On the Shrinkage Exponent for Read-Once Formulae

Johan Håstad
Royal Institute of Technology
100 44 Stockholm, SWEDEN

Alexander Razborov*
Steklov Mathematical Institute
Vavilova 42, 117966, GSP–1, Moscow, RUSSIA

Andrew Yao†
Princeton University
Princeton, NJ 08544, USA

Abstract

We prove that the size of any read-once de Morgan formula reduces in average by a factor of at least $p^{\alpha-o(1)}$ when all but a fraction p of the input variables are randomly assigned to $\{0,1\}$ (here $\alpha \rightleftharpoons 1/\log_2(\sqrt{5}-1) \approx 3.27$). This resolves in the affirmative a conjecture of Paterson and Zwick. The bound is shown to be tight up to a polylogarithmic factor for all $p \geq n^{-1/\alpha}$.

1. Introduction

Assume that we randomly assign all but a fraction p of variables in a de Morgan formula of size s. What will be the expected formula size of the induced function? The obvious answer of course is that this size will be at most ps.

*This research was done while the last two authors were visiting the Royal Institute of Technology and Mittag Leffler, Stockholm.

†The research of this author was supported in part by the US National Science Foundation under grant NSF-8813283 and by a Guggenheim Fellowship.

Subbotovskaya [13] was the first to observe that actually formulae shrink more. Namely she established an upper bound

$$O\left(p^{1.5}s+1\right) \tag{1}$$

on the expected formula size of the induced function. This result allowed her to derive an $\Omega\left(n^{1.5}\right)$ lower bound on the de Morgan formula size of the parity function.

This latter bound was superseded by Khrapchenko [14, 15] who, using a different method, proved a tight $\Omega\left(n^2\right)$ lower bound for the parity function. His result implied that the parity function shrinks by a factor $\theta\left(p^2\right)$, and provided an upper bound $\Gamma \leq 2$ on the *shrinkage exponent* Γ, defined as the least upper bound of all γ that can replace 1.5 in (1).

The study of shrinking properties of constant depth circuits led in [5, 1, 11, 6] to good lower bounds for such circuits.

The new impetus for research on the expected size of the reduced formula was given by Andreev [12] who, based upon Subbotovskaya's result, derived an $n^{2.5-o(1)}$ lower bound on the de Morgan formula size for a function in P. A close inspection of the proof reveals that his method actually gives for the same function the bound $n^{\Gamma+1-o(1)}$.

New improvements of the lower bound on Γ followed. Nisan and Impagliazzo [8] proved that $\Gamma \geq \frac{21-\sqrt{73}}{8} \approx 1.55$. Paterson and Zwick [9], complementing the technique from [8] by very clear and natural arguments, pushed this bound further to $\Gamma \geq \frac{5-\sqrt{3}}{2} \approx 1.63$. Combined with Andreev's work this gives the currently best known lower bound $\Omega\left(n^{2.63}\right)$ on the de Morgan formula size for functions in NP.

It is generally believed that $\Gamma = 2$ (see e.g. [8, 9]). A natural starting point to prove this conjecture is to investigate the special case of read-once formulae. Note that Khrapchenko's example of parity function does not provide a shrink-resistant instance in this case, and the only upper bound known so far on the shrinkage exponent Γ^* for read-once formulae was proved by Paterson and Zwick in [9]. In that paper a sequence of read-once functions in n variables was presented so that the expected size of the induced functions is at least $\Omega\left(pn^{1/\alpha}\right)$ where $\alpha \rightleftharpoons 1/\log_2(\sqrt{5}-1) \approx 3.27$. With $p = Cn^{-1/\alpha}$ this gives the upper bound $\Gamma^* \leq \alpha$. Paterson and Zwick conjectured that this bound is tight, that is $\Gamma^* = \alpha$.

The main purpose of this paper is to prove their conjecture. More precisely, we show that the expected formula size of the function resulted from a read-once formula in n variables after assigning in it all but a fraction p of the variables at random is at most $O\left(p^\alpha\left(\log\frac{1}{p}\right)^{\alpha-1}n+(\log n)^{-1}\right)$ (Theorem 2.1). If the original formula is balanced then

the factor $\left(\log \frac{1}{p}\right)^{\alpha-1}$ can be omitted (Corollary 2.4). We also improve, in the range $p = p(n) \geq n^{-1/\alpha}$, the upper bound of Paterson and Zwick by presenting an example of read-once functions in n variables with the expected size of the induced functions $\Omega\left(p^{\alpha} n\right)$ (Theorem 2.5). This shows that our lower bounds are tight up to a polylogarithmic factor.

In the heart of our approach lie various links between the shrinkage properties of a function and its behavior under random restrictions assigning *all* variables. These links allow us to apply to our problem the strong machinery developed by Valiant [10] and Boppana [2].

More generally, our proofs are assembled from several independent pieces. It seems that many of these auxiliary statements have a scope of application much broader than the original task they were designed for. We hope that at least some of them will be useful for attacking the general case.

The paper is organized as follows. In Section 2 we introduce the necessary notations and state our main results. In Section 3 we exhibit our collection of auxiliary lemmas. By the reasons explained above we prefer to gather them in one place and formulate with reasonable generality. After that it is comparatively easy to prove our main result, which we do in section 4. Section 5 contains a simpler proof of the slightly better lower bound for the case of balanced formulae. In the last section 6 we present an example showing that our bounds are tight up to a polylogarithmic factor.

