Electronic Colloquium on Computational Complexity, Report No. 57 (2019)

Proof Complexity of Symmetry Learning in QBF

Joshua Blinkhorn and Olaf Beyersdorff

Institut für Informatik, Friedrich-Schiller-Universität Jena, Germany

Abstract. For quantified Boolean formulas (QBF), a resolution system with a symmetry rule was recently introduced by Kauers and Seidl (Inf. Process. Lett. 2018). In this system, many formulas hard for QBF resolution admit short proofs.
Kauers and Seidl apply the symmetry rule on symmetries of the original formula. Here we propose a new formalism where symmetries are dynamically recomputed during the proof on restrictions of the original QBF. This presents a theoretical model for the potential use of symmetry learning as an inprocessing rule in QCDCL solving.
We demonstrate the power of symmetry learning by proving an exponential separation between Q-resolution with the symmetry rule and Q-resolution with our new symmetry learning rule. In fact, we show that bounding the number of symmetry recomputations gives rise to a hierarchy of QBF resolution systems of strictly increasing strength.

1 Introduction

The last decade has seen tremendous advances in our understanding and algorithmic handling of quantified Boolean formulas (QBF), both theoretically and practically. QBF solving has emerged as a powerful technique to apply to hard problems from many application domains (e.g. [6,14,15]). To theoretically model and analyse the success of QBF solving, a number of QBF proof systems have been developed and analysed, yielding a surge in QBF proof complexity research (e.g. [4, 7, 8]). Ideally, this interaction works in both directions: QBF resolution systems aim to model central solving features and lower bounds to proof size in these systems translate to lower bounds for solver running time. Conversely, new theoretical models can also stimulate practical improvements.

This paper explores the power of symmetry learning for QBF from a theoretical proof-complexity perspective. It is well known that many combinatorial principles exhibit symmetry properties [22], both propositionally and in QBF. Breaking these symmetries is an effective technique for SAT and QBF that can significantly reduce the search space and speed up search.

In SAT solving, symmetry breaking is done both statically [12] – as a preprocessing technique – as well as dynamically during the search [1,13,23]. Part of the work on static symmetry breaking was lifted to the more complex setting of QBF [2, 3], while dynamic symmetry breaking has not yet been realised in the QBF domain (cf. [17] for a recent overview of symmetry breaking in QBF).

On the proof-theoretic side, the propositional resolution system – underlying CDCL solving [5, 21] – has been augmented with symmetry rules of different strengths [19, 24, 25]. In the most basic setting of [19], a symmetry rule is added that from a clause C allows to derive $\sigma(C)$ for each symmetry σ of the input CNF. Already this yields a quite powerful proof system, which e.g. admits short proofs of the pigeonhole principle [25].

Recently, the system of [19] was generalised to QBF by Kauers and Seidl [16]. This proof system Q-Res+S builds on Q-resolution (Q-Res [18]) and again augments it with a symmetry rule for symmetries of the original formula. In [16] the power of this proof system was demonstrated by the construction of short proofs for the formulas of Kleine Büning et al. [18] and of the parity formulas [8], two notorious examples of QBFs hard for Q-resolution.

In this paper we continue the proof-theoretic study of proof systems exploiting QBF symmetries. Our contributions can be summarised as follows.

1. QBF resolution systems for symmetry learning. We introduce a framework for symmetry learning during proof search. While the system Q-Res+S only allows to use symmetries of the input CNF, our new system Q-Res+SL additionally exploits symmetries of restrictions of the input formula. These restrictions correspond to certain partial assignments (called linear assignments here, Definition 5) that arise during runs of QCDCL algorithms. During such a run of QCDCL,

ISSN 1433-8092

we allow to recompute symmetries for the restricted formula currently considered by the solver. These 'newly learned' symmetries then replace the existing set of symmetries. When the QCDCL solver backtracks and unassigns variables, symmetries based on the corresponding restrictions of these variables are discarded as well.

