KTH Computer Science and Communication

DD2442 Proof Complexity: Problem Set 2

Due: Monday October 31, 2016, at 23:59 AoE. Submit your solutions as a PDF file by e-mail to `jakobn at kth dot se` with the subject line `Problem set 2: ⟨your full name⟩`. Name the PDF file `PS2_⟨YourFullName⟩.pdf` with your name written in CamelCase without blanks and in ASCII without national characters. State your name and e-mail address at the very top of the first page. Solutions should be written in LaTeX or some other math-aware typesetting system with reasonable margins on all sides (at least 2.5 cm). Please try to be precise and to the point in your solutions and refrain from vague statements. *Write so that a fellow student of yours can read, understand, and verify your solutions.* In addition to what is stated below, the general rules stated on the course webpage always apply.

Collaboration: Discussions of ideas in groups of two people are allowed—and indeed, encouraged—but you should write down your own solutions individually and understand all aspects of them fully. You should also acknowledge any collaboration. State at the very top of the first page of your problem set solutions if you have been collaborating with someone and if so with whom. *Note that collaboration is on a per problem set basis, so you should not discuss different problems on the same problem set with different people.*

Reference material: Some of the problems are "classic" and hence it might be easy to find solutions on the Internet, in textbooks or in research papers. It is not allowed to use such material in any way unless explicitly stated otherwise. Anything said during the lectures or in the lecture notes should be fair game, though, unless you are specifically asked to show something that we claimed without proof in class. All definitions should be as given in class and cannot be substituted by versions from other sources. It is hard to pin down 100% watertight formal rules on what all of this means—when in doubt, ask the main instructor.

About the problems: Some of the problems are meant to be quite challenging and you are not necessarily expected to solve all of them. On the contrary, you can chooose to solve just a subset of the problems and still get a top grade. A total score of around 80 points should be enough for grade E, 110 points for grade D, 140 points for grade C, 170 points for grade B, and 200 points for grade A on this problem set. Any corrections or clarifications will be given at `piazza.com/kth.se/fall2016/dd2442/` and any revised versions will be posted on the course webpage `www.csc.kth.se/DD2442/semteo16/`.

1 (10 p) Suppose that for a proof system $\mathcal{P}$ we can prove a lower bound for the length of refutations of uniformly randomly sampled 3-XOR formulas. Prove that this implies at least as good a lower bound in $\mathcal{P}$ for the length of refutations of uniformly randomly sampled 3-CNF formulas.

Hint: Note that you only know that one kind of random formulas sampled according to one distribution are hard, and now you want to show that another kind of formulas sampled from a slightly different distribution are also hard. Please make sure to get the formal argument correct.

2 (10 p) Prove for DNF formulas $G_1, \ldots, G_s, H$ that if $G_1, \ldots, G_s$ strongly imply H in the sense of Lecture 7, then $G_1, \ldots, G_s$ imply H in the usual sense (i.e., $G_1 \wedge \cdots \wedge G_s \vDash H$). Then show that the opposite direction does not hold. Please make sure to motivate your line of reasoning.

3 (20 p) The purpose of this problem is to make more formal the claim that "restrictions preserve refutations" (which will be true for all proof systems we encounter in this course). In what follows below, let $\rho : \mathit{Vars}(F) \to \{0, 1, *\}$ denote a restriction, i.e., a partial assignment, to the variables of an unsatisfiable CNF formula F.

3a (10 p) Prove that if π is a d-DNF resolution refutation of F, then $\pi\restriction_\rho$ can easily be made into a refutation of $F\restriction_\rho$ in at most the same length. Is it in fact the case that $\pi\restriction_\rho$ is already a syntactically legal d-DNF resolution refutation even without any modification? Here $\pi\restriction_\rho$ is defined so that any proof line containing a d-DNF formula with a satisfied term is removed, and in all other lines falsified terms are removed, after which satisfied literals are removed.

