Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 77 (2014)

Size of Sets with Small Sensitivity: a Generalization of Simon's Lemma*

Andris Ambainis and Jevgēnijs Vihrovs

Faculty of Computing, University of Latvia, Raiņa bulv. 19, Rīga, LV-1586, Latvia

Abstract. We study the structure of sets $S \subseteq \{0,1\}^n$ with small sensitivity. The well-known Simon's lemma says that any $S \subseteq \{0,1\}^n$ of sensitivity s must be of size at least 2^{n-s}. This result has been useful for proving lower bounds on sensitivity of Boolean functions, with applications to the theory of parallel computing and the "sensitivity vs. block sensitivity" conjecture.

In this paper, we take a deeper look at the size of such sets and their structure. We show an unexpected "gap theorem": if $S \subseteq \{0,1\}^n$ has sensitivity s, then we either have $|S| = 2^{n-s}$ or $|S| \geq \frac{3}{2}2^{n-s}$. This is shown via classifying such sets into sets that can be constructed from low-sensitivity subsets of $\{0,1\}^{n'}$ for $n' < n$ and *irreducible* sets which cannot be constructed in such a way and then proving a lower bound on the size of irreducible sets.

This provides new insights into the structure of low sensitivity subsets of the Boolean hypercube $\{0,1\}^n$.

1 Introduction

The complexity of computing Boolean functions (for example, in the decision tree model of computation) is related to a number of combinatorial quantities, such as sensitivity and block sensitivity of the function, its certificate complexity and the degree of polynomials that represent the function (exactly or approximately) [4]. Study of these quantities has resulted in both interesting results and longstanding open problems.

For example, it has been shown that decision tree complexity (in either deterministic or probabilistic or quantum model of computation) is polynomially related to a number of these quantities: certificate complexity, block sensitivity and the minimum degree of polynomials that represent or approximate f [8, 3]. This result, in turn, implies that deterministic, probabilistic and quantum

* The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under projects QALGO (Grant Agreement No. 600700) and RAQUEL (Grant Agreement No. 323970) and ERC Advanced Grant MQC. Part of this work was done while Andris Ambainis was visiting Institute for Advanced Study, Princeton, supported by National Science Foundation under agreement No. DMS-1128155. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

ISSN 1433-8092

decision tree complexities are polynomially related — which is very interesting because similar result is not known in Turing machine world (and, for deterministic vs. quantum complexity, is most likely false because of Shor's factoring algorithm).

The question about the relation between the sensitivity of a function and the other quantities is, however, a longstanding open problem, known as the "sensitivity vs. block sensitivity" question. (Since the other quantities are all polynomially related, showing a polynomial relation between the sensitivity and one of them would imply a polynomial relation between the sensitivity and all of them). This question has attracted quite much attention (since being first posed by Nisan in 1991 [7]) but there has been quite little progress and the gap between the best upper and lower bounds remains huge. The biggest separation between the two quantities is $bs(f) = \Omega(s^2(f))$ [9, 11, 2] (here, $bs(f)$ and $s(f)$ denote the block sensitivity and the sensitivity of f, respectively) while the best upper bound on $bs(f)$ in terms of $s(f)$ is exponential: $bs(f) \leq s(f)2^{s(f)-1} - s(f) + 1$ [6, 1].

In this paper, we study the following question: assume that a subset S of Boolean hypercube $\{0,1\}^n$ has a low sensitivity (that is, for every $x \in S$, there are at most s indices $i \in \{1, \ldots, n\}$ such that changing x_i to the opposite value results in $y \in S$). What can we say about this set?

Simon's lemma [10] says that a set $S \subset \{0,1\}^n$ with sensitivity s must contain at least 2^{n-s} input vectors $x \in S$. Simon [10] then used this result to show that $s(f) \geq \frac{1}{2}\log_2 n - \frac{1}{2}\log_2\log_2 n + \frac{1}{2}$ for any Boolean function that depends on n variables. Since $bs(f) \leq n$ (trivially), this implies $bs(f) \leq s(f)4^{s(f)}$. This was the first upper bound on $bs(f)$ in terms of $s(f)$.

Since then, this upper bound was improved, first by Kenyon and Kutin [6] to $bs(f) \leq C\sqrt{s(f)}e^{s(f)}$ (where C is constant) and then by Ambainis et al. [1], to $bs(f) \leq s(f)2^{s(f)-1} - s(f) + 1$. But the best currently known upper bound by [1] is based on isoperimetric inequality for Boolean hypercube which is closely related to Simon's lemma (and a bound of $bs(f) \leq s(f)2^{s(f)-1}$ which is only weaker by only a tiny amount can be derived by using Simon's lemma instead of isoperimetric inequality).

