Width lower-bounds for Resolution proof system

Why do we care about width lower bounds

The width lower bounds for resolution are important because due to the following theorem, we can transform the width lower bound into a size lower bound.

Theorem 1 (Ben-Sasson, Wigderson)**.** For all k and $k-$CNF $\mathcal{F}$ the following holds:

1. $w(\mathcal{F} \vdash 0) \leq k + log\left(S_{tRes}(\mathcal{F} \vdash 0)\right)$
2. $w(\mathcal{F} \vdash 0) \leq k + O\left(\sqrt{nlog(S_{Res}(\mathcal{F} \vdash 0))}\right)$.

Proof strategy

All lower bounds on width follow the same strategy:

1. Define a complexity measure $\mu\colon \{Clauses\} \to \mathbb{N}$ such that $\mu(Axiom) \leq 1$.
2. Prove $\mu(\emptyset)$ is large.
3. Infer that in any refutation there is some clause C with medium size $\mu(C)$.
4. Prove that if $\mu(C)$ is medium then $w(C)$ is large.

We shall now formalize and explain this strategy. First, we will define a measure that will satisfy condition 1.-3.

Definition 2 ($\mu_{\mathcal{A}}$)**.** For f a boolean function, let $Vars(f)$ denote the set of variables appearing in f. Let $\alpha \in \{0,1\}^{Vars(f)}$ be an assignment to f. We say that α satisfies f, if $f(\alpha) = 1$. For C a clause and Γ a set of boolean functions, let $V = Vars(f) \cup Vars(C)$. We say that Γ implies C, denoted by $\Gamma \models C$, if every assignment satisfying every function $\gamma \in \Gamma$ satisfies C as well.

Let $\mathcal{A}$ be an unsatisfiable set of boolean functions, i.e. $\mathcal{A} \models 0$, and let C be a clause. We define $\mu_{\mathcal{A}}(C) := \min\{|\mathcal{A}'|; \mathcal{A}' \subseteq \mathcal{A}, \mathcal{A}' \models C\}$.

$\mu_{\mathcal{A}}$ is a sub-additive complexity measure with respect to resolution steps:

Exercise 3. Suppose D was inferred from B, C by a single resolution step. Then for any set of boolean functions $\mathcal{A}$: $\mu_{\mathcal{A}}(D) \leq \mu_{\mathcal{A}}(B) + \mu_{\mathcal{A}}(C)$.

In order to assure condition 1 of the strategy, we want $\mu(Axiom)$ to be small:

Definition 4 (Compatibility)**.** For $\mathcal{F}$ a non-satisfiable CNF we say that $\mathcal{A}$ is compatible with $\mathcal{F}$ if $\mathcal{A} \models 0$ and $\forall C \in \mathcal{F}\ \mu(C) \leq 1$.

We always pick a compatible $\mathcal{A}$ and use it to define $\mu = \mu_{\mathcal{A}}$. Note that part 2 of the strategy puts another requirements on $\mathcal{A}$, namely that no "small" subset of it is contradictory. However, this would be intuitively easy to achieve with "hard" tautologies.

We now claim that part 3 can be deduced from the definitions:

Exercise 5. If $\mathcal{A}$ is compatible with $\mathcal{F}$ then in every refutation of $\mathcal{F}$ there must be a clause C with $\mu(\emptyset)/3 \leq \mu(C) \leq 2\mu(\emptyset)/3$.

Definition 6. A boolean function f is called Sensitive if any two distinct falsifying assignments $\alpha, \beta \in f^{-1}(0)$, have Hamming distance greater than 1. An example of a Sensitive function is PARITY.

For $\mathcal{A}$ a set of boolean functions, and $f \in \mathcal{A}$, a Critical Assignment for f is an assignments $\alpha \in \{0,1\}^{Vars(\mathcal{A})}$ such that $g(\alpha) = \begin{cases} 0 & g = f \\ 1 & g \neq f, g \in \mathcal{A} \end{cases}$

For $\alpha, \beta \in \{0,1\}^{Vars(\mathcal{A})}$, we say that β is the result of flipping α on the variable x, if $\beta(y) = \begin{cases} 1 - \alpha(y) & y = x \\ \alpha(y) & otherwise \end{cases}$

We shall now define the expansion of a CNF formula in terms of its minimal boundary:

Definition 7 (Boundary)**.** For f a boolean function and x a variable, we say that f is dependent on x if there is some assignment α such that $f(\alpha) = 0$, but flipping α on x satisfies f.

For $\mathcal{A}$ a set of boolean functions, the Boundary of $\mathcal{A}$, denoted $\partial\mathcal{A}$, is the set of variables x such that there is a unique function $f \in \mathcal{A}$ that is dependent on x.

Exercise 8 (Sanity check)**.** Check that a critical assignment to a sensitive function can be changed to a satisfying assignment, by flipping a boundary variable. Formally:
If $f \in \mathcal{A}$ is Sensitive, α is a Critical Assignment for f, and $x \in Vars(f) \cap \partial\mathcal{A}$ then flipping α on x yields an assignment β that satisfies $\mathcal{A}$.

We define the expansion of $\mathcal{F}$ to be the minimal boundary of a medium size sub-formula of $\mathcal{A}$:

Definition 9 (Expansion)**.** For $\mathcal{A} \models 0$, let $k = \mu_{\mathcal{A}}(\emptyset)$. We define the Expansion of $\mathcal{A}$ to be:

$$e(\mathcal{A}) := \mathbf{min}\{|\partial\mathcal{A}'|; \mathcal{A}' \subseteq \mathcal{A}, 1/3 \cdot k \leq |\mathcal{A}'| \leq 2/3 \cdot k\}.$$

The main tool, used in proving most lower bounds on width, presents the connection between width and expansion:

Theorem 10. For $\mathcal{F}$ an unsatisfiable CNF:

$$w(\mathcal{F} \vdash 0) \geq \max e(\mathcal{A}),$$

where the maximum is taken over all sets $\mathcal{A}$ of sensitive functions, compatible with $\mathcal{F}$.

Proof. Fix some $\mathcal{A}$ that is compatible with $\mathcal{F}$, and let $\mu_{\mathcal{A}}(\emptyset) = k$. By Exercise 5 there must exist some clause C such that $k/3 \leq \mu_{\mathcal{A}}(C) \leq 2k/3$. Let $\mathcal{A}' \subset \mathcal{A}$ be a minimal set such that $\mathcal{A}' \models C$. We claim that any variable $x \in \partial\mathcal{A}'$ must appear in C. To see this, notice that for every $f \in \mathcal{A}'$ there is some assignment α_f such that $\alpha_f(C) = \alpha_f(f) = 0$ and $\alpha_f(g) = 1$ for all $g \in \mathcal{A}, g \neq f$. This follows from the minimality of $\mathcal{A}'$, for otherwise $\mathcal{A}' \setminus f \models C$. Suppose, for the sake of contradiction, that $x \in \partial\mathcal{A}' \cap Vars(f)$ but $x \notin C$. By Exercise 8, flipping α_f on x satisfies $\mathcal{A}'$, but the new assignment agrees with α_f on $Vars(C)$. Hence $\mathcal{A}' \not\models C$, contradiction. □

Main results for Tseitin formulas and PHP

Definition 11 ($TSE_{G,f}$)**.** Let $G = ([n], E)$ be an undirected, acyclic graph. Let $f\colon [n] \to \{0,1\}$ be a function assigning to each vertex 0 or 1. Consider the set of equations in $\mathbb{F}_2$ in the variables x_e, where $e = \{i,j\}$ are the edges from E:

$$\bigoplus_{j:\{i,j\}\in E} x_{i,j} = f(i), \text{ for each } i \in [n].$$

$TSE_{G,f}$ is a formula saying that these equations hold together.

Exercise 12 ($TSE_{G,f}$)**.** Assuming d is a maximum degree of G, write $TSE_{G,f}$ as d-CNF and give an upper bound on number of clauses and number of variables.

Fact 13. Assume $\sum_{i\in[n]} f(i) = 1$ in $\mathbb{F}_2$. Then $TSE_{G,f}$ is unsatisfiable.

Definition 14 (Expansion for graphs)**.** For G a finite connected graph, the Expansion of G is

$$e(G) := \min\{|E(V', V \setminus V')|; V' \subseteq V, |V|/3 \leq |V'| \leq 2|V|/3\}.$$

Theorem 15. For G a finite, connected, undirected, acyclic graph and f an odd-weight function on $V(G)$, $w(TSE_{G,f} \vdash 0) \geq e(G)$.

Proof. For $v \in V(G)$, we define $PARITY_v := \left(\bigoplus_{v\in e, e\in E} x_e = f(v)\right)$ where the equations are over $\mathbb{F}_2$. Set $\mathcal{A}_V = \{PARITY_v; v \in V(G)\}$ and denote $\mu(C) = \mu_{\mathcal{A}_V}(C)$. Every axiom C is one of the defining axioms of $PARITY_v$. Clearly, for this very same v, $PARITY_v \models C$. Hence for any axiom C, $\mu(C) = 1$.

So far we have shown that $\mathcal{A}_V$ is compatible for $TSE_{G,f}$. Next, we claim that $\mu(\emptyset) = |V(G)|$, because for any $|V'| < |V(G)|$, $\mathcal{A}_{V'}$ is satisfiable. This latter claim is seen by the following reasoning: Let v be some vertex in $V \setminus V'$. Look at the formula $TSE_{G,f'}$ for $f'(u) = \begin{cases} 1 - f(u) & u = v \\ f(u) & otherwise \end{cases}$. By Fact 13, $TSE_{G,f'}$ is satisfiable. $\mathcal{A}_V$ is a sub-formula of $TSE_{G,f'}$, and hence satisfiable as well. $\mathcal{A}_{V(G)}$ is a collection of $PARITY$ functions which are Sensitive. Finally, for $V' \subseteq V$, $\partial\mathcal{A}_{V'} = \{x_e; e \in E(V', V \setminus V')\}$. This is true because if $e = \{v, u\}, v \in V', u \in V \setminus V'$ then $PARITY_v$ is the only function in $\mathcal{A}_{V'}$ dependent on x_e. Hence $e(\mathcal{A}_V) \geq e(G)$ and we apply Theorem 9 to complete the proof. □

Corollary 16. For G a 3-connected expander (i.e. $e(G) = \Omega(|V|))$n and f an odd-weight function on $V(G)$, $S(TSE_{G,f}) = 2^{\Omega(|TSE_{G,f}|)}$.

Definition 17 (Our version of PHP)**.** The pigeonhole principle PHP^n_{n-1} is the CNF formula over variables x_{ij} with $1 \leq i \leq n$, $1 \leq j \leq n-1$ consisting of:

- **Pigeon axioms:**
$$P_i = \bigvee_{j=1}^{n-1} x_{ij},$$
asserting that each pigeon goes to some hole.
- **Hole axioms:**
$$x_{i_1 j} \vee x_{i_2 j} \qquad (i_1 \neq i_2),$$
asserting that no two pigeons occupy the same hole.

Clearly PHP^n_{n-1} is unsatisfiable, since no injective function $[n] \to [n-1]$ exists.

Theorem 18. There is a constant $c > 1$ such that any resolution refutation of PHP^n_{n-1} requires size c^n.

Proof. Here we call a truth assignment α *i-critical* if it falsifies exactly the pigeon axiom P_i. It is *critical* if it is i-critical for some i.

Given a clause C, let $C(\alpha)$ denote its truth value under α.

Claim 1. Let C be any clause and C^+ the clause obtained by replacing every negated literal $\neg x_{ij}$ by the subclause $X_{ij} = \bigvee_{i' \neq i} x_{i'j}$. Then for any critical assignment α we have $C^+(\alpha) = C(\alpha)$.

Exercise 19. Prove Claim 1.

Claim 2. Every resolution refutation of PHP^n_{n-1} must contain a clause where the width of C^+ satisfie $w(C^+) \geq 2n^2/9$.

Proof of Claim 2. Given a clause C let $Pigeon(C) \subseteq [n]$ be the set of pigeons where there is some i-crirical assignment α for which $C(\alpha) = 0$. Let $\mu(C) = |Pigeon(C)|$. We observe that

1. $\mu(P_i) = 1$ for each pigeon axiom P_i,
2. $\mu(\emptyset) = n$ since every assignment falsifies $\emptyset$,
3. if clause C was derived from clauses A, B using the one resolution step then $\mu(C) \leq \mu(A) + \mu(B)$. This is because every assignment falsifying C must falsify A or B.

Therefore, we can conclude from the conditions above that if R was a resolution refutation of PHP^n_{n-1} there is some clause C in R where $n/3 \leq \mu(C) \leq 2n/3$. Now we want to argue that for that clause C, we have $w(C^+) \geq 2n^2/9$.

Fix $i \in \mathit{Pigeon}(C)$ and some $j \notin \mathit{Pigeon}(C)$ so that α is an i-critical assignment that falsifies C. Without loss of generality, assume that α sent a pigeon j to a hole k. By flipping α on x_{ik} and x_{jk} we obtain a j-critical assignment α'. Furthermore, $C(\alpha) = C^+(\alpha) = 0$ and $C(\alpha') = C^+(\alpha') = 1$ by construction and the previous claim. Therefore, since C^+ contains only positive literals and only variable whose truth value was switched from 0 to 1 is x_{ik}, we conclude that C^+ must contain x_{ik}.

Repeating the argument for all pairs (i, j) such that $i \in \mathit{Pigeon}(C)$ and $j \notin \mathit{Pigeon}(C)$, and since all pairs must yield a distinct variable in C^+, we can conclude that the width of C^+ is at least $s(n - s)$ where $s = |\mathit{Pigeon}(C)|$. By assumption that $n/3 \leq |\mathit{Pigeon}(C)| \leq 2n/3$, we get that $w(C^+) \geq 2n^2/9$. □

Now we will use that width lower bound to get the size lower bound. We let R be a resolution refutation of PHP^n_{n-1} and let R^+ be the positive version where each clause $C \in R$ is replaced by its positive version C^+. Let $\varepsilon > 0$ be a constant whose exact value we will choose later.

Definition 20. A clause in R^+ is ε-wide if its width satisfies $w(C^+) \geq \varepsilon n^2$.

Let S be the number of wide clauses in the refutation R^+. We can conclude that there is some variable x_{ij} appearing in at least εS wide clauses. by counting the number of variables in wide clauses. Therefore, we can define a restriction that sets $x_{ij} = 1$, all $x_{i'j} = 0$ for $i' \neq i$, and all $x_{ij'} = 0$ for $j' \neq j$.

Notice that once we apply the restriction to R (and hence to R^+), R is now a resolution refutation of PHP^{n-1}_{n-2} and there are now at most $(1-\varepsilon)S$ wide clauses in R^+. Therefore, we can inductively apply k restrictions for $k = ln(S)/\varepsilon$ so that we get a proof of PHP^{n-k}_{n-k-1} where the positive version R^+ contains no wide clauses since there are at most $S(1-\varepsilon)^k$ wide clauses and $S(1-\varepsilon)^k < Se^{-k\varepsilon} \leq 1$ by the choice of ε and k.

However, from the previous claim, we know that in any refutation of PHP^{n-k}_{n-k-1} there is some clause C where $w(C^+) \geq 2(n-k)^2/9$ but on the other hand, we

have produced a refutation of PHP^{n-k}_{n-k-1} with no wide clauses by restricting the proof of PHP^n_{n-1}. Therefore, the inequality

$$\frac{2(n-k)^2}{9} = \frac{2(n - \frac{ln(S)}{\varepsilon})}{9} \leq w(C^+) < \varepsilon n^2$$

must be satisfied. After some algebra, we can conclude that $ln(S) \geq \varepsilon n - \varepsilon\sqrt{\frac{9\varepsilon}{2}}n$, so picking $\varepsilon = 8/81$ to maximize the bound means that the resolution refutation of PHP^n_{n-1} must contain at least

$$S \geq exp\left(\frac{8n}{243}\right) \geq 1.033^n$$

wide clauses, which finishes the proof. □