QBF Resolution Systems and their Proof Complexities

Complementary material

Valeriy Balabanov[1], Magdalena Widl[2], and Jie-Hong R. Jiang[1]

[1]National Taiwan University, [2]Vienna University of Technology
{balabasik@gmail.com, widl@kr.tuwien.ac.at, jhjiang@ntu.edu.tw}

1 Proof of Proposition 1 and discussion on mandatory universal reduction

Below we restate Proposition 1 from Section 3.1 and present a proof.

Proposition 1. *Given a false QBF $\Phi = Q_1 v_1 \,..\, Q_k v_k.\ C_1 \wedge .. \wedge C_j \wedge .. \wedge C_n$ over the set V of variables. Let $\Phi^* = Q_1 v_1 \,..\, Q_k v_k.\ C_1 \wedge .. \wedge C_j^* \wedge .. \wedge C_n$, where clause $C_j^* = (C_j \cup \{\mathsf{lit}(v)\})$ for some $v \in V$. If Φ^* is false, then the shortest* Q*-resolution proof for Φ^* is at least as large as that for Φ.*

Proof. Consider the resolution step $C = \mathsf{resolve}((p, q, C_1), p, (\overline{p}, C_2)) = (q, C_1, C_2)$ that is part of some Q-resolution proof Π. Observe that regardless of the quantification of the variables, eliminating the literal q from the left clause updates the resolution to $C = \mathsf{resolve}((p, C_1), p, (\overline{p}, C_2)) = (C_1, C_2)$. Hence it decreases or leaves unchanged the number of literals in C, and *does not affect* any resolutions that precede C in Π. Similarly, eliminating the literal p from the left clause allows us to directly assign $C = (q, C_1)$, hence decreasing or leaving unchanged the number of literals in C, and *decreasing* the number of resolutions preceding C in Π by one.

Now consider the clause $C = \mathsf{reduce}((C_1, p, q)) = (C_1, p)$ derived from universal reduction of variable q. Elimination of p results in a smaller clause $C = C_1$. Elimination of q leaves C unchanged. Neither of the two elimination cases affects the number of resolutions preceding C in Π.

We refer to the above construction as *rebuilding rules* applied to the given Q-resolution step whenever an elimination occurs. Rebuilding rules can be easily extended to a simultaneous elimination of several variables.

Consequently, if Π^* is the shortest Q-resolution proof for Φ^*, then after the literal $\mathsf{lit}(v)$ is eliminated from the clause C_j^*, we can apply the described rebuilding rules to Π^* iteratively, eventually leading to a new Q-resolution proof that is at most as long as Π^*. Hence there exists a Q-refutation proof for Φ that is at most as long as Π^*. □

Please note that the proof of Proposition 1 is not restricted to the $\mathsf{resolve}_\exists$-rule only. Therefore it can be extended to QU-resolution, LQ-resolution, LQU-resolution, and LQU+-resolution as the following proposition states.

Proposition 4. *Given two false QBFs Φ and Φ^* as described in Proposition 1, let Π and Π^* be their respective shortest* QU*-refutations (*LQ*-refutations,* LQU*-refutations, or* LQU+*-refutations). Then $|\Pi^*| \geq |\Pi|$.*

For all the Q-resolution based proof systems in our paper (i.e. {Q, QU, LQ, LQU, LQU+}-resolution), we have followed the assumption that universal reduction is performed whenever possible. If one allows postponing the reduction arbitrarily (as in the definition of universal reduction in reference [19]), it will generalize the aforementioned proof systems and allow a larger number of sound refutations. In the sequel we call a refutation where the reduction of at least one universal variable has been postponed a *postponed refutation.* Postponing, however, cannot lead to shorter refutations in terms of the number of resolutions for any of the {Q, QU, LQ, LQU, LQU+}-resolution proof systems, as the following corollary from Propositions 1 and 4 states.

Corollary 3. *Given a false QBF Φ, let Π be its shortest* QU*-refutation ({*Q*,* LQ*,* LQU*,*LQU+*}-refutation), and let Π^* be its shortest postponed* QU*-refutation ({*Q*,* LQ*,*LQU*,*LQU+*}-refutation). Then $|\Pi^*| \geq |\Pi|$.*

Proof. The corollary directly follows from Propositions 1 and 4 once we apply the rebuilding rules to the clauses in Π^* where postponing occurs. □

In the light of Corollary 3, Theorem 1 holds for postponed QU-refutations as well.

2 Proofs for invariants in Theorem 2

In the following we prove the invariants used in the proof of Theorem 2.

Proposition 5. *Given any* LQ*-resolution proof Π of a formula* KBKF[t]*, any clause $C \in \Pi$ has at most one positive existential literal.*

Proof. First, the statement holds for any clause in the original clause set of KBKF[t]. Now consider any resolution step $C = \mathsf{resolve}((C_1, p), p, (C_2, \overline{p})) = (C_1, C_2)$. If the statement holds for the clauses (C_1, p) and $(C_2, \overline{p})$, then C_1 has no positive existential literals, and C_2 has at most one. Thus C also has at most one positive existential literal. By induction, any clause C has at most one positive existential literal. □

Lemma 1 (Invariant 1). *Given any* LQ*-resolution proof Π of a formula* KBKF[t]*, the following holds for any clause $C \in \Pi$. For all $i \in [1..t]$, if $f_i \in C$ then $\mathsf{lit}(x_i) \in C$, and if $\overline{f}_i \in C$ then for any $j \in [i..t]$ either $\overline{f}_j \in C$ or $\mathsf{lit}(x_j) \in C$.*

Proof. First, observe that the invariant holds for any clause in the original clause set of KBKF[t]. Let C be a clause derived from C' by exactly one derivation step, such that $f_i \in C$ and $f_i \in C'$. If $\mathsf{lit}(x_i) \in C'$ then it must hold that $\mathsf{lit}(x_i) \in C$, because resolution on universal variables is forbidden and the presence of f_i

disallows the universal reduction of $\mathsf{lit}(x_i)$ in both C' and C. Thus by induction it holds for any clause C that if $f_i \in C$ then $\mathsf{lit}(x_i) \in C$.