Because of that, we think that Simon's lemma is quite important for a better understanding of the "sensitivity vs. block sensitivity" problem and related questions. In this paper, we study the question: can one improve the bound of Simon's lemma?

Initially, it looks like Simon's lemma is exactly optimal. If we let S be a subcube of the hypercube $\{0,1\}^n$ obtained by fixing s of variables x_i (that is, S is the set of all $x = (x_1, \ldots, x_n)$ that satisfy $x_{i_1} = a_1$, ..., $x_{i_s} = a_s$ for some choice of $i_1, \ldots, i_k \in \{1, \ldots, n\}$ (which are all distinct) and $a_1, \ldots, a_s \in \{0,1\}$), then every $x \in S$ is sensitive to changing s bits $x_{i_1}, \ldots, x_{i_k}$ and $|S| = 2^{n-s}$.

In this paper, we show that any S with sensitivity s that is not a subcube must be substantially larger. To do that, we study the structure of sets S with sensitivity s by classifying them into two types:

1. sets S that are contained in a subcube $S' \subset \{0,1\}^n$ obtained by fixing one or more of values x_i;
2. sets S that are not contained in any such subcube.

There is one-to-one correspondence between the sets of the first type and low-sensitivity subsets of $\{0,1\}^{n-k}$ for $k \in \{1,\ldots,s\}$.[1] In contrast, the sets of the second type do not reduce to low-sensitivity subsets of $\{0,1\}^{n-k}$ for $k > 0$. Therefore, we call them *irreducible*.

Our main result is that any irreducible $S \subseteq \{0,1\}^n$ must be of size $|S| \geq 2^{n-s+1} - 2^{n-2s}$ (almost twice as large as a subcube obtained by fixing s variables) and this bound is tight.

As a consequence, we obtain a surprising result: if $S \subseteq \{0,1\}^n$ has sensitivity s, then either $|S| = 2^{n-s}$ or $|S| \geq \frac{3}{2}2^{n-s}$. That is, such a set S cannot have size between 2^{n-s} and $\frac{3}{2}2^{n-s}$!

2 Preliminaries

In this section we give the basic definitions used in the paper. Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function of n variables, where the i-th variable is denoted by x_i. We use $x = (x_1,\ldots,x_n)$ to denote a tuple consisting of all input variables x_i.

Definition 1. *The* sensitivity complexity $s(f,x)$ *of f on an input x is defined as $|\{i \mid f(x) \neq f(x^{(i)})\}|$, where $x^{(i)}$ is an input obtained from x by flipping the value of the i-th variable. The* sensitivity $s(f)$ *of f is defined as*

$$s(f) = \max\{s(f,x) \mid x \in \{0,1\}^n\}. \tag{1}$$

The c-sensitivity $s_c(f)$ *of f is defined as*

$$s_c(f) = \max\{s(f,x) \mid x \in \{0,1\}^n, f(x) = c\}. \tag{2}$$

In this paper, we will look at $\{0,1\}^n$ as a set of vertices for a graph Q_n (called *n-dimensional Boolean cube* or *hypercube*) in which we have an edge (x,y) whenever $x = (x_1,\ldots,x_n)$ and $y = (y_1,\ldots,y_n)$ differ in exactly one position. We look at subsets $S \subseteq \{0,1\}$ as subgraphs (induced by the subset of vertices S) in this graph.

Definition 2. *We define an m-dimensional* subcube *or m-*subcube *of Q_n to be a cube induced by the set of all vertices that have the same bit values on $n-m$ positions $x_{i_1},\ldots,x_{i_{n-m}}$ where i_j are all different.*

If a subcube can be obtained by fixing some continuous sequence b of starting bits, we denote such subcube by G_b. For example, G_0 and G_1 can be obtained by fixing a single first bit, and G_{01} can be obtained by fixing first two bits to 01. By G_{*10} we denote a cube obtained by fixing the second and the third bit to 10.

[1] If a set S of sensitivity s is contained in a subcube S' obtained by fixing $x_{i_1},\ldots,x_{i_k}$, removing the variables that have been fixed gives us a set $S'' \subseteq \{0,1\}^{n-k}$ of sensitivity $s-k$.