3b (10 p) Prove that if π is a cutting planes refutation of F, then $\pi\restriction_\rho$ can easily be made into a refutation of $F\restriction_\rho$ in at most the same length. Is it in fact the case that $\pi\restriction_\rho$ is a legal cutting planes refutation even without any modification? Here $\pi\restriction_\rho$ is defined so that in a linear inequality $\sum_i A_i x_i \geq B$ the values $\rho(x_i) \in \{0, 1\}$ are substituted for assigned variables x_i, after which all constants are moved to the right-hand side to yield $\sum_{i,\rho(x_i)=*} A_i x_i \geq B - \sum_{i,\rho(x_i)=1} A_i$.

4 (30 p) In the first lecture, we briefly discussed the proof system polynomial calculus. The purpose of this problem is to establish that polynomial calculus is an implicationally complete proof system for CNF formulas and to compare it to resolution. Recall that if we write any clause C as $C = C^+ \vee C^-$, where C^+ contains all positive (unnegated) literals of C and C^- contains all negative (negated) literals, then in polynomial calculus we translate the axiom clauses $C \in F$ to polynomials $p(C) = \prod_{x \in C^+} x \cdot \prod_{\overline{y} \in C^-} (1 - y)$ expanded out as linear combinations of monomials in some (fixed) field $\mathbb{F}$. We also have Boolean axioms $x^2 - x$ for all variables x. The size of a polynomial is the number of monomials when the polynomial is written as a linear combination of monomials with coefficients in $\mathbb{F}$, and the size of a polynomial calculus refutation (i.e., a derivation of the multiplicative identity 1 in the field $\mathbb{F}$) is the sum of the sizes of all polynomials in it (including any Boolean axioms used).

4a (20 p) Show that if a resolution refutation uses the resolution rule

$$\frac{B \vee x \qquad C \vee \overline{x}}{B \vee C} \ ,$$

then polynomial calculus can derive $p(B \vee C)$ from $p(B \vee x)$ and $p(C \vee \overline{x})$. Explain why it follows that polynomial calculus can refute any unsatisfiable CNF formula F.

4b (10 p) Is it true that the simulation outlined in Problem 4a shows that in fact polynomial calculus polynomially simulates resolution? Fill in the necessary details to prove that this is so, or explain why this fails to be a polynomial simulation.

5 (30 p) In Lecture 8 we saw the theorem that if a d-DNF resolution refutation $\pi = (H_1, H_2, \ldots, H_L)$ of a h-CNF formula F has the property that all lines H_i can be strongly represented by decision trees of height h, then there is a resolution refutation of F where all clauses have width at most $2h$. We proved this by arguing that a decision tree of height h can be represented by a CNF formula of width h, and then showing that the sets of CNF formulas corresponding to each H_i could be stringed together into a resolution refutation with a bit of extra work.

It is a natural question whether we really need the decision trees here, or whether they just happen to be a by-product of the switching lemma used for d-DNF resolution and we would be equally fine in this particular part of our overall lower bound construction if we only knew that the lines in the d-DNF resolution refutation could be represented efficiently as CNF formulas. Your task is to shed light on this.

To formalize the question we want to understand, suppose that F is a h-CNF formula such that there is a d-DNF resolution refutation $\pi = (H_1, H_2, \ldots, H_L)$ of F where each line H_i can be represented as a h-CNF formula. Under this assumption, can the proof we did in class be adapted to show that F has a resolution refutation in width $\mathrm{O}(h)$?

If your answer is yes, please explain clearly how to modify the proof given in class to establish this claim. You do not need to give all the details, but a fellow student of yours who understood the proof of the simulation of decision trees by resolution should be able to piece together a proof of the new claim from your description. If your answer is no, then you do not necessarily have to show unconditionally that the claim is false, but you need to argue clearly which parts of the proof seem challenging or impossible to extend to the new setting and why.

6 (30 p) Suppose that $A(\mathbf{p}, \mathbf{q}) \wedge B(\mathbf{p}, \mathbf{r})$ is an unsatisfiable CNF formula over disjoint sets of variables $\mathbf{p}, \mathbf{q}, \mathbf{r}$. In class we showed monotone feasible interpolation theorems first in the case when all the $\mathbf{p}$-variables appear only positively in $A(\mathbf{p}, \mathbf{q})$ for resolution, and then in the case when all the $\mathbf{p}$-variables appear only negatively in $B(\mathbf{p}, \mathbf{r})$ for cutting planes. In this problem we want you do the proof of the monotone interpolation theorems for the cases that were not covered in class.

