Shrinkage of de Morgan formulae under restriction *

Michael S. Paterson
Department of Computer Science
University of Warwick
Coventry CV4 7AL, England

Uri Zwick
Department of Computer Science
Tel Aviv University
Tel Aviv 69978, Israel

Abstract

It is shown that a random restriction leaving only a fraction ε of the input variables unassigned reduces the expected de Morgan formula size of the induced function by a factor of $O(\varepsilon^{\frac{5-\sqrt{3}}{2}}) = O(\varepsilon^{1.63})$. (A de Morgan, or unate, formula is a formula over the basis $\{\wedge, \vee, \neg\}$.)

This improves a long-standing result of $O(\varepsilon^{1.5})$ by Subbotovskaya and a recent improvement to $O(\varepsilon^{\frac{21-\sqrt{73}}{8}}) = O(\varepsilon^{1.55})$ by Nisan and Impagliazzo.

The new exponent yields an increased lower bound of $n^{\frac{7-\sqrt{3}}{2}-o(1)} = \Omega(n^{2.63})$ for the de Morgan formula size of a function in P defined by Andreev. This is the largest formula size lower bound known, even for functions in NP.

1 Introduction

In [11], Subbotovskaya introduced the random restriction method and used it to derive an $\Omega(n^{1.5})$ lower bound on the de Morgan formula size of the parity function. Khrapchenko [5],[6] (see also [3],[13],[14]) then used a different method to improve the lower bound for the parity function to a tight $\Omega(n^2)$ bound. Quite a few years then passed before Andreev [1] observed that the random restriction method used by Subbotovskaya could be used to derive an $\Omega(n^{2.5-o(1)})$ lower bound on the de Morgan formula size of a rather simple function in P. This function is easily shown to have formula size $o(n^3)$.

*This research was partially supported by the ESPRIT BRA Programme of the EC under contracts # 3075 (ALCOM) and # 7141 (ALCOM II).

The following basic observation is used in Subbotovskaya's proof that a random restriction leaving only a fraction ε of the variables of a function unassigned reduces the expected formula size of the induced function by a factor of at least $\varepsilon^{1.5}$. If a leaf in a de Morgan formula is set once to 0 and once to 1, then under one of these assignments the size of the formula decreases by at least one and under the other by at least two, giving a total of at least 3 and an average of at least 1.5.

Any formula in which all the leaves are arranged in pairs, or *twigs* as we shall call them, and in which the gates in the lowest two levels alternate, shows that the preceding argument, in its present form, cannot be strengthened.

Nisan and Impagliazzo [8] found an ingenious way to improve Subbotovskaya's exponent from 1.5 to about 1.55. They argue that, even if some formula is in a form in which only small savings can be expected, subsequent reductions are likely to transform its structure into one in which greater savings are expected. In order to capture this property they assign to every formula a *weight* that takes into account not only its size but also the likelihood of further size reductions in subsequent random assignments. A formula with a greater potential for savings will get a smaller weight, reflecting the fact that if this formula is obtained progress has been made, even if only a small reduction of size was obtained.

Nisan and Impagliazzo's method leads to technical complications which hinder further development. In the sequel we follow the general approach in [8], but with a simpler and more natural weight function. Our weight function is simply the formula size plus a supplement for each twig in the formula. Another significant difference between our work and [8] is that we use a top-down approach. The resulting simplifications allow us to improve Nisan and Impagliazzo's results while retaining a more lucid proof structure which may pave the way for further improvements.

We shall define the *shrinkage exponent* Γ to be largest possible exponent for ε in the expected formula size reduction after a random restriction of the kind considered above. A precise definition will be given in the next section. Subbotovskaya showed that $\Gamma \geq 1.5$ and Nisan and Impagliazzo increased this bound to 1.55. The parity function is an example showing that $\Gamma \leq 2$.

Our main result is Theorem 2.5, which improves an analogous result in [8]. The shrinkage exponent bound and the improved lower bound for Andreev's function then follow using arguments similar to those in [8], Andreev's paper [1] or Dunne's book [3]. Our lower bound for this function is the largest known bound for the unate formula size of explicitly given functions.

Finally we consider shrinkage of *read-once* formulae, i.e., those in which each variable occurs only once. Clearly any read-once formula is of minimal size. Intuitively, one might expect that read-once formulae should be particularly resistant to shrinkage. We show that this is unlikely by analysing a set of formulae which we conjecture to have optimal shrink-resistance among read-once formulae.

A preliminary version of this paper appeared in [9].

2 Shrinkage results

Definitions. A *de Morgan formula* is a binary tree in which each leaf is labelled by a literal from the set $\{x_1, \ldots, x_n, \overline{x}_1, \ldots, \overline{x}_n\}$ and each internal node functions as either an AND ($\wedge$) or an OR ($\vee$) gate. The *size* of a formula T is defined to be the number of leaves in it, and it is denoted by $L(T)$. A *twig* is a subtree with just two leaves. The number of twigs contained in a formula T is denoted by $\tau(T)$. The *weight* of a formula T is defined to be $w(T) = L(T) + \alpha \cdot \tau(T)$ where α is a parameter. We will take $\alpha = \sqrt{3} - 1$.

Every formula T computes in a natural way a function from $\{0,1\}^n$ to $\{0,1\}$, denoted by $f(T)$. For a function f we denote by $L(f)$ the minimal size of a formula computing f and by $w(f)$ the minimal weight of a formula computing f. If T is a literal, $L(T) = w(T) = 1$, and if T is a constant, $L(T) = w(T) = 0$. If $L(f) \le 1$, we say f is a *trivial* function.

The complexity measures L and w differ by at most a constant factor. It is easy to see that $w(f)/(1 + \frac{\alpha}{2}) \le L(f) \le w(f)$, for all f.