Definition 3. *Two m-dimensional subcubes of Q_n are* adjacent*, if the fixed $n - m$ positions of both subcubes are the same and their bit values differ in exactly one position.*

Each Boolean function f can be represented as a set of vertices $V(f) = \{x \mid f(x) = 1\}$, thus each function of n variables represents a single subgraph $G(f)$ of Q_n induced by $V(f)$. Note that for an input $x \in V(f)$, the sensitivity $s(f, x)$ is equal to the number of vertices not in $V(f)$ and connected to x with an edge in Q_n. Thus the sensitivity of $V(f)$ is equal to $s_1(f)$.

For a Boolean function f, the minimum degree $\delta(G(f))$ corresponds to $n - s_1(f)$, and the minimum degree of a graph induced by $\{0,1\}^n \setminus V$ corresponds to $n - s_0(f)$.

In the rest of this paper, we phrase our results in terms of subgraphs of Q_n.

Definition 4. *Let X and Y be subgraphs of Q_n. By $X \cap Y$ we denote the* intersection *graph of X and Y, which is the graph $(V(X) \cap V(Y), E(X) \cap E(Y))$. By $X \setminus Y$ denote the* complement *of Y in X, which is the graph induced by the vertex set $V(X) \setminus V(Y)$ in X.*

We also denote the degree of a vertex v in a graph G by $\deg(v, G)$.

The main focus of the paper is on the *irreducible* class of subgraphs:

Definition 5. *We call a subgraph $G \subset Q_n$* reducible*, if it is a subgraph of some graph $S \subset Q_n$, where $V(S)$ can be obtained by fixing one or more of values x_i. Conversely, other subgraphs we call* irreducible*.*

Another way to define the irreducible graphs is to say that each such graph contains at least one vertex in each of the $(n-1)$-subcubes of Q_n (so there are no redundant dimensions for this graph).

3 Simon's Lemma

In this section, we present a theorem proved by Simon [10] along with the proof. In the next section, we will build on the ideas from this proof to prove our results.

Theorem 1. *Let $G = (V, E)$ be a non-empty subgraph of Q_n ($n \geq 0$) of minimum cardinality among the subgraphs with $\delta(G) = d$ ($d \geq 0$). Then G is a d-dimensional subcube of Q_n and*

$$|V| = 2^d. \tag{3}$$

Proof. We use an induction on n.

1. Base case, $n = 0$. Then $G = Q_n$, and $|V| = 1$.

2. In the induction step we assume the claim holds for $n-1$, and prove it for n. If $G_j \cap G$ is empty for some $j \in \{0,1\}$, then G is a subgraph of G_{1-j} thus by the claim G is a d-subcube of G_{1-j}, and $|V| = 2^d$. Otherwise the graphs $G_0 \cap G$ and $G_1 \cap G$ are not empty, and their minimum degree is at least $d-1$, since each vertex in G_0 has exactly one neighbour in G_1. Then by the claim $|V| = |V(G_0 \cap G)| + |V(G_1 \cap G)| \geq 2 \cdot 2^{d-1} = 2^d$, and the minimum is achieved only if $G_0 \cap G$ and $G_1 \cap G$ both are $(d-1)$-subcubes of Q_n. Since $\delta(G) = d$, $G_0 \cap G$ is a subcube adjacent to $G_1 \cap G$, together forming a single d-subcube, and $|V| = 2^d$.

□

This theorem implies:

Corollary 1. *Let $f(x)$ be a Boolean function of n variables. If $f(x)$ is not always 0, then*

$$|\{x \mid f(x) = 1\}| \geq 2^{n-s_1(f)}, \tag{4}$$

and the minimum is obtained iff some $s_1(f)$ positions hold the same bit values for all $x : f(x) = 1$.

Proof. Let G be a subgraph of Q_n induced by the set of vertices $V = \{x \mid f(x) = 1\}$. The minimum degree of G is $\delta(G) = n - s_1(f)$. Then by Theorem 1 $|V| \geq 2^{n-s_1(f)}$. The minimum is obtained iff G is an $(n - s_1(f))$-subcube of Q_n. That means that then it is defined by some bits fixed in $s_1(f)$ positions. □

4 Smallest Irreducible Subgraphs

In this section we prove the following theorem.