We proof-theoretically model this algorithmic approach using the framework of Lonsing, Egly, and Seidl [20], who propose a general approach with additional axioms (corresponding to learned clauses) for a proof-theoretic modelling of inprocessing techniques in QCDCL. This framework has also been previously used for dependency learning in QBF [10]. This gives rise to our new QBF resolution system Q-Res+SL, where at each point in the proof, the current symmetries are made available to the symmetry rule. Using the approach of [10, 20] we show soundness of the system (Theorem 9); completeness follows as the calculus extends Q-Res.

We can parameterise the system Q-Res+SL by keeping track of the maximal number d of symmetry recomputations on a QCDCL branch without any backtracking. We call this number d the degree of the Q-Res+SL proof. Restricting proofs in Q-Res+SL to degree 0 yields the system Q-Res+S.

2. Exponential separations. Our main results assess the proof complexity of the new calculi for symmetry learning. We show that Q-Res+SL is exponentially stronger than Q-Res+S, and in fact, the subsystem of Q-Res+SL restricted to degree d proofs is exponentially stronger than Q-Res+SL restricted to degree $d-1$ proofs (Theorem 18). Thus allowing to successively learn more symmetries corresponds to a hierarchy of proof systems of strictly increasing strength.

To show this result we start by noticing that also the equality formulas – known to be hard in Q-Res and even stronger systems such as QBF cutting planes and polynomial calculus [7] – admit short proofs in Q-Res+S (Theorem 4).

We then devise a *symmetry blocker* (Definition 11): a simple gadget that when applied to an arbitrary QBF yields a 'blocked' QBF without any symmetries (Proposition 12). Such blocking can also be repeated, and applying it d times results in the d-blocked version of the formula. However, the symmetries can be unlocked again with symmetry learning, and in particular the d-blocked versions of the equality formulas have short proofs in Q-Res+SL of degree d (Lemma 17).

The main technical difficulty lies in showing that lower bounds for QBFs Φ_n for Q-Res lift to lower bounds for the d-blocked versions of Φ_n in the subsystem of Q-Res+SL restricted to degree $d-1$ proofs (Lemma 13). In combination, these upper and lower bounds yield the strictness of the hierarchy for symmetry learning proof system (Theorem 18).

To ease the technical presentation, our model of QCDCL assignments used to define the system Q-Res+SL does neither incorporate unit propagation nor pure literal elimination (in contrast to the model of [20]). However, we argue (Section 6) that all our hardness results can be lifted to the practically more realistic setting where unit propagation and pure literal elimination is by default built into the proof system.

2 Preliminaries

Conjunctive normal form. A *literal* is a Boolean variable z or its negation $\bar{z}$. The *complement* $\bar{a}$ of a literal a is z if $a = \bar{z}$, and $\bar{z}$ if $a = z$. A *clause* is a finite disjunction of literals, and a *conjunctive normal form* formula (CNF) is a finite conjunction of clauses. We denote clauses as sets of literals and CNFs as sets of clauses. A clause is *tautological* if it contains some literal and its complement, otherwise it is *non-tautological*.

The variable of a literal $a = z$ or $a = \bar{z}$ is $\mathrm{var}(a) := z$. The variable set of a clause is $\mathrm{vars}(C) := \{\mathrm{var}(a) : a \in C\}$, and the variable set of a CNF is $\mathrm{vars}(f)$, the union of the variable sets of the clauses in f.

An assignment to a set $Z \subseteq \mathrm{vars}(f)$ of variables is a mapping $\alpha : Z \to \{0, 1\}$, typically represented as a set of literals $\{a_1, \ldots, a_k\}$, where literals $\bar{z}$ and z represent the respective assignments $z \mapsto 0$ and $z \mapsto 1$. The *restriction* of f by α, denoted $f[\alpha]$, is obtained from f by removing any clause containing a literal in α, and removing the complementary literals $\bar{a}_1, \ldots, \bar{a}_k$ from the remaining clauses.