For any function f, we denote by $f_{x=0}, f_{x=1}$ the functions obtained from f by assigning to x the values 0,1 respectively. We define $\sigma_{x=c}(f) = w(f) - w(f_{x=c})$ for $c = 0, 1$. Also $\sigma_x(f) = \sigma_{x=0}(f) + \sigma_{x=1}(f)$ and $\sigma(f) = \sum_{x \in f} \sigma_x(f)$, where the summation is over all the variables on which f depends. These quantities measure the *savings* resulting from basic restrictions.

For any function $f = f(x_1, \ldots, x_n)$, let $\hat{f}$ denote the random function $f_{x_i=c}$ obtained by choosing $i \in \{1, 2, \ldots, n\}$ and $c \in \{0, 1\}$ independently and with equal probabilities.

The following theorem lies at the heart of our improved bounds. The proof is given in Section 5.

Theorem 2.1 *For any non-trivial function f,*

$$\frac{\sigma(f)}{w(f)} \ge 2\gamma \quad \textit{where} \quad \gamma = \tfrac{3+2\alpha}{2+\alpha} = \tfrac{5-\sqrt{3}}{2} \simeq 1.63\,.$$

The following corollary immediately follows from Theorem 2.1. The expectation operator is denoted by E.

Corollary 2.2 *If f is non-trivial then*

$$E[w(\hat{f})] \leq \left(1 - \frac{\gamma}{n}\right) w(f).$$

Proof :

$$\begin{aligned} E[w(\hat{f})] &= \frac{1}{2n}\sum_{i=1}^{n}(w(f_{x_i=0}) + w(f_{x_i=1})) \\ &= \frac{1}{2n}\sum_{i=1}^{n}(2w(f) - \sigma_{x_i}(f)) \;=\; w(f) - \frac{\sigma(f)}{2n} \;\leq\; \left(1 - \frac{\gamma}{n}\right) w(f). \end{aligned}$$ □

For any $\varepsilon > 0$, define R^ε to be the random assignment which gives values independently to each variable x_i, for $i \geq 1$, as follows: $(x_i \leftarrow 0)$ with probability $\frac{1-\varepsilon}{2}$, $(x_i \leftarrow 1)$ with probability $\frac{1-\varepsilon}{2}$, and x_i remains unchanged with probability ε. We define f^ε to be the result of applying R^ε to the function f.

There is a similar random assignment $\hat{R}_{n,m}$, where $1 \leq m < n$, which selects uniformly at random a subset of $\{x_1, \ldots, x_n\}$ of size m, and assigns 0 or 1 with equal independent probabilities to each variable not in the chosen subset. The following lemma shows that $R^{m/n}$ and $\hat{R}_{n,m}$ are interchangeable for our purposes.

Lemma 2.3 *Let $\varepsilon = m/n$, where $1 \leq m < n$ and m may depend on n, and let $\gamma \geq 1$ be a constant. If m_ε is the random variable giving the number of variables not set by R^ε, then*

$$\begin{aligned} E[m_\varepsilon^\gamma] &= \Theta(m^\gamma) \quad \textit{as } n \to \infty, \\ &\sim m^\gamma \quad \textit{if } m \to \infty \textit{ as } n \to \infty. \end{aligned}$$

Proof : The result follows from simple estimations of the tails of the distribution of m_ε away from its mean of m. □

Definition 2.4 We define *the shrinkage exponent* Γ to be the least upper bound for constants γ such that there exist constants c, d for which, for all functions $f = f(x_1, \ldots, x_n)$ and for all $\varepsilon > 0$,

$$E[L(f^\varepsilon)] \leq c \cdot \varepsilon^\gamma L(f) + d.$$

This definition is relative to the unate basis, $\{\wedge, \vee, \neg\}$. For the full basis B_2, the parity function demonstrates that the corresponding shrinkage exponent is 1.

Our main result is the following lower bound for Γ which improves the analogous result of $\Gamma \geq \frac{21-\sqrt{73}}{8} \simeq 1.55$ in [8].

Theorem 2.5 $\Gamma \geq \frac{5-\sqrt{3}}{2} \simeq 1.63$.

Proof : Corollary 2.2 has the condition that f is non-trivial. If $w(f) = 1$, we have only $E[w(\hat{f})] \leq \left(1 - \frac{1}{n}\right) w(f)$, since the single variable from $x_1, \ldots, x_n$ upon which f depends is selected only with probability $1/n$. To cope with the cases where $w(f) \leq 1$, we use the weaker inequality

$$E[w(\hat{f}) - d] \leq \left(1 - \frac{\gamma}{n}\right)(w(f) - d)$$

where $d = 1 - 1/\gamma$, which is valid in all cases.

For any $f(x_1, \ldots, x_n)$, we define the sequence of random functions $f_n, \ldots, f_m$, by $f_n = f$ and $f_k = \hat{f}_{k+1}$ for $n > k \geq m$. Note that the random restrictions at each stage are chosen independently of each other. For each k, f_k is considered as a function of the k remaining variables, irrespective of whether it actually depends on all of them. From the inequality given above, we have

$$E[w(f_{k-1}) - d] \leq \left(1 - \frac{\gamma}{k}\right) E[w(f_k) - d]$$

for $n \geq k > m$, and so

$$E[w(\hat{f_m}) - d] \leq (1 - \tfrac{\gamma}{m+1}) \ldots \left(1 - \tfrac{\gamma}{n}\right)(w(f) - d).$$

Since

$$\left(1 - \frac{\gamma}{m+1}\right) \ldots \left(1 - \frac{\gamma}{n}\right) \leq \left(\frac{m}{n}\right)^{\gamma},$$

we have

$$E[w(\hat{f_m})] < \left(\frac{m}{n}\right)^{\gamma} w(f) + d.$$

Since $w(f)$ and $L(f)$ agree to within a constant factor, there exist constants c', d' such that

$$E[L(\hat{f_m})] \leq c' \left(\frac{m}{n}\right)^{\gamma} L(f) + d'.$$

Lemma 2.3 ensures the same form of inequality for f^{ε} where $\varepsilon = m/n$. This completes the proof of the theorem. $\square$

3 The formula size lower bound

The principal application of our theorem on shrinkage is to provide lower bounds on formula size. In [1] (see also [3]), Andreev derived an $\Omega(n^{2.5-o(1)})$ lower bound on the formula size of a particular function. This was improved by Nisan and Impagliazzo in [8] to $\Omega(n^{2.55})$, and we can now increase this further to $\Omega(n^{2.63})$.