Theorem 2. *Let $G = (V, E)$ be a non-empty subgraph of Q_n ($n \geq 1$) with the properties:*

1. *G is irreducible.*
2. *The minimum degree is $\delta(G) = d$ ($d \geq 0$).*

Then, the smallest possible cardinality of V is

$$|V| = \left\lceil 2^{d+1} - 2^{2d-n} \right\rceil \tag{5}$$

In the language of Boolean functions, this theorem corresponds to:

Corollary 2. *Let $f(x)$ be a Boolean function of n variables. If $\forall i \in [n] \forall j \in \{0,1\} \exists x (x_i = j, f(x) = 1)$, then*

$$|\{x \mid f(x) = 1\}| \geq 2^{n-s_1(f)+1} - 2^{n-2s_1(f)}, \tag{6}$$

Theorem 2 also implies the following generalization of Simon's lemma:

Theorem 3. *Let $G = (V, E)$ be a non-empty subgraph of Q_n ($n \geq 0$) with $\delta(G) = d$. Then either $|V| = 2^d$ or $|V| \geq \frac{3}{2} \cdot 2^d$, with $V = |2^d|$ achieved if an only if V is a d-subcube.*

Equivalently, if G has the sensitivity s, then either $|V| = 2^{n-s}$ or $|V| \geq \frac{3}{2}2^{n-s}$. Thus, there is a gap between the possible values for $|V|$ — which we find quite surprising.

In the next two subsections we prove Theorem 2, and in the last two subsections we show how it implies Corollary 2 and Theorem 3.

4.1 Sufficiency

In this section we prove that the given number of vertices is sufficient. We distinguish three cases:

1. $n = 1$. The only valid graph satisfying the properties is $G = Q_n$ with $d = 1$. Then $|V| = 2$.
2. $n > 1$, $2d < n$. Since $2^{2d-n} < 1$, $|V|$ should be 2^{d+1}. We take

$$S_j = \{x \mid \forall i \in [n-d](x_i = j)\} \tag{7}$$

   for $j \in \{0, 1\}$, and $V = S_0 \cup S_1$. Let G be the graph induced by V in Q_n. Then G consists of two d-subcubes of Q_n with no common vertices. Since $n - d > 1$, no edge connects any two vertices between these subcubes, thus $\delta(G) = d$. For the irreducibility, suppose that some $(n-1)$-subcube H is defined by fixing $x_i = j$. If $i \leq n - d$, then $H \cap S_j \neq \varnothing$. If $i > n - d$, then $H \cap S_j \neq \varnothing$ for any j. Then $|V| = 2 \cdot 2^d = 2^{d+1}$.
3. $n > 1$, $2d \geq n$. Then $|V|$ should be $2^{d+1} - 2^{2d-n}$. We take

$$S_l = \{x \mid \forall i \in [n-d](x_i = 1)\}, \tag{8}$$

$$S_r = \{x \mid \forall i \in [n-d+1; 2(n-d)](x_i = 1)\} \tag{9}$$

   and $V = S_l \cup S_r$. Let G be the graph induced by V in Q_n. Graphs induced by S_l and S_r are d-dimensional subcubes of Q_n. Since they are not adjacent, $\delta(G) = d$. For the irreducibility, observe that any bit position i is not fixed for at least one of S_l or S_r. Then $(n-1)$-subcube H obtained by fixing x_i holds at least one of vertices of G. Since $S_l \cap S_r = \{x \mid \forall i \in [2(n-d)](x_i = 1)\}$, it follows that

$$|V| = 2 \cdot 2^d - 2^{n-2(n-d)} = 2^{d+1} - 2^{2d-n}. \tag{10}$$

4.2 Optimality

In this section we prove that there are no such graphs with a number of vertices less than $\lceil 2^{d+1} - 2^{2d-n} \rceil$.

We use an induction on n. As the base case we take $n \leq 2$. From the fact that each $(n-1)$-subcube contains at least one vertex of G it follows that $|V| \geq 2$. This proves the cases $n = 1$, $d = 1$ and $n = 2$, $d = 0$ (and the case $n = 1$, $d = 0$

is not possible). Suppose $n = 2$, $d = 1$: if there were 2 vertices in G, then either some of the 1-subcubes would contain none of vertices of G or there would be a vertex of G with degree 0 (which is in contradiction with $d = 1$). Thus, in this case $|V| \geq 3 = 2^{1+1} - 2^{2-2}$. Suppose $n = 2$, $d = 2$. Then $G = Q_n$ and $|V| = 4 = 2^{2+1} - 2^{4-2}$.

