 6), \\ \pi &= (Y, N, N, Y, N, Y, Y, N), \end{aligned}$$

so that $\rho \star \pi = (5, 1, 2, 3)$ and $(\rho \star \pi) \circ f(\cdot, z)$ is the function obtained from $f(\cdot, z)$ by applying the switch operator on co-ordinates in the order $5, 1, 2, 3$. Now we are ready to state the lemma.

Lemma 4.1. *For $\sigma \in \Sigma^{|S|}$,*

$$g_{S,\rho}(\sigma, z) = ((\rho \star \pi(\sigma)) \circ f)(\varphi(\sigma), z).$$

Proof. Since the input z is "auxiliary" and just "floats around", we can drop it from the notation. Equivalently, we can assume that $S = \{1, \ldots, n\}$ is the full set. Also, we can assume w.l.o.g. that the permutation ρ is the identity permutation, i.e. the ordered list $(1, 2, \ldots, n)$. Thus the function $g_{S,\rho}$ is the function obtained from f by applying the split operator on co-ordinates 1 through n, in that order. We write $g = g_{S,\rho}$ and drop S and ρ from the notation. Further, for $\pi \in \{Y, N\}^n$, by denoting $\pi \circ f$ as the function obtained from f by considering the co-ordinates 1 through n in that order, and applying the switch operator on j^{th} co-ordinate if and only if $\pi_j = Y$, the lemma amounts to saying

$$\forall\, \sigma \in \Sigma^n, \qquad g(\sigma) = (\pi(\sigma) \circ f)(\varphi(\sigma)). \tag{7}$$

Using the short-form $\nabla_{[1,\ldots,i]}(f)$ to denote the "prefix"

$$\nabla_i(\nabla_{i-1}(\ldots \nabla_2(\nabla_1(f)) \ldots)),$$

we note that $g = \nabla_{[1,\ldots,n]}(f)$. Also, $\nabla_{[1,\ldots,i]}(f)$ is a function

$$\nabla_{[1,\ldots,i]}(f) : \Sigma^i \times \{0, 1\}^{n-i} \mapsto \{0, 1\}.$$

We prove by induction on i that

$$\forall\, \sigma \in \Sigma^i,\ x \in \{0,1\}^{n-i}, \qquad \nabla_{[1,\ldots,i]}(f)(\sigma, x) = (\pi(\sigma) \circ f)(\varphi(\sigma), x), \tag{8}$$

and the lemma follows from the case $i = n$. We note that in the inductive statement above, $\sigma \in \Sigma^i$, $\pi(\sigma) \in \{Y, N\}^i$ and $\pi(\sigma) \circ f$ denotes the function obtained from f by considering the co-ordinates 1 through i in that order, and applying the switch operator on j^{th} co-ordinate if and only if $\pi_j = Y$.

For $i = 0$, there is nothing to prove as the statement is $f(x_1, \ldots, x_n) = f(x_1, \ldots, x_n)$. Assume the statement (8) for some $0 \leqslant i \leqslant n-1$ and for convenience, write $x = (x_{i+1}, y)$ where $y = (y_{i+2}, \ldots, y_n)$ is a vector of formal boolean variables. Thus the inductive hypothesis is that:

$$\nabla_{[1,\ldots,i]}(f)(\sigma, x_{i+1}, y) = (\pi(\sigma) \circ f)(\varphi(\sigma), x_{i+1}, y).$$

By applying the split operator ∇_{i+1} to the L.H.S.

$$\nabla_{[1,\ldots,i+1]}(f)(\sigma, \sigma_{i+1}, y) \;=\; \begin{cases} \nabla_{[1,\ldots,i]}(f)(\sigma, 0, y) & \sigma_{i+1} = 00 \\ \min\left\{\nabla_{[1,\ldots,i]}(f)(\sigma, 0, y), \nabla_{[1,\ldots,i]}(f)(\sigma, 1, y)\right\} & \sigma_{i+1} = 01 \\ \max\left\{\nabla_{[1,\ldots,i]}(f)(\sigma, 0, y), \nabla_{[1,\ldots,i]}(f)(\sigma, 1, y)\right\} & \sigma_{i+1} = 10 \\ \nabla_{[1,\ldots,i]}(f)(\sigma, 1, y) & \sigma_{i+1} = 11. \end{cases}$$

By inductive hypothesis, the R.H.S. can be replaced as

$$\nabla_{[1,\ldots,i+1]}(f)(\sigma, \sigma_{i+1}, y) \;=\; \begin{cases} (\pi(\sigma) \circ f)(\varphi(\sigma), 0, y) & \sigma_{i+1} = 00 \\ \min\{(\pi(\sigma) \circ f)(\varphi(\sigma), 0, y), (\pi(\sigma) \circ f)(\varphi(\sigma), 1, y)\} & \sigma_{i+1} = 01 \\ \max\{(\pi(\sigma) \circ f)(\varphi(\sigma), 0, y), (\pi(\sigma) \circ f)(\varphi(\sigma), 1, y)\} & \sigma_{i+1} = 10 \\ (\pi(\sigma) \circ f)(\varphi(\sigma), 1, y) & \sigma_{i+1} = 11. \end{cases}$$

The R.H.S. can be further replaced, by the definition of the switch operator, as

$$\nabla_{[1,\ldots,i+1]}(f)(\sigma, \sigma_{i+1}, y) \;=\; \begin{cases} (\pi(\sigma) \circ f)(\varphi(\sigma), 0, y) & \sigma_{i+1} = 00 \\ S_{i+1}((\pi(\sigma) \circ f)(\varphi(\sigma), 0, y) & \sigma_{i+1} = 01 \\ S_{i+1}((\pi(\sigma) \circ f)(\varphi(\sigma), 1, y) & \sigma_{i+1} = 10 \\ (\pi(\sigma) \circ f)((\varphi(\sigma), 1, y) & \sigma_{i+1} = 11. \end{cases}$$

The R.H.S. can be written succinctly as

$$((\pi(\sigma), N) \circ f)(\varphi(\sigma), \varphi(\sigma_{i+1}), y) \quad \text{for } \sigma_{i+1} = 00, 11,$$

$$((\pi(\sigma), Y) \circ f)(\varphi(\sigma), \varphi(\sigma_{i+1}), y) \quad \text{for } \sigma_{i+1} = 01, 10,$$

which can be written even more succinctly as

$$((\pi(\sigma), \pi(\sigma_{i+1})) \circ f)(\varphi(\sigma), \varphi(\sigma_{i+1}), y),$$

completing the inductive step. □

Using Lemma 4.1, it is now easy to lower bound the distance of $g_{S,\rho}(\cdot, z)$ to strong monotonicity by the average distance of $(\rho \star \pi) \circ f(\cdot, z)$ to monotonicity, averaged over a random choice of $\pi \in \{Y, N\}^{|S|}$.