Quantified Boolean formulas. A *quantified Boolean formula* (QBF) $F := Q \cdot f$ consists of a *quantifier prefix* $Q = Q_1 z_1 \cdots Q_n z_n$, in which each Q_i is a quantifier in $\{\exists, \forall\}$, and a CNF f called the *matrix*, for which $\text{vars}(f) = \{z_1, \ldots, z_n\}$. The variable set of a QBF is $\text{vars}(F) := \text{vars}(f)$. The prefix Q defines a total order $<_Q$ on $\text{vars}(F)$ such that $z_i <_Q z_j$ holds whenever $i < j$, in which case we say that z_i *is left of* z_j and z_j *is right of* z_i. Variables in the first and last blocks are termed *leftmost* and *rightmost*, respectively.

The *restriction* of F by an assignment α is $F[\alpha] := Q[\alpha] \cdot f[\alpha]$, where $Q[\alpha]$ is obtained from Q by removing the variables $\text{vars}(f) \setminus \text{vars}(f[\alpha])$ along with their associated quantifiers.

A model g for a QBF $F := Q \cdot f$ is a set $\{G_x : x \in \text{vars}_\exists(F)\}$ for which (a) each G_x is a Boolean circuit over the universla variables left of x, and (b) the simultaneous substitution of each G_x for x in f is a tautology, i.e. a circuit computing the constant function 1. A QBF is true iff it has a model, otherwise it is false.

Proof systems. Given any literal p, a clause $C = C_1 \cup C_2$ is called a *resolvent* of $C_1 \cup \{p\}$ and $C_2 \cup \{\bar{p}\}$. Given any prefix Q, a clause C is a *Q-reduction* of a clause $C \cup R$ if every variable in R is universally quantified and right of every variable in C (with respect to Q).

A Q-resolution (Q-Res) [18] refutation of a QBF $Q \cdot f$ is a sequence $C_1, \ldots, C_k$ of non-tautological clauses in which (a) C_k is the empty clause, and (b) each clause either belongs to f, or is a resolvent or Q-reduction of preceding clauses.

A proof system P *p-simulates* a proof system Q if there exists a polynomial-time computable function that takes a Q-proof to a P-proof of the same formula [11].

3 Static Symmetries in Q-resolution

In this section, we provide the necessary background on Q-Res+S and its proof complexity.

Definition 1 (symmetry). *A* symmetry σ *for a QBF* $Q \cdot f$ *is a bijection on its literals for which* (a) *applying* σ *to every literal in every clause preserves* f*, and* (b) *for each literal* $a \in \text{dom}(\sigma)$*,* $\sigma(a), \sigma(\bar{a})$ *are complementary literals and* $\text{var}(a), \text{var}(\sigma(a))$ *belong to the same block of* Q*.*

The set of all symmetries of a QBF forms a group under composition, and is denoted $\mathcal{S}(F)$.

Symmetries are incorporated into Q-Res with the addition of a single inference rule: any symmetry of the input QBF can be applied to a derived clause. This rule is labelled 'S' in the following definition.

Definition 2 (Kauers and Seidl [17]). *A* Q-Res+S *refutation of a QBF* $Q \cdot f$ *is a sequence of non-tautological clauses* $\pi := C_1, \ldots, C_k$ *in which* C_k *is the empty clause and one of the following holds for each* $i \in [k]$*:*

A ***axiom:*** C_i *is a clause in the matrix* f*;*
R ***resolution:*** C_i *is the resolvent of two preceding clauses;*
U ***universal reduction:*** C_i *is a* Q*-reduction of a preceding clause;*
W ***weakening:*** C_i *is subsumed by a preceding clause;*
S ***symmetry:*** C_i *is the image of a preceding clause under a symmetry of* $Q \cdot f$*.*

The size of π *is* $|\pi| = k$*.*

It was shown in [16] that Q-Res+S is exponentially stronger than Q-Res, the separation was demonstrated by two different QBF families. Here, we briefly point out that the separation is also performed by another QBF family, namely the equality formulas [7]. The particulars of the linear-size Q-Res+S refutations of equality are used later on in Section 5.

