

Extension Variables in QBF Resolution

Olaf Beyersdorff¹ and Leroy Chew¹ and Mikoláš Janota²

¹School of Computing, University of Leeds, United Kingdom

²Microsoft Research, Cambridge, United Kingdom

Abstract

We investigate two QBF resolution systems that use extension variables: *weak extended Q-resolution*, where the extension variables are quantified at the innermost level, and *extended Q-resolution*, where the extension variables can be placed inside the quantifier prefix. These systems have been considered previously by (Jussila et al. 2007), who give experimental evidence that extended Q-resolution is stronger than weak extended Q-resolution.

Here we prove an exponential separation between the two systems, thereby confirming the conjecture of (Jussila et al. 2007). Conceptually, this separation relies on showing strategy extraction for weak extended Q-resolution by bounded-depth circuits. In contrast, we show that this strong strategy extraction result fails in extended Q-resolution.

Introduction

Using extension variables to abbreviate possibly complex formulas is a well known and powerful concept in proof complexity and solving. In Tseitin transformations, extension variables are used to encode arbitrary propositional formulas in CNF. More generally, allowing the extension rule in proofs is known to shorten proof size drastically for many examples. This makes extension variables also very interesting in the context of solving, and indeed modern proof checking formats such as RAT for SAT-solvers (Heule, Jr., and Wetzler 2013) and QRAT for QBF solvers (Heule, Seidl, and Biere 2014) incorporate the use of extension variables.

When augmenting the classical resolution system with the extension rule, allowing to introduce a new variable v to abbreviate a disjunction $\neg x \vee \neg y$, we arrive at *extended resolution* (Tseitin 1968), cf. also (Kullmann 1999). Although resolution itself is considered a weak proof systems with many lower bounds known (cf. Segerlind 2007), extended resolution is an extremely powerful system. Extended resolution is equivalent to extended Frege systems (Cook and Reckhow 1979) (Krajíček and Pudlák 1989), one of the strongest proof systems considered today. Showing any non-trivial lower bounds for extended Frege constitutes an extremely challenging problem with even hard candidate formulas currently lacking.

In QBF solving and proof complexity, using extension variables was first considered by Jussila et al. (2007), where they augment Q-resolution, a QBF analogue of resolution (Kleine Büning, Karpinski, and Flögel 1995), to an extended Q-resolution system. In comparison to the propositional case, extension variables in QBF present one additional challenge. We cite the relevant passage from (Jussila et al. 2007):

“In adapting the extension rule to the QBF setting, the crucial question is where to ‘put’ new variables, e.g., how defined variables are ordered with respect to variables that already occur in the formula. It seems intuitive that new variables can only be existential. It is also clear that they cannot be moved further out than the innermost variable on which they depend. In the experimental section we show that we actually need this freedom to move defined variables as far out as possible. Keeping them in the innermost existential scope, as for instance in the Tseitin encoding of a non-CNF QBF formula, is insufficient.”

The main contribution of this work is to underline this experimental observation with a rigorous theoretical argument. For this we consider two systems: weak extended Q-resolution, where extension variables are quantified at the innermost level, and extended Q-resolution, where extension variables are quantified immediately after the innermost variable on which they depend.

Our main result is an exponential separation between these two versions. We show that QPARITY formulas recently introduced in (Beyersdorff, Chew, and Janota 2015) have short proofs in extended Q-resolution, but require exponential-size proofs in weak extended Q-resolution.

The lower bound uses the strategy extraction technique, recently introduced for Q-resolution in (Beyersdorff, Chew, and Janota 2015) and further developed for stronger systems in (Beyersdorff, Bonacina, and Chew 2015). Here we show that weak extended Q-resolution admits strategy extraction in AC^0 , i.e., from each refutation of a false QBF we can extract a winning strategy for the universal player that can be computed by constant-depth circuits. This allows to transfer Håstad’s circuit lower bound for parity (Håstad 1987) to a proof size lower bound for the QPARITY formulas in weak extended Q-resolution.

Preliminaries

We recall some definitions on QBFs and proof systems.