In the induction step we assume the result holds for all $n' < n$, and prove it for n. Suppose that each of $(n-2)$-subcubes of Q_n contains at least one of vertices of G, then $G \cap G_0$ and $G \cap G_1$ are irreducible. The minimum degrees of $G \cap G_0$ and $G \cap G_1$ are at least $d-1$, since each vertex of $G \cap G_0$ can have at most one neighbour in G_1 (and conversely). By applying the inductive assumption to cubes G_0 and G_1, we obtain that

$$|V| \geq 2 \cdot \left\lceil 2^{(d-1)+1} - 2^{2(d-1)-(n-1)} \right\rceil = \tag{11}$$

$$= 2 \cdot \left\lceil 2^{d} - 2^{2d-n-1} \right\rceil \geq \tag{12}$$

$$\geq \left\lceil 2^{d+1} - 2^{2d-n} \right\rceil. \tag{13}$$

Now suppose that there is some $(n-2)$-subcube without vertices of G. WLOG assume it is G_{00}, i.e. $G \cap G_{00} = \varnothing$. We prove two lemmas.

Lemma 1. *Denote by $S(n, d)$ the minimum possible number of vertices for a graph that conforms to the conditions of Theorem 2. Let $G = (V, E)$ be a non-empty subgraph of Q_n ($n \geq 0$) with $\delta(G) = d$ ($d \geq 0$). Then either $|V| = 2^d$ or $|V| \geq \min_{i=d+1}^{n} S(i, d)$.*

Proof. We use an induction on n. Base case: $n = 0$. Then $G = Q_n$, $d = 0$ and $|V| = 1 = 2^{0-0}$. In the induction step we assume the Lemma holds for $n-1$, and prove it for n. If $n = d$, then $G = Q_n$, and $|V| = 2^n = 2^d$. Otherwise $n > d$. If each $(n-1)$-subcube of Q_n contains vertices of G, then $|V| \geq S(n, d)$ by the definition of S. Otherwise there is an $(n-1)$-subcube of Q_n that does not contain any of vertices of G. Then by induction the other $(n-1)$-subcube contains either 2^d or at least $\min_{i=d+1}^{n-1} S(i, d)$ of vertices of G. Combining the two cases together gives us the result. □

Lemma 2. *Let $G = (V, E)$ be a subgraph of Q_n ($n \geq 1$). Let $G' = G \cap G_0$. If G' is not empty and $\min_{v \in G'} \deg(v, G) \geq d$, then*

$$|V| \geq 2^d. \tag{14}$$

Note that this lemma is also a stronger version of Simon's result. Here we require the lower bound for the minimum degree only for vertices of G in one $(n-1)$-subcube of Q_n.

Proof. We use an induction on n.

1. Base case, $n = 1$. Since G' is non-empty, $G' = G_0$. If $d = 0$, $|V| \geq 1 = 2^0$. If $d = 1$, then $G = Q_n$ and $|V| = 2 = 2^1$.

2. In the induction step we assume the Lemma holds for $n-1$, and prove it for n. If $G_{0j} \cap G'$ is empty for some $j \in \{0,1\}$, then $G' \subseteq G_{0(1-j)}$. Thus by the Lemma $|V(G_{*(1-j)})| \geq 2^d$. Otherwise both G_{00} and G_{01} contain some of vertices of G. Since each of vertices of $G_{0j} \cap G$ has at most one neighbour in $G_{0(1-j)} \cap G$, it follows that $\min_{v \in G_{0j} \cap G} \deg(v, G_{*j}) \geq d-1$ for any $j \in \{0,1\}$. By applying the Lemma for $G_{*j} \cap G$ in cube G_{*j} for each j, we obtain that $|V| \geq 2 \cdot 2^{d-1} = 2^d$.

□

We observe that $\delta(G \cap G_{01}) \geq d-1$ and $\delta(G \cap G_{10}) \geq d-1$, since both $G \cap G_{01}$ and $G \cap G_{10}$ are adjacent to G_{00} (which does not contain vertices of G) and G_{11} (which may contain vertices of G). Now we distinguish two cases:

1. $|V(G \cap G_{01})| \neq 2^{d-1}$ and $|V(G \cap G_{10})| \neq 2^{d-1}$.
Cube G_{01} has $n-2$ dimensions and $\delta(G_{01} \cap G) \geq d-1$. By Lemma 1,

$$|V(G_{01} \cap G)| \geq \min_{i=(d-1)+1}^{n-2} S(i, d-1) = \min_{i=d}^{n-2} S(i, d-1). \tag{15}$$

From the inductive assumption, it follows that

$$|V(G_{01} \cap G)| \geq \min_{i=d}^{n-2} \left\lceil 2^{(d-1)+1} - 2^{2(d-1)-i} \right\rceil. \tag{16}$$

The minimum is achieved when i is the smallest, $i = d$. Thus $|V(G_{01} \cap G)| \geq \left\lceil 2^d - 2^{d-2} \right\rceil$. Similarly we prove that $|V(G_{10} \cap G)| \geq \left\lceil 2^d - 2^{d-2} \right\rceil$.
It remains to estimate the number of vertices of G in G_{11}. We deal with two cases:

(a) Some $(n-3)$-subcube of Q_n in G_{11} does not contain vertices of G. WLOG we assume it is G_{110}, i.e., $G \cap G_{110} = \varnothing$. We again distinguish two cases:

i. One of the subcubes G_{010} and G_{100} does not contain vertices of G. WLOG assume it is G_{010}, i.e., $G \cap G_{010} = \varnothing$. Then for the subcube G_{011} holds $\min_{v \in G \cap G_{011}} \deg(v, G \cap G_{*11}) \geq d$, since $G \cap G_{001} = \varnothing$ (because $G_{001} \subset G_{00}$), $G \cap G_{010} = \varnothing$ and G_{111} may contain vertices of G. Applying Lemma 2 to $G \cap G_{011}$ in G_{*11}, we get $|V(G \cap G_{*11})| \geq 2^d$. Similarly we prove that $|V(G \cap G_{10*})| \geq 2^d$. That gives us

$$|V| \geq 2 \cdot 2^d = 2^{d+1} \geq \left\lceil 2^{d+1} - 2^{2d-n} \right\rceil \tag{17}$$

and the case is done.

ii. Both of the subcubes G_{010} and G_{100} contain vertices of G. Then for the subcube G_{010} holds $\min_{v \in G \cap G_{010}} \deg(v, G \cap G_{01*}) \geq d$, since $G \cap G_{000} = \varnothing$, $G \cap G_{110} = \varnothing$, and G_{011} may contain vertices of G. Applying Lemma 2 to $G \cap G_{010}$ in G_{01*}, we get $|V(G \cap G_{01*})| \geq 2^d$. Similarly we prove that $|V(G \cap G_{10*})| \geq 2^d$. That gives us

$$|V| \geq 2 \cdot 2^d = 2^{d+1} \geq \left\lceil 2^{d+1} - 2^{2d-n} \right\rceil \tag{18}$$

and this case also is done.

(b) Each $(n-3)$-subcube of Q_n in G_{11} contains vertices of G. Since G_{11} is adjacent to G_{01} and G_{10}, $\delta(G \cap G_{11}) \geq d-2$. From the inductive assumption it follows that

$$|V(G \cap G_{11})| \geq 2^{(d-2)+1} - 2^{2(d-2)-(n-2)} = 2^{d-1} - 2^{2d-n-2}. \quad (19)$$

Thus

$$\begin{aligned} |V| &= |V(G \cap G_{01})| + |V(G \cap G_{10})| + |V(G \cap G_{11})| \geq & (20)\\ &\geq 2 \cdot \left\lceil 2^d - 2^{d-2} \right\rceil + \left\lceil 2^{d-1} - 2^{2d-n-2} \right\rceil \geq & (21)\\ &\geq \left\lceil 2 \cdot (2^d - 2^{d-2}) + 2^{d-1} - 2^{2d-n-2} \right\rceil = & (22)\\ &= \left\lceil 2^{d+1} - 2^{d-1} + 2^{d-1} - 2^{2d-n-2} \right\rceil = & (23)\\ &= \left\lceil 2^{d+1} - 2^{2d-n-2} \right\rceil, & (24) \end{aligned}$$

which is not less than $\left\lceil 2^{d+1} - 2^{2d-n} \right\rceil$, and this case is complete.

2. $|V(G \cap G_{01})| = 2^{d-1}$ or $|V(G \cap G_{10})| = 2^{d-1}$. WLOG assume that this holds for G_{01}.