Lemma 4.2.

$$\delta(g_{S,\rho}(\cdot, z)) \geqslant \mathop{\mathbb{E}}_{\pi \in \{Y,N\}^{|S|}} [\varepsilon((\rho \star \pi) \circ f(\cdot, z))].$$

Proof. As before, we may drop z, assume that $S = \{1, \ldots, n\}$, that ρ is the identity permutation, so that $g_{S,\rho} = g = \nabla_{[1,\ldots,n]}(f) : \Sigma^n \mapsto \{0, 1\}$ and prove that

$$\delta(g) \geqslant \mathop{\mathbb{E}}_{\pi \in \{Y,N\}^n} [\varepsilon(\pi \circ f)].$$

Let $h : \Sigma^n \mapsto \{0, 1\}$ be a strongly monotone function that is $\delta(g)$-close to g. Identify each $\pi \in \{Y, N\}^n$ with the sub-cube $V = V_\pi \subseteq \Sigma^n$ defined as

$$V = \prod_{\pi[i]=Y} \{01, 10\} \times \prod_{\pi[i]=N} \{00, 11\}.$$

We look at the restriction of functions g, h, φ to the sub-cube V. The map $\varphi|_V : V \mapsto \{0, 1\}^n$ is a bijection and Lemma 4.1 (or rather the specialized statement (7)) implies that

$$g|_V(\sigma) = (\pi \circ f)(\varphi|_V(\sigma)).$$

Since $h : \Sigma^n \mapsto \{0, 1\}^n$ is strongly monotone, i.e. monotone under the ordering $00 \leqslant 01 \leqslant 10 \leqslant 11$ on Σ, it follows that $h|_V(\varphi|_V^{-1}(\cdot))$ viewed as a function on $\{0, 1\}^n$ is monotone. Thus

$$\varepsilon(\pi \circ f) \leqslant \Delta(\pi \circ f,\ h|_V(\varphi|_V^{-1}(\cdot))) = \Delta((\pi \circ f)(\varphi|_V(\cdot)),\ h|_V(\cdot)) = \Delta(g|_V, h|_V).$$

Now taking expectations over the choice of π, we get as desired

$$\mathop{\mathbb{E}}_{\pi \in \{Y,N\}^n} [\varepsilon(\pi \circ f)] \leqslant \mathop{\mathbb{E}}_{\pi \in \{Y,N\}^n} [\Delta(g|_{V_\pi}, h|_{V_\pi})] = \delta(g).$$

$\square$

4.2.2 The Telescoping Argument

We are now ready to complete the proof of Theorem 1.6. Combining inequality (6) and Lemma 4.2, we get

$$C \cdot \mathbb{E}_x \left[\sqrt{I_f^-(x)}\right] \geqslant \mathop{\mathbb{E}}_{z} \left[\mathop{\mathbb{E}}_{\pi} [\varepsilon\left((\rho \star \pi) \circ f(\cdot, z)\right)]\right].$$

Note that on the R.H.S, it is understood that there is a underlying set S of co-ordinates, $f(\cdot, z)$ is regarded as a function of co-ordinates in S, ρ is a permutation of S and $\pi \in \{Y, N\}^{|S|}$. Using Definition 3.6, we can replace $\varepsilon(\cdot)$ by $\gamma(\cdot)$ and write

$$2\,C \cdot \mathbb{E}_x \left[\sqrt{I_f^-(x)}\right] \geqslant \mathop{\mathbb{E}}_{z} \left[\mathop{\mathbb{E}}_{\pi} [\gamma\left((\rho \star \pi) \circ f(\cdot, z)\right)]\right]. \tag{9}$$

By definition of $\gamma(\cdot)$, denoting by τ a random permutation on S, we have

$$\gamma\left((\rho \star \pi) \circ f(\cdot, z)\right) = \mathbb{E}_\tau \left[\Delta\left((\rho \star \pi) \circ f(\cdot, z),\ \tau \circ (\rho \star \pi) \circ f(\cdot, z)\right)\right],$$

which by triangle inequality is at least

$$\mathbb{E}_\tau \left[\Delta\left(f(\cdot,z),\ \tau\circ(\rho\star\pi)\circ f(\cdot,z)\right)\right] \ - \ \Delta\left(f(\cdot,z),\ (\rho\star\pi)\circ f(\cdot,z)\right).$$

Now we look at the function $\tau\circ(\rho\star\pi)\circ f(\cdot,z)$ closely. Here, after applying the switch operator according to $\rho\star\pi$, one applies the switch operator again according to τ. The co-ordinates on which the switch operator was applied in the first phase, i.e. according to $\rho\star\pi$, are already "monotonized" and these co-ordinates do not really participate when the switch operator is applied in the second phase, i.e. according to τ. The combined effect of the two phases is same as that of applying the switch operator according to some permutation λ that depends on τ,ρ,π. As an illustrative example, suppose $S=\{1,2,3,4,5,6,7,8\}$, $\rho=(5,4,8,1,7,2,3,6)$, $\pi=\{Y,N,N,Y,N,Y,Y,N\}$ so that $\rho\star\pi=(5,1,2,3)$. Suppose $\tau=(2,8,6,1,4,5,3,7)$. In that case

$$\lambda=\lambda(\tau,\rho,\pi)=(5,1,2,3,8,6,4,7).$$

Clearly, if τ,ρ,π were selected uniformly at random, then λ is also distributed as a uniformly random permutation of S. Taking the expectation of inequality (9) over a random choice of ρ, we thus see that $2\,C\cdot\mathbb{E}_x\left[\sqrt{I_f^-(x)}\right]$ is lower bounded by

$$\mathbb{E}_z\left[\,\mathbb{E}_\lambda\left[\Delta\left(f(\cdot,z),\ \lambda\circ f(\cdot,z)\right)\right]\,\right] \ - \ \mathbb{E}_z\left[\,\mathbb{E}_{\rho,\pi}\left[\Delta\left(f(\cdot,z),\ (\rho\star\pi)\circ f(\cdot,z)\right)\right]\,\right].$$

Finally, we take expectation of the lower bound above over the choice of a random set $S\subseteq\{1,\ldots,n\}$ where each co-ordinate is in S with probability p independently. It is easily seen that we now get a lower bound

$$2\,C\cdot\mathbb{E}_x\left[\sqrt{I_f^-(x)}\right]\geqslant\Psi_f(p)-\Psi_f\left(\frac{p}{2}\right),$$

where $\Psi_f(p)$ denotes the expected change in f by considering the co-ordinates 1 through n in a random order and applying the switch operator on each co-ordinate with probability p independently. Considering this lower bound for $p=1,\frac{1}{2},\frac{1}{4},\ldots,\frac{1}{2^s}$ where $s=\lceil 5\log n+5\log(1/\varepsilon(f))\rceil$, and using telescoping sum, we get that

$$(s+1)\cdot 2\,C\cdot\mathbb{E}_x\left[\sqrt{I_f^-(x)}\right]\geqslant\Psi_f(p=1)-\Psi_f\left(p=\frac{1}{2^{s+1}}\right)\geqslant\varepsilon(f)-\frac{1}{2}\cdot\varepsilon(f)\geqslant\frac{1}{2}\cdot\varepsilon(f),$$

and hence

$$\mathbb{E}_x\left[\sqrt{I_f^-(x)}\right]\geqslant\Omega\left(\frac{\varepsilon(f)}{\log n+\log(1/\varepsilon(f))}\right),$$

proving Theorem 1.6. We noted that $\Psi_f(p=1)=\gamma(f)\geqslant\varepsilon(f)$ and for the choice of s, with probability $1-\frac{1}{2}\cdot\varepsilon(f)$, no switch is applied at all (not even on a single co-ordinate), so $\Psi_f\left(p=\frac{1}{2^{s+1}}\right)\leqslant\frac{1}{2}\cdot\varepsilon(f)$.