Definition 3 (equality family [7]). *The* equality family *is the QBF family whose* n^{th} *instance is* $\text{EQ}_n := \exists x_1 \cdots x_n \forall u_1 \cdots u_n \exists z_1 \cdots z_n \cdot \text{eq}_n$*, where*

$$\text{eq}_n := \left(\bigcup_{i \in [n]} \{\{\bar{x}_i \bar{u}_i \bar{z}_i\}, \{x_i u_i \bar{z}_i\}\}\right) \cup \{\{z_1 \cdots z_n\}\}.$$

Theorem 4. *The equality family has linear-size refutations in* Q-Res+S*.*

Proof. For $i \in [n]$, let σ_i be the bijection on the literals of EQ_n that sends literals in x_i and u_i to their complements, and is the identity everywhere else. It is easy to see that each σ_i is a symmetry of EQ_n. In n resolution steps, resolving $\{z_1 \cdots z_n\}$ against each $\{\bar{x}_i \bar{u}_i \bar{z}_i\}$ over pivot variable z_i, we obtain the clause $\{\bar{x}_1 \cdots \bar{x}_n \bar{u}_1 \cdots \bar{u}_n\}$. By universal reduction, we then obtain $\{\bar{x}_1 \cdots \bar{x}_n\}$. From any $\{\bar{x}_1 \cdots \bar{x}_i\}$, we obtain $\{\bar{x}_1 \cdots \bar{x}_{i-1}\}$ by application of σ_i and resolution over x_i. □

4 A Theoretical Model of Symmetry Learning

In this section we introduce the proof system Q-Res+SL, a theoretical model for QBF symmetry learning.

Our model is built on the foundation of 'Q-Resolution with Generalised Axioms' due to Lonsing et al. [20], whose primary motivation was to model the integration of preprocessing techniques into the QCDCL search process itself; in other words, to reformulate preprocessing as inprocessing. The authors, however, noted that their setup offers a more general 'interface to Q-resolution' capable of modelling 'the direct combination of orthogonal solving techniques' [20].

This direction was subsequently taken up in the paper [10]. The setup was reformulated to interface with QBF dependency schemes, thereby modelling the integration of dependency-awareness into QCDCL search. An important development there was the introduction of *proof referencing*. A refutation becomes a nested structure of 'subrefutations' of restrictions of the instance, allowing the work behind the interface to contribute to the overall proof size. This addition is key for proof complexity, since without it, every QBF has a trivial refutation (cf. [20]).

Here, in much the same spirit, we propose to interface with the computation of symmetries, thereby modelling the integration of symmetry techniques into QCDCL search. This offers a new possibility: otherwise unidentified symmetries of the current formula can be learned at arbitrary search nodes. Learned symmetries can be applied to learned clauses, thereby strengthening the knowledge base.

The trail. Central to the model of Lonsing et al. is a particular kind of QBF assignment (QCDCL assignment [20, p. 437]), intended to represent the current assignment, or *trail*, of the solver at an arbitrary search node. Here, we have chosen to omit constraint propagation, and work instead with *linear assignments*. Propagation can be safely detached from the proof complexity discussion, with considerable simplification of technical content – this is explained in greater detail in Section 6.

Definition 5 (linear assignment). *A partial assignment $\{a_1, \ldots, a_n\}$ to a QBF F is* linear *if, for each i in $[n]$, $\mathrm{var}(a_i)$ is in the first block of $F[\{a_1, \ldots, a_{i-1}\}]$.*

The proof system Q-Res+SL. Derivations in Q-Res+SL are defined recursively by *degree*, using the proof referencing method of [10].

Definition 6 (Q-Res+SL). *Given a QBF F,* Q-Res+SL *refutations of F are defined inductively by degree:*

- *a degree-0 refutation is a* Q-Res+S *refutation of F;*
- *for $d \in \mathbb{N}$, a degree-d refutation is a sequence $\pi_0 \circ \rho_1 \circ \cdots \circ \rho_k$ satisfying*
  - *π_0 is a* Q-Res+S *refutation of F with extra axioms $A_1, \ldots, A_k$;*
  - *each A_i is the negation of a linear assignment α_i to F;*
  - *each ρ_i is a* Q-Res+SL *refutation of $F[\alpha_i]$;*
  - *the maximum degree of the ρ_i is $d - 1$.*

The size of a Q-Res+SL *refutation is the number of clauses in the sequence.*

The extra axioms $A_1, \ldots, A_k$ are said to *reference* the lower-degree refutations $\rho_1, \ldots, \rho_k$. We illustrate the system, and the use of proof referencing, with the following example.