Now let C be a clause derived from C' by exactly one derivation step, such that $\overline{f}_i \in C$ and $\overline{f}_i \in C'$. If $\mathsf{lit}(x_j) \in C'$ for some $j \in [i..t]$, then $\mathsf{lit}(x_j) \in C$ for the same reasons as above. If $\overline{f}_j \in C'$ for some $j \in [i..t]$, then either $\mathsf{lit}(x_j) \in C$ (in the case where f_j is the pivot variable, i.e., $C = \mathsf{resolve}(C', f_j, C'')$ with $f_j, \mathsf{lit}(x_j) \in C''$ by the above discussion), or $\overline{f}_j \in C$ (in any other case). Thus by induction it holds for any clause C that if $\overline{f}_i \in C$ then for any $j \in [i..t]$ either $\overline{f}_j \in C$ or $\mathsf{lit}(x_j) \in C$. □

Lemma 2 (Invariant 2). *Given any* LQ*-resolution proof* Π *of a formula* KBKF[t] *the following holds for any clause* $C \in \Pi$*. For all* $i \in [1..t]$*, if* $\mathsf{lit}(d_i) \in C$ *or* $\mathsf{lit}(e_i) \in C$ *then* $f_j \notin C$ *for any* $j \in [1..t]$*.*

Proof. First, the invariant holds for any clause in the original clause set of KBKF[t]. Now let $C = \mathsf{resolve}(C_1, p, C_2)$, where $\mathsf{lit}(e_i) \in C$ or $\mathsf{lit}(d_i) \in C$, and $\mathsf{lit}(e_i) \in C_1$ or $\mathsf{lit}(d_i) \in C_1$ for some $i \in [1..t]$.

If $\mathsf{lit}(e_k) \in C_2$ or $\mathsf{lit}(d_k) \in C_2$ for some $k \in [1..t]$ then by inductive hypothesis it holds that for all $j \in [1..t]$ $f_j \notin C_1$ and $f_j \notin C_2$. Therefore, by the definition of resolve, it holds that for all $j \in [1..t]$ $f_j \notin C$.

Else, $\mathsf{lit}(e_i) \notin C_2$ and $\mathsf{lit}(d_i) \notin C_2$, thus we are left with $p = f_k$ for some $k \in [1..t]$. By inductive hypothesis, $f_j \notin C_1$ for all $j \in [1..t]$, therefore $\overline{f}_k \in C_1$ and $f_k \in C_2$. By Proposition 5 it holds that $f_j \notin C_2$ for all $j \in [1..t]$ with $j \neq k$. Thus for all $j \in [1..t]$ it holds that $f_j \notin C$.

Therefore, by induction it holds for any clause C and for all $i \in [1..t]$ that if $\mathsf{lit}(d_i) \in C$ or $\mathsf{lit}(e_i) \in C$ then $f_j \notin C$ for any $j \in [1..t]$. □

Lemma 3 (Invariant 3). *Given any* LQ*-resolution proof* Π *of a formula* KBKF[t] *the following holds for any clause* $C \in \Pi$*. For all* $i \in [1..t]$ *it holds that if* $\mathsf{lit}(d_i) \in C$ *or* $\mathsf{lit}(e_i) \in C$ *then for any* $j \in [1..i-1]$ *either* $\overline{f}_j \in C$ *or* $\mathsf{lit}(x_j) \in C$*.*

Proof. First, note that the invariant holds for any clause of the original clause set of KBKF[t]. Now, let C be a clause retrieved from C' by one derivation step, such that $\mathsf{lit}(e_i) \in C'$ or $\mathsf{lit}(d_i) \in C'$, and $\mathsf{lit}(e_i) \in C$ or $\mathsf{lit}(d_i) \in C$. If for some $j \in [1..i-1]$ it holds that $\mathsf{lit}(x_j) \in C'$, then $\mathsf{lit}(x_j) \in C$ for the same reasons as in the proof of Invariant 1 (recall that $\mathsf{lev}(e_i) = \mathsf{lev}(d_i) > \mathsf{lev}(x_j)$ for $j \in [1..i-1]$, therefore disallowing universal reduction of $\mathsf{lit}(x_j)$ in the presence of either $\mathsf{lit}(e_i)$ or $\mathsf{lit}(d_i)$). If $\overline{f}_j \in C'$ for some $j \in [1..i-1]$, then either $\mathsf{lit}(x_j) \in C$ (in the case where f_j is the pivot variable, i.e., $C = \mathsf{resolve}(C', f_j, C'')$ with $\{f_j, \mathsf{lit}(x_j)\} \in C''$ by Invariant 1), or $\overline{f}_j \in C$ (in any other case).

Therefore by induction it holds for any clause C and for all $i \in [1..t]$ that if $\mathsf{lit}(d_i) \in C$ or $\mathsf{lit}(e_i) \in C$ then for any $j \in [1..i-1]$ either $\overline{f}_j \in C$ or $\mathsf{lit}(x_j) \in C$. □