Let the selection function *sel* be defined by

$$\mathrm{sel}^{(r)}(x_0,\ldots,x_{r-1};y_0,\ldots,y_{2^r-1}) = y_{x^*}$$

where $x^* = \sum_{i=0}^{r-1} x_i 2^i$, i.e., the vector $\vec{x}$ is interpreted as a binary number and the value of the function is the bit in $\vec{y}$ indexed by this number. In a sense, *sel* is a 'universal function' since any function in B_r may be obtained as a restriction of $sel^{(r)}$ by setting $\vec{y}$ to the truth-table of that function. The parity function $\mathcal{P}_n$ is defined by

$$\mathcal{P}_n(x_1,\ldots,x_n) = x_1 \oplus \ldots \oplus x_n.$$

Following Andreev [1], we define the function $\mathcal{A}$ by

$$\mathcal{A}_{r,s}(\vec{x}_1,\ldots,\vec{x}_r,\vec{y}) = \mathrm{sel}(\mathcal{P}_s(\vec{x}_1),\ldots,\mathcal{P}_s(\vec{x}_r);\vec{y})$$

where each $\vec{x}_i$ is an s-tuple and $\vec{y}$ is a 2^r-tuple. We have introduced the parameters r, s in order to demonstrate the optimisation involved. For any $n > 0$, let $\mathcal{A}_n(z_1,\ldots,z_n) = \mathcal{A}_{r,s}(z_1,\ldots,z_{rs+2^r})$, where $r = \lfloor \lg n \rfloor - 1$ and $s = \lfloor 2^r/r \rfloor$. Here lg denotes the logarithm to base 2.

It is easy to construct formulae for $\mathcal{A}_n$ of size $O(s^2 2^r) = O(n^3/(\lg n)^2)$. Our principal application provides a corresponding lower bound.

Theorem 3.1 *The formula size of the function $\mathcal{A}_n$ satisfies*

$$L(\mathcal{A}_n) = \Omega(n^{\frac{7-\sqrt{3}}{2}-o(1)}) = \Omega(n^{2.63}).$$

Proof : Classical results of Riordan and Shannon [10] and Lupanov [7] establish the existence of a sequence of functions $f_1, f_2, \ldots$ such that $f_k \in B_k$ and $L(f_k) \sim 2^k/\lg k$. Since f_r is some restriction of $sel^{(r)}$, we have

$$L(\mathrm{sel}^{(r)}) \geq L(f_r) \sim 2^r/\lg r.$$

Similarly, if we define

$$g_r(\vec{x}_1,\ldots,\vec{x}_r) = f_r(\mathcal{P}_s(\vec{x}_1),\ldots,\mathcal{P}_s(\vec{x}_r)),$$

we have

$$L(\mathcal{A}_{r,s}) \geq L(g_r).$$

If $\varepsilon = (\lg r)/s$ then, with high probability, the random assignment R^ε leaves at least one variable in each $\vec{x}_i$ unset. Hence with high probability,

$$L(g_r^\varepsilon) \geq L(f_r) \sim 2^r/\lg r,$$

but by Theorem 2.1, we also have

$$E[L(g_r^\varepsilon)] = O(\varepsilon^\gamma)L(g_r).$$

The combination of these results yields

$$L(\mathcal{A}_{r,s}) = \Omega(\varepsilon^{-\gamma})2^r/\lg r = \Omega\left(\frac{s^\gamma 2^r}{(\lg r)^{\gamma+1}}\right).$$

Hence

$$L(\mathcal{A}_n) = \Omega\left(\frac{n^{\gamma+1}}{(\lg n)^\gamma(\lg\lg n)^{\gamma+1}}\right) = \Omega(n^{\gamma+1-o(1)}).$$

□

4 Preliminary results

For the proof of Theorem 2.1 it is useful to extend the definition of a formula by allowing leaves to be labelled by constants. The size of a formula containing constant leaves is defined to be the number of non-constant leaves in it. The weight of such a formula is obtained by adding the constant α for every twig, even if one or both of its leaves are constant.

We first show that this extension does not change the weights that were assigned to functions.

Lemma 4.1 *For any formula T computing a non-constant function, there exists a formula T' with no constant leaves that computes the same function and for which $w(T') \leq w(T)$.*

Proof : Let ℓ be a leaf in T labelled by a constant c. Let m be the parent of ℓ and let S be the other subtree of m. If the function computed at m is constant, then remove ℓ and S from T making m a leaf labelled by the appropriate constant. This action does not increase the number of twigs (at most one new twig, involving m, is formed, but at least one twig is removed) and clearly does not increase the number of non-constant leaves in T. The weight therefore can only decrease and the new tree still computes the same function.

If the function computed at m is not constant, then it is equal to the function computed by S. We can therefore replace the subtree of ℓ, m and S in T by S alone. Again the function computed is unchanged and its weight does not increase.

By repeating the above process we eventually reach a formula with the required properties. □

Corollary 4.2 *Let T be a formula and let f' be a function computed by T when k of its non-constant leaves are assigned constant values. Then, $w(f') \leq w(T) - k$.* $\square$

This corollary will be used in the next section to show that if f depends on $x_1, \ldots, x_k$ and $f' = f_{x_1=c_1,\ldots,x_k=c_k}$ then $w(f') \leq w(f) - k$.