Quantified Boolean formulas. A (closed prenex) *quantified Boolean formula* (QBF) is a formula in quantified propositional logic where each variable is quantified at the beginning of the formula, using either an existential or universal quantifier. We denote such formulas as $\mathcal{Q}. \phi$, where ϕ is a propositional Boolean formula in conjunctive normal form (CNF), called *matrix*, and \mathcal{Q} is its *quantifier prefix*. The *quantification level* $\text{lv}(y)$ of a variable y in $\mathcal{Q}. \phi$ is the number of alternations of quantifiers y has on its left in the quantifier prefix of $\mathcal{Q}. \phi$. Given a variable y , we will sometimes refer to the variables with quantification level lower than $\text{lv}(y)$ as variables *left of* y ; analogously the variables with quantification level higher than $\text{lv}(y)$ will be *right of* y .

Often it is useful to think of a QBF $\mathcal{Q}_1 X_1 \dots \mathcal{Q}_k X_k. \phi$ as a *game* between the *universal* and the *existential player*. In the i -th step of the game, the player \mathcal{Q}_i assigns values to all the variables X_i . The existential player wins the game iff the matrix ϕ evaluates to 1 under the assignment constructed in the game. The universal player wins iff the matrix ϕ evaluates to 0. Given a universal variable u with index i , a *strategy for* u is a function from all variables of index $< i$ to $\{0, 1\}$. A QBF is false iff there exists a *winning strategy* for the universal player, i.e. if the universal player has a strategy for all universal variables that wins any possible game (Goultiaeva, Van Gelder, and Bacchus 2011)(Arora and Barak 2009, Sec. 4.2.2)(Papadimitriou 1994, Chap. 19).

Classical resolution. *Resolution*, introduced by Blake (1937) and Robinson (1965), is a refutational proof system manipulating unsatisfiable CNFs as sets of clauses. The only inference rule is $\frac{C \vee x \quad D \vee \neg x}{C \cup D}$ where C, D denote clauses and x is a variable. A resolution refutation derives the empty clause \square .

QBF resolution calculi. *Q-resolution* by Kleine Büning, Karpinski, and Flögel (1995) is a resolution-like calculus that operates on QBFs in prenex form where the matrix is a CNF. It uses the propositional resolution rule above with the side conditions that variable x is existential and if $z \in C$, then $\neg z \notin D$. We do this to prevent tautological clauses, which we also forbid from being introduced in the proof system via axiom. In addition Q-resolution has a universal reduction rule

$$\frac{C \vee u}{C} \quad \frac{C \vee \neg u}{C} \text{ (\forall-Red)}$$

where variable u is universal and all other existential variables $x \in C$ are left of u in the quantifier prefix.

Q-resolution is not the only QBF resolution system; for further calculi cf. e.g. (Beyersdorff, Chew, and Janota 2014).

Two versions of extended Q-resolution

Extended resolution for propositional resolution (Tseitin 1968), enables adding clauses expressing the equality $v \Leftrightarrow (\neg x \vee \neg y)$, for a fresh variable v . We follow this idea in the context of Q-resolution. Here, we need to decide the position of the fresh variable in the prefix. Two versions are considered. A weak one a general one.

$$\frac{}{(\neg x \vee \neg y \vee \neg v), (x \vee v), (y \vee v)}$$

where:

v is a fresh variable,

v is inserted into prefix as existentially quantified,

weak extension: v is inserted at the end of the prefix

general extension: v is inserted after x and y in the prefix

Figure 1: Two versions of *extension rule*

Figure 1 defines the two forms of the extension rule, which gives us two flavors of extended Q-resolution.

Definition 1. *Weak extended Q-resolution is the calculus of Q-resolution enhanced with the extension rule in its weak form.*

Definition 2. *Extended Q-resolution is the calculus of Q-resolution enhanced with the extension rule in its general form.*

The QPARITY formulas

We recall the QPARITY formulas from (Beyersdorff, Chew, and Janota 2015). Let $\text{xor}(o_1, o_2, o)$ be the set of clauses $\{\neg o_1 \vee \neg o_2 \vee \neg o, o_1 \vee o_2 \vee \neg o, \neg o_1 \vee o_2 \vee o, o_1 \vee \neg o_2 \vee o\}$, which defines o to be equal to $o_1 \oplus o_2$. Define QPARITY_n as $\exists x_1, \dots, x_n \forall z \exists t_2, \dots, t_n.$

$$\text{xor}(x_1, x_2, t_2) \cup \bigcup_{i=3}^n \text{xor}(t_{i-1}, x_i, t_i) \cup \{z \vee t_n, \neg z \vee \neg t_n\}.$$

