Logic and Proof Hilary 2024

Lower Bounds for Resolution

James Worrell

No polynomial-time algorithm is known for the SAT problem. However if a propositional formula is satisfiable then there is a short and easily checkable certificate of this fact—namely a satisfying valuation. An important question is whether there likewise exist short certificates of unsatisfiability. One candidate for such a certificate would be a resolution refutation, i.e., we can ask whether there exists a polynomial upper bound on the length of the shortest refutation of an unsatisfiable formula. In this lecture we give a negative answer to this question by exhibiting a family of formulas with long (i.e., exponential in the formula size) refutations.

We will give a lower bound on the length of resolution proofs of the pigeonhole principle:

If n pigeons are placed in $n-1$ boxes then some box contains at least two pigeons.

Fix $n \in \mathbb{N}$. For $i \in \{1, \ldots, n\}$ and $j \in \{1, \ldots, n-1\}$, let propositional variable $x_{i,j}$ denote that pigeon i is in box j. We consider the following formulas:

$$P_i := \bigvee_{j=1}^{n-1} x_{i,j} \quad \text{“pigeon } i \text{ is in some box”}$$

$$\begin{aligned}\mathrm{CRIT}_n := & \bigwedge_{j=1}^{n-1} \bigvee_{i=1}^{n} x_{i,j} \quad \text{“every box contains some pigeon”} \\ & \wedge \bigwedge_{j=1}^{n-1} \bigwedge_{1 \le i < i' \le n} (\neg x_{i,j} \vee \neg x_{i',j}) \quad \text{“no box contains two different pigeons”} \\ & \wedge \bigwedge_{i=1}^{n} \bigwedge_{1 \le j < j' \le n-1} (\neg x_{i,j} \vee \neg x_{ij'}) \quad \text{“no pigeon is in two different boxes”}\end{aligned}$$

A valuation that satisfies CRIT_n is said to be *critical.* Such a valuation corresponds to a bijective assignment of $n-1$ out of n pigeons to $n-1$ boxes, with one pigeon left unassigned. We formalise the pigeonhole principle for n pigeons to be the statement that $\mathrm{PHP}_n := \mathrm{CRIT}_n \wedge \bigwedge_{i=1}^{n} P_i$ is unsatisfiable. The rest of the lecture is devoted to a proof of the following result.

Theorem 1. Every resolution refutation of PHP_n has length at least $2^{n/21}$.

We lay the groundwork for the proof by introducing some key concepts. We say that a sequence of monotone clauses (i.e., clauses with only positive literals) $C_1, \ldots, C_m$ is a *pseudo refutation* of PHP_n if C_m is the empty clause and for all $1 \le i \le m$ either:

PR1 $\mathrm{CRIT}_n \wedge P_j \models C_i$ for some $1 \le j \le n$, or

PR2 $\mathrm{CRIT}_n \wedge C_j \wedge C_k \models C_i$ for some $j, k < i$.

We say that $W \subseteq \{1, \ldots, n\}$ is a *witness* of a clause C if $\mathrm{CRIT}_n \wedge \bigwedge_{i \in W} P_i \models C$ (i.e., every critical assignment that houses all pigeons in W satisfies C). Every clause in a pseudo refutation has a witness: a clause that follows by rule PR1 has a singleton witness, while a clause that follows from PR2 has as witness the union of the witnesses of its two antecedents under PR2. We further define the *weight* of C to be the minimum cardinality of any witness of C.

For a clause C, let C^* be the clause in which each negative literal $\neg x_{i,j}$ is replaced by $\bigvee_{i' \neq i} x_{i',j}$. Observe that $\mathrm{CRIT}_n \models C \leftrightarrow C^*$—that is, for all critical assignments, pigeon i is not in box j iff some other pigeon is in box j. It follows that if $C_1, \ldots, C_m$ is a resolution refutation of PHP_n then $C_1^*, \ldots, C_m^*$ is a pseudo refutation of PHP_n. It thus suffices to prove a lower bound on the length of pseudo refutations of PHP_n.

Proposition 2. Every pseudo refutation of PHP_n contains a clause with at least $\frac{2n^2}{9}$ variables.