5 Proof of Theorem 1.9

Now that we have proved Theorem 1.6, we observe that the proof of Theorem 1.9 follows in essentially the same manner. We only point out the minor differences. Let $\mathsf{col}:\mathcal{S}_f^-\mapsto\{\mathsf{red},\mathsf{blue}\}$ be any coloring of the negatively sensitive edges of f. We intend to lower bound

$$\mathbb{E}_x\left[\sqrt{I_{f,\mathsf{red}}^-(x)}\right]+\mathbb{E}_x\left[\sqrt{I_{f,\mathsf{blue}}^-(x)}\right].$$

Similar to the beginning of Section 4.2, thinking of the input x as $x = (w, z)$ w.r.t. an underlying partition of the co-ordinates as $S \cup \overline{S}$, we get the lower bound

$$\mathbb{E}_z \left[\mathbb{E}_w \left[\sqrt{I^-_{f(\cdot,z),\mathsf{red}}(w)} \right] + \mathbb{E}_w \left[\sqrt{I^-_{f(\cdot,z),\mathsf{blue}}(w)} \right] \right].$$

The reason is the same as before, that we are now ignoring edges along the co-ordinates in $\overline{S}$. We now consider the function $g_{S,\rho}$ as before. Using Lemma 3.18, there is a coloring of negatively sensitive edges of $g_{S,\rho}(\cdot, z)$ so that we get a lower bound

$$\mathbb{E}_z \left[\mathbb{E}_\sigma \left[\sqrt{I^-_{g_{S,\rho}(\cdot,z),\mathsf{red}}(\sigma)} \right] + \mathbb{E}_\sigma \left[\sqrt{I^-_{g_{S,\rho}(\cdot,z),\mathsf{blue}}(\sigma)} \right] \right].$$

By Lemma 3.14 and 3.16, we get a lower bound, up to a constant factor,

$$\mathbb{E}_z \left[\delta(g_{S,\rho}((\cdot,\cdot), z)) \right].$$

Now Sections 4.2.1 and 4.2.2 only concern a lower bound on this last quantity (and taking expectation over the choice of S, ρ and telescoping) and we are done. We note that even though the proof in Section 4.2.2 is written with reference to the quantity $\mathbb{E}_x \left[\sqrt{I^-_f(x)} \right]$, one only needs to replace it with the analogue $\mathbb{E}_x \left[\sqrt{I^-_{f,\mathsf{red}}(x)} \right] + \mathbb{E}_x \left[\sqrt{I^-_{f,\mathsf{blue}}(x)} \right]$ everywhere.

6 A Combinatorial Fact

Definition 6.1. *A bipartite graph $G(V, W, E)$ is called right-(K, d)-good if $|W| \geqslant K$, degree of every vertex in W is in the range $[d, 2d]$ and degree of every vertex in V is at most $2d$.*

Definition 6.2. *A bipartite graph $G(V, W, E)$ is called left-(K, d)-good if $|V| \geqslant K$, degree of every vertex in V is in the range $[d, 2d]$ and degree of every vertex in W is at most $2d$.*

Definition 6.3. *A bipartite graph $G(V, W, E)$ is called (K, d)-good if it is either right-(K, d)-good or left-(K, d)-good.*

Definition 6.4. *A bipartite graph $G(V, W, E)$ is called L-robust if for any 2-coloring of its edges* $\mathsf{col} : E \mapsto \{\mathsf{red}, \mathsf{blue}\}$*, we have*

$$\sum_{v \in V} \sqrt{D_{\mathsf{red}}(v)} + \sum_{w \in W} \sqrt{D_{\mathsf{blue}}(w)} \geqslant L.$$

Here $D_{\mathsf{red}}(v)$ denotes the number of red edges incident on v and $D_{\mathsf{blue}}(w)$ denotes the number of blue edges incident on w.

Lemma 6.5. *If a bipartite graph $G(V, W, E)$ is L-robust and all its vertices have degree less than 2^s, then it has a subgraph that is (K, d)-good with $K\sqrt{d} \geqslant \frac{L}{8s}$.*

Proof. We search for the desired good subgraph by looking at a decreasing sequence of subgraphs of $G(V, W, E)$. Let $G_0(V_0, W_0, E_0) = G(V, W, E)$ be the starting graph. For $j = 0, 1, 2, \ldots$, we assume that $G_j(V_j, W_j, E_j)$ is $\left(L - j \cdot \frac{L}{2s}\right)$-robust and all its vertices have degree less than 2^{s-j}. Then we show that either $G_j(V_j, W_j, E_j)$ has a subgraph that is (K, d)-good with $K\sqrt{d} \geqslant \frac{L}{8s}$ (in which case we are done and we can stop) or else has a subgraph $G_{j+1}(V_{j+1}, W_{j+1}, E_{j+1})$ that is $\left(L - (j+1) \cdot \frac{L}{2s}\right)$-robust and has all degrees less than 2^{s-j-1} and we resume the next iteration. This iterative process must find a good subgraph since otherwise for $j = s$, the graph G_s will have no edges and still be $\frac{L}{2}$-robust.

We prove the iterative claim. In the graph $G_j(V_j, W_j, E_j)$, let $A \subseteq V_j, B \subseteq W_j$ be the sets of vertices whose degree is in the range $[2^{s-j-1}, 2^{s-j})$. Let $V_{j+1} = V_j \backslash A$, $W_{j+1} = W_j \backslash B$, and $G_{j+1}(V_{j+1}, W_{j+1}, E_{j+1})$ be the induced subgraph of $G_j(V_j, W_j, E_j)$ on vertex set (V_{j+1}, W_{j+1}). Note that the degrees of vertices in G_{j+1} are less than 2^{s-j-1}.

Let $H(A, W_j, E_H)$ be the induced subgraph of $G_j(V_j, W_j, E_j)$ on the vertex set (A, W_j). With $d = 2^{s-j-1}$, note that in the graph $H(A, W_j, E_H)$, every vertex in A has degree in the range $[d, 2d]$ and every vertex in W_j has degree at most $2d$. Thus H is left-$(|A|, d)$-good and if $|A|\sqrt{d} \geqslant \frac{L}{8s}$, we are done. Hence we may assume that $|A|\sqrt{d} \leqslant \frac{L}{8s}$. Similarly, we may assume that $|B|\sqrt{d} \leqslant \frac{L}{8s}$. Now we prove that $G_{j+1}(V_{j+1}, W_{j+1}, E_{j+1})$ is $\left(L - (j+1) \cdot \frac{L}{2s}\right)$-robust.

Consider any $\{\mathsf{red}, \mathsf{blue}\}$-coloring of edges of $G_{j+1}(V_{j+1}, W_{j+1}, E_{j+1})$. Extend this to a coloring of edges of $G_j(V_j, W_j, E_j)$ by coloring all edges between A and W_{j+1} as red, all edges between B and V_{j+1} as blue, and coloring edges between A and B arbitrarily. Using the fact that G_j is $\left(L - j \cdot \frac{L}{2s}\right)$-robust, we get (degrees below are degrees in G_j):

$$\begin{aligned}
L - j \cdot \frac{L}{2s} &\leqslant \sum_{v \in V_j} \sqrt{D_{\mathsf{red}}(v)} + \sum_{w \in W_j} \sqrt{D_{\mathsf{blue}}(w)} \\
&= \sum_{v \in V_{j+1}} \sqrt{D_{\mathsf{red}}(v)} + \sum_{v \in A} \sqrt{D_{\mathsf{red}}(v)} + \sum_{w \in W_{j+1}} \sqrt{D_{\mathsf{blue}}(w)} + \sum_{w \in B} \sqrt{D_{\mathsf{blue}}(w)} \\
&\leqslant \sum_{v \in V_{j+1}} \sqrt{D_{\mathsf{red}}(v)} + |A| \cdot \sqrt{2d} + \sum_{w \in W_{j+1}} \sqrt{D_{\mathsf{blue}}(w)} + |B| \cdot \sqrt{2d}.
\end{aligned}$$