We note that Lemma 4.1 and Corollary 4.2 are not valid for the weight function of Nisan and Impagliazzo. These properties make our weight function easier to work with.

If $f = f_1 \vee f_2$ (or $f = f_1 \wedge f_2$) where $w(f) = w(f_1) + w(f_2)$ and $f_{1|x=c}$ and $f_{2|x=c}$ are not both literals then $\sigma_{x=c}(f) \geq \sigma_{x=c}(f_1) + \sigma_{x=c}(f_2)$, i.e., the savings behave super-additively. The hard cases in the proof of Theorem 2.1 will be those in which both f_1 and f_2 reduce to literals under certain assignments. We therefore need some understanding of the circumstances under which this can happen.

The *solo structure* of a function f is the relation '$\Rightarrow$' on literals defined by $x \Rightarrow y$ if $f_{x=1} = y$, or in other words if $f = xy \vee \overline{x}g$ for some function g. For example, if $f_{x=0} = y$ then this is expressed by writing $\overline{x} \Rightarrow y$. The next lemma classifies all possible solo structures.

Lemma 4.3 *For any non-trivial function f, the solo structure takes one of the following forms:*

1. *the empty relation,*
2. $\{x \Rightarrow y, y \Rightarrow x, \overline{x} \Rightarrow \overline{y}, \overline{y} \Rightarrow \overline{x}\}$ *and* $f = xy \vee \overline{x}\,\overline{y}$,
3. $\{x \Rightarrow y, \overline{x} \Rightarrow z\}$ *and* $f = xy \vee \overline{x}z$,
4. $\{x \Rightarrow y, y \Rightarrow x\}$ *and* $f = xy \vee \overline{x}\,\overline{y}g$,
5. $\{x \Rightarrow y, \overline{y} \Rightarrow \overline{x}\}$ *and* $f = xy \vee \overline{x}\,\overline{y} \vee \overline{x}yg$,
6. $\{x_1 \Rightarrow y, \ldots, x_k \Rightarrow y\}$ *for some* $k \geq 1$ *and* $f = (x_1 \vee \ldots \vee x_k)y \vee \overline{x}_1 \ldots \overline{x}_k g$.

In cases (4) and (5) the function g does not depend on x or y, in case (4) $g \neq 1$, in case (5) $g \neq 0$, and in case (6) the function g does not depend on $x_1, \ldots, x_k$ and if $k = 1$ then $g \neq 0$.

Proof : Let $x \Rightarrow y$ be in the solo structure of f. If $x \Rightarrow y$ is the only pair in the solo structure of f then f is of the form (6) with $k = 1$. Suppose therefore that $x' \Rightarrow y'$ is also in the solo structure of f.

If $x' = x$ then $y = y'$ and therefore $x' \Rightarrow y'$ is not a new pair. If $x' = \overline{x}$ then either $y' = \overline{y}$ in which case f is of the form (2), or $y' = z$ in which case f is of the form (3). The case $\{x \Rightarrow y, \overline{x} \Rightarrow y\}$ is impossible since then $f = y$ and f is trivial.

If $x' = y$ then $f = xy \vee \overline{x}g = yy' \vee \overline{y}h$ and $f_{x=1,y=1} = y'_{x=1} = 1$ and therefore $y' = x$. It follows that $f = xy \vee \overline{x}\,\overline{y}g$ and therefore f is either of the form (2) if $g = 1$ or of the form (4) otherwise.

If $x' = \overline{y}$ then similarly we get that $y' = \overline{x}$ and f is either of the form (2) or (5).

If $x' \neq x, \overline{x}, y, \overline{y}$ then by assigning $x \leftarrow 1, x' \leftarrow 1$ simultaneously we get that $y = y'$. The function is therefore of the form (6) with $k \geq 2$. □

As examples, note that xy is of the form (4) with $g = 0$, $x \oplus y = x\overline{y} \vee \overline{x}y = x\overline{y} \vee \overline{x}\,\overline{\overline{y}}$ is of the form (2), and $x \vee y = \overline{x}y \vee \overline{\overline{x}}\,\overline{y} \vee \overline{\overline{x}}y$ is of the form (5) with $g = 1$.

5 Proof of main technical result

For convenience we restate the result here.

Theorem 2.1 *For any non-trivial function f,*

$$\frac{\sigma(f)}{w(f)} \geq 2\gamma \quad \textit{where} \quad \gamma = \tfrac{3+2\alpha}{2+\alpha} = \tfrac{5-\sqrt{3}}{2} \simeq 1.63\,.$$

Proof : The proof is by induction on the weight of f. As the basis of the induction we verify the theorem directly for all non-trivial functions with size at most 4. These are exactly the non-trivial functions with weight at most $4 + 2\alpha \simeq 5.46$, or equivalently with weight less than $5 + \alpha \simeq 5.73$. Details are given in Table 5.1. These cases motivate the choice of values for α and γ such that $2\gamma = \frac{6+4\alpha}{2+\alpha} = \frac{10+3\alpha}{3+\alpha}$.

Suppose now that $L(f) > 4$ (therefore $w(f) > 5.7$) and that the theorem holds for all non-trivial functions with less weight than f. By definition, there exist two functions f_1, f_2 such that $f = f_1 \vee f_2$ or $f = f_1 \wedge f_2$, and $w(f) = w(f_1) + w(f_2)$. Without loss of generality we will assume that $f = f_1 \vee f_2$, the other case being dual.

If one of f_1 or f_2 is trivial then the required inequality for f follows from the induction hypothesis and from Lemma 5.1 below.

