KTH Computer Science and Communication

DD2442 Proof Complexity: Problem Set 3

Due: Monday January 23, 2017, 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 3: ⟨your full name⟩`. Name the PDF file `PS3_⟨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) Let $k(C)$ denote the minimum Hamming weight of a satisfying assignment for a monotone circuit C. Prove the following "self-improvement" property of monotone circuit that was needed for the Alekhnovich-Razborov non-automatizability result we covered in class: For any fixed $d \in \mathbb{N}^+$ there is a polynomial-time computable function f_d that maps monotone circuits to monotone circuits in such a way that $k\big(f_d(C)\big) = (k(C))^d$.

2 (20 p) Recall that a *bridge* in an undirected graph $G = (V, E)$ is an edge $e \in E$ such that the number of (maximal) connected components in $G \setminus \{e\} = (V, E \setminus \{e\})$ is larger than that in G. In the lower bounds we studied for Tseitin formulas we implicitly used that the set of bridges of a graph G is independent of the order in which the bridges are identified. Prove formally that this is indeed so. That is, prove that assuming that e_1 is a bridge in $G = (V, E)$, then $e_2 \in E \setminus \{e_1\}$ is a bridge in $G \setminus \{e_1\}$ if and only if e_2 is a bridge in G.

3 (20 p) Let $\mathcal{F}$ denote a Frege proof system over $\{\vee, \neg\}$ and let d be a positive integer. We proved in class that for large enough n any depth-d refutation of PHP_n^{n+1} in $\mathcal{F}$ requires size at least $\exp\big(n^{6^{-d}}\big)$. For weaker proof systems such as resolution and k-DNF resolution we have seen earlier in the course that the lower bound proofs are quite robust in that they hold even for cn pigeons being mapped into n holes for any constant $c > 1$. In contrast, even for just depth-2 Frege it has been shown that PHP_n^{2n} can be refuted in at most quasi-polynomial size $\exp\big((\log n)^{O(1)}\big)$.

Thus, there must be one place (or more) in the lower bound for bounded-depth Frege refutations of formulas PHP_n^{n+1} where the argument critically fails if we try to adapt it to PHP_n^{2n}. Your task is to point out clearly where and why.

4 (30 p) Suppose that G is an undirected, connected graph and that $\chi : V(G) \to \{0, 1\}$ is a charge function. Recall that χ is said to have odd charge if $\sum_{v \in V(G)} \chi(v)$ is an odd number.

4a Prove that for any two odd-charge functions χ and χ' the Tseitin formulas $Ts_{G,\chi}$ and $Ts_{G,\chi'}$ are equivalent in the formal sense that any resolution refutation $\pi : Ts_{G,\chi} \vdash \bot$ can be transformed into a resolution refutation $\pi' : Ts_{G,\chi'} \vdash \bot$ of exactly the same length by simple syntactic manipulations.

4b Prove that $Ts_{G,\chi}$ is unsatisfiable if and only if χ has odd charge.

5 (40 p) As discussed in class, a *projection* ρ is a generalization of a restriction where each variable x can be mapped by ρ to not only 0 or 1, but also y or $\overline{y}$ for some variable y. We made the somewhat handwavy claim in class that projections work mostly like restrictions in that they preserve refutations in well-behaved proof systems. The purpose of this problem is to make this claim slightly more precise.

Recall that to any resolution refutation $\pi : F \vdash \bot$ we can associate a DAG G_π with vertices labelled by the clauses in the refutation and with edges from resolved clauses to resolvents. We say that π is *tree-like* if G_π is a tree. The refutation π is said to be *regular* if along any path in G_π from a source (i.e., axiom clause) to the sink (i.e., the empty clause $\bot$, which can be assumed to be the only sink of G_π without loss of generality) every variable is resolved over at most once. *Tree-like resolution* and *regular resolution* are the two subsystems of resolution where refutations are restricted to be tree-like or regular, respectively. In *general resolution* there are no restrictions on G_π. We want to understand how projections affect refutations in these three flavours of resolution.

5a Is it true for any projection ρ and any tree-like resolution refutation π of a CNF formula F that $\pi\restriction_\rho$ is a tree-like resolution refutation of $F\restriction_\rho$?

5b Is it true for any projection ρ and any regular resolution refutation π of F that $\pi\restriction_\rho$ is a regular refutation of $F\restriction_\rho$?

5c Is it true for any projection ρ and any general resolution refutation π of F that $\pi\restriction_\rho$ is a resolution refutation of $F\restriction_\rho$?

For each of the subproblems above, either give a formal proof that $\pi\restriction_\rho$ is a refutation of $F\restriction_\rho$ of the required form (possibly using the weakening rule), or explain clearly why this is not necessarily the case.