2. Preliminaries

A *de Morgan formula* is a binary tree in which each leaf is labeled by a literal from the set $\{x_1, \ldots, x_n, \bar{x}_1, \ldots, \bar{x}_n\}$ and each internal node v is labeled by an operation $o(v)$ which is either $\wedge$ or $\vee$. The *size* of a formula F is defined as the number of leaves and is denoted by $L(F)$. The *depth* $D(F)$ is the depth of the underlying tree. The *size* and the *depth* of a Boolean function f are, respectively, the minimal size and depth of any de Morgan formula computing f in the natural sense. For convenience we define the size and depth of a constant function to be 0.

A de Morgan formula is *read-once* if for each input variable x_i there exists exactly one leaf labeled by x_i or $\bar{x}_i$. We will always assume that leaves of a read-once formula are numbered in such a way that the l-th leaf is labeled by x_l or $\bar{x}_l$. A Boolean function is *read-once* if it can be computed by a read-once formula. For a read-once function f, $L(f)$ equals the number of variables f essentially depends on.

A de Morgan formula is *balanced* if the underlying tree is balanced that is all branches

have the same length. A Boolean function is *read-once balanced* if it can be computed by a de Morgan formula which is both read-once and balanced. Clearly,[1] $D(f) = \log L(f)$ for read-once balanced functions f.

A de Morgan formula is *monotone* if it contains no negated literals from $\{\bar{x}_1, \ldots, \bar{x}_n\}$. Monotone formulae compute monotone (in the natural sense) Boolean functions.

A *restriction* is an element of $\{0, 1, *\}^n$. For $p \in [0, 1]$ let $\boldsymbol{\rho}_p$ be the random restriction, in which we set randomly and independently each variable to $*$ with probability p and to $0, 1$ with equal probabilities $\frac{1-p}{2}$. A restriction ρ naturally takes a function f of n variables into a function of the variables given the value $*$ by ρ. We will denote this function by $\rho(f)$. A probability distribution on restrictions together with a fixed function f gives a probability distribution on functions, and we denote $\boldsymbol{\rho}_p(f)$ by $\boldsymbol{f}_p$. Let

$$\begin{aligned} E_f(p) &\rightleftharpoons \mathbf{E}\left[L(\boldsymbol{f}_p)\right], \\ N_f(p) &\rightleftharpoons \mathbf{P}\left[\boldsymbol{f}_p \not\equiv \text{const}\right], \\ C_f^{\epsilon}(p) &\rightleftharpoons \mathbf{P}\left[\boldsymbol{f}_p \equiv \epsilon\right] \ (\epsilon \in \{0, 1\}). \end{aligned}$$

Clearly,

$$N_f(p) + C_f^0(p) + C_f^1(p) = 1. \tag{2}$$

The *shrinkage exponent* Γ^* *for read-once formulae* [9] is defined as the least upper bound for those constants γ for which the bound $E_f(p) \le O\left(p^\gamma L(f) + 1\right)$ holds uniformly for all read-once functions f. Let $\alpha \rightleftharpoons 1/\log(\sqrt{5} - 1) \approx 3.27$. Paterson and Zwick [9, Theorem 6.5] came up with an example of read-once functions $f(x_1, \ldots, x_n)$ such that $E_f(p) \ge \Omega\left(pn^{1/\alpha}\right)$. This showed $\Gamma^* \le \alpha$. We have the following:

Theorem 2.1. *For any read-once function* $f(x_1, \ldots, x_n)$,

$$E_f(p) \le O\left(p^\alpha \left(\log \frac{1}{p}\right)^{\alpha-1} n + (\log n)^{-1}\right).$$

Hence $\Gamma^* = \alpha$.

Remark 2.2. For the case when $p^\alpha n$ is small, we have the better bound matching the lower bound of Paterson and Zwick. It will be explicitly stated in Theorem 4.1.

If f is, in addition, known to be balanced, we can do slightly better.

[1] all logarithms and exponents in this paper are to base 2

Theorem 2.3. *For read-once functions f,*

$$E_f(p) \leq O\left(p^\alpha 2^{D(f)} + 1\right).$$

Corollary 2.4. *For read-once balanced functions $f(x_1, \ldots, x_n)$,*

$$E_f(p) \leq O\left(p^\alpha n + 1\right).$$

We also have the following lower bound.

Theorem 2.5. *For each real-valued function $p(d)$ with $2^{-d/\alpha} \leq p(d) \leq 1$ there exists a sequence of read-once functions f_d such that $D(f_d) \leq d$ and*

$$E_{f_d}(p(d)) \geq \Omega\left(p^\alpha 2^d\right).$$

This shows that the bound of Theorem 2.3 is tight (in the interval $p^\alpha 2^{D(f)} \geq 1$) and the bound of Theorem 2.1 is tight up to a polylogarithmic factor.