Using the upper bound on $|A|$ and $|B|$ as above, we get

$$\sum_{v \in V_{j+1}} \sqrt{D_{\mathsf{red}}(v)} + \sum_{w \in W_{j+1}} \sqrt{D_{\mathsf{blue}}(w)} \geqslant L - (j+1) \cdot \frac{L}{2s}.$$

Noting that for vertices in V_{j+1} (resp. W_{j+1}) their red-degree (resp. blue-degree) is same in the graphs G_j and G_{j+1}, the claim follows. □

7 Algorithm for Monotonicity Testing

In this section, we present our monotonicity testing algorithm. Let $f : \{0,1\}^n \mapsto \{0,1\}$ be the given function that is $\varepsilon(f)$-far from being monotone. The tester is a simple path tester that picks a distance parameter $\tau \in \{1, \ldots, \lfloor\sqrt{n}\rfloor\}$ from a certain distribution and tries to detect a violation to monotonicity on a pair (x, y) of inputs such that y is monotonically above x at distance τ. Using Theorem 1.9, we deduce existence of a good subgraph of violated edges and then the tester is analyzed with reference to this good subgraph (note that the good subgraph is used for purposes of the analysis only).

As we will see, we only need to consider the case when $\varepsilon(f) \geqslant \frac{1}{\sqrt{n}}$, hence we make this assumption henceforth and in particular, $\log(1/\varepsilon(f))$ is $O(\log n)$. Secondly, we can ignore "atypical" inputs, i.e. those whose Hamming weight is outside the range $\frac{n}{2} \pm 10\sqrt{n \log n}$. The fraction of these inputs is at most, say $\frac{1}{n^{10}}$, and ignoring them does not affect the analysis. One can get around this issue in a formal manner also, as follows. One may pretend that instead of a query access to the given function f, we have a query access to the function $\tilde{f}$ that is same as f except that on inputs of Hamming weight less than $\frac{n}{2} - 10\sqrt{n \log n}$, $\tilde{f} = 0$ and on inputs of Hamming weight greater than $\frac{n}{2} + 10\sqrt{n \log n}$, $\tilde{f} = 1$. It is easily seen that the atypical inputs do not participate in any violating pair (x, y) for $\tilde{f}$. One may carry out the analysis for $\tilde{f}$ giving a lower bound on the rejection probability. Since the rejection probability amounts to detecting a violating pair and any violating pair for $\tilde{f}$ is also a violating pair for f, the same lower bound on rejection probability applies to f.

7.1 Existence of a Good Subgraph

We recall that $G_f^-(V, W, E)$ denotes 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.

Theorem 1.9 amounts to saying that the graph $G_f^-(V, W, E)$ is $\Omega(2^n \cdot \frac{\varepsilon(f)}{\log n})$-robust. We note that the degrees of vertices in G_f^- are at most n, hence less than 2^s with $s = \lceil \log n \rceil + 1$. Applying Lemma 6.5, we get that $G_f^-(V, W, E)$ has a subgraph $G_{\mathsf{good}}(A, B, E_{AB})$ that is (K, d)-good where $K\sqrt{d} \geqslant \Omega(\frac{2^n \cdot \varepsilon(f)}{\log^2 n})$.

We assume w.l.o.g. that G_{good} is a right-(K, d)-good graph with $|B| = K$. By definition of a good graph, every vertex in A has degree at most $2d$ and every vertex in B has degree in the range $[d, 2d]$. By deleting some edges if necessary, we may assume that every vertex in B has degree exactly d. By deleting some vertices from B if necessary, we may assume that $K\sqrt{d} = \Theta(\frac{2^n \cdot \varepsilon(f)}{\log^2 n})$. Finally, we write $K = \sigma \cdot 2^n$ for convenience. We summarize the conclusion below.

Lemma 7.1. *$G_f^-(V, W, E)$ contains a subgraph $G_{\mathsf{good}}(A, B, E_{AB})$ such that $|B| = \sigma \cdot 2^n$, vertices in B have degree exactly d, vertices in A have degree at most $2d$ and $\sigma\sqrt{d} = \Theta(\frac{\varepsilon(f)}{\log^2 n})$.*

7.2 The Tester

Let p be the largest integer such that $2^p \leqslant \sqrt{\frac{n}{\log n}}$. Our tester works as follows:

1. Pick an integer $k \in \{0, 1, 2, \ldots, p\}$ uniformly at random and let $\tau = 2^k$.
2. Pick an input $x \in \{0,1\}^n$ uniformly at random.
3. Let $S \subseteq \{1, \ldots, n\}$ be the set of 0-co-ordinates of x. Pick a random subset $T \subseteq S, |T| = \tau$ and obtain z by changing co-ordinates of x in T to 1.
4. Reject if and only if $f(x) = 1$ and $f(z) = 0$.

Remark 7.2. *(i) In Step (3), as remarked before, we may assume that the Hamming weight of x is in the range $\frac{n}{2} \pm 10\sqrt{n \log n}$, so x does have enough 0-co-ordinates. (ii) The description of the tester and the analysis are written assuming that the good subgraph, as in Lemma 7.1, is a right-good subgraph. If it were a left-good subgraph, the tester (and the analysis) would work in the anti-symmetric manner, by letting the input y to be monotonically below x by a distance τ and then rejecting if $f(x) = 0, f(y) = 1$. Formally, the tester would pick one of the two options with probability $\frac{1}{2}$ each and proceed.*

7.3 The Analysis of the Tester

We intend to show that the tester rejects with probability $\tilde{\Omega}(\frac{\varepsilon(f)^2}{\sqrt{n}})$. We first make a series of preliminary observations and then proceed to the main analysis. We note first that the tester, when $\tau = 1$, actually runs the edge tester and this happens with probability $\Theta(\frac{1}{\log n})$ over the choice of τ. Thus our tester already includes, implicitly, the edge tester.

The Lower Bound $\varepsilon(f) \geqslant \frac{1}{\sqrt{n}}$:

The edge tester has rejection probability $\Omega(\frac{\varepsilon(f)}{n})$ as shown by Goldreich et al [11] and when $\varepsilon(f) \leqslant \frac{1}{\sqrt{n}}$, it also qualifies as a tester with rejection probability $\Omega(\frac{\varepsilon(f)^2}{\sqrt{n}})$. Thus we assume henceforth that $\varepsilon(f) \geqslant \frac{1}{\sqrt{n}}$.

The Upper Bound $I_f \leqslant 6\sqrt{n}$:

If the total negative influence I_f^- is large, i.e. if $I_f^- \geqslant \sqrt{n}$, then the edge tester rejects with probability $\frac{I_f^-}{n}$ which is $\Omega(\frac{1}{\sqrt{n}})$ and we are done. Thus we assume henceforth that $I_f^- \leqslant \sqrt{n}$. By Theorem 9.1, we then have an upper bound on the total influence, i.e. $I_f \leqslant \max\{3 \cdot I_f^-, 6\sqrt{n}\} = 6\sqrt{n}$.

The Lower Bound $\sigma \geqslant \frac{\log n}{\sqrt{n}}$:

We know that $\sigma\sqrt{d} = \Theta(\frac{\varepsilon(f)}{\log^2 n})$. If $\sigma \leqslant \frac{\log n}{\sqrt{n}}$, we get that

$$\sigma \cdot d \;=\; \frac{(\sigma\sqrt{d})^2}{\sigma} \;\geqslant\; \Omega\left(\frac{\varepsilon(f)^2 \cdot \sqrt{n}}{\log^5 n}\right).$$

The graph $G_{\mathsf{good}}(A, B, E_{AB})$ has exactly $|B| \cdot d = \sigma \cdot 2^n \cdot d$ edges and all these are violating edges. Thus $I_f^- \geqslant \sigma \cdot d \geqslant \tilde{\Omega}(\varepsilon(f)^2 \cdot \sqrt{n})$. As before, the edge tester rejects with probability $\frac{I_f^-}{n}$ which is $\tilde{\Omega}(\frac{\varepsilon(f)^2}{\sqrt{n}})$ and we are done. Thus we assume henceforth that $\sigma \geqslant \frac{\log n}{\sqrt{n}}$.