6a (15 p) Under the assumption that $\mathbf{p}$-variables appear only negatively in $B(\mathbf{p}, \mathbf{r})$, show that if $A(\mathbf{p}, \mathbf{q}) \wedge B(\mathbf{p}, \mathbf{r})$ has a resolution refutation of length L, then there exists a monotone Boolean interpolating circuit $I(\mathbf{p})$ of size $\mathrm{O}(L)$.

6b (15 p) Under the assumption that $\mathbf{p}$-variables appear only positively in $A(\mathbf{p}, \mathbf{q})$, show that if $A(\mathbf{p}, \mathbf{q}) \wedge B(\mathbf{p}, \mathbf{r})$ has a cutting planes refutation of length L, then there exists a monotone real interpolating circuit $I(\mathbf{p})$ of size $\mathrm{O}(nL)$, where n is the number of variables in the formula.

Remark: Note that there is no need to reproduce the entire proofs covered in class—you can assume that they are all known. Instead, just explain in a precise manner exactly which parts of the proofs need to be modified and how, and then motivate (potentially briefly, but clearly and to the point) why the modified proofs work.

7 (50 p) Let $F = \bigwedge_{i=1}^{m} E_i$ be an XOR formula with linear equations E_i and let G_F be the constraint-variable incidence graph as defined in class (i.e., G_F is a bipartite graph with left vertices labelled by constraints, right vertices labelled by variables, and edges corresponding to variable occurrences in constraints). Let F' denote the canonical encoding of F as a CNF formula. The purpose of this problem is to prove that if G_F is an (r, c)-boundary expander for some $c > 0$, then F' requires refutation width $W(F' \vdash \bot) \geq rc/2$ in resolution.

7a (20 p) For any clause C, define the complexity measure

$$\mu(C) = \min\left\{|I| \,:\, I \subseteq [m],\, \textstyle\bigwedge_{i \in I} E_i \vDash C\right\}$$

to be the minimal size of any subset of linear equations $\{E_i \mid i \in I\} \subseteq F$ such that any assignment α that satisfies $\bigwedge_{i \in I} E_i$ must also satisfy C. Prove that $\mu(C) = 1$ for any clause C in the CNF formula F' encoding the XOR formula F and that $\mu(\bot) \geq r$.

7b (10 p) Prove that the measure μ is *subadditive*, i.e., that for any clause $B \vee C$ derived by resolving clauses $B \vee x$ and $C \vee \overline{x}$ it holds that $\mu(B \vee C) \leq \mu(B \vee x) + \mu(C \vee \overline{x})$. Then show that this implies that every resolution refutation $\pi : F' \vdash \bot$ must contain a clause D such that $r/2 \leq \mu(D) \leq r$.

7c (20 p) Prove that for any clause D such that $\mu(D) \leq r$ it holds that $W(D) \geq c \cdot \mu(D)$, and use this to show that $W(F' \vdash \bot) \geq cr/2$.

8 (70 p) In our proofs of the length lower bounds for d-DNF refutations of PHP_n^{cn} and random 3-XOR formulas we claimed that refutations of these formulas can be assumed to be in different kinds of *normal form* without loss of (too much) generality. The purpose of this problem is to formalize and establish these claims.

8a (30 p) Prove that if the formula PHP_n^{cn} has a d-DNF resolution refutation in length L, then it also has a d-DNF resolution refutation in length $2L$ where all lines are in pigeon-normal form. Recall that we say that a term t is in *pigeon-normal form* if it does not contain two positive literals $x_{i,j}$ and $x_{i',j}$ for $i \neq i'$, and that a DNF formula is in pigeon-normal form if all terms in it are in pigeon-normal form.

8b (40 p) Let $F = \bigwedge_{i=1}^{\Delta n} E_i$ be a 3-XOR formula with Δn linear constraints E_i over n variables and let F' be the canonical encoding of F in 3-CNF. Suppose that the constraint-variable incidence graph G_F is an (r, c)-boundary expander for some $c > 1/2$.