3. Lemmas

Let $\boldsymbol{\rho}^q$ be the random restriction which assigns independently each variable to 1 with probability q and to 0 with probability $(1-q)$. It may be helpful for the reader to think of p as being small and q as being around $1/2$. We will denote $\boldsymbol{\rho}^q(f)$ by $\boldsymbol{f}^q$. Note that, unlike $\boldsymbol{f}_p$, $\boldsymbol{f}^q$ is always a constant. Let

$$A_f(q) \rightleftharpoons \mathbf{P}\left[\boldsymbol{f}^q \equiv 1\right].$$

The following remarkable result of Boppana [2] lies in the heart of our approach:

Proposition 3.1. (Boppana) *If f is a read-once monotone function and $q \in (0,1)$ then*

$$A_f'(q) \leq L(f)^{1/\alpha} \cdot \frac{H(A_f(q))}{H(q)},$$

where $H(q) = -q \log q - (1-q)\log(1-q)$.

Our first lemma is an easy exercise in mathematical calculus:

Lemma 3.2. *Let f be as in Proposition* 3.1, $p \in (0,1)$ *and*

$$x \geq \exp\left(-\frac{1}{2}p^{-1}L(f)^{-1/\alpha}\right). \tag{3}$$

Then

$$A_f\left(\frac{1+p}{2}\right) \leq A_f(1/2)\left(1 + pL(f)^{1/\alpha}\log\frac{1}{x}\right) + O(x). \tag{4}$$

Proof. Because of the term $O(x)$ in (4) we may assume that

$$A_f\left(\frac{1+p}{2}\right) \geq x \tag{5}$$

and x is arbitrarily small. By (3) we may assume now that p is also arbitrarily small.

By the mean value theorem and Proposition 3.1,

$$\Delta \rightleftharpoons A_f\left(\frac{1+p}{2}\right) - A_f(1/2) = \frac{p}{2}\cdot A_f'(q) \leq \frac{p}{2}\cdot L(f)^{1/\alpha}\cdot\frac{H(A_f(q))}{H(q)},\ q \in \left[\frac{1}{2}, \frac{1+p}{2}\right]. \tag{6}$$

We are going to prove that

$$\Delta \leq \frac{2}{3}pL(f)^{1/\alpha}A_f\left(\frac{1+p}{2}\right)\log\frac{1}{x}. \tag{7}$$

Consider two cases.

Case 1. $A_f\left(\frac{1+p}{2}\right) \geq 0.01$. Since p and x are arbitrarily small,

$$\Delta \leq \text{ (by (6))}\ p\cdot L(f)^{1/\alpha} \leq \frac{2}{3}pL(f)^{1/\alpha}\cdot 0.01\cdot\log\frac{1}{x}$$

which proves (7) in this case.

Case 2. $A_f\left(\frac{1+p}{2}\right) \leq 0.01$. Since f is monotone, $A_f(q) \leq A_f\left(\frac{1+p}{2}\right)$. Along with the assumption of Case 2, this gives $H(A_f(q)) \leq H\left(A_f\left(\frac{1+p}{2}\right)\right) \leq 1.3\cdot A_f\left(\frac{1+p}{2}\right)\log\left(\frac{1}{A_f\left(\frac{1+p}{2}\right)}\right)$ and allows us to continue the chain of inequalities (6) as follows:

$$\Delta \leq \frac{2}{3}p\cdot L(f)^{1/\alpha}\cdot A_f\left(\frac{1+p}{2}\right)\log\left(\frac{1}{A_f\left(\frac{1+p}{2}\right)}\right) \leq \text{ (by (5))}\ \frac{2}{3}p\cdot L(f)^{1/\alpha}\cdot A_f\left(\frac{1+p}{2}\right)\log\frac{1}{x}.$$

So, (7) is proved. It implies $A_f(1/2) \geq \left(1 - \frac{2p}{3} L(f)^{1/\alpha} \log \frac{1}{x}\right) A_f\left(\frac{1+p}{2}\right)$ and then

$$A_f\left(\frac{1+p}{2}\right) \leq A_f(1/2)\left(1 + pL(f)^{1/\alpha} \log \frac{1}{x}\right)$$

since $pL(f)^{1/\alpha} \log \frac{1}{x} \leq 1/2$ by (3).■

Our main tool is the following lemma which for monotone f allows us to express $E_f(p), C_f^\epsilon(p), N_f(p)$ in terms of $A_f(p)$:

Lemma 3.3. *If f is monotone then:*

a) $C_f^1(p) = A_f\left(\frac{1-p}{2}\right)$,

b) $C_f^0(p) = 1 - A_f\left(\frac{1+p}{2}\right)$,

c) $N_f(p) = A_f\left(\frac{1+p}{2}\right) - A_f\left(\frac{1-p}{2}\right)$.