Proof. Consider a pseudo refutation $\rho := C_1, \ldots, C_m$ of PHP_n. Since C_m (which is the empty clause) has weight n, there exists a first clause C in ρ with weight $\geq n/3$. But then C has weight at most $2n/3$ (at most the sum of the weight of its antecedents under rule PR2), i.e., its weight lies between $n/3$ and $2n/3$.

We now argue that the clause C, identified above, contains at least $\frac{2n^2}{9}$ variables. Let W be a minimal witness for C, with $n/3 \leq |W| \leq 2n/3$. For each $i_1 \in W$ we exhibit $n - |W|$ different variables in C of the form $x_{i_1,j}$. We conclude that C contains at least $|W|(n - |W|) \geq \frac{2n^2}{9}$ variables.

Fix $i_1 \in W$. By the minimality of W as a witness, there exists a critical assignment $\mathcal{A}$ that leaves out pigeon i_1 and does not satisfy C. Now let $i_2 \notin W$ and suppose that $\mathcal{A}$ assigns i_2 to box j_2. Define an assignment $\mathcal{A}'$ by $\mathcal{A}'[\![x_{i_1,j_2}]\!] = 1$, $\mathcal{A}'[\![x_{i_2,j_2}]\!] = 0$, and otherwise $\mathcal{A}'$ agrees with $\mathcal{A}$. (That is, $\mathcal{A}'$ assigns pigeon i_1 to box j_2 and makes i_2 the unassigned pigeon.) Then $\mathcal{A}'$ satisfies $\mathrm{CRIT}_n \wedge \bigwedge_{i \in W} P_i$ and hence $\mathcal{A}'$ satisfies C. But the fact that $\mathcal{A}'$ satisfies C while $\mathcal{A}$ does not satisfy C entails that x_{i_1,j_2} is mentioned in C. This completes the proof. □

Proof of Theorem 1. Let $\rho := C_1, \ldots, C_m$ be a pseudo refutation of PHP_n. Say that a clause is *long* if it contains at least $\frac{n^2}{8}$ variables. Suppose that ρ has ℓ long clauses. By double counting (i.e., using the fact that the sum of the number of variables in each long clause equals the sum of the number of long clauses that each variable belongs to) we see that some variable is mentioned in as least $\ell/8$ long clauses. By renaming variables if necessary, we can assume that the aforementioned variable is $x_{n,n-1}$. Now we transform ρ by *"assigning pigeon n to box $n-1$"*—formally we delete any clause containing $x_{n,n-1}$ and then delete from the remaining clauses every variable $x_{i,j}$ with either $i = n$ or $j = n-1$. Then the resulting sequence $C_1', \ldots, C_{m'}'$ is a pseudo refutation of PHP_{n-1} with at most $\frac{7}{8}\ell$ long clauses.[1] Repeating this process $n/4$ times we arrive at pseudo refutation of $\mathrm{PHP}_{3n/4}$ with at most $\left(\frac{7}{8}\right)^{n/4} \ell$ long clauses. But by Proposition 2, every pseudo refutation of $\mathrm{PHP}_{3n/4}$ contains a clause with $\frac{2}{9}(3n/4)^2 = \frac{n^2}{8}$ variables, i.e., a long clause. We deduce that $\left(\frac{7}{8}\right)^{n/4} \ell \geq 1$ and hence $\ell \geq \left(\frac{8}{7}\right)^{n/4} \geq 2^{n/21}$. But this means that there are more than $2^{n/21}$ clauses in ρ. □

The proof technique above has been called a *bottleneck counting argument.* The bottlenecks are the long clauses. The proof combines a lower bound on bottlenecks (that every pseudo refutation of $\mathrm{PHP}_{3n/4}$ contains at least one long clause) and an upper bound on bottlenecks (if a pseudo refutation of PHP_n has ℓ long clauses then $\mathrm{PHP}_{3n/4}$ has a pseudo refutation with at most $(7/8)^{n/4}\ell$ long clauses).

[1] You are asked to give a formal proof of this in Exercise Sheet 2. The substitution lemma for propositional logic will prove helpful.