Topics in proof complexity (Phd Open 2021)
Problem set
Due: March 8, 2021

Update March 8. I was just alerted to an issue with Problem 4, in which there was a bug. The definition of Res(2) was missing the information that the system is also allowed to use axioms, which have the form $x, \overline{x}$ for a variable x. This is now corrected. If you were working on Problem 4 and need a bit more time to take this correction into account, let me know.

Note. I think that none of the 8 problems below is "trivial", but many are "routine", for instance because they can be solved by slightly tweaking one of the techniques we discussed. In general, if you want grade n, I would like you to solve $n-1$ problems more or less correctly. If you want to protect yourself against the effect of unexpected major gaps/errors, feel free to send solutions of more problems or to do something about the "bonus points" questions below some of the problems. (The bonus questions have no specified value in problem-equivalents. How much they are worth depends on the question at hand and what you do with it.)

In upper bound proofs, do not be obsessively detailed, but do provide a good enough sketch that I can determine whether your argument is correct. If I am in doubt about that, I may contact you.

Please send your solutions to `lak@mimuw.edu.pl`. If you find any bugs or have questions, send them to the same address. I will try to maintain an as-debugged-as-possible version of the problem set at the link where you found this.

By the way: for those of you who attended or watched the class on Saturday, the problem I got stuck on at the end has a very easy solution once you recall that the variables are partially disjoint. You should be able to watch a recording of an explanation at the PhD Open subpage for these lectures.

Problem 1. The CNF formula $\mathrm{CPLS}_{n,m}$ (the letters stand for *Coloured Polynomial Local Search*) claims to describe an $n \times n$ matrix in which each cell can have one of m colours, and in which each cell in row $i < n$ has at least one distinguished "teacher" among the cells in row $i+1$. The formula states that each cell in row n has some colour, and that if a teacher of a cell has a colour, then the cell also has it, but that cell $(1,1)$ has no colour.

The variables of $\mathrm{CPLS}_{n,m}$ are x_{ijk} for $i, j \in [n]$, $k \in [m]$ ("cell (i,j) has colour k") and $y_{ijj'}$ for $i \in [n-1], j, j' \in [n]$ ("cell $(i+1, j')$ is the teacher of (i,j)"). The clauses are:

$$\begin{array}{ll}
\bigvee_{k=1}^{m} x_{njk} & \text{for each } j \in [n], \\
\overline{x_{11k}} & \text{for each } k \in [m], \\
\bigvee_{j'=1}^{n} y_{ijj'} & \text{for each } i \in [n-1], j \in [n], \\
\overline{y_{ijj'}} \vee \overline{x_{(i+1)j'k}} \vee x_{ijk} & \text{for each } i \in [n-1], j, j' \in [n], k \in [m].
\end{array}$$

Prove that $\mathrm{CPLS}_{n,n}$ has a poly(n)-size refutation in Resolution.
(*Remark.* $\mathrm{CPLS}_{n,n}$ is, in some sense, "the hardest" formula that is easy for Resolution. This is a bit difficult to state as a formal result because if we defined a notion of reduction that would turn this into a completeness theorem, the class of formulas easy for Resolution would not be closed under it.)

Problem 2. Prove that $\mathrm{CPLS}_{n,1}$ has a $\mathrm{poly}(n)$-size treelike refutation in Resolution. (*Bonus points.* Prove that $\mathrm{CPLS}_{n,n}$ does not.)

Problem 3. The *relativized ordering principle* rOP_n is a CNF that claims to describe a relation on n elements with the property that, for some nonempty set $U \subseteq [n]$, the relation is a linear order such that each element is strictly smaller than another element of U.