The auxiliary variables t_i express the prefix sums $x_1 \oplus \dots \oplus x_i$, and hence t_n computes the parity $x_1 \oplus \dots \oplus x_n$ of the x variables. Therefore, the formulas express that there exists an assignment to the x variables such that $x_1 \oplus \dots \oplus x_n$ is neither 0 nor 1, an obvious contradiction. The crucial feature of QPARITY_n is that the only strategy of the universal player to win on this formula is to play $z = x_1 \oplus \dots \oplus x_n$, i.e., the strategy has to compute parity.

In (Beyersdorff, Chew, and Janota 2015) the QPARITY formulas are shown to require exponential-size Q-Res refutations. Here we will show that they in fact provide an exponential separation of weak extended Q-resolution and extended Q-resolution.

Short proofs for QPARITY in extended Q-resolution

We show that QPARITY is easy for extended Q-resolution.

Theorem 3. *The formulas QPARITY_n have linear-size proofs in extended Q-resolution.*

Proof. We define extension variable $s_2 = (x_1 \vee x_2) \wedge (\bar{x}_1 \vee \bar{x}_2)$. In clausal form this introduces the following clauses:

$$s_2 \vee \bar{x}_1 \vee x_2 \quad s_2 \vee x_1 \vee \bar{x}_2 \quad \bar{s}_2 \vee x_1 \vee x_2 \quad \bar{s}_2 \vee \bar{x}_1 \vee \bar{x}_2$$

Note that in our definition of extension variables we only allow extension variables defined as a conjunction of two literals. Here and later in this proof, we will ignore the strictness of the requirement. Notice that, if we were more strict,

s_2 can be defined as $s_2 = \neg a \wedge \neg b$ where $a = \neg x_1 \wedge \neg x_2$ and $b = x_1 \wedge x_2$. The above clauses can be derived easily from the extension rules and the extension variables a, b literals can be cut out.

For $2 < i \leq n$, We define extension variables $s_i = s_{i-1} \oplus x_i = (s_{i-1} \vee x_i) \wedge (\bar{s}_{i-1} \vee \bar{x}_i)$.

In clausal form this introduces the following clauses:

$$s_i \vee \bar{x}_i \vee s_{i-1} \quad s_i \vee x_i \vee \bar{s}_{i-1} \quad \bar{s}_i \vee x_i \vee s_{i-1} \quad \bar{s}_i \vee \bar{x}_i \vee \bar{s}_{i-1}$$

The main part of this short proof is to show we can easily substitute s_i for t_i . We will show this by induction on i . The shortness comes from the fact that s_n literals are left of z . This allows us to reduce z literals ‘early’ and get a contradiction.

Induction Claim: Clauses $s_i \vee \bar{t}_i$ and $t_i \vee \bar{s}_i$ (that show $s_i = t_i$) are provable in $O(i)$ size proofs.

Base Case: Let $i = 2$,

$$\frac{\frac{s_2 \vee \bar{x}_1 \vee x_2 \quad \bar{t}_2 \vee \bar{x}_1 \vee \bar{x}_2}{s_2 \vee \bar{t}_2 \vee \bar{x}_1} \quad \frac{s_2 \vee x_1 \vee \bar{x}_2 \quad \bar{t}_2 \vee x_1 \vee x_2}{s_2 \vee \bar{t}_2 \vee x_1}}{s_2 \vee \bar{t}_2}$$

$$\frac{\frac{t_2 \vee \bar{x}_1 \vee x_2 \quad \bar{s}_2 \vee \bar{x}_1 \vee \bar{x}_2}{t_2 \vee \bar{s}_2 \vee \bar{x}_1} \quad \frac{s_2 \vee x_1 \vee \bar{x}_2 \quad \bar{t}_2 \vee x_1 \vee x_2}{s_2 \vee \bar{t}_2 \vee x_1}}{s_2 \vee \bar{t}_2}$$

Inductive Step: we assume from induction hypothesis that we have for some $i \leq n$, $s_{i-1} \vee \bar{t}_{i-1}$ and $t_{i-1} \vee \bar{s}_{i-1}$