Choosing τ that "Works" and Persistence of B-vertices:

Since τ takes values that are powers of 2, we can fix a value of τ such that

$$\tau \leqslant \sigma\sqrt{\frac{n}{\log n}} \leqslant 2\,\tau.$$

Using the bounds $\frac{\log n}{\sqrt{n}} \leqslant \sigma \leqslant 1$, it holds that $\frac{1}{2} \cdot \sqrt{\log n} \leqslant \tau \leqslant \sqrt{\frac{n}{\log n}}$. We will show that the tester "works" when this specific value of τ is chosen by the tester, which happens with probability $\Theta(\frac{1}{\log n})$.

We call a vertex $y \in B$ "$(\tau-1)$-persistent" if changing $\tau - 1$ of its co-ordinates at random from 0 to 1 changes the value of the function from 0 to 1 with probability at most $\frac{1}{10}$. By Lemma 9.3, the fraction of vertices $y \in \{0,1\}^n$ that are not $(\tau-1)$-persistent is bounded by $O(\frac{I_f}{n} \cdot \tau)$, which is using upper bounds $I_f \leqslant 6\sqrt{n}$ and $\tau \leqslant \sigma\sqrt{\frac{n}{\log n}}$, is upper bounded by $\frac{\sigma}{100}$. Since B constitutes σ fraction of vertices in $\{0,1\}^n$, at least a $\frac{99}{100}$ fraction of the vertices in B are $(\tau-1)$-persistent. We retain only the persistent vertices in B and assume henceforth that all vertices in B are $(\tau-1)$-persistent, redefining the parameter σ to reflect the new, slightly reduced, size of B.

Main Analysis

We are now ready to present the main argument in our analysis. Let $G_{\mathsf{good}}(A, B, E_{AB})$ be the good graph of violated edges. Let τ be the specifically chosen parameter as above. The tester picks an input $x \in \{0,1\}^n$ uniformly at random and then picks input z at random that is monotonically above x by a distance τ. We consider the probability of the following event $\mathcal{R}$. We note that in case of event $\mathcal{R}$, the tester does detect a violation of monotonicity and rejects. Naturally, the probability of event $\mathcal{R}$ is a lower bound on the rejection probability of the tester, however, one needs to be careful to avoid "double-counting" as explained towards the end of the analysis.

The event $\mathcal{R}$

- $x \in A$ and hence $f(x) = 1$.
- There is a unique y such that $(x, y) \in E_{AB}$ (i.e. $x \leqslant y$ and $y \in B$) and moreover $y \leqslant z$.
- $f(y) = f(z) = 0$.

Fix some $x \in A$. Let $s = \mathsf{deg}_A(x) \leqslant 2d$ be the degree of x in the graph $G_{\mathsf{good}}(A, B, E_{AB})$ and let $\{y_1, \ldots, y_s\}$ be the set of neighbors of x in this graph. We note that $d \leqslant n$ and

$$s \cdot \tau \ \leqslant\ 2d \cdot \sigma\sqrt{\frac{n}{\log n}} \ =\ O(\sigma\sqrt{d}) \cdot \sqrt{\frac{dn}{\log n}} \ \leqslant\ O\left(\frac{1}{\log^2 n}\right) \cdot \frac{n}{\sqrt{\log n}}.$$

Consider the choice of a random z that is monotonically above x by a distance τ. This amounts to changing, at random, τ of the co-ordinates of x from 0 to 1. The probability that $y_i \leqslant z$ for some $i \in \{1, \ldots, s\}$ is $\Omega(\frac{s \cdot \tau}{n})$. This reflects the probability that while changing τ of the co-ordinates from $\approx \frac{n}{2}$ 0-co-ordinates of x, one of the s relevant co-ordinates corresponding to its neighbors $y_1, \ldots, y_s$ gets changed. It is important that $s \cdot \tau \ll n$ for this argument to work. On the other hand, the probability that $y_i \leqslant z$ for two or more indices $i \in \{1, \ldots, s\}$ is at most $O(s^2 \cdot \frac{\tau^2}{n^2})$, which is negligible compared to the probability that there is at least one such index. In other words, whenver $y_i \leqslant z$ for some $i \in \{1, \ldots, s\}$, such an index i is likely unique. Moreover, for any fixed $i \in \{1, \ldots, s\}$, conditioning the choice of z so that $y_i \leqslant z$, the choice of z amounts to changing, at random, $\tau - 1$ of the co-ordinates of y_i from 0 to 1. Since $y_i \in B$ is $(\tau - 1)$-persistent, with probability $\frac{9}{10}$ over the choice of z (conditional on $y_i \leqslant z$), it holds that $f(z) = f(y_i) = 0$.

The discussion above shows that for a fixed $x \in A$, the probability of event $\mathcal{R}$ over the choice of z is $\Omega(\frac{\mathsf{deg}_A(x) \cdot \tau}{n})$. Hence the overall probability of event $\mathcal{R}$ can be lower bounded as:

$$\mathsf{Pr}[\mathcal{R}] \ =\ \frac{1}{2^n} \sum_{x \in A} \Omega\left(\frac{\mathsf{deg}_A(x) \cdot \tau}{n}\right) \ =\ \Omega\left(\frac{\tau}{n}\right) \cdot \frac{1}{2^n} \sum_{x \in A} \mathsf{deg}_A(x) \ =\ \Omega\left(\frac{\tau}{n}\right) \cdot \frac{|B| \cdot d}{2^n} \ =\ \Omega\left(\frac{\tau}{n}\right) \cdot \sigma \cdot d,$$

which by substituting $\tau = \Theta(\sigma\sqrt{\frac{n}{\log n}})$ and $\sigma^2 d = (\sigma\sqrt{d})^2 \geqslant \Omega((\frac{\varepsilon(f)}{\log^2 n})^2)$, gives a lower bound of $\tilde{\Omega}(\frac{\varepsilon(f)^2}{\sqrt{n}})$ on $\mathsf{Pr}[\mathcal{R}]$. We now note that whenever the event $\mathcal{R}$ occurs, the tester detects that (x, z) is a violating pair. Moreover, the pair (x, z) uniquely determines the edge $(x, y_i) \in E_{AB}$ with $x \leqslant y_i \leqslant z$ and the violating pair (x, z) can be "credited" to the edge $(x, y_i) \in E_{AB}$. The sets of violating pairs credited to different edges in E_{AB} are disjoint and there is no "double-counting". This shows that $\mathsf{Pr}[\mathcal{R}]$ is also a lower bound on the rejection probability of the tester.

8 A Lower Bound for Monotonicity Testing

In this section, we prove Theorem 1.12. For a subset $S \subseteq \{1, 2, \ldots, 4n\}, |S| = 2n$ and an index $j \in \{1, 2, \ldots, 4n\}$, $j \notin S$, define a function $f_{S,j} : \{0,1\}^{4n} \mapsto \{0,1\}$ as follows: for input $x \in \{0,1\}^{4n}$, denoting by x_S, its restriction to co-ordinates in S,

- If x_S has Hamming weight less than n, $f(x) = 0$.
- If x_S has Hamming weight larger than n, $f(x) = 1$.
- If x_S has Hamming weight exactly n, $f(x) = 1 - x_j$.