Example 7. The QBF

$$F := \exists a x_1 \forall u_1 \exists z_1 \cdot \{\{a x_1 u_1 z_1\}\{a \bar{x}_1 \bar{u}_1 z_1\}\{a \bar{z}_1\}\{\bar{a}_1\}\}$$

has a degree-1 Q-Res+SL refutation $\pi \circ \rho$, where π and ρ are the sequences:

$$\begin{array}{lll} \pi := 1 & \{a\} & \text{extra axiom} \\ 2 & \{\bar{a}\} & \text{axiom} \\ 3 & \square & \text{resolution} \end{array} \qquad \begin{array}{lll} \rho := 1 & \{x_1 u_1 z_1\} & \text{axiom} \\ 2 & \{\bar{z}_1\} & \text{axiom} \\ 3 & \{x_1 u_1\} & \text{resolution} \\ 4 & \{x_1\} & \text{universal reduction} \\ 5 & \{\bar{x}_1\} & \text{symmetry} \\ 6 & \square & \text{resolution} \end{array}$$

Consider the linear assignment $\alpha := \{\bar{a}\}$ to F. Notice that $F[\alpha] = \mathrm{EQ}_1$, and that the sequence ρ is the Q-Res+S refutation of EQ_1 described in the proof of Theorem 4. In line 1 of π, the clause $\{a\}$, being the negation of α, can be introduced as an extra axiom, referencing the degree-0 refutation ρ. The refutation is concluded in π by resolution against the unit clause $\{\bar{a}\}$.

Notice that the application of the symmetry rule in line 5 of ρ would not be allowed in π, since σ_1, which is a symmetry of the restricted formula EQ_1, is not a symmetry of F itself. □

Modelling symmetry learning. We provide some high-level intuition on how symmetry learning would work in practice, and its connection to proof referencing in Q-Res+SL.

In our model, we consider the symmetry groups of the input formula and its restrictions; derived clauses and free axioms do not contribute to these symmetry groups. This is analogous to the solving approach in which learned clauses are the target of the new learned symmetries, and not the source. While there are other possibilities, the current approach is perhaps the most straightforward for a first theoretical model.

The degree of a refutation can be understood as the maximum depth of nested symmetry recomputations. When recomputation takes place, a pointer to the new set of symmetries is placed on the trail. The current symmetry set can be applied to any newly learned clauses until it is either replaced by recomputation, or it is removed from the trail by backtracking step. In the latter case, the solver reverts to an earlier set of symmetries, following the highest level pointer that remains on the trail. In this way, higher degree refutations are associated with symmetries learned deeper into the search, under increasingly larger restrictions of the input formula. Degree-0 refutations, which coincide with Q-Res+S, represent the traditional setting in which no recomputation takes place.

One might ask whether the symmetries of the parent formula should be available to referenced refutations. However, a simple example demonstrates that this is not sound, even at the propositional level. Consider the true QBF

$$F := \exists ax \cdot \{\{ax\}\{\bar{a}\bar{x}\}\},$$

and the symmetry σ that sends both literals a and x to their complements. Applying σ to either of

$$F[\{\bar{a}\}] = \exists x \cdot \{\{x\}\}, \quad F[\{a\}] = \exists x \cdot \{\{\bar{x}\}\}$$

permits a refutation, but $F[\{\bar{a}\}]$ and $F[\{a\}]$ are both true QBFs. Moreover, the unit clauses $\{a\}, \{\bar{a}\}$ can be introduced as extra axioms, refuting F itself.

There is a subtle point here: application and restriction of clauses do not associate. In practice, symmetries must be applied to the whole learned clause, not merely to its restriction under the current trail assignment. The restriction, if performed first, may remove literals which would have been satisfied under the symmetry – this is exactly the issue with the foregoing example.

Soundness. It is already known from [10, 20] that the method of proof referencing admits a soundness proof by induction on degree. We follow this method to prove the soundness of Q-Res+SL. The following lemma constitutes the chief observation. It is the analogue of [20, Theorem 2] and [10, Lemma 14].