Firstly we use the clauses that express $s_{i-1} = t_{i-1}$ to substitute s_{i-1} for t_{i-1} .

$$\frac{\bar{t}_i \vee \bar{x}_i \vee \bar{t}_{i-1} \quad t_{i-1} \vee \bar{s}_{i-1} \quad t_i \vee x_i \vee \bar{t}_{i-1}}{\bar{t}_i \vee \bar{x}_i \vee \bar{s}_{i-1} \quad t_i \vee x_i \vee \bar{s}_{i-1}}$$

$$\frac{\bar{t}_i \vee x_i \vee t_{i-1} \quad s_{i-1} \vee \bar{t}_{i-1} \quad t_i \vee \bar{x}_i \vee t_{i-1}}{\bar{t}_i \vee x_i \vee s_{i-1} \quad t_i \vee \bar{x}_i \vee s_{i-1}}$$

We can use that the clauses that define s_i and t_i are identical to derive $s_i = t_i$.

$$\frac{\frac{s_i \vee \bar{x}_i \vee s_{i-1} \quad \bar{t}_i \vee \bar{x}_i \vee \bar{s}_{i-1}}{t_i \vee \bar{s}_{i-1} \vee \bar{x}_i} \quad \frac{s_i \vee x_i \vee \bar{s}_{i-1} \quad \bar{t}_i \vee x_i \vee s_{i-1}}{s_2 \vee \bar{t}_2 \vee x_1}}{s_i \vee \bar{t}_i}$$

$$\frac{\frac{t_i \vee \bar{x}_i \vee s_{i-1} \quad \bar{s}_i \vee \bar{x}_i \vee \bar{s}_{i-1}}{t_i \vee \bar{s}_i \vee \bar{x}_i} \quad \frac{t_i \vee x_i \vee \bar{s}_{i-1} \quad \bar{s}_i \vee x_i \vee s_{i-1}}{t_i \vee \bar{s}_i \vee x_i}}{t_i \vee \bar{s}_i}$$

Since we only add a constant number of clauses, the induction argument allows us to keep a $O(i)$ size proof up until $s_n \vee \bar{t}_n$ and $t_n \vee \bar{s}_n$.

$$\frac{s_n \vee \bar{t}_n \quad z \vee t_n}{s_n \vee z} \quad \frac{t_n \vee \bar{s}_n \quad \bar{z} \vee \bar{t}_n}{\bar{s}_n \vee \bar{z}_n}$$

Because s_n is defined only on x_i variables and so universal reduction is available

$$\frac{s_n \vee z \quad \bar{s}_n \vee \bar{z}_n}{s_n} \quad \frac{\bar{s}_n \vee \bar{z}_n}{\bar{s}_n}$$

⊥

This completes the short refutation in extended Q-resolution. \square

Strategy extraction and hardness of QPARITY in weak extended Q-resolution

We now complement the upper bound from the previous section with a lower bound for the same formulas in weak extended Q-resolution. The lower bound argument rests on strategy extraction, which is a widely used paradigm in QBF solving and proof systems. Strategy extraction as a lower bound technique was introduced in (Beyersdorff, Chew, and Janota 2015) and further developed in (Beyersdorff, Bonacina, and Chew 2015).

The idea is to show that from refutations of false formulas it is possible to efficiently extract winning strategies for the universal player. For Q-resolution it is known (Balabanov and Jiang 2012) that strategies can be extracted in a very simple model, namely decision lists.

Definition 4 ((Rivest 1987)). A decision list is a finite sequence of pairs (t_i, c_i) where t_i is a term and $c_i \in \{0, 1\}$ is a Boolean constant. Additionally, the last term is the empty term, semantically equivalent to true. For an assignment μ , a decision list $D = (t_1, c_1), \dots, (t_n, c_n)$ evaluates to c_i if i is the least index such that $\mu \models t_i$. We say that (t_i, c_i) triggers under μ if this condition is satisfied.

We remark that it is easy to transform decision lists into bounded-depth Boolean circuits.

Lemma 5 ((Beyersdorff, Chew, and Janota 2015)). If the function f can be represented as a polynomial-size decision list D , then f can be computed by polynomial-size circuits of depth 3.