Lemma 5.1 *If $f = x \vee g$ for some g, then $\sigma(f)/w(f) \geq 2\gamma$.*

Proof : Let k be the number of different assignments under which g reduces to a literal. Note that since $w(f) = 1 + w(g)$, the function g is independent of x.

weight	size	function	$\sigma(f)/w(f)$
$2+\alpha$	2	$x \vee y$	$\frac{6+4\alpha}{2+\alpha} = 2\gamma$
$3+\alpha$	3	$x \vee yz$	$\frac{10+3\alpha}{3+\alpha} = 2\gamma$
$3+\alpha$	3	$x \vee y \vee z$	$\frac{12+3\alpha}{3+\alpha} > 2\gamma$
$4+\alpha$	4	$x \vee yzw$	$\frac{17+4\alpha}{4+\alpha} > 2\gamma$
$4+\alpha$	4	$x \vee y(z \vee w)$	$\frac{15+2\alpha}{4+\alpha} > 2\gamma$
$4+\alpha$	4	$x \vee y \vee zw$	$\frac{16+2\alpha}{4+\alpha} > 2\gamma$
$4+\alpha$	4	$x \vee y \vee z \vee w$	$\frac{20+4\alpha}{4+\alpha} > 2\gamma$
$4+2\alpha$	4	$xy \vee zw$	$\frac{12+8\alpha}{4+2\alpha} = 2\gamma$
$4+2\alpha$	4	$xy \vee \overline{x}\,\overline{y}$	$\frac{12+8\alpha}{4+2\alpha} = 2\gamma$
$4+2\alpha$	4	$xy \vee \overline{x}z$	$\frac{14+8\alpha}{4+2\alpha} > 2\gamma$

Table 5.1: Saving factors for all functions with $2 \le L(f) \le 4$.

If $k \le 4$ then the induction hypothesis for g yields

$$\begin{aligned}\sigma(x \vee g) &\ge 1 + w(f) + \sigma(g) - 4\alpha \\ &\ge 1 + w(f) + 2\gamma(w(f) - 1) - 4\alpha \\ &\ge (1+2\gamma)w(f) - (4\alpha + 2\gamma - 1) \\ &\ge 2\gamma \cdot w(f)\end{aligned}$$

since $w(f) > 5.7 > (4\alpha + 2\gamma - 1) \simeq 5.2$.
If $k \ge 5$ then

$$\begin{aligned}\sigma(x \vee g) &\ge 1 + w(f) + 5(w(f) - 2 - \alpha) \\ &= 6w(f) - (9 + 5\alpha) \ge 2\gamma \cdot w(f)\end{aligned}$$

since $w(f) > 5.7 > (9+5\alpha)/(6-2\gamma) \simeq 4.7$. The proof of the lemma is complete. □

We can now assume that both f_1 and f_2 are non-trivial. The required inequality for f follows from the induction hypothesis and from Lemma 5.2 below.

Lemma 5.2 *If $f = f_1 \vee f_2$ then $\sigma(f)/w(f) \ge 2\gamma$.*

Proof : Since $w(f) = w(f_1) + w(f_2)$, it is enough to prove that $\sigma(f) \ge \sigma(f_1) + \sigma(f_2)$.

Let k be the number of different assignments under which both f_1 and f_2 simplify to literals based on different variables. It is clear that $\sigma(f) \ge$

$\sigma(f_1) + \sigma(f_2) - k\alpha$, but this inequality does not take into account extra savings that arise in the cases where $w(f_{x=c}) < w(f_{1|x=c}) + w(f_{2|x=c})$. We will show that these extra savings sum to more than $k\alpha$. Unless $k > 0$, there is nothing to prove. We consider two cases according to the value of k.

Case 1. $1 \leq k \leq 2$.
Let $x \leftarrow 1$ be one of the special assignments. Then $f_1 = xy \vee \overline{x}g$ and $f_2 = xz \vee \overline{x}h$ for some functions g, h independent of x where the literals x, y, z are all from distinct variables. Without loss of generality we may assume that x, y, z *are* variables.

Consider the effect on f of setting y or z to 1 :

$$\begin{array}{rcccc} f & = & (xy \vee \overline{x}g) & \vee & (xz \vee \overline{x}h) \\ f_{y=1} & = & (x \vee g_{y=1}) & \vee & (xz \vee \overline{x}h_{y=1}) \\ f_{z=1} & = & (xy \vee \overline{x}g_{z=1}) & \vee & (x \vee h_{z=1}). \end{array}$$

Since $f_{1|y=1} = x \vee g_{y=1}$, we have $f_{y=1} = f_{1|y=1} \vee f_{2|y=1,x=0}$. If $h_{y=1} \neq z$, then $f_{2|y=1} = xz \vee \overline{x}h_{y=1}$ depends on x and therefore any formula for $f_{2|y=1}$ contains at least one occurrence of x. Corollary 4.2 therefore implies that $w(f_{2|y=1,x=0}) \leq w(f_{2|y=1}) - 1$, which gives an extra saving of at least 1 when $y \leftarrow 1$. Similarly, if $g_{z=1} \neq y$, there is again an extra saving of at least 1 when $z \leftarrow 1$.

If both $h_{y=1} \neq z$ and $g_{z=1} \neq y$ hold then these two extra savings add up to give an extra credit of 2, which is greater than $k\alpha$.

Suppose now that $h_{y=1} = z$ but $g_{z=1} \neq y$. (The case where $h_{y=1} \neq z$ but $g_{z=1} = y$ is similar.) We then have $f_2 = (x \vee y)z \vee \overline{x}\,\overline{y}h'$ for some h'. Consider again the effect of setting z to 1 :

$$\begin{array}{rcccc} f & = & (xy \vee \overline{x}g) & \vee & ((x \vee y)z \vee \overline{x}\,\overline{y}h') \\ f_{z=1} & = & (xy \vee \overline{x}g_{z=1}) & \vee & (x \vee y \vee h'_{z=1}). \end{array}$$

Using the same argument as before, we get that $f_{z=1} = f_{1|z=1,y=0,x=0} \vee f_{2|z=1}$. The function $f_{1|z=1} = xy \vee \overline{x}g_{z=1}$ depends both on y (as seen with $x \leftarrow 1$) and on x (since we assumed that $g_{z=1} \neq y$). Using Corollary 4.2 we get extra savings of at least 2 from this case alone. Again this is enough since $k \leq 2$.

The last subcase to consider is when both $h_{y=1} = z$ and $g_{z=1} = y$. We then have $f_1 = (x \vee z)y \vee \overline{x}\,\overline{z}g'$ and $f_2 = (x \vee y)z \vee \overline{x}\,\overline{y}h'$, for some g' and h'. The effect of setting y or z to 1 is now

$$\begin{array}{rcccc} f & = & (x \vee z)y \vee \overline{x}\,\overline{z}g' & \vee & (x \vee y)z \vee \overline{x}\,\overline{y}h' \\ f_{y=1} & = & x \vee z \vee g'_{y=1} & \vee & z \\ f_{z=1} & = & y & \vee & x \vee y \vee h'_{z=1}. \end{array}$$

If $g'_{y=1} = 1$ then $f_{y=1} = 1$ and $f = y \vee f'$ for some f', and the result follows from Lemma 5.1. The case when $h'_{z=1} = 1$ is similar. Hence we can assume that $f_{1|y=1}$ depends on z and $f_{2|z=1}$ depends on y. Therefore we get an extra saving of 1 when y is set to 1, and an additional such saving when z is set to 1. The total savings of 2 are sufficient.

Case 2. $k \geq 3$.
The functions f_1 and f_2 both reduce to literals for at least three different assignments. According to Lemma 4.3, f_1 and f_2 are either both of the form $xy \vee \overline{x}\,\overline{y}$ (or one of its duals), in which case the condition $w(f) = w(f_1) + w(f_2)$ is clearly violated, or else all k assignments involve literals $x_1, \ldots, x_k$ of different variables. In this case, $f_1 = Xy \vee \overline{X}g$ and $f_2 = Xz \vee \overline{X}h$, where $X = x_1 \vee \ldots \vee x_k$, the functions g and h do not depend on $x_1, \ldots, x_k$ and y, z are literals of distinct variables, distinct also from the variables of X. Again, without loss of generality, we may assume that $x_1, \ldots, x_k, y, z$ are (distinct) variables.

If $h_{y=1} \neq z$ then as in the previous case we get an extra saving of at least k when y is set to 1 since all the occurrences of $x_1, \ldots, x_k$ in $f_{2|y=0}$ can be replaced by 0. A similar situation arises if $g_{z=1} \neq y$.

The only case left to consider is therefore when $f_1 = (X \vee z)y \vee \overline{X}\overline{z}g'$ and $f_2 = (X \vee y)z \vee \overline{X}\overline{y}h'$ for some functions g' and h'. In this case we prove directly that $\sigma(f)/w(f) \geq 2\gamma$. Note that $\sigma_{x_i=1}(f) = w(f) - (2+\alpha)$ for $1 \leq i \leq k$, $\sigma_{y=1}(f) \geq 1 + (w(f_2) - 1) = w(f_2)$ and $\sigma_{z=1}(f) \geq 1 + (w(f_1) - 1) = w(f_1)$. The first relation holds because f reduces to a twig when any one of the x_i's is set to 1. The second relation (and similarly the third) holds because $\sigma_{y=1}(f_1) \geq 1$, and because f_2 is reduced to a literal when y is set to 1 and therefore $\sigma_{y=1}(f_2) = w(f_2) - 1$. Corollary 4.2 implies that $\sum_v \sigma_{v=0}(f) \geq L(f) \geq w(f)/(1+\alpha/2)$. Combining all this with the fact that $k \geq 3$, we get that

$$\sigma(f) \geq (4 + \tfrac{1}{1+\alpha/2})w(f) - 3(2+\alpha) \geq 2\gamma \cdot w(f)$$

since

$$w(f) > 5.7 > \tfrac{3(2+\alpha)^2}{4} \simeq 5.6\,.$$

This finishes the proof of the second lemma. □

The proof of Theorem 2.1 is now complete. □

6 Shrink-resistant formulae

As observed by Nisan and Impagliazzo [8], the parity function $\mathcal{P}_n$ on n arguments has good resistance to shrinkage.

Theorem 6.1 $\Gamma \leq 2$.

Proof : It is well known that $L(\mathcal{P}_n) = \Theta(n^2)$ (see [5, 6]). By Lemma 2.3, $E[L(\mathcal{P}_n^\varepsilon)] = \Theta(\varepsilon^2 n^2) = \Theta(\varepsilon^2 L(\mathcal{P}_n))$. $\square$

One would expect that for *read-once* formulae, i.e., those formulae in which each variable occurs only once, the analysis of shrinkage should be particularly tractable. We define a set of such functions which appears to have certain extremal properties. We conjecture that these have optimal shrink-resistance among read-once formulae. It is interesting to note the resemblance of our construction to that used by Valiant [12] to construct short monotone formulae for the majority function. The exponent $1/\log_2(\sqrt{5}-1) \simeq 3.27$ is the exponent in his first stage which 'amplifies' $(\psi - \frac{1}{n}, \psi + \frac{1}{n})$ to $(\frac{1}{4}, \frac{3}{4})$, where $\psi = \frac{\sqrt{5}-1}{2}$, and which was shown to be optimal by Boppana [2].

Let $\{r_n\}$ be a sequence of bits. Define two sequences of read-once formulae $\{T_{n,0}\}$ and $\{T_{n,1}\}$ as follows :

$$\begin{aligned}
T_{1,0} &= x_1 \wedge x_2 \\
T_{1,1} &= x_1 \vee x_2 \\
T_{n,i} &= T_{n-1,1-i} \boxed{\text{NAND}}\, T_{n-1,r_n} \quad \text{for } n \geq 2 \text{ and } i = 0, 1.
\end{aligned}$$

In defining $T_{n,i}$ for $n \geq 2$, *disjoint sets of variables* are to be used for $T_{n-1,1-i}$ and for T_{n-1,r_n}. NAND gates are used for convenience. Each of these formulae, or its complement, is equivalent to a complete binary tree with alternating levels of AND and OR gates, except possibly at the level next to the leaves, where both types of gates may appear.

Let $p_{n,i}$ (respectively $q_{n,i}$) be the probability that $T_{n,i}$ takes the value 1 (respectively 0) when each leaf is independently assigned a random value 0 or 1 with equal probabilities.

For $n \geq 2$ we clearly have

$$\begin{aligned}
p_{n,i} &= 1 - p_{n-1,1-i} \cdot p_{n-1,r_n} \\
q_{n,i} &= 1 - p_{n,i} = p_{n-1,1-i} \cdot p_{n-1,r_n} .
\end{aligned}$$

We choose a sequence $\{r_n\}$ for which the probabilities $p_{n,i}$ and $q_{n,i}$ converge to a limit. The next lemma shows that this is possible.

Lemma 6.2 *There exists a sequence $\{r_n\}$ for which*

1. $p_{n,0}, p_{n,1} \to \psi = \frac{\sqrt{5}-1}{2} \simeq 0.62$,
2. $p_{n,0} < \psi < p_{n,1}$ *for every $n \geq 1$,*
3. $p_{n,1} - p_{n,0} < c \cdot \psi^n$ *for every $n \geq 1$ and a fixed $c > 0$.*

Proof : Note that $p_{1,0} = \frac{1}{4} < \psi < p_{1,1} = \frac{3}{4}$. Suppose that we have already chosen $r_2, \ldots, r_{n-1}$ in such a way that $p_{n-1,0} < \psi < p_{n-1,1}$. If $1 - p_{n-1,1}p_{n-1,0} < \psi$ then choose $r_n = 0$, so that $p_{n,0} = 1 - p_{n-1,1}p_{n-1,0} < \psi$ and $p_{n,1} = 1 - p_{n-1,0}p_{n-1,0} > 1 - \psi^2 = \psi$. If $1 - p_{n-1,1}p_{n-1,0} > \psi$ (note that equality cannot hold since ψ is irrational) then choose $r_n = 1$, and the required inequalities are again easily verified.

Next we want to show that $p_{n,0}$ and $p_{n,1}$ do indeed converge to ψ. To that end we observe that

$$\begin{aligned} p_{n,1} - p_{n,0} &= (1 - p_{n-1,0} \cdot p_{n-1,r_n}) - (1 - p_{n-1,1} \cdot p_{n-1,r_n}) \\ &= p_{n-1,r_n}(p_{n-1,1} - p_{n-1,0}). \end{aligned}$$

In particular $p_{n,1} - p_{n,0}$ is a decreasing sequence. It is easy to check that $r_2 = r_3 = 1$ and that $p_{2,0} = \frac{7}{16}$, $p_{2,1} = \frac{13}{16}$, $p_{3,0} = \frac{87}{256}$ and $p_{3,1} = \frac{165}{256}$. Since $p_{n,1} - p_{n,0} \leq p_{3,1} - p_{3,0} = \frac{39}{128}$ for $n \geq 3$, we also get then that $p_{n,i} < \psi + (p_{n,1} - p_{n,0}) < \frac{15}{16}$. Thus $p_{n,1} - p_{n,0}$ decreases exponentially to 0 and $p_{n,0}, p_{n,1} \to \psi$.

Having established the exponential convergence and using the fact that for any $\alpha < 1$, $\prod_{k=1}^{n}(1 + \alpha^k)$ converges, we get that $p_{n,1} - p_{n,0} < c \cdot \psi^n$ for some constant c. $\square$

Assuming now that the formulae $T_{n,i}$ are constructed using the sequence $\{r_n\}$ defined in the proof of Lemma 6.2, we investigate their shrinkage under the random assignment R^ε defined in Section 2.

Under R^ε, each leaf in $T_{n,i}$ is assigned the value 0 or 1 with probability $\frac{1-\varepsilon}{2}$. In the complementary probability ε, the leaf is left unassigned. Denote by $p^*_{n,i}$ (respectively $q^*_{n,i}$) the probability that the random restriction $T^\varepsilon_{n,i}$ is equivalent to the constant function 1 (respectively 0). It is easy to check that

$$\begin{aligned} p^*_{n,i} &= 1 - (1 - q^*_{n-1,1-i})(1 - q^*_{n-1,r_n}) \\ q^*_{n,i} &= p^*_{n-1,1-i} \cdot p^*_{n-1,r_n} \end{aligned}$$

for $n > 1$, and

$$\begin{aligned} p^*_{1,0} &= q^*_{1,1} = \left(\tfrac{1-\varepsilon}{2}\right)^2 \\ p^*_{1,1} &= q^*_{1,0} = 1 - \left(\tfrac{1+\varepsilon}{2}\right)^2. \end{aligned}$$

For all $n \geq 1$ and $i = 0, 1$, we have

$$p^*_{n,i} < p_{n,i} \quad \text{and} \quad q^*_{n,i} < q_{n,i}\,.$$

These inequalities follow immediately from the observation that the random total restriction corresponding to the unstarred variables is a further restriction from R^ε.

Denote by $E_{n,i}$ the expected size of the formula $T^{\varepsilon}_{n,i}$, let $e_{n,i} = E_{n,i}/L(T_{n,i}) = E_{n,i}/2^n$, and finally $e_n = \min\{e_{n,0}, e_{n,1}\}$.

Lemma 6.3 *For $n \geq 2$ we have $e_n \geq p_{n-1,0} \cdot e_{n-1}$.*

Proof : For $n \geq 2$ we have

$$\begin{aligned} e_n &= e_{n,i} \\ &= \tfrac{e_{n-1,1-i}}{2}(1 - q^*_{n-1,r_n}) + \tfrac{e_{n-1,r_n}}{2}(1 - q^*_{n-1,1-i}) \\ &\geq \tfrac{e_{n-1}}{2}\left((1 - q_{n-1,r_n}) + (1 - q_{n-1,1-i})\right) \\ &= \tfrac{e_{n-1}}{2}(p_{n-1,r_n} + p_{n-1,1-i}) \\ &\geq e_{n-1} \cdot p_{n-1,0} \,. \end{aligned}$$

□

From this is it easy to establish

Theorem 6.4 *$E[L(T^{\varepsilon}_{n,i})] \geq c \cdot \psi^n \varepsilon \cdot L(T_{n,i})$ for some $c > 0$.*

If we choose $\varepsilon = 2^{-n+o(n)}$ (which is the minimal value of ε for which a non-constant number of unassigned variables is expected to remain), then for $f = T_{n,i}$ we have

$$E[L(f^{\varepsilon})] = \Omega(\varepsilon^{1-\log_2 \psi - o(1)})L(f) = \Omega(\varepsilon^{1.70})L(f).$$

Although this inequality may seem to suggest that our lower bound of 1.63 is close to optimal, it does not yield a useful upper bound for Γ, since for such a choice of ε we have $E[L(f^{\varepsilon})] = o(1)$. To outweigh the constant d which appears in our definition of the shrinkage exponent requires $E[L(f^{\varepsilon})] = \Omega(1)$. The best that we can get therefore is the following

Theorem 6.5 *If $\varepsilon = (2\psi)^{-n}$ and $f = T_{n,i}$ then*

$$E[L(f^{\varepsilon})] = \Omega\left(\varepsilon^{1/\log_2(2\psi)} \cdot L(f) + 1\right).$$

Therefore $1/\log_2(2\psi) \simeq 3.27$ is an upper bound for the shrinkage exponent Γ^* for read-once formulae.

7 Possible improvements

A natural extension of our proof method in Theorem 2.1 would be to define new weight functions giving credits or debits for other small subtrees. Such an approach should yield an increasing sequence of lower bounds on Γ.

A challenging gap remains between our lower bound and the obvious upper bound. Is parity the most shrink-resistant function?

Conjecture 1 (Nisan and Impagliazzo)

$$\Gamma = 2\,.$$

Are read-once functions so much more shrinkable?

Conjecture 2 For read-once formulae

$$\Gamma^* = 1/\log_2 2\psi \simeq 3.27\,.$$

Read-once formulae can be regarded as a simple special case of monotone formulae.

Conjecture 3 The shrinkage exponent for monotone formulae is the same as for read-once formulae.

Acknowledgements

The authors thank Noam Nisan and Russell Impagliazzo for making available to them a preprint of their paper.

Note added in proof

Recently Håstad, Razborov and Yao [4] have proved Conjecture 2.

References

[1] A.E. Andreev, *On a method for obtaining more than quadratic effective lower bounds for the complexity of π-schemes. Vestnik Moskov. Univ. Mat. 42(1)(1987) 70-73* (in Russian). English translation in *Moscow Univ. Math. Bull. 42(1)(1987) 63-66.*

[2] R.B. Boppana, *Amplification of probabilistic Boolean formulas.* In *Advances in Computer Research, Vol.5: Randomness and Computation* (JAI Press, Greenwich, Conn., to appear).

[3] P.E. Dunne, *The complexity of Boolean networks. Academic Press, 1988.*

[4] J. Håstad, A. Razborov, A. Yao, *On the shrinkage exponent for read-once formulae.* Manuscript.

[5] V.M. Khrapchenko, *Complexity of the realization of a linear function in the class of π-circuits. Mat. Zametki 9 (1971) 35-40* (in Russian). English translation in *Math. Notes Acad. Sciences USSR 9 (1971) 21-23.*

[6] V.M. Khrapchenko, *A method of determining lower bounds for the complexity of* Π*-schemes. Mat. Zametki 10 (1972) 83-92* (in Russian). English translation in *Math. Notes Acad. Sciences USSR 10 (1971) 474-479*.

[7] O.B. Lupanov, *Complexity of formula realization of functions of logical algebra. Prob. Kibernetiki 3 (1962) 61-80* (in Russian). English translation in *Prob. Cyb. 3 (1962) 782-811*.

[8] N. Nisan, R. Impagliazzo, *The effect of random restrictions on formulae size.* Preprint. A revised version appears in *Random Structures and Algorithms*, this issue.

[9] M.S. Paterson, U. Zwick, *Shrinkage of de Morgan formulae under restriction. Proc. 32nd Annual IEEE Symposium on Foundations of Computer Science. (1991), pp. 324-333.*

[10] J. Riordan, C.E. Shannon, *The number of two-terminal series-parallel networks. J. Math. and Physics 21 (1942) 155-171.*

[11] B. A. Subbotovskaya, *Realizations of linear functions by formulas using* $+, *, -$. *Doklady Akademii Nauk SSSR 136 (1961) 553-555* (in Russian). English translation in *Soviet Mathematics Doklady 2 (1961) 110-112.*

[12] L.G. Valiant, *Short monotone formulae for the majority function. J. Algorithms 5 (1984) 363-366.*

[13] I. Wegener, *The complexity of Boolean functions. Wiley-Teubner Series in Computer Science, 1987.*

[14] U. Zwick, *An extension of Khrapchenko's theorem. Information Processing Letters 37 (1991) 215-217.*