In short, $f_{S,j}(x)$ is the majority function on x_S except in the "middle layer" where it is the anti-dictatorship of co-ordinate j. It is easily seen that $\varepsilon(f_{S,j}) = \Theta(\frac{1}{\sqrt{n}})$ (to turn $f_{S,j}$ into a monotone function, the best strategy, up to a constant factor, is to make it 0 in the middle layer). Consider the family $\mathcal{F}$ of functions $f_{S,j}$ over all choices of $S, |S| = 2n$ and $j \notin S$.

We show that for any "reasonable" pair of inputs (x, y) such that y is monotonically above x, the probability that (x, y) is a violating pair for the function $f_{S,j}$ is $O(1/n^{3/2})$ over a randomly chosen function $f_{S,j} \in \mathcal{F}$. It then follows that any pair tester that queries $o(n^{3/2})$ pairs, has only $o(1)$ rejection probability on some function in $\mathcal{F}$, completing the proof. A pair (x, y) is "reasonable" if Hamming weights of x, y are in the range $2n \pm O(\sqrt{n \log n})$. The fraction of inputs that participate in an unreasonable pair is polynomially small in n and hence we may ignore unreasonable pairs without affecting our argument.

Fix any reasonable pair (x, y) where y is monotonically above x by a distance τ. Note that the pair is violating pair for the function $f_{S,j}$ if and only if $x_S = y_S$ with Hamming weight exactly n, $x_j = 0$ and $y_j = 1$. Denoting by $D_{x,y}$ the set of co-ordinates where x and y differ, $|D_{x,y}| = \tau$ and this is a violating pair if and only if

$$x_S \text{ has Hamming weight } n \qquad \text{and} \qquad S \cap D_{x,y} = \phi \qquad \text{and} \qquad j \in D_{x,y}.$$

The probability that the above event happens for a random choice of S, j is at most $O(\frac{1}{\sqrt{n}} \cdot 2^{-\Omega(\tau)} \cdot \frac{\tau}{n})$, where the three events indicated above have probabilities roughly $\frac{1}{\sqrt{n}}$, $2^{-\Omega(\tau)}$ and $\frac{\tau}{n}$ respectively and are nearly independent. This probability is maximized when $\tau = 1$ and is at most $O(1/n^{3/2})$.

9 Some Useful Results and Additional Observations

9.1 Bound on Total Influence

Recall that for a function $f : \{0,1\}^n \mapsto \{0,1\}$, $\mathcal{S}_f$ denotes the set of sensitive edges and $\mathcal{S}_f^- \subseteq \mathcal{S}_f$ denotes the set of negatively sensitive edges. The total influence is $I_f = \frac{|\mathcal{S}_f|}{2^n}$ and the total negative influence is $I_f^- = \frac{|\mathcal{S}_f^-|}{2^n}$. We show that if $I_f^- \ll I_f$ then I_f is at most $O(\sqrt{n})$. This upper bound on the total influence is very useful in the analysis of our monotonicity tester. Formally:

Theorem 9.1.

$$I_f^- \leqslant \frac{1}{3} \cdot I_f \implies I_f \leqslant 6\sqrt{n}.$$

Proof. For co-ordinate $i \in \{1, \ldots, n\}$, let (x_{-i}, b) denote an input that equals b in the i^{th} co-ordinate and $x_{-i} \in \{0,1\}^{n-1}$ in the remaining co-ordinates. Let $\mathbf{1}_{(\cdot)}$ denote the indicator of an event. One observes that the Fourier coefficient $\hat{f}(\{i\})$ is, by definition:

$$2 \cdot \hat{f}(\{i\}) = \mathbb{E}[f(x_{-i}, 0) - f(x_{-i}, 1)] = \mathbb{E}[\, 2 \cdot \mathbf{1}_{f(x_{-i},0) > f(x_{-i},1)} - \mathbf{1}_{f(x_{-i},0) \neq f(x_{-i},1)}].$$

Rearranging and summing over $i \in \{1, \ldots, n\}$ gives

$$\sum_{i=1}^{n} \mathbb{E}[\mathbf{1}_{f(x_{-i},0) \neq f(x_{-i},1)}] \;=\; 2 \cdot \sum_{i=1}^{n} \mathbb{E}[\mathbf{1}_{f(x_{-i},0) > f(x_{-i},1)}] - 2 \cdot \sum_{i=1}^{n} \hat{f}(\{i\}).$$

Observing that the first and the second sums are precisely I_f and I_f^- respectively and taking absolute values,

$$I_f \;\leqslant\; 2 \cdot I_f^- + 2 \cdot \sum_{i=1}^{n} |\hat{f}(\{i\})|.$$

Now if $I_f^- \;\leqslant\; \frac{1}{3}\, I_f$, it follows that I_f is at most $6 \cdot \sum_{i=1}^{n} |\hat{f}(\{i\})|$, which by Cauchy-Schwartz is at most $6\sqrt{n}\sqrt{\sum_{i=1}^{n} \hat{f}(\{i\})^2} \leqslant 6\sqrt{n}$. $\square$

9.2 Bound on Fraction of Non-persistent Inputs

Fix an integer parameter $\tau \in \left[1, \sqrt{\frac{n}{\log n}}\right]$ and a function $f : \{0,1\}^n \mapsto \{0,1\}$. For an input $x \in \{0,1\}^n$, a random input y that is monotonically above x at distance τ is picked by changing τ co-ordinates of x from 0 to 1, picked uniformly at random from the set of 0-co-ordinates of x. In the following, we assume that all inputs under consideration have Hamming weight in the range $\frac{1}{2}n \pm O(\sqrt{n \log n})$, since the fraction of remaining "atypical" inputs is at most, say $\frac{1}{n^{10}}$, and ignoring these atypical inputs has no effect on our arguments.

Definition 9.2. *An input $x \in \{0,1\}^n$ is called τ-persistent if over the choice of a random input y that is monotonically above x at distance τ,*

$$\mathsf{Pr}_y[f(x) \neq f(y)] \;\leqslant\; \frac{1}{10}.$$

The following lemma bounds the fraction of non-persistent inputs in terms of the total influence of the function. Here is an intuitive argument. For a non-persistent input, changing τ of its co-ordinates changes the value of the function with a constant probability. Thus, roughly speaking, changing one of its co-ordinates changes the value of the function with probability $\Omega(1/\tau)$ and hence, the total influence is at least the fraction of non-persistent inputs multiplied by $\Omega(n/\tau)$. However, a formal proof offers some subtleties.