Proof. Let $S = \{i \mid (t_i, 1) \in D\}$ be the indices of all pairs in D with 1 as the second component. Observe that f evaluates to 1 under μ iff one of the t_i with $i \in S$ triggers under μ . For each t_i with $i \in S$ construct a function $f_i = t_i \wedge \bigwedge_{l=1}^{i-1} \neg t_l$. Construct a circuit for the function $f_S = \bigvee_{i \in S} f_i$. The function f_S is equal to f_D and can be computed in depth 3 as all t_i are just terms. \square

Balabanov and Jiang (2012) showed that Q-resolution allows strategy extraction in decision lists. Here we show that the same remains true in weak extended Q-resolution. In comparison to (Balabanov and Jiang 2012) we provide a simplified proof.

Theorem 6. Given a refutation π of QBF ϕ in weak extended Q-resolution, there exists a winning strategy for the universal player for ϕ such that for each universal variable u of ϕ the winning strategy can be represented as a Boolean function f_u that is expressible as a decision list whose size is polynomial in $|\pi|$.

Proof. We first review the proof of the strategy extraction theorem for Q-resolution and then explain at the end how it applies to weak extended Q-resolution.

Let $\pi = (L_1, \dots, L_\ell)$ be a resolution refutation of the false QBF $\mathcal{Q}.\phi$ and let

$$\pi_i = \begin{cases} \emptyset & \text{if } i = \ell, \\ (L_{i+1}, \dots, L_\ell) & \text{otherwise.} \end{cases}$$

We show, by downward induction on i , that from π_i it is possible to construct in linear time (w.r.t. $|\pi_i|$) a winning strategy σ^i for the universal player for the QBF formula $\mathcal{Q} \cdot \phi_i$, where

$$\phi_i = \begin{cases} \phi & \text{if } i = 0, \\ \phi \wedge L_1 \wedge \dots \wedge L_i & \text{otherwise,} \end{cases}$$

such that for each universal variable u in $\mathcal{Q} \cdot \phi$, there exists a decision list D_u^i computing σ_u^i as a function of the variables in \mathcal{Q} left of u , having size $O(|\pi_i|)$.

The statement of the theorem corresponds to the case when $i = 0$. The base case of the induction is for $i = \ell$. In this case σ^ℓ is trivial since ϕ_ℓ contains the line $L_\ell = \perp$, and we can define all the D_u^ℓ as $u \leftarrow 0$.

We show now how to construct σ_u^{i-1} and D_u^{i-1} from σ_u^i and D_u^i :

- If L_i is derived by resolution, then for each universal variable u we set $\sigma_u^{i-1} = \sigma_u^i$ and $D_u^{i-1} = D_u^i$.
- Suppose L_i is the result of an application of a \forall red rule on clause L_j , that is $L_j = L_i \vee u^c$ with $c \in \{0, 1\}$, where u is the rightmost variable in L_j , and u^0 stands for $\neg u$ and u^1 for u . Let $\vec{x}_{u'}$ denote the variables on the left of u' in the quantifier prefix of $\mathcal{Q} \cdot \phi$. Then we define

$$\sigma_{u'}^{i-1}(\vec{x}_{u'}) = \begin{cases} \sigma_{u'}^i(\vec{x}_{u'}) & \text{if } u' \neq u, \\ 1 - c & \text{if } u' = u \text{ and } L_i(\vec{x}_u) = 0, \\ \sigma_u^i(\vec{x}_u) & \text{if } u' = u \text{ and } L_i(\vec{x}_u) = 1. \end{cases}$$

Moreover for each $u' \neq u$ we set $D_{u'}^{i-1} = D_{u'}^i$, and we set D_u^{i-1} as follows:

$$\text{if } \neg L_i(\vec{x}_u) \text{ then } u \leftarrow 1 - c; \\ \text{else } D_u^i(\vec{x}_u).$$

We now check that for each u' , $\sigma_{u'}^{i-1}$ respects all the properties of the inductive claim.

It is clear that $\sigma_{u'}^{i-1}$ and $D_{u'}^{i-1}$ are well defined and constructed in linear time w.r.t. $|\pi_{i-1}|$. Also, by construction $D_{u'}^{i-1}$ computes $\sigma_{u'}^{i-1}$.