The formula rOP_n has variables u_i for $i \in [n]$ ("the i-th element is in U"), x_{ij} for pairwise distinct $i, j \in [n]$ ("the i-th element is below the j-th in the ordering"), and z_{ij} for pairwise distinct $i, j \in [n]$ ("the j-th element is a designated element above the i-th"). The clauses are:

$$\begin{array}{ll}
\bigvee_{i=1}^{n} u_i, & \\
\overline{u_i} \vee \overline{u_j} \vee x_{ij} \vee x_{ji} & \text{for distinct } i, j \in [n], \\
\overline{u_i} \vee \overline{u_j} \vee \overline{x_{ij}} \vee \overline{x_{ji}} & \text{for distinct } i, j \in [n], \\
\overline{u_i} \vee \overline{u_j} \vee \overline{u_k} \vee \overline{x_{ij}} \vee \overline{x_{jk}} \vee x_{ik} & \text{for pairwise distinct } i, j, k \in [n], \\
\bigvee_{j \neq i} z_{ij} & \text{for each } i \in [n], \\
\overline{u_i} \vee \overline{z_{ij}} \vee u_j & \text{for distinct } i, j \in [n], \\
\overline{u_i} \vee \overline{z_{ij}} \vee x_{ij} & \text{for each } i \in [n].
\end{array}$$

Prove that rOP_n requires size $2^{\Omega(n)}$ to refute in Resolution.

Problem 4. Res(2) is a system that operates like Resolution except that each line is a set (understood as a disjunction) of formulas that are either literals or conjunctions of 2 literals. In addition to the resolution rule and weakening, the rules are as follows:

$$\frac{C, \ell_1 \qquad D, \ell_2}{C, D, \ell_1 \wedge \ell_2}\ (\wedge\text{-introduction}) \qquad\qquad \frac{C, \ell_1 \wedge \ell_2 \qquad D, \overline{\ell_1}, \overline{\ell_2}}{C, D}\ (\text{cut})$$

(We could reformulate this so that resolution would be a special case of cut.) It is also allowed to use lines of the form $x, \overline{x}$ for a variable x as axioms.

For a CNF formula F in variables $x_1, \ldots, x_n$, let $F(2)$ be the CNF formula obtained by adding, for each pair $\{\ell_1, \ell_2\}$ of literals over the x variables, new variables y_{ℓ_1,ℓ_2}, and adding the following clauses to the clauses of F:

$$\overline{\ell_1}, \overline{\ell_2}, y_{\ell_1,\ell_2} \qquad\qquad \overline{y_{\ell_1,\ell_2}}, \ell_1 \qquad\qquad \overline{y_{\ell_1,\ell_2}}, \ell_2$$

Prove that for unsatisfiable F, the size of the smallest Res(2) refutation of F and the size of the smallest Resolution refutation of $F(2)$ are polynomially bounded in terms of one another. (*Bonus points.* Show how to conclude from this that Res(2) is weakly automatable iff Resolution is weakly automatable.)

Problem 5. Show that the formula rOP_n from Problem 3 has a $\mathrm{poly}(n)$-size refutation in Res(2).

Problem 6. The *bijective pigeonhole principle* bijPHP_n^{n+1} (also known as the *functional and onto* PHP or the *perfect matching principle* on $[n+1] \sqcup [n]$) is a CNF that claims to describe a bijection between a set of $n+1$ pigeons and a set of n holes. The variables of bijPHP_n^{n+1} are x_{ij} for $i \in [n+1], j \in [n]$ ("pigeon i goes to hole j").

We rewrite this formula as a set of polynomials in Polynomial Calculus with Resolution (recall that a line p in PCR is intended to say that $p = 0$). In that context, bijPHP_n^{n+1} consists of the following polynomials:

$$\begin{array}{ll}
\prod_{j=1}^{n} \overline{x_{ij}} & \text{for each } i \in [n+1], \\
\prod_{i=1}^{n+1} \overline{x_{ij}} & \text{for each } j \in [n], \\
x_{ij}x_{i'j} & \text{for distinct } i, i' \in [n+1] \text{ and each } j \in [n], \\
x_{ij}x_{ij'} & \text{for each } i \in [n+1] \text{ and distinct } j, j' \in [n].
\end{array}$$

Prove that bijPHP_n^{n+1} has a $\mathrm{poly}(n)$-size refutation in PCR over $\mathbb{F}_2$. (You can assume that the size of a refutation is defined to be the total number of monomials appearing in the refutation.)
(*Hints.* It may be useful to refute the polynomials $1 - \sum_{j=1}^{n} x_{ij}$ and $1 - \sum_{i=1}^{n+1} x_{ij}$ and then derive them from the axioms of bijPHP_n^{n+1}. In the latter step of the construction, it may also be useful to think of the refutation as a Prover strategy, with the Prover's knowledge at each point of the game being $p = 1$ (that is $p \neq 0$) for some polynomial p.)
(*Bonus points.* Possibly using yet another variant of bij-PHP, give an example of a family of k-CNF's for some fixed k that have constant-degree refutations in PCR over $\mathbb{F}_2$, but require non-constant width to refute in Resolution.)