Lemma 8. *Let A be the negation of a linear assignment α to a QBF $Q \cdot f$. If $Q \cdot f$ is false under α, then $Q \cdot f \vDash Q \cdot \{A\}$.*

It is known that the rules of Q-Res+S preserve the models of the input QBF [16]. Therefore, we are given the soundness of degree-0 refutations for free. Moreover, if we can show that models of the input QBF also satisfy the extra axioms, then we prove the soundness of refutations at the next degree. This is merely the contrapositive statement of Lemma 8, since extra axioms reference refutations of lower degree, which refute false formulas by the inductive hypothesis.

Theorem 9. *If a QBF has a* Q-Res+SL *refutation, then it is false.*

Since our setting differs at the technical level from both [20] and [10], we provide a full proof of the core lemma.

Proof (of Lemma 8). The lemma is trivially true if F is false, so we assume otherwise.

Let $A := \{a_1 \cdots a_k\}$ be the negation of a linear assignment α to $F := Q \cdot f$. For each $0 \leq i \leq k$, let A_i be the first i literals of A, let α_i be its negation, and define

$$E_i := \mathrm{vars}_\exists(F[\alpha_{i-1}]) \setminus \mathrm{vars}_\exists(F[\alpha_i])\,,$$
$$U_i := \mathrm{vars}_\forall(F[\alpha_{i-1}]) \setminus \mathrm{vars}_\forall(F[\alpha_i])\,.$$

Now, let g be a model for F. Set $g_0 := g$, and for each i in $[k]$, obtain g_i from g_{i-1} by discarding the circuits for variables in E_i and restricting the rest by the assignment

$$\beta_i := \{\bar{a}_i\} \cup \{\bar{z} : z \in U_i \setminus \{\mathrm{var}(a_k)\}\}\,.$$

Lastly, let T_i be the circuit obtained by substituting the circuits in g_i for the existential variables in the CNF $f[\alpha_i]$.

By induction on k, we prove the following: if g models F but not $Q \cdot A_k$, then g_k models $F[\alpha_k]$. The base case $k = 0$ is trivial. For the inductive step, let $k \geq 1$. Suppose that g models F but not $Q \cdot \{A_k\}$. Then g does not model $Q \cdot \{A_{k-1}\}$, so g_{k-1} models $F[\alpha_{k-1}]$ by the inductive hypothesis, as the negation of A_{k-1} is a linear assignment. Hence T_{k-1} is a tautology.

On the other hand, $T_k = T_{k-1}[\beta_i]$, so T_k is a tautology, and g_k models $F[\alpha_k]$. Indeed, given $f[\alpha_{k-1}]$, applying the substitution based on g_{k-1} followed the restriction β_k has the same effect as applying the restriction β_k first, followed by substitution of the restricted circuits g_k. To see this, one must note the following: if $\{\bar{a}_k\}$ represents an existential assignment, say $x \mapsto b$, then the circuit for x in g_{k-1} computes the constant b, so substituting that circuit for x is the same as applying the assignment $\{\bar{a}_k\}$. □

5 Proof Complexity of Symmetry Learning Systems

In this section we show that degree-d Q-Res+SL refutations may be exponentially shorter than refutations of degree $d-1$, for each natural number d.

We make use of the product operation [9] on CNFs and QBFs. The product of two CNFs f and g is $f \otimes g := \{C \cup D : C \in f, D \in g\}$, and the product of two QBFs $F := Q \cdot f, G := R \cdot g$ is $F \otimes G := QR \cdot f \otimes g$.

Provided that the concatenation of prefixes does not create a longer block in the middle, taking a product has the natural effect on the symmetry groups.

Proposition 10. *Let F and G be variable-disjoint QBFs. If the rightmost block of F and the leftmost block of G are oppositely quantified, then*

$$\mathcal{S}(F \otimes G) = \{\sigma \cup \tau : \sigma \in \mathcal{S}(F), \tau \in \mathcal{S}(G)\}\,.$$

Proof. It is clear that $\sigma \cup \tau$ is a symmetