Width and size of regular resolution proofs

Alasdair Urquhart

Banff Proof Complexity Workshop

October 2011

An irregular refutation

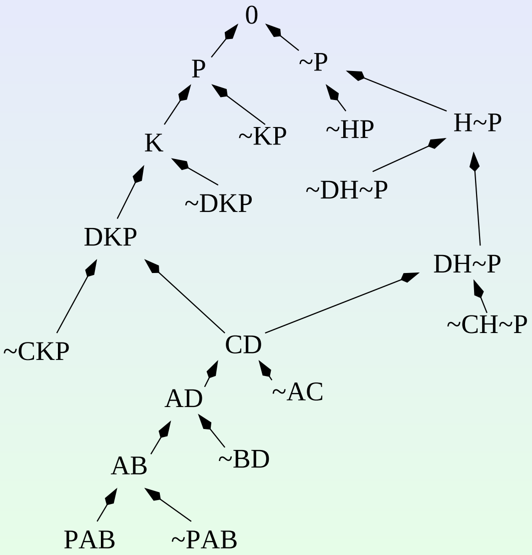

Are minimal refutations regular?

This refutation is interesting because it is the first example discovered of a set of clauses where the minimal refutation is necessarily irregular [Wenqi Huang and Xiangdong Yu 1987].

Prior to this discovery, several researchers had attempted to show that there is always a regular refutation of minimal size. This is true for tree resolution, but definitely not true when refutations are presented as directed acyclic graphs.

Tseitin on regularity

Tseitin (1966) makes the following remarks about the heuristic interpretation of the regularity restriction:

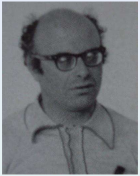

The regularity condition can be interpreted as a requirement for not proving intermediate results in a form stronger than that in which they are later used (if A and B are disjunctions such that $A \subseteq B$, then A may be considered to be the stronger assertion of the two); if the derivation of a disjunction containing a variable ξ involves the annihilation of the latter, then we can avoid this annihilation, some of the disjunctions in the derivation being replaced by "weaker" disjunctions containing ξ.

A superpolynomial separation

The first superpolynomial separation between regular and general resolution was proved by Andreas Goerdt in 1993.

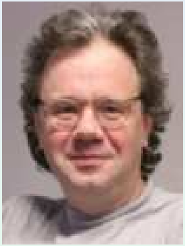

His proof is rather complicated and depends on a modified version of the propositional pigeonhole principle.

An exponential separation

The first exponential separation was proved by Alekhnovich, Johannsen, Pitassi and Urquhart in 2002 [STOC 2002, Theory of Computing 2007].

The paper contains two sets of examples providing exponential separations between regular and general resolution.

First example

Let GT_n be the set of clauses saying that there is a linear ordering of $\{1, \ldots, n\}$ with no last element. This example has size $O(n^3)$, but requires tree resolution refutations of size $2^{\Omega(n)}$.

Gunnar Stålmarck

Krishnamurthy [1985] conjectured that GT_n requires superpolynomial size resolution refutations, but this was refuted by Stålmarck, who showed that they in fact have linear size resolution refutations [1996].

Misha Alekhnovich's trick

The linear-size refutations of Stålmarck are regular, so there is no hope of a separation by using these clause sets directly.

However, Misha Alekhnovich thought of a trick that converts GT_n into a set of clauses that is hard, not just for tree resolution, but also for regular resolution. The basic idea is to replace a single clause by a pair of clauses

$$C \longmapsto \{C \vee x, C \vee \overline{x}\},$$

where x is a variable chosen in a particular way (more on this later).

Misha's examples

Misha Alekhnovich
1978 – 2006

Using Misha's trick, we can convert the set of clauses GT_n (that has linear-size regular resolution refutations) into a set of clauses GT_n^* that requires regular refutations with size $2^{\Omega(n)}$.

Second example

The second family of examples is constructed from a family of pebbling formulas.

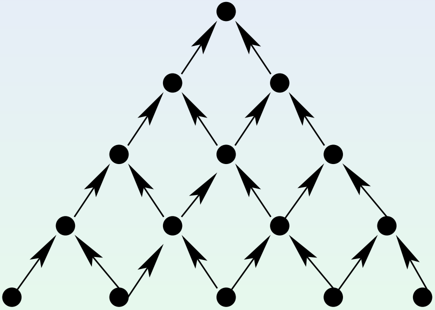

This graph has pebbling number 6. More generally, the pyramid graph with n source vertices has size $O(n^2)$ and pebbling number $n+1$.

Pebbling formulas

Let G be a directed acyclic graph with a unique sink node. The pebbling formula $\mathrm{Peb}(G)$ is a set of clauses that says:

- Any source node can be pebbled.
- If all predecessors of a node can be pebbled, then the node itself can be pebbled.
- The sink node cannot be pebbled.

The set of clauses $\mathrm{Peb}(G)$ has a resolution refutation that is linear in the size of G. However, any resolution refutation of $\mathrm{Peb}(G)$ requires depth bounded below by the pebbling number of G. This last property is the key feature of the pebbling clauses that allows us to separate both tree resolution and regular resolution from general resolution.

Construction of second set of examples

The second set of examples (producing the $2^{\Omega(\sqrt[4]{R}/(\log R)^3)}$ separation), are constructed from a directed acyclic graph G, and can be understood as asserting the following claims.

- There is a non-empty set of pebbles, each of which is red or blue (but not both).
- Every node in the graph G has a pebble on it.
- If all predecessors of a node are pebbled with a red pebble, so is the node.
- The sink node is pebbled with a blue pebble.

Separating general from regular width

Both sets of examples separate general width from regular width. That is, the proofs showing the size separation between general and regular resolution also show that the examples have small general width, but large regular width (any regular refutation of them must have large width).

This suggests a possible generalization of a theorem of Ben-Sasson and Wigderson.

The width-size tradeoff theorem (Ben-Sasson and Wigderson 1999)

Let Σ be a contradictory set of clauses with an underlying set of variables V, $w(\Sigma)$ the maximum number of literals in a clause in Σ, and $w(\Sigma \vdash 0)$ the maximum width of a resolution refutation of Σ. Then

$$S(\Sigma) = \exp\left(\Omega\left(\frac{[w(\Sigma \vdash 0) - w(\Sigma)]^2}{|V|}\right)\right).$$

Could there be a similar theorem for "regular width" and "regular size"?

NO!

The width-size tradeoff theorem (Ben-Sasson and Wigderson 1999)

Let Σ be a contradictory set of clauses with an underlying set of variables V, $w(\Sigma)$ the maximum number of literals in a clause in Σ, and $w(\Sigma \vdash 0)$ the maximum width of a resolution refutation of Σ. Then

$$S(\Sigma) = \exp\left(\Omega\left(\frac{[w(\Sigma \vdash 0) - w(\Sigma)]^2}{|V|}\right)\right).$$

Could there be a similar theorem for "regular width" and "regular size"?

NO!

Graphs with high pebbling number

Paul, Celoni and Tarjan [1977] constructed a sequence G_i of directed graphs, with pebbling number $cn/\log n$, where $n = n(i)$ is the number of nodes in the graph.

Adding random literals (iterating Misha's trick)

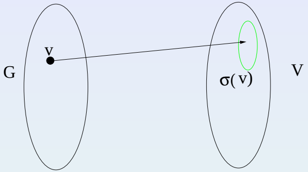

For each $v \in G$, let $\sigma(v)$ be a sequence of variables with size $p = \lceil \log^5 n \rceil$. Then let the set $\text{Clauses}^\sigma(v)$ be the set of all clauses having the form

$$\text{Clause}(v) \vee \pm\sigma_1(v) \vee \cdots \vee \pm\sigma_p(v),$$

where $\pm r$, for $r \in V$, is either r or $\neg r$. Then define

$$\text{Peb}^\sigma(G) = \bigcup \{\text{Clauses}^\sigma(v) | v \in \text{Peb}(G)\}.$$

A combinatorial condition

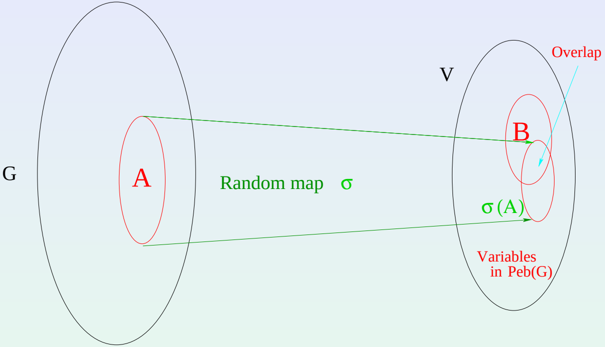

We need a map σ from clauses to variables so that the image under ρ of a large set of vertices has a large overlap with a large set of variables. A set of clauses is "large" if it contains $\Theta(n/\log n)$ elements, similarly for a set of variables. A "large" overlap contains $\Theta(n/\log n)$ elements. We can prove the existence of σ by a probabilistic construction.

Main Result

There is an infinite sequence $\Sigma_1, \Sigma_2, \ldots, \Sigma_i, \ldots$ of contradictory sets of clauses and a corresponding list of parameters $n(1), n(2), \ldots, n(i), \ldots$ so that (abbreviating $n(i)$ as n):

1. Each clause set Σ_i contains $n - 1$ variables and $n^{O(log^4 n)}$ clauses with width $O(\log^5 n)$;
2. Σ_i has a regular tree refutation with size $n^{O(log^4 n)}$;
3. Any regular refutation of Σ_i must contain a clause with width $\Omega(n/\log n)$.

Proof: Let $\Sigma_i = \mathrm{Peb}^{\sigma}(G_i)$.

Open Problems

1. What is the complexity of determining the minimum regular width of a set of clauses? (Conjecture: PSPACE-complete).
2. (Moshe Vardi) What is the complexity of determining the resolution width of a set of clauses? (Conjecture: EXPTIME-complete).
3. Prove or disprove: The Tseitin graph tautologies always have a regular proof with minimal size. Same question for the pigeonhole principle.


THE
END