Proof. a) First note that a monotone function is identically 1 iff it takes the value 1 on the all zero input. Let $\boldsymbol{\tau}$ be the restriction which assigns to 0 all variables set to $*$ by $\boldsymbol{\rho}_p$. Clearly $\boldsymbol{\tau}\boldsymbol{\rho}_p$ has the same distribution as $\boldsymbol{\rho}^{\left(\frac{1-p}{2}\right)}$ and hence by the above observation we have $C_f^1(p) = \mathbf{P}\left[\boldsymbol{f}_p \equiv 1\right] =$ (since f and hence $\boldsymbol{f}_p$ are monotone) $\mathbf{P}\left[\boldsymbol{\tau}(\boldsymbol{f}_p) \equiv 1\right] = \mathbf{P}\left[\boldsymbol{\rho}^{\left(\frac{1-p}{2}\right)}(f) \equiv 1\right] = A_f\left(\frac{1-p}{2}\right)$.

b) is dual to a) and is proved in the same way.

c) immediately follows from a), b) and (2).■

Assume now that F is a formula, l is a leaf and $v_1, v_2, \ldots, v_{d(l)} = l$ is the path leading from the root ($= v_1$) to this leaf. For $1 \leq i \leq d(l) - 1$ consider the subtree rooted at the brother node of v_{i+1}. Let $f_i(l)$ be the function computed at the root of this subtree,

$$K_{\wedge,l}(F) \rightleftharpoons \bigwedge_{\substack{1 \leq i \leq d(l)-1 \\ o(v_i)=\wedge}} f_i(l),$$

$$K_{\vee,l}(F) \rightleftharpoons \bigvee_{\substack{1 \leq i \leq d(l)-1 \\ o(v_i)=\vee}} f_i(l),$$

$$K_l(F) \rightleftharpoons K_{\wedge,l}(F) \wedge (\neg K_{\vee,l}(F)).$$

If F is a de Morgan formula computing a function f, we say that a restriction ρ *kills a variable x_l in F* if and only if $\rho(f)$ does not depend on x_l. Our next lemma describes in terms of the function $K_l(F)$ the killing relation for the case when F is read-once:

Lemma 3.4. *Let F be a read-once formula, l be a leaf and ρ be a restriction. Then ρ kills x_l in F if and only if $\rho(x_l) \neq *$ or $\rho(K_l(F)) \equiv 0$.*

Proof. Obvious from the construction of $K_l(F)$.■

We derive now from Lemma 3.4 two extremely useful formulas.

Lemma 3.5. *If f is the function computed by a read-once formula F, then*

$$E_f(p) = p \cdot \sum_{l=1}^{n} \left(1 - C^0_{K_l(F)}(p)\right).$$

Proof. $L(\boldsymbol{f}_p)$ equals the number of leaves not killed by $\boldsymbol{\rho}_p$. Therefore

$$E_f(p) = \sum_{l=1}^{n} \mathbf{P}\left[\boldsymbol{\rho}_p \text{ does not kill } x_l\right] = \text{ (by Lemma 3.4)}$$
$$\sum_{l=1}^{n} \mathbf{P}\left[\boldsymbol{\rho}_p(x_l) = *, \boldsymbol{\rho}_p(K_l(F)) \not\equiv 0\right] = p \cdot \sum_{l=1}^{n} \left(1 - C^0_{K_l(F)}(p)\right)$$

since x_l does not occur in $K_l(F)$ and hence the events $\boldsymbol{\rho}_p(x_l) = *$ and $\boldsymbol{\rho}_p(K_l(F)) \equiv 0$ are independent.■

Lemma 3.6. *If f is the function computed by a monotone read-once formula F, then*

$$A'_f(q) = \sum_{l=1}^{n} A_{K_l(F)}(q).$$

Proof. Let $\boldsymbol{\rho}^{q_1,\ldots,q_n}$ be the random restriction which independently assigns x_l to 1 with probability q_l and to 0 with probability $1 - q_l$ and let

$$\mathcal{A}_f(q_1, \ldots, q_n) \rightleftharpoons \mathbf{P}\left[\boldsymbol{\rho}^{q_1,\ldots,q_n}(f) \equiv 1\right].$$

Since $A_f(q) = \mathcal{A}_f(q, \ldots, q)$, we have

$$A'_f(q) = \sum_{l=1}^{n} \left.\frac{\partial \mathcal{A}_f}{\partial q_l}\right|_{(q,\ldots,q)}.$$

So we only have to show that for each fixed leaf l,

$$\left.\frac{\partial \mathcal{A}_f}{\partial q_l}\right|_{(q,\ldots,q)} = A_{K_l(F)}(q). \tag{8}$$

To do this, denote by $\boldsymbol{\rho}$ the restriction which assigns x_l to $*$ and independently assigns all remaining variables to 0 with probability $1-q$ and to 1 with probability q. By Lemma 3.4,

$$\mathbf{P}\left[\boldsymbol{\rho} \text{ does not kill } x_l\right] = \mathbf{P}\left[\boldsymbol{\rho}(K_l(F)) \equiv 1\right] = A_{K_l(F)}(q). \tag{9}$$

On the other hand,

$$\mathcal{A}_f\left(q,\ldots,q_l,\ldots,q\right) = \mathbf{P}\left[\boldsymbol{\rho}^{q_l}\boldsymbol{\rho}(f)=1\right] = \mathbf{P}\left[\boldsymbol{\rho}(f)\equiv 1\right] + q_l \cdot \mathbf{P}\left[\boldsymbol{\rho} \text{ does not kill } x_l\right]$$

since f is monotone. Hence

$$\left.\frac{\partial \mathcal{A}_f}{\partial q_l}\right|_{(q,\ldots,q)} = \mathbf{P}\left[\boldsymbol{\rho} \text{ does not kill } x_l\right]. \tag{10}$$

(9) and (10) together imply the desired equality (8).■

Our last lemma is a tool to handle unbalanced formulae.

Lemma 3.7. *Let F be a formula of size at least s where $s \geq 1$ is an integer. Then there exists a subformula H of F such that $L(H) \geq L(F) - s + 1$ and for the representation*

$$F \equiv G(x_1,\ldots,x_r,H) \tag{11}$$

with

$$L(F) = L(G) + L(H) - 1 \tag{12}$$

we have either $L(G) > s/2$ or $H \equiv H_1 \circ H_2$ where $L(H_1), L(H_2) \geq s/2$.

Proof. F contains subformulae H of size at least $L(F) - s + 1$; for example $H \equiv F$. Let us choose a minimal subformula H with this property and consider the corresponding representation (11). If $L(G) > s/2$ we are done. Otherwise $L(H) \geq L(F) - s/2 + 1$ by (12) which along with assumption $L(F) \geq s$ implies that H can not be a single variable. Hence $H \equiv H_1 \circ H_2$, where $\circ \in \{\wedge, \vee\}$. If, say, $L(H_1) < s/2$ then we would have $L(H_2) \geq L(F) - s + 1$ which would contradict the choice of H. Hence $L(H_1) \geq s/2$ and similarly we prove $L(H_2) \geq s/2$.■

4. Proof of Theorem 2.1

We will denote $L(f)$ by n throughout the section. We divide the analysis according to whether $pn^{1/\alpha} \leq \frac{1}{2\log n}$ or not.

4.1. $p^\alpha n$ is small

In the case $pn^{1/\alpha} \leq \frac{1}{2\log n}$ the claim of Theorem 2.1 is clearly implied by the following statement:

Theorem 4.1. *For any read-once function $f(x_1, \ldots, x_n)$ and any p such that*

$$pn^{1/\alpha} \leq \frac{1}{2\log n} \tag{13}$$

we have the bound $E_f(p) \leq O\left(pn^{1/\alpha}\right)$.

Proof. Let F be the read-once formula computing f. Replacing all occurrences of $\bar{x}_l$ in F by x_l, we may assume w.l.o.g. that F is monotone. Let l be a leaf. Since $K_{\wedge,l}(F)$ and $K_{\vee,l}(F)$ have disjoint sets of variables,

$$1 - C^0_{K_l(F)}(p) = \left(1 - C^0_{K_{\wedge,l}(F)}(p)\right) \cdot \left(1 - C^1_{K_{\vee,l}(F)}(p)\right)$$

and similarly

$$A_{K_l(F)}(1/2) = A_{K_{\wedge,l}(F)}(1/2) \cdot \left(1 - A_{K_{\vee,l}(F)}(1/2)\right). \tag{14}$$

Since $K_{\wedge,l}(F)$ and $K_{\vee,l}(F)$ are monotone we may apply Lemma 3.3 to conclude

$$1 - C^0_{K_l(F)}(p) = A_{K_{\wedge,l}(F)}\left(\frac{1+p}{2}\right) \cdot \left(1 - A_{K_{\vee,l}(F)}\left(\frac{1-p}{2}\right)\right).$$

Substituting this to Lemma 3.5 we get

$$E_f(p) = p \cdot \sum_{l=1}^{n} A_{K_{\wedge,l}(F)}\left(\frac{1+p}{2}\right) \cdot \left(1 - A_{K_{\vee,l}(F)}\left(\frac{1-p}{2}\right)\right) \tag{15}$$

and similarly from (14) and Lemma 3.6 we have

$$A'_f(1/2) = \sum_{l=1}^{n} A_{K_{\wedge,l}(F)}(1/2) \cdot \left(1 - A_{K_{\vee,l}(F)}(1/2)\right). \tag{16}$$

Apply now Lemma 3.2 with $f \equiv K_{\wedge,l}(F)$ and $x = n^{-1}$ (note that (3) follows from (13)). We derive $A_{K_{\wedge,l}(F)}\left(\frac{1+p}{2}\right) \leq O\left(A_{K_{\wedge,l}(F)}(1/2) + n^{-1}\right)$ and, by dual arguments, $1 - A_{K_{\vee,l}(F)}\left(\frac{1-p}{2}\right) \leq O\left(\left[1 - A_{K_{\vee,l}(F)}(1/2)\right] + n^{-1}\right)$. Substituting these two bounds to (15) we have

$$E_f(p) \leq p \cdot O\left(\sum_{l=1}^{n} A_{K_{\wedge,l}(F)}(1/2) \cdot \left(1 - A_{K_{\vee,l}(F)}(1/2)\right) + 1\right) = \text{ (by (16))}$$

$$p \cdot O\left(A'_f(1/2) + 1\right) \leq \text{ (by Proposition 3.1) } O\left(pn^{1/\alpha}\right). \blacksquare$$

4.2. $p^\alpha n$ is large

Without loss of generality, we can assume that $0 < p < 10^{-2}$. Let $s = s_p$ be the maximum integer n satisfying (13). Clearly, $s \geq 12$ and (13) holds for all $n \leq s$. By Theorem 4.1 we have

$$E_f(p) \leq \frac{C}{\log s} \tag{17}$$

for an arbitrary read-once $f(x_1, \ldots, x_n)$ with $n \leq s$ where C is an absolute constant. Since $s = \theta((p \log(1/p))^{-\alpha})$, in order to complete the proof of Theorem 2.1 it suffices to establish the following bound:

Lemma 4.2. *Let $f(x_1, \ldots, x_n)$ be a read-once function, $n \geq s/4$. Then*

$$E_f(p) \leq \frac{C(12n - 2s)}{s \log s}.$$

Proof. Induction on n.

Base $s/4 \leq n \leq s$ follows from (17).

Inductive step. Let $n > s$ and F be the read-once formula computing f. Replace in Lemma 3.7 s by $s/2$ and apply it in this form to the formula F. We will get a representation of the form (11). Let g, h be the functions computed by the formulas G, H respectively. Renaming variables we may assume w.l.o.g. that g depends on the variables $x_1, \ldots, x_r, y$ whereas h depends on $x_{r+1}, \ldots, x_n$; $r + 1 \leq s/2$. Let g^ϵ be the function obtained from g by setting y to ϵ ($\epsilon \in \{0, 1\}$). The crucial observation is that

$$E_f(p) \leq E_{g^0}(p) + E_{g^1}(p) + E_h(p). \tag{18}$$

We prove (18) locally i.e. we show that for any *fixed* restriction $\rho \in \{0, 1, *\}^n$,

$$L(\rho(f)) \leq L(\rho(g^0)) + L(\rho(g^1)) + L(\rho(h)). \tag{19}$$

Extend ρ by setting $\rho(y) \rightleftharpoons *$. If ρ kills y in G or ρ reduces h to a constant, $\rho(f)$ coincides with $\rho(g^\epsilon)$ for some $\epsilon \in \{0,1\}$ and (19) becomes obvious. Otherwise $L(\rho(f)) = L(\rho(g)) + L(\rho(h)) - 1$. Moreover we can set y to a constant ϵ so that this does not produce any extra killings in $\rho(g)$. Which means $L(\rho(g^\epsilon)) = L(\rho(g)) - 1$ and again implies (19). (19) and hence (18) are proved.

Now, since $r+1 \le s/2$, we may apply (17) and derive from (18) that $E_f(p) \le \frac{2C}{\log s} + E_h(p)$. Let us recall from Lemma 3.7 that additionally we have either $r+1 > s/4$ or $H \equiv H_1 \circ H_2$ where $L(H_1), L(H_2) \ge s/4$. In the first case we apply the inductive assumption to H (note that $L(H) \ge L(F) - s/2 + 1 \ge s/2$) to conclude

$$E_f(p) \le \frac{2C}{\log s} + \frac{C(12(n-r)-2s)}{s \log s} \le \frac{2C}{\log s} + \frac{C(12n-5s+12)}{s \log s} \le \frac{C(12n-2s)}{s \log s}.$$

In the second case the inductive assumption can be applied to both H_1 and H_2 and we have

$$E_f(p) \le \frac{2C}{\log s} + E_{h_1}(p) + E_{h_2}(p) \le$$
$$\frac{2C}{\log s} + \frac{C(12 \cdot L(H_1) - 2s)}{s \log s} + \frac{C(12 \cdot L(H_2) - 2s)}{s \log s} \le \frac{C(12n-2s)}{s \log s}.$$

In either case the inductive step is completed and this also completes the proofs of Lemma 4.2 and Theorem 2.1.■

5. Proof of Theorem 2.3

The proof of Theorem 4.1 is not "analytic" in the sense that it does not provide us with an analytic bound on $E_f(p)$ provable by induction on $L(f)$. Instead it requires a rather non-trivial analysis essentially involving the tree structure of the read-once formula computing f. We do not know whether this proof can be smoothed in the general (unbalanced) case. In this section we show how to do this for balanced formulae by proving Theorem 2.3. As a reward, we get rid of the factor $\left(\log \frac{1}{p}\right)^{\alpha-1}$ (see Corollary 2.4).

Denote by C the constant assumed in the term $O(x)$ in (4). Let

$$a \rightleftharpoons \log\left[5(C+1)\right], \; D \rightleftharpoons \alpha \log\left(\frac{1}{p}\right) - a - 6.$$

Lemma 5.1. *For a read-once function $f(x_1, \ldots, x_n)$ such that $D(f) \leq D$ we have the bound $E_f(p) \leq O(1)$.*

Theorem 2.3 follows easily from Lemma 5.1. In fact, any read-once function f with $D(f) \geq D$ can be decomposed as $f = g(h_1, \ldots, h_l)$ where $l \leq O\left(p^\alpha 2^{D(f)}\right)$ and h_i are read-once functions of depth at most D. This gives us immediately $E_f(p) \leq \sum_{i=1}^{l} E_{h_i}(p) \leq O(l) \leq O\left(p^\alpha 2^{D(f)}\right)$, and hence Theorem 2.3.

Proof of Lemma 5.1. We can assume $D \geq 0$. For $d \leq D$ set

$$\delta_d \rightleftharpoons \prod_{k=0}^{d-1} \left(1 + p2^{k/\alpha}(D + a - k)\right)$$

and

$$\Delta_d \rightleftharpoons Cp2^{d+d/\alpha-D}.$$

We are going to prove that for any $d \leq D$ and any monotone read-once function f of depth $D(f) \leq d$,

$$E_f(p) \leq \delta_d \left(pA'_f(1/2) + \Delta_d\right). \tag{20}$$

With the help of $1 + x \leq e^x$, it is easy to see that $\delta_D \leq O(1)$ and $\Delta_D \leq O(1)$. It is also easy to derive from Proposition 3.1 that $A'_f(1/2) \leq p^{-1}$. Hence (20) would suffice to finish the proof of Lemma 5.1.

We prove (20) by induction on d.

Base $d = 0$ is obvious since $\delta_0 = 1$, $E_f(p) = p$ and $A'_f(1/2) = 1$.

Inductive step. Let $D(f) = d + 1$, $d \leq D - 1$. Assume that $f \equiv g \wedge h$ where $D(g), D(h) \leq d$ and that (20) is already established for g and h.

The subfunction g contributes $L(\rho(g))$ to the size of $\rho(f)$ unless $\rho(h) \equiv 0$ and the same holds for the contribution of $L(\rho(h))$. This implies

$$\left.\begin{aligned} &E_f(p) = E_g(p) \cdot \left(1 - C_h^0(p)\right) + E_h(p) \cdot \left(1 - C_g^0(p)\right) = \text{ (by Lemma 3.3)} \\ &E_g(p) \cdot A_h\left(\frac{1+p}{2}\right) + E_h(p) \cdot A_g\left(\frac{1+p}{2}\right). \end{aligned}\right\} \tag{21}$$

Since $A_f(q) = A_g(q) \cdot A_h(q)$,

$$A'_f(1/2) = A'_g(1/2)A_h(1/2) + A_g(1/2)A'_h(1/2). \tag{22}$$

We are going to apply Lemma 3.2 with $x = x_d = 2^{d-D-a}$ to $A_g\left(\frac{1+p}{2}\right), A_h\left(\frac{1+p}{2}\right)$ in (21). For this we should first check (3). In our situation this inequality becomes $\alpha \log \frac{1}{p} - d - 6 \leq \frac{1}{2}p^{-1}2^{-d/\alpha}$ or $\alpha \log y - 6 \leq \frac{1}{2}y$ where $y \rightleftharpoons p^{-1}2^{-d/\alpha}$. This is easily checked by finding the maximum of the function $\alpha \log y - \frac{1}{2}y$.

Applying Lemma 3.2 allows us to continue the chain of inequalities (21) as follows:

$$
\begin{aligned}
&E_f(p) \leq E_g(p) \cdot (A_h(1/2) + Cx_d)\left(1 + p2^{d/\alpha}(D + a - d)\right) + \\
&E_h(p) \cdot (A_g(1/2) + Cx_d)\left(1 + p2^{d/\alpha}(D + a - d)\right) \leq \text{ (by inductive assumption)} \\
&\delta_d\left(1 + p2^{d/\alpha}(D + a - d)\right) \cdot \\
&\left[\left(pA'_g(1/2) + \Delta_d\right) \cdot (A_h(1/2) + Cx_d) + (pA'_h(1/2) + \Delta_d) \cdot (A_g(1/2) + Cx_d)\right] = \\
&\delta_{d+1}\left[\left(pA'_g(1/2) + \Delta_d\right) \cdot (A_h(1/2) + Cx_d) + (pA'_h(1/2) + \Delta_d) \cdot (A_g(1/2) + Cx_d)\right] \leq \\
&\text{(by (22))}\ \delta_{d+1}\left[pA'_f(1/2) + 2\Delta_d + Cx_d\left(2\Delta_d + pA'_g(1/2) + pA'_h(1/2)\right)\right] \leq \\
&\text{(by Proposition 3.1)}\ \delta_{d+1}\left[pA'_f(1/2) + 2\Delta_d + 2Cx_d\left(Cp2^{d/\alpha} + p2^{d/\alpha}\right)\right] \leq \\
&\delta_{d+1}\left[pA'_f(1/2) + \Delta_{d+1}\right].
\end{aligned}
$$

The case $f \equiv g \vee h$ can be treated similarly.

The inductive step is completed. This also completes the proofs of (20) and Lemma 5.1.■

6. Proof of Theorem 2.5

The example of Paterson and Zwick [9] shows that the bound of Theorem 4.1 is tight. Extending their argument we prove that also Theorem 2.3 is tight and hence Theorem 2.1 is tight up to a polylogarithmic factor. Quite fortunately, Lemma 3.3 allows us to use results of the computations already performed in [10, 2].

Proposition 6.1. *For each $p \leq 1$ there exists a monotone read-once function f_p of depth at most $\alpha \log \frac{1}{p} + O(1)$ such that*

$$A_{f_p}(1/2) \leq 1/10,\ A_{f_p}\left(\frac{1+p}{2}\right) \geq 9/10.$$

Proof. Combining [2, Theorem 1.2 (a)] (with $p := 1/2$ and $m := 2/p$) and [2, Theorem 1.2 (b)] (with $m := 4$) we get the desired function f_p of size $O\left(p^{-\alpha}\right)$. A close inspection of the constructions involved in the proof shows that moreover $D(f_p) \leq \log\left[O\left(p^{-\alpha}\right)\right] = \alpha \log \frac{1}{p} + O(1)$.■

Proof of Theorem 2.5. Suppose $2^{-d/\alpha} \leq p \leq 1$. Let d_p be the depth of the function f_p defined above. We may assume that $d \geq d_p$ since otherwise [9, Theorem 6.5] applies. Define f_d by a formula which is an AND-OR tree with a $\wedge$ closest to the inputs of depth $d - d_p$ and where input number l is replaced by a copy of f_p, called f_p^l. For $i = 1, 3, \ldots, d - d_p$ let r_i be the probability that an input to an $\wedge$-gate on level i is forced by $\boldsymbol{\rho}_p$ to 0. By Lemma 3.3 and the construction, $r_1 \leq 1/10$ and it is easy to see that $r_{i+2} \leq 4r_i^2$ which implies

$$r_i \leq \frac{1}{4} \cdot \left(\frac{2}{5}\right)^{2^{\frac{i-1}{2}}}.$$

In a similar way let s_i be the probability that an input to an $\vee$-gate on level i is forced to 1. By Lemma 3.3 and construction we have $s_2 \leq 1/100$ and $s_{i+2} \leq 4s_i^2$ which gives us

$$s_i \leq \frac{1}{4} \cdot \left(\frac{1}{25}\right)^{2^{\frac{i-2}{2}}}.$$

Similarly to Lemma 3.5 we have

$$E_{f_d}(p) = \sum_l E_{f_p^l}(p) \mathbf{P}\left[R_l\right]$$

where R_l is the event that f_d is not made independent of the leaf l by other fixings. Now by construction $E_{f_p^l}(p) \geq N_{f_p^l}(p) \geq 4/5$ and by the bounds on the s_i and r_i,

$$\mathbf{P}[R_l] \geq 1 - \sum_{\text{odd } i} r_i - \sum_{\text{even } i} s_i \geq \Omega(1).$$

The theorem now follows.

Note added in proof: There has been a number of recent papers in the related area. Dubiner and Zwick [3] established the bounds of Boppana for other functions than H. Using these bounds would decrease the exponent of the $\log \frac{1}{p}$-factor. In a different paper Dubiner and Zwick [4], using related methods to ours remove this factor totally. For the general case Håstad [7] has established that the shrinkage exponent is 2.

References

[1] M. Ajtai. Σ_1^1-formulae on finite structures. *Annals of Pure and Applied Logic*, 24:1–48, May 1983.

[2] R. B. Boppana. Amplification of probabilistic Boolean formulas. In S. Micali, editor, *Advances in Computer Research, vol. 5: Randomness and Computation*, pages 27–46. JAI Press, Greenwich, CT, 1989.

[3] M. Dubiner and U. Zwick Amplification and Percolation. In *Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science*, pages 258–267, 1992.

[4] M. Dubiner and U. Zwick How do read-once formulae shrink?. Manuscript, 1993.

[5] M. Furst, J. B. Saxe, and M. Sipser. Parity, circuits and the polynomial time hierarchy. *Math. Syst. Theory*, 17:13–27, 1984.

[6] J. Håstad. *Computational limitations on Small Depth Circuits.* PhD thesis, Massachusetts Institute of Technology, 1986.

[7] J. Håstad. The shrinkage exponent is 2. In *Proceedings of the 34th IEEE Symposium on Foundations of Computer Science*, pages 114–123, 1993.

[8] N. Nisan and R. Impagliazzo. The effect of random restrictions on formulae size. *Random Structures and Algorithms*, Vol 4, No 2. 1993 pp 121–134.

[9] M. S. Paterson and U. Zwick. Shrinkage of de Morgan formulae under restriction. *Random Structures and Algorithms*, Vol 4, No 2. 1993 pp 135–150.

[10] L. G. Valiant. Short monotone formulae for the majority function. *Journal of Algorithms*, 5:363–366, 1984.

[11] A. Yao. Separating the polynomial-time hierarchy by oracles. In *Proceedings of the 26th IEEE FOCS*, pages 1–10, 1985.

[12] А. Е. Андреев. О методе получения более чем квадратичных нижних оценок для сложности π-схем. *Вестник МГУ, сер. матем и механ.*, т. 42, в. 1, 1987, стр. 70-73. (A.E. Andreev, On a method for obtaining more than quadratic effective lower bounds for the complexity of π-schemes, *Moscow Univ. Math. Bull* 42(1)(1987), 63-66).

[13] Б. А. Субботовская. О реализации линейных функций формулами в базисе $\&, \vee, -$. *ДАН СССР*, т. 136, в. 3, 1961, стр. 553-555. (B.A.Subbotovskaya, Realizations of linear functions by formulas using $+, *, -$, *Soviet Mathematics Doklady* 2(1961), 110-112).

[14] В. М. Храпченко. О сложности реализации линейной функции в классе П-схем. *Матем. заметки*, т. 9, в. 1, 1971, стр. 35-40. (V.M. Khrapchenko, Complexity of the realization of a linear function in the class of π-circuits, *Math. Notes Acad. Sciences USSR* 9(1971), 21-23).

[15] В. М. Храпченко. Об одном методе получения нижних оценок сложности П-схем. *Матем. заметки*, т. 10, в. 1, 1971, стр. 83-92. (V.M. Khrapchenko, A method of determining lower bounds for the complexity of П-schemes, *Math. Notes Acad. Sciences USSR* 10(1971), 474-479).