Lemma 9.3. *Let $\tau \in \left[1, \sqrt{\frac{n}{\log n}}\right]$. The fraction of τ-non-persistent inputs is at most $O(\frac{I_f}{n} \cdot \tau)$.*

Proof. Let α be the fraction of non-persistent inputs. Consider the random process that picks input x uniformly at random and then picks, at random, input y that is monotonically above x by distance τ. We consider the probability that $f(x) \neq f(y)$. Clearly,

$$\mathsf{Pr}_{x,y}[f(x) \neq f(y)] \;=\; \frac{1}{2^n} \cdot \sum_{x \in \{0,1\}^n} \mathsf{Pr}_y[f(x) \neq f(y)] \;\geqslant\; \alpha \cdot \frac{1}{10}, \tag{10}$$

since for a non-persistent input x, the probability $\mathsf{Pr}_y[f(x) \neq f(y)]$ is at least $\frac{1}{10}$. Let $x = x^0, x^1, \ldots, x^\tau = y$ denote the sequence of inputs, starting with x and changing one co-ordinate at a time (from '0' to '1') till one reaches y. The order in which the co-ordinates are changed is itself random. Clearly,

$$\mathsf{Pr}_{x,y}[f(x) \neq f(y)] \;\leqslant\; \sum_{\ell=0}^{\tau-1} \mathsf{Pr}[f(x^\ell) \neq f(x^{\ell+1})]. \tag{11}$$

We note that while moving from x^ℓ to $x^{\ell+1}$ the co-ordinate that is changed is itself random among the '0'-co-ordinates of x^ℓ. The distribution of the input x^ℓ is not necessarily uniform (but close to it, as we will see). We observe that the lemma follows immediately if we were to assume that the distribution of x^ℓ is uniform. Indeed, supposing so, the distribution of the pair $(x^\ell, x^{\ell+1})$ is same as the distribution of the pair (u, v) where u is uniformly random and v is obtained by changing a random '0' co-ordinate of u to '1'. Thus, we would have, combining Equations (10), (11),

$$\frac{\alpha}{10} \;\leqslant\; \sum_{\ell=0}^{\tau-1} \mathsf{Pr}_{u,v}[f(u) \neq f(v)] = \tau \cdot \frac{I_f}{n},$$

completing the proof of the lemma. We now show how to handle the issue that the distribution of x^ℓ is not necessarily uniform. We show that for any fixed input $z \in \{0,1\}^n$ with Hamming weight in the range $\frac{1}{2}n \pm O(\sqrt{n \log n})$,

$$\mathsf{Pr}[z = x^\ell] \leqslant C \cdot \mathsf{Pr}[z = u],$$

for some constant C, i.e. the probability of sampling z from the distribution corresponding to that of x^ℓ is at most a constant times the probability of sampling z from the uniform distribution. Indeed, suppose the Hamming weight of z is $\frac{1}{2}n + k$ for some $k \in [-O(\sqrt{n \log n}), O(\sqrt{n \log n})]$. We note that if $z = x^\ell$, then x must have Hamming weight $\frac{1}{2}n + k - \ell$ and since x is uniformly random,

$$\mathsf{Pr}[z = x^\ell] \;=\; \left(\frac{1}{2^n} \cdot \binom{n}{\frac{1}{2}n + k - \ell}\right) \cdot \frac{1}{\binom{n}{\frac{1}{2}n+k}}.$$

Here the first factor reflects the probability that x has Hamming weight $\frac{1}{2}n + k - \ell$ and the second factor reflects the probability that z happens to be one specific input among all inputs with Hamming weight $\frac{1}{2}n+k$. By Lemma 9.4, $\mathsf{Pr}[z = x^\ell]$ is at most a constant times $\frac{1}{2^n} = \mathsf{Pr}[z = u]$ as claimed. Combining Equations (10), (11) as before,

$$\frac{\alpha}{10} \;\leqslant\; \sum_{\ell=0}^{\tau-1} \mathsf{Pr}[f(x^\ell) \neq f(x^{\ell+1})] \;\leqslant\; \sum_{\ell=0}^{\tau-1} C \cdot \mathsf{Pr}[f(u) \neq f(v)] \;=\; \tau \cdot C \cdot \frac{I_f}{n},$$

completing the proof of the lemma. $\square$

Lemma 9.4. *Consider integers* $\ell \in \left[0, \sqrt{\frac{n}{\log n}}\right]$ *and* $k \in [-O(\sqrt{n \log n}), O(\sqrt{n \log n})]$. *Then layers* $\frac{1}{2}n + k$ *and* $\frac{1}{2}n + k - \ell$ *of the* n*-dimensional hypercube have the same number of vertices up to a constant factor. That is:*

$$\frac{\binom{n}{\frac{1}{2}n+k-\ell}}{\binom{n}{\frac{1}{2}n+k}} = \Theta(1).$$

Proof.

$$\frac{\binom{n}{\frac{1}{2}n+k-\ell}}{\binom{n}{\frac{1}{2}n+k}} = \frac{(\frac{1}{2}n+k)!(\frac{1}{2}n-k)!}{(\frac{1}{2}n+k-\ell)!(\frac{1}{2}n-k+\ell)!} = \frac{(\frac{1}{2}n+k)(\frac{1}{2}n+k-1)...(\frac{1}{2}n+k-\ell+1)}{(\frac{1}{2}n-k+\ell)(\frac{1}{2}n-k+\ell-1)...(\frac{1}{2}n-k+1)} =$$
$$= \prod_{j=0}^{\ell-1}\left(1+\frac{2k-\ell}{\frac{1}{2}n-k+\ell-j}\right) = \left(1 \pm O\left(\sqrt{\frac{\log n}{n}}\right)\right)^{\ell} = \Theta(1),$$

since $|2k-\ell|$ and $|k-\ell+j|$ are $O(\sqrt{n\log n})$ and ℓ is at most $\sqrt{\frac{n}{\log n}}$. $\square$

9.3 Theorem 1.9 $\Longrightarrow$ Theorem 1.7

We show how to derive Theorem 1.7 from Theorem 1.9 (up to the poly-log factor). For a function $f : \{0,1\}^n \mapsto \{0,1\}$, Let $G_f^-(V,W,E)$ denote the bipartite graph of negatively sensitive edges as before (i.e. $E = \mathcal{S}_f^-$). We note that $I_f^- = \frac{|\mathcal{S}_f^-|}{2^n}$ and $\Gamma^-_{f,\mathsf{matching}}$ is the maximum size of a matching in G_f^- divided by 2^n. We intend to show that

$$I_f^- \cdot \Gamma^-_{f,\mathsf{matching}} \geqslant \tilde{\Omega}(\varepsilon(f)^2).$$

Towards this end, for subsets $A \subseteq V,\ B \subseteq W$, let (A,B) be a minimum vertex cover in G_f^- so that its size $|A|+|B|$ is also the maximum size of a matching. Color the edges of G_f^- with two colors so that all edges incident on A are colored red and all remaining edges (which must be incident on B) blue. Applying Theorem 1.9, we get that

$$\sum_{x\in A}\sqrt{I^-_{f,\mathsf{red}}(x)} + \sum_{y\in B}\sqrt{I^-_{f,\mathsf{blue}}(x)} \geqslant 2^n \cdot \tilde{\Omega}(\varepsilon(f)).$$

Using Cauchy-Schwartz, the L.H.S. is upper bounded by

$$\sqrt{|A|+|B|}\cdot\sqrt{\sum_{x\in A} I^-_{f,\mathsf{red}}(x) + \sum_{y\in B} I^-_{f,\mathsf{blue}}(x)} = \sqrt{\Gamma^-_{f,\mathsf{matching}}\cdot 2^n}\cdot\sqrt{|\mathcal{S}_f^-|}.$$

Combining these two observations and squaring both sides of the inequality gives the desired result.

9.4 Undirected Theorems from Directed Theorems

In this section, we show that Theorems 1.1, 1.2, 1.3, 1.8 follow from the corresponding directed versions of these theorems, namely, Theorems 1.4, 1.5, 1.6, 1.9 respectively (possibly up-to poly-log factor). Thus the directed versions are indeed generalizations of the undirected versions (as ought to be the case).