Problem 7. Let Res^+ be a system that has the same lines as Resolution, but has the following stregthened resolution rule (with an arbitrary number of premises):

$$\frac{C, \ell_1, \ell_2, \ldots, \ell_k \qquad D, \overline{\ell_1} \qquad D, \overline{\ell_2} \qquad \cdots \qquad D, \overline{\ell_k}}{C, D}\ (\mathrm{res}^+)$$

In other words, we can resolve w.r.t. many literals at once. We also have weakening.

Assume that the CNF formula F has a treelike Resolution refutation of size s. Prove that F also has a Res^+ refutation of *height* $O(\log s)$: that is, the length of the longest path from an initial clause to $\perp$ in that Res^+ refutation is $O(\log s)$.

Problem 8. Let ψ be a first-order sentence of the form

$$\forall x_1 \ldots x_k\, (\eta(\vec{x}) \to \exists y_1 \ldots y_\ell\, \delta(\vec{x}, \vec{y})),$$

where η, δ are quantifier-free formulas in the signature with one binary relation symbol R. Additionally, we assume that δ is just a disjunction of atoms and negated atoms.

The n-th propositional translation $\langle\psi\rangle_n$ of ψ is a propositional formula that is satisfiable exactly if ψ has a model with universe $\{1, \ldots, n\}$. For each $i, j \in [n]$, the formula $\langle\psi\rangle_n$ has a variable p_{ij} to state that $R(i, j)$ holds, so for any assignment $\vec{x} \mapsto \vec{a}, \vec{y} \mapsto \vec{b}$ the statement $\eta(\vec{a}) \to \delta(\vec{a}, \vec{b})$ can be rewritten as a propositional formula simply by replacing each $R(i, j)$ appearing in it by p_{ij}. The existential quantification over $\vec{y}$ is then rewritten as a disjunction, and the universal quantification over $\vec{x}$ becomes a conjunction.

Note that $\langle\psi\rangle_n$ can be stated as a $\mathrm{poly}(n)$-size CNF: we can assume that η is written as a DNF, so the procedure described above turns each $\neg\eta(\vec{a})$ into a constant-size CNF and thus each $\eta(\vec{a}) \to \exists\vec{y}\, \delta(\vec{a}, \vec{y})$ into a constant-size $\bigwedge$ of polysize $\bigvee$'s; we then take a polysize conjunction of such formulas.

Assume that ψ is false in all finite structures but has an infinite model. Prove that the formulas $\langle\psi\rangle_n$ do not have poly(n)-size refutations in treelike Resolution.
(*Hints.* It may be useful (and is allowed, even if you do not do Problem 7) to use the statement of Problem 7 as a lemma. As a guiding example, you may think of the ordering principles OP_n we considered in class: modulo cosmetic details, OP_n is $\langle\psi\rangle_n$ for ψ saying "R is a linear order with no smallest element". But whether this is helpful is probably up to individual taste.)
(*Bonus points.* Think about the formulas $\mathrm{CPLS}_{n,1}$; cf. Problem 2. Clearly, the model-theoretic criterion from the present problem could be generalized to ψ using more relations of various arities, and just as clearly, one could define an infinite "matrix" with rows and columns indexed by say $1, 2, 3, \ldots, \ldots, -4, -3, -2, -1$ in that order (so that -1 is "n"), in which all cells with positive row indices do not have the unique colour, all cells with negative row indices do, and the teacher relation connects cells in the same column in neighbouring rows. Given this, explain why the criterion does not apply to $\mathrm{CPLS}_{n,1}$.)