6 (40 p) Let $\mathbf{p}$, $\mathbf{q}$, and $\mathbf{r}$, be disjoint set of variables and let $F = A(\mathbf{p},\mathbf{q}) \wedge B(\mathbf{p},\mathbf{r})$ be an unsatisfiable CNF formula such that the $\mathbf{p}$-variables occur only positively in A. Consider the Karchmer-Wigderson game where Alice is given an assignment $\alpha : \mathbf{p} \to \{0,1\}$ such that $A(\alpha,\mathbf{q})$ is satisfiable, Bob is given an assignment $\alpha' : \mathbf{p} \to \{0,1\}$ such that $B(\alpha',\mathbf{r})$ is satisfiable, and their task is to communicate to find an index i such that $\alpha_i = 1$ and $\alpha'_i = 0$.

Prove from first principles (i.e., without using any theorems stated on the board during Pavel Pudlák's guest lecture) that if F has a resolution refutation in depth d (i.e., such that the longest path in the refutation DAG G_π has length d), then there is a deterministic two-party protocol with communication $\mathrm{O}(d)$ that solves this Karchmer-Wigderson game. Recall that such a protocol can be described as a binary tree where Alice and Bob start at the root node, where every node is labelled by whose turn it is to send a bit b, where the bit b sent by Alice (or Bob) is a function only of Alice's (Bob's) input and of the bits sent so far, where there are directed edges labelled 0 and 1 such that Alice and Bob follow the branch labelled b, and where every leaf is labelled by an answer that is correct for any pair of inputs to Alice and Bob that reach that leaf in the protocol tree. Such a protocol achieves a communication upper bound c if the longest path in the tree has length at most c.

Hint: Depending on how one thinks about this, the argument might become slightly more straightforward if one instead considers the modified formula $F' = A(\mathbf{p},\mathbf{q}) \wedge B(\mathbf{p}',\mathbf{r}) \wedge \bigwedge_i (\overline{p}_i \vee p'_i)$ where every occurrence of p_i in $B(\mathbf{p}',\mathbf{r})$ is replaced by p'_i in $B(\mathbf{p}',\mathbf{r})$ and where Bob is given an assignment $\alpha' : \mathbf{p}' \to \{0,1\}$ such that $B(\alpha',\mathbf{r})$ is satisfiable, and then studies resolution refutations of this formula. However, if this hint confuses you more than it helps you, then you should feel perfectly free to ignore it.

7 (40 p) Returning to the bounded-depth Frege lower bound for PHP formulas over a set $V = P \dot\cup H$ with $|P| = n+1$ and $|H| = n$, let η denote a k-evaluation for a set of formulas Γ and let α be a partial matching in V. Prove formally the claim that was hand-waved in class that $\eta\restriction_\alpha$ is a k-evaluation for the set $\Gamma\restriction_\alpha$. Recall that for a nontrivial restricted formula $F\restriction_\alpha \in \Gamma\restriction_\alpha$ we define $\eta\restriction_\alpha (F\restriction_\alpha)$ as $\eta(F)\restriction_\alpha$, where $\eta(F)\restriction_\alpha$ denotes the restriction over trees as defined in lectures 13–16.

In particular, this means that you have to prove the following properties of matching decision trees T under restrictions (with notation as used in the relevant lectures).

7a If T is a complete tree for the set of vertices V, then $T\restriction_\alpha$ is complete for $V\restriction_\alpha$. Recall that $V\restriction_\alpha$ is the set of all vertices $v \in V$ not matched by α.

7b $\mathit{Disj}(T)\restriction_\alpha = \mathit{Disj}(T\restriction_\alpha)$.

7c $T^c\restriction_\alpha = (T\restriction_\alpha)^c$. Recall that T^c denotes the operation over trees that consists of changing the labels of the leaves of T from 0 to 1 and vice versa.

7d If T represents a matching disjunction G, then $T\restriction_\alpha$ represents $G\restriction_\alpha$.

8 (100 p) Recall the formula $\tau(C, \mathcal{A}, \mathbf{F}, \mathbf{f})$ in the Alekhnovich-Razborov non-automatizability result as described in lectures 22–24, where $C(p_1, \ldots, p_n)$ is a monotone circuit, $\mathcal{A} \subseteq \{0,1\}^m$ is a set of vectors, and $F_j : \{0,1\}^s \to \mathcal{A}$ and $f_i : \{0,1\}^s \to [r]$ are onto functions, where functions f_i are possibly partial. Recall also the notational shorthands $[\mathrm{col}_j = \mathbf{a}]$ denoting "$F_j\big(x_j^1, \ldots, x_j^s\big) = \mathbf{a}$" and $[\mathrm{ctrl}_i = c]$ denoting "$f_i\big(y_i^1, \ldots, y_i^s\big)$ is defined and equals c", using which the formula $\tau(C, \mathcal{A}, \mathbf{F}, \mathbf{f})$ could be described as consisting of encodings into CNF of the following conditions:

$$\big(y_i^1, \ldots, y_i^s\big) \in \mathrm{Dom}(f_i) \qquad [i \in [m]] \tag{1a}$$

$$\big([\mathrm{col}_j = \mathbf{a}] \wedge [\mathrm{ctrl}_i = c]\big) \to z_{i,p_j}^c \qquad [c \in [r],\, j \in [n],\, \mathbf{a} \in \mathcal{A} \text{ and } i \in [m] \text{ s.t. } \mathbf{a}_i = 1] \tag{1b}$$

$$\big([\mathrm{ctrl}_i = c] \wedge (z_{i,u}^c \circ z_{i,v}^c)\big) \to z_{i,w}^c \qquad [c \in [r],\, i \in [m],\, \text{gate } w = u \circ v \text{ for } \circ \in \{\wedge, \vee\}] \tag{1c}$$

$$[\mathrm{ctrl}_i = c] \to \overline{z}_{i,v_{\mathrm{out}}}^c \qquad [c \in [r],\, i \in [m],\, v_{\mathrm{out}} \text{ output gate of } C] \tag{1d}$$

We refer to the lecture notes for a more detailed discussion of exactly what the notation above means and of the combinatorial principle encoded.