Definition 9.5. *For input $x \in \{0,1\}^n$, let $x^* \in \{0,1\}^n$ denote the input with all the bits of x flipped. For a function $f : \{0,1\}^n \mapsto \{0,1\}$, let $f^* : \{0,1\}^n \mapsto \{0,1\}$ denote the function defined so that*

$$\forall x \in \{0,1\}^n, \quad f(x) = f^*(x^*).$$

Lemma 9.6. *For a function $f : \{0,1\}^n \mapsto \{0,1\}$,*

$$\varepsilon(f) + \varepsilon(f^*) \geqslant \Omega(\mathsf{var}(f)).$$

Proof. Let h be a monotone function that is nearest to f so that $\varepsilon(f) = \Delta(f, h)$. Let $p = \mathbb{E}[f]$ and $q = \mathbb{E}[h]$ so that $|p - q| \leqslant \varepsilon(f)$. Thus

$$\mathsf{var}(h) = q(1-q) \geqslant p(1-p) - 2\,|p-q| \geqslant \mathsf{var}(f) - 2\,\varepsilon(f).$$

If $\varepsilon(f) \geqslant \frac{1}{4} \cdot \mathsf{var}(f)$ we are done. Otherwise, we get $\mathsf{var}(h) \geqslant \frac{1}{2} \cdot \mathsf{var}(f)$. Now let g be a monotone function that is nearest to f^* so that $\varepsilon(f^*) = \Delta(f^*, g) = \Delta(f, g^*)$. Since h is monotone and g^* is anti-monotone, using Lemma 3.12,

$$\varepsilon(f) + \varepsilon(f^*) = \Delta(f, h) + \Delta(f, g^*) \geqslant \Delta(h, g^*) \geqslant \mathsf{var}(h) \geqslant \frac{1}{2} \cdot \mathsf{var}(f),$$

and we are done. □

We make some preliminary observations and then all the implications will follow immediately. Let $G(V, W, E)$ denote the graph of sensitive edges of f, i.e. $f(V) \equiv 1$, $f(W) \equiv 0$, E is precisely the set of sensitive edges of f and every vertex in $V \cup W$ has at least one sensitive edge of f incident on it. By definition:

$$I_f = \frac{|E|}{2^n}, \qquad \Gamma_f = \frac{|V|}{2^n}, \qquad \mathbb{E}_x\left[\sqrt{I_f(x)}\right] = \frac{1}{2^n} \cdot \sum_{x \in V} \sqrt{\mathsf{deg}_G(x)}.$$

$G_1(V_1, W_1, E_1)$ be the graph of negatively sensitive edges of f. This is a subgraph of G induced by precisely the negatively sensitive edges, i.e. $(x, y) \in E_1$ if and only if $(x, y) \in E$ and $x \leqslant y$. Also, $V_1 \subseteq V$, $W_1 \subseteq W$ are precisely the subsets of vertices that have at least one negatively sensitive edge on them. By definition:

$$I_f^- = \frac{|E_1|}{2^n}, \qquad \Gamma_f^- = \frac{|V_1|}{2^n}, \qquad \mathbb{E}_x\left[\sqrt{I_f^-(x)}\right] = \frac{1}{2^n} \cdot \sum_{x \in V_1} \sqrt{\mathsf{deg}_{G_1}(x)}.$$

Let $E_2 = E \setminus E_1$ and let $G_2(V_2, W_2, E_2)$ be the subgraph of G induced by edges in E_2. Thus V_2, W_2 are subsets of vertices that have at least one edge in E_2 incident on them. We may write

$$G(V, W, E) \;=\; G_1(V_1, W_1, E_1) \;\cup\; G_2(V_2, W_2, E_2),$$

noting that $E_1 \cap E_2 = \phi$, but V_1, V_2 and similarly W_1, W_2 need not be disjoint.

Consider the mapping $\varphi : x \mapsto x^*$. This is clearly an isomorphism of the hypercube and let us denote the isomorphic copy of $G_2(V_2, W_2, E_2)$ under this isomorphism as $G_2^*(V_2^*, W_2^*, E_2^*)$. The main observation is that $G_2^*(V_2^*, W_2^*, E_2^*)$ is precisely the graph of negatively sensitive edges of the function f^*. Indeed, an edge $e = (x, y) \in E_2$ satisfies $f(x) = 1, f(y) = 0, x \geqslant y$. So its isomorphic copy $e^* = (x^*, y^*) \in E_2^*$ satisfies $f^*(x^*) = 1, f^*(y^*) = 0, x^* \leqslant y^*$ and hence is a negatively sensitive edge of f^*. The same argument works backwards.

The isoperimetric parameters for f^* concerning the graph $G_2^*(V_2^*, W_2^*, E_2^*)$ can now be expressed in terms of its isomorphic copy $G_2(V_2, W_2, E_2)$ as:

$$I_{f^*}^- = \frac{|E_2|}{2^n}, \qquad \Gamma_{f^*}^- = \frac{|V_2|}{2^n}, \qquad \mathbb{E}_x\left[\sqrt{I_{f^*}^-(x)}\right] = \frac{1}{2^n} \cdot \sum_{x \in V_2} \sqrt{\mathsf{deg}_{G_2}(x)}.$$

Theorem 1.1 now follows from Theorem 1.4 and Lemma 9.6 as:

$$I_f = \frac{|E|}{2^n} = \frac{|E_1|}{2^n} + \frac{|E_2|}{2^n} = I_f^- + I_{f^*}^- \geqslant \Omega(\varepsilon(f)) + \Omega(\varepsilon(f^*)) \geqslant \Omega(\mathsf{var}(f)).$$

Theorem 1.2 follows from Theorem 1.5 as:

$$I_f \cdot \Gamma_f = (I_f^- + I_{f^*}^-) \cdot \frac{|V|}{2^n} \geqslant (I_f^- + I_{f^*}^-) \cdot \frac{1}{2} \cdot \left(\frac{|V_1|}{2^n} + \frac{|V_2|}{2^n} \right) = (I_f^- + I_{f^*}^-) \cdot \frac{1}{2} \cdot \left(\Gamma_f^- + \Gamma_{f^*}^- \right),$$

which is lower bounded, up to the factor $\frac{1}{2}$, by

$$I_f^- \cdot \Gamma_f^- + I_{f^*}^- \cdot \Gamma_{f^*}^- \geqslant \Omega(\varepsilon(f)^2) + \Omega(\varepsilon(f^*)^2) \geqslant \Omega((\varepsilon(f) + \varepsilon(f^*))^2) \geqslant \Omega(\mathsf{var}(f)^2).$$

Theorem 1.3 follows from Theorem 1.6 (up to a poly-log factor) as:

$$\mathbb{E}_x \left[\sqrt{I_f(x)} \right] = \frac{1}{2^n} \cdot \sum_{x \in V} \sqrt{\mathsf{deg}_G(x)} \geqslant \frac{1}{2} \cdot \frac{1}{2^n} \cdot \left(\sum_{x \in V_1} \sqrt{\mathsf{deg}_{G_1}(x)} + \sum_{x \in V_2} \sqrt{\mathsf{deg}_{G_2}(x)} \right),$$

which is same, up to the factor $\frac{1}{2}$, as

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

Similarly, Theorem 1.8 follows from Theorem 1.9 (up to a poly-log factor), where a coloring of edges of E induces a coloring of edges of E_1 and E_2. We omit the straightforward proof.

10 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.

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] Shahar Fattal and Dana Ron. Approximating the distance to monotonicity in high dimensions. *ACM Trans. Algorithms*, 6(3):Art. 52, 37, 2010.

[9] 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.

[10] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre. Correlation inequalities on some partially ordered sets. *Comm. Math. Phys.*, 22:89–103, 1971.

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

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

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

[14] 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.

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

ECCC ISSN 1433-8092

http://eccc.hpi-web.de