Prove that if F' has a d-DNF resolution refutation in length L for $d \leq \min\left\{\frac{r}{4}, \log n\right\}$, then F' has a d-DNF resolution refutation in length $L \cdot n^{\mathrm{O}(1)}$ where all lines are in XOR-normal form. We say that a term t is in *XOR-normal form* if $t \wedge \bigwedge_{i \in supp(t)} E_i$ is satisfiable (where we refer to the lecture notes for the definition of $supp(t)$), and that a DNF formula is in XOR-normal form if all its terms are in XOR-normal form.

Hint: It might be helpful to recall that any literal a is locally consistent since $supp(a) = \emptyset$.

9 (90+ p) Given an undirected graph $G = (V, E)$ and a parameter $k \in \mathbb{N}^+$, we can define a formula $\mathit{Clique}(G, k)$ that encodes the claim that G has a k-clique in the following way. The variables of $\mathit{Clique}(G, k)$ are $x_{i,v}$ for $i \in [k]$ and $v \in V$, where the intended meaning is that $x_{i,v}$ is true if vertex v is the ith member of the k-clique, and the clauses are

$$\bigvee_{v \in V} x_{i,v} \qquad i \in [k]; \tag{1}$$

$$\overline{x}_{i,u} \vee \overline{x}_{i,v} \qquad i \in [k] \text{ and } u, v \in V,\ u \neq v; \tag{2}$$

$$\overline{x}_{i,u} \vee \overline{x}_{j,v} \qquad i, j \in [k],\ i \neq j, \text{ and } u, v \in V,\ (u, v) \notin E; \tag{3}$$

where clauses (1) and (2) say that the ith member of the clique is some unique vertex $v_i \in V$ for $i \in [k]$ and clauses (3) say that two vertices u and v cannot both be members of the clique if there is no edge between them. (Note that in contrast to the clique-colouring formulas, here there is a concrete graph G hardcoded into the formula by the clauses of type (3)).

A reasonably well-established hypothesis in computational complexity theory is that for any constant $k \in \mathbb{N}^+$ deciding whether a graph $G = (V, E)$ with $|V| = n$ contains a k-clique cannot be done faster than $n^{\Omega(k)}$ (i.e., there is no fundamentally better way than to exhaustively check all $\binom{n}{k}$ k-clique candidates). It is an interesting question whether such a lower bound can be established *unconditionally* for weak computational models such as conflict-driven clause learning (CDCL) SAT solvers and the resolution proof system in which these solvers search for proofs, and this is the topic of this problem.

Hint: In what follows, you are encouraged to use the fact mentioned in class that general resolution is equivalent to Pudlák's Prosecutor-Defendant game. You can also use the equivalence proven in problem set 1 between tree-like resolution and decision trees. In this context it might be helpful to think of the decision tree as a Prosecutor who never forgets anything and the values of the queries as Defendant answers (although if you want to claim a formal equivalence you need to prove it).

9a (40 p) Consider formulas $\mathit{Clique}(G_n, k)$ defined over complete $(k-1)$-partite graphs, i.e., graphs over $n = (k-1)n'$ vertices with $V = V_1 \dot\cup V_2 \dot\cup \cdots \dot\cup V_{k-1}$ for $|V_i| = n'$ and with edge set $E = \bigcup_{1 \leq i < j \leq k-1} \{(v_i, v_j) \mid v_i \in V_i, v_j \in V_j\}$. A moment of thought reveals that such graphs do not contain k-cliques. Prove that tree-like resolution proofs require length $n^{\Omega(k)}$ to establish this fact.

9b (50 p) Prove that formulas $\mathit{Clique}(G_n, k)$ defined over complete $(k-1)$-partite graphs are easy to refute for general resolution in that they only require refutations of length $n^{\mathrm{O}(1)}$.

9c (300+ p) **Open problem:** Find some family of graphs $\{G_n\}_{n=1}^{\infty}$ on n vertices that do not contain k-cliques but are such that the refutation length of formulas $\mathit{Clique}(G_n, k)$ in general resolution grows like $n^{\Omega(k)}$, or indeed like $n^{\omega_k(1)}$ for any arbitrarily slowly growing but unbounded function $\omega_k(1)$ of k. If it helps, you can also omit the clauses (2) from the formulas (removing these clauses can only make the formulas harder to refute, although it is clear that they are still unsatisfiable).