The purpose of this problem is to consider other possible encodings of the same combinatorial principle, and what would happen to the upper and lower bounds we proved in class if these other encodings are used instead. For all of the subproblems below, your task is to answer the follow questions for both the upper and the lower bound shown in class:

1. Does the bound still hold for the new suggested encoding with essentially the same proof as we did in class (except possibly for minor, obvious fixes)?
2. If your answer to question 1 is yes, explain briefly what small fixes are needed, if any, and why the proof goes through.
3. If your answer to question 1 is no, then point out where the proof breaks.
4. In case of an answer no, is it possible to give a different argument that can recover all of, or parts of, the bound, or do there seem to be more fundamental problems making it hard to see how such a result could be obtained?

8a Consider an encoding where for conditions (1b)–(1d) we insist on equivalence $\leftrightarrow$ instead of implication $\to$ (so that the circuit evaluation is forced to be correct and cannot as before introduce spurious 1s along the way).

8b Consider an encoding where we eliminate the double indexing over rows i and circuit copies c, and instead just specify that every row i gets evaluated in its own circuit copy c as indicated by the y-variables. That is, we add conditions that every row i gets its unique, own circuit copy c not used by any other row, and then remove indices i from variables $z_{i,v}^c$ in conditions (1b)–(1d).

8c Consider an encoding where there are no circuit copies, and hence no clauses (1a). Instead, the value $z_{i,v}$ of a circuit gate v when evaluated on row i is encoded redundantly by a function $g_{i,v} : \{0,1\}^s \to \{0,1\}$ (for instance, one concrete encoding could be to to set $g_{i,v}(\sigma) = \bigoplus_i \sigma_i$, so that $z_{i,v}$ is the exclusive or of the bits in the input to $g_{i,v}$).

We then introduce the shorthands

$$[z_{i,v}] = \bigwedge_{\sigma \in g_{i,v}^{-1}(0)} \left(\left(z_{i,v}^1\right)^{1-\sigma_1} \vee \left(z_{i,v}^2\right)^{1-\sigma_2} \vee \cdots \vee \left(z_{i,v}^s\right)^{1-\sigma_s} \right)$$

and

$$[\overline{z}_{i,v}] = \bigwedge_{\sigma \in g_{i,v}^{-1}(1)} \left(\left(z_{i,v}^1\right)^{1-\sigma_1} \vee \left(z_{i,v}^2\right)^{1-\sigma_2} \vee \cdots \vee \left(z_{i,v}^s\right)^{1-\sigma_s} \right)$$

to denote the possible values of gate v when evaluated on row i, and, for instance, the encoding of (1c) would become simply $\left([z_{i,u}] \circ [z_{i,v}]\right) \to [z_{i,v}]$ expanded out in CNF.

Remark: It goes without saying that there is some flexibility as to exactly how the questions above can be interpreted, and thus the emphasis in the grading will be on assessing how convincingly, and in how appropriate a level of detail, you present your arguments for the different encodings, rather than on your exact wording of the final answers to the questions.

As a general rule, you should not expect to have to write pages and pages of detailed arguments to answer to the questions above in the cases where you believe the proofs still work or can be adapted to work. Also, when it seems that a proof does not work, and perhaps cannot even be fixed, you do not have to prove beyond all doubt that no way of formalizing an argument along similar lines can possibly work in any universe. It is enough to point out, briefly but concretely, what technical difficulties arise, and, when applicable, why they seem hard to circumvent.

Hint: You should expect to have to study the handwritten notes for lectures 22–24 in detail in order to be able to solve this problem, including (and especially) the part of the notes that we did not have time to cover in detail in class.

9 (150 p) Prove that Tseitin formulas over $n \times n$ rectangular grids require resolution refutations of length $\exp(\Omega(n))$. For partial credit, you can instead explain what the problem is if one tries a straightforward adaptation of the the approach in Lecture 3 to yield lower bounds for grid graphs.

Hint: Use projections.

More hints: For this problem, an additional hint in the form of an outline for how to prove the lower bound can be purchased at a cost of 30 points. In this way, you can configure yourself whether you want this problems to be more creative and open-ended, where a lot might depend on finding the right idea, or whether you want it to be more of a guided exercise providing a useful work-out. If you do not solve this problem, there is no charge for the hint (i.e., it is not deducted from the score on other problems). Contact the main instructor via Piazza if you want to buy a hint.