To verify that σ^{i-1} is a winning strategy for $\mathcal{Q} \cdot \phi_{i-1}$ we fix an assignment ρ to the existential variables of ϕ . Let τ_i be the complete assignment to existential and universal variables, constructed in response to ρ under the strategy σ^i . By induction hypothesis τ_i falsifies ϕ_i . We need to show that τ_{i-1} falsifies ϕ_{i-1} . To show this we distinguish again two cases.

If L_i is derived by the resolution rule, then $\sigma^{i-1} = \sigma^i$ and $\tau_{i-1} = \tau_i$. Hence by induction hypothesis, τ_i falsifies a conjunct from ϕ_i . To argue that τ_{i-1} also falsifies a conjunct from ϕ_{i-1} we only need to look at the case when the falsified conjunct is L_i . As L_i is false under τ_i and L_i is derived by resolution, by soundness of the resolution rule one of the parent formulas of L_i in the application of the resolution rule must be falsified as well. Hence τ_{i-1} falsifies ϕ_{i-1} .

Let now L_i be derived by \forall red from $L_j = L_i \vee u^c$ for some $j < i$. In this case, our strategy σ^{i-1} changes the assignment τ_i only when τ_i made the universal player win by

falsifying L_i . As we set u to $1 - c$, the modified assignment τ_{i-1} falsifies L_j . Otherwise, if τ_i does not falsify L_i we keep $\tau_{i-1} = \tau_i$ and hence falsify one of the conjuncts of ϕ_{i-1} by induction hypothesis.

Let us now explain how the above applies to weak extended Q-resolution. Consider a refutation π of a QBF $\mathcal{Q}\vec{x} \cdot \phi(\vec{x})$ in weak extended Q-resolution. We can view π as a Q-resolution refutation of the QBF $\mathcal{Q}\vec{x}\exists\vec{y} \cdot \phi(\vec{x}) \wedge \psi(\vec{x}, \vec{y})$, where suitable extension variables \vec{y} together with their definitions $\psi(\vec{x}, \vec{y})$ have been added.

We apply the argument above to construct decision lists computing a winning strategy for $\mathcal{Q}\vec{x}\exists\vec{y} \cdot \phi(\vec{x}) \wedge \psi(\vec{x}, \vec{y})$. By definition of the \forall red rule, no \vec{y} variables are present in \forall red steps in π . Hence the decision list will also not use any of the extension variables and therefore in fact compute a winning strategy for the original formula $\mathcal{Q}\vec{x} \cdot \phi(\vec{x})$. \square

This result enables us to show lower bounds for formulas that require hard strategies. An example of such formulas are the QPARITY formulas. As observed earlier, the only winning strategy of the universal player on the QPARITY formulas is to actually compute the PARITY function. However, PARITY is the classic example of a function hard for bounded-depth circuits (and hence by Lemma 5 for decision lists).

Theorem 7 (Furst, Saxe, Sipser (1984), Håstad (1987)). *Every non-uniform family of bounded-depth circuits computing PARITY is of exponential size.*

Using strategy extraction we can now immediately transfer this circuit lower bound to an exponential lower bound in weak extended Q-resolution.

Theorem 8. *Any refutation of QPARITY_n in weak extended Q-resolution is of exponential size.*

Proof. The unique winning strategy for the variable z in QPARITY_n is to compute $x_1 \oplus \dots \oplus x_n$. By Theorem 6, there is a polynomial-time algorithm for constructing a decision list D_n from any refutation of QPARITY_n in weak extended Q-resolution. Such decision list can be converted in polynomial time into a depth-3 circuit by Lemma 5. Hence, the refutation must be of exponential size due to Theorem 7. \square

Theorem 8 together with the upper bound from the previous section imply the following exponential separation.

Corollary 9. *Weak extended Q-resolution does not simulate extended Q-resolution.*

This also has consequences for strategy extraction in extended Q-resolution.

Theorem 10. *Extended Q-resolution has strategy extraction in P, but does not admit strategy extraction in AC⁰.*

Proof. Inspecting the proof of Theorem 6 it is clear that extended Q-resolution still has strategy extraction in polynomial time, and in fact for the original formula not involving extension variables: for this it suffices to perform the construction of the decision lists as in Theorem 6 and then replace extension variables by their definition. To keep the

size polynomial, this requires reusing definitions of extensions variables and hence instead of formulas will lead to polynomial-size circuits. However, the substitutions will increase the depth and not result in AC^0 circuits.