By Theorem 1 it follows that $G \cap G_{01}$ is a $(d-1)$-dimensional subcube of Q_n. WLOG we can assume that the set of its vertices is

$$V(G \cap G_{01}) = \{x \mid x_1 = 0, \forall i \in [2; n-d+1](x_i = 1)\}. \quad (25)$$

Observe that $\deg(v, G \cap G_{01}) = d-1$ for all $v \in G \cap G_{01}$. Since $\delta(G) = d$, each $x \in V(G \cap G_{01})$ has $x^{(1)}$ as a neighbour in G. Then $\{x^{(1)} \mid x \in V(G \cap G_{01})\} = V(G \cap G_{11})$, and $G \cap G_{11}$ contains a $(d-1)$-subcube of Q_n, adjacent to $G \cap G_{01}$. Thus $G \cap G_{*1}$ contains a d-subcube of Q_n, with $\{x \mid \forall i \in [2; n-d+1](x_i = 1)\} \subseteq V(G \cap G_{*1})$.

Now we need to estimate the number of vertices of G in G_1 that do not belong to this subcube. We prove the following lemma.

Lemma 3. *Let $G = (V, E)$ be a subgraph of Q_n with the property that $G \cap G_0 \neq \varnothing$. Let G_l be an l-subcube of Q_n ($n \geq 1$) such that $G_l \subseteq G_1$. If $G_l \subset G$ and $\min_{v \in G \cap G_0} \deg(v, G) \geq d$, then*

$$|V(G \setminus G_l)| \geq 2^d - 2^{d-(n-l)}. \quad (26)$$

Note that the subcube G_l is defined by fixing $n-l$ bits.

Proof. We use an induction on $n-l$. Base case: $n-l=1$. Then $G_l = G_1$ and $\delta(G \cap G_0) = d-1$. By Theorem 1 it follows that $|V(G \cap G_0)| \geq 2^{d-1}$, and $2^{d-1} = 2^d - 2^{d-1}$.

In the induction step we assume the Theorem holds for $n-l-1$, and prove it for $n-l$. WLOG we can assume that the value of each fixed bit defining G_l is 1. We distinguish two cases:

(a) $G \cap G_{10} = \varnothing$. Then $\min_{v \in G \cap G_{00}} \deg(v, G) \geq d$. By applying Lemma 2 to $G \cap G_0$ and G_0, we obtain

$$|V(G \setminus G_l)| \geq |V(G \cap G_0)| \geq 2^d > 2^d - 2^{d-(n-l)}. \quad (27)$$

(b) $G \cap G_{10} \neq \varnothing$. Then $\min_{v \in G \cap G_{10}} \deg(v, G) \geq d-1$, since G_{00} may contain vertices of G. Then we apply the induction to $G \cap G_1$ ($G \cap G_1$ as the graph, G_1 as the binary cube, G_l remains the same) and obtain

$$|V(G \cap G_1 \setminus G_l)| \geq 2^{d-1} - 2^{d-1-(n-1-l)} = 2^{d-1} - 2^{d-(n-l)}. \quad (28)$$

On the other hand, $\delta(G \cap G_0) \geq d-1$, since G_1 contains vertices of G. We apply Theorem 1 to $G \cap G_0$ in G_0 and we obtain $|V(G \cap G_0)| \geq 2^{d-1}$. Then

$$|V(G \setminus G_l)| = |V(G \cap G_1 \setminus G_l)| + |V(G \cap G_0)| \geq \quad (29)$$

$$\geq 2^{d-1} - 2^{d-(n-l)} + 2^{d-1} = 2^d - 2^{d-(n-l)} \quad (30)$$

and we are done.

□

Consider the graph $G \cap G_1$. It contains the $(d-1)$-subcube induced by the set $\{x \mid \forall i \in [1; n-d+1](x_i = 1)\}$, denote it by D. Since $G \cap G_{00} = \varnothing$, the subcube G_{10} must contain vertices of G, so the $(n-1)$-subcube G_{*0} contains vertices of G. Since $G \cap G_{00} = \varnothing$, it follows that $\min_{v \in G \cap G_{10}} \deg(v, G \cap G_1) \geq d$. Thus we can apply Lemma 3 and get

$$|V(G_1 \setminus D)| \geq 2^d - 2^{d-((n-1)-(d-1))} = 2^d - 2^{2d-n}. \quad (31)$$

Since $G \cap G_0$ is a $(d-1)$-subcube of Q_n, $|V(G \cap G_0)| = 2^{d-1}$. Then

$$|V| = |V(G \cap G_0)| + |V(D)| + |V(G_1 \setminus D)| \geq \quad (32)$$

$$\geq 2^{d-1} + 2^{d-1} + (2^d - 2^{2d-n}) = \quad (33)$$

$$= 2^{d+1} - 2^{2d-n} \quad (34)$$

This completes the proof of Theorem 2. □

4.3 Application for Boolean Functions

The result lets us prove Corollary 2.

Proof. Let G be a subgraph of Q_n induced by the set of vertices $V = \{x \mid f(x) = 1\}$. The minimum degree of G is $\delta(G) = n - s_1(f)$. The given constraint means that G is irreducible. Then, by Theorem 2,

$$|V| \geq 2^{(n-s_1(f))+1} - 2^{2(n-s_1(f))-n} = 2^{n-s_1(f)+1} - 2^{n-2s_1(f)}. \quad (35)$$

□

4.4 Improvement of Simon's Lemma

We can use the obtained result to prove Theorem 3, which is a stronger version of Simon's lemma (Theorem 1).

Proof. We substitute $\lceil 2^d - 2^{2d-n} \rceil$ instead of $S(n,d)$ in Lemma 1. Then, in

$$\min_{i=d+1}^{n} S(i,d) = \min_{i=d+1}^{n} \lceil 2^d - 2^{2d-i} \rceil \tag{36}$$

the minimum is obtained for $i = d+1$. Thus either $|V| = 2^d$ or $|V| \geq 3 \cdot 2^{d-1}$. □

5 Conclusion

In this paper, we have shown two results on the structure of low sensitivity subsets of Boolean hypercube:

- a tight lower bound on the size of irreducible low sensitivity sets $S \subseteq \{0,1\}^n$ (that is, sets S that are not contained in any subcube of $\{0,1\}^n$ obtained by fixing one or more variables x_i);
- a gap theorem that shows that $S \subseteq \{0,1\}^n$ of sensitivity s must either have $|S| = 2^{n-s}$ or $|S| \geq \frac{3}{2} 2^{n-s}$.

Th gap theorem follows from the first result, by classifying $S \subseteq \{0,1\}^n$ into irreducible sets and sets that are constructed from irreducible subsets $S' \subseteq \{0,1\}^{n-k}$ for some $k \in \{1,2,\ldots,s\}$ and then using the first result for each of those categories.

We find this gap theorem quite surprising. A gap theorem of a similar type is known for the spectral norm of Boolean functions [5]: the spectral norm of a Boolean function is either equal to 1 or is at least 3/2. Both results have the constant 3/2 appearing in them and there is some resemblance between the constructions of optimal sets/functions but the proof methods are quite different and it is not clear to us if there is a more direct connection between the results.

Both results contribute to understanding the structure of low-sensitivity subsets of Boolean hypercube and we hope that they will find applications in resolving the "sensitivity vs. block sensitivity" conjecture or other open problems that involve the sensitivity of Boolean functions.

References

1. A. Ambainis, M. Bavarian, Y. Gao, J. Mao, X. Sun, and S. Zuo. New decision tree complexity upper bounds in terms of sensitivity. *ICALP'2014, to appear.*
2. A. Ambainis and X. Sun. New separation between $s(f)$ and $bs(f)$. *CoRR*, abs/1108.3494, 2011.
3. R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001.

4. H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: a survey. *Theor. Comput. Sci.*, 288(1):21–43, 2002.
5. B. Green and T. Sanders. Boolean functions with small spectral norm. *Geometric and Functional Analysis*, 18(1):144–162, 2008.
6. C. Kenyon and S. Kutin. Sensitivity, block sensitivity, and ℓ-block sensitivity of Boolean functions. *Information and Computation*, 189(1):43 – 53, 2004.
7. N. Nisan. CREW PRAMs and decision trees. *SIAM Journal on Computing*, 20:999–1007, 1991.
8. N. Nisan and M. Szegedy. On the degree of Boolean functions as real polynomials. *Computational Complexity*, 4(4):301–313, 1994.
9. D. Rubinstein. Sensitivity vs. block sensitivity of Boolean functions. *Combinatorica*, 15(2):297–299, 1995.
10. H.-U. Simon. A tight $\Omega(\log\log N)$-bound on the time for parallel RAM's to compute nondegenerated Boolean functions. In *Proceedings of the 1983 International FCT-Conference on Fundamentals of Computation Theory*, pages 439–444, London, UK, 1983. Springer-Verlag.
11. M. Virza. Sensitivity versus block sensitivity of Boolean functions. *Inf. Process. Lett.*, 111(9):433–435, 2011.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de