To argue that extended Q-resolution does not admit AC^0 strategy extraction we use again the QPARITY formulas, which have short proofs in extended Q-resolution by Theorem 3, but require exponential-size strategies by Theorem 7. \square

Acknowledgments. This work was partially supported by CMU-Portugal grant AMOS (CMUP-EPB/TIC/0049/2013), FCT grant POLARIS (PTDC/EIA-CCO/123051/2010), INESC-ID's multiannual PIDDAC funding PEst-OE/EEI/LA0021/2011, grant no. 48138 from the John Templeton Foundation, and EPSRC grant EP/L024233/1. The second author was supported by a Doctoral Training Grant from EPSRC.

References

- Arora, S., and Barak, B. 2009. *Computational Complexity – A Modern Approach*. Cambridge University Press.
- Balabanov, V., and Jiang, J.-H. R. 2012. Unified QBF certification and its applications. *Formal Methods in System Design* 41(1):45–65.
- Beyersdorff, O.; Bonacina, I.; and Chew, L. 2015. Lower bounds: from circuits to QBF proof systems. In *ITCS'16*. technical report available at ECCC, 22:152.
- Beyersdorff, O.; Chew, L.; and Janota, M. 2014. On unification of QBF resolution-based calculi. In *MFCS*, 81–93.
- Beyersdorff, O.; Chew, L.; and Janota, M. 2015. Proof complexity of resolution-based QBF calculi. In *Proc. Symposium on Theoretical Aspects of Computer Science*, 76–89. LIPIcs series.
- Blake, A. 1937. *Canonical expressions in boolean algebra*. Ph.D. Dissertation, University of Chicago.
- Cook, S. A., and Reckhow, R. A. 1979. The relative efficiency of propositional proof systems. *The Journal of Symbolic Logic* 44(1):36–50.
- Furst, M. L.; Saxe, J. B.; and Sipser, M. 1984. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory* 17(1):13–27.
- Goultiaeva, A.; Van Gelder, A.; and Bacchus, F. 2011. A uniform approach for generating proofs and strategies for both true and false QBF formulas. In *IJCAI*, 546–553.
- Håstad, J. 1987. *Computational Limitations of Small Depth Circuits*. Cambridge: MIT Press.
- Heule, M.; Jr., W. A. H.; and Wetzler, N. 2013. Verifying refutations with extended resolution. In *Automated Deduction - CADE-24 - 24th International Conference on Automated Deduction*, 345–359.
- Heule, M.; Seidl, M.; and Biere, A. 2014. A unified proof system for QBF preprocessing. In *IJCAR*, 91–106.
- Jussila, T.; Biere, A.; Sinz, C.; Kröning, D.; and Wintersteiger, C. M. 2007. A first step towards a unified proof checker for QBF. In Marques-Silva, J., and Sakallah, K. A., eds., *SAT*, volume 4501, 201–214. Springer.
- Kleine Büning, H.; Karpinski, M.; and Flögel, A. 1995. Resolution for quantified Boolean formulas. *Inf. Comput.* 117(1):12–18.
- Krajíček, J., and Pudlák, P. 1989. Propositional proof systems, the consistency of first order theories and the complexity of computations. *The Journal of Symbolic Logic* 54(3):1063–1079.
- Kullmann, O. 1999. On a generalization of extended resolution. *Discrete Applied Mathematics* 96-97:149–176.
- Papadimitriou, C. H. 1994. *Computational Complexity*. Addison-Wesley.
- Rivest, R. L. 1987. Learning decision lists. *Machine Learning* 2(3):229–246.
- Robinson, J. A. 1965. A machine-oriented logic based on the resolution principle. *Journal of the ACM* 12:23–41.
- Segerlind, N. 2007. The complexity of propositional proofs. *Bulletin of Symbolic Logic* 13(4):417–481.
- Tseitin, G. S. 1968. On the complexity of derivations in the propositional calculus. *Studies in Constructive Mathematics and Mathematical Logic Part II*, ed. A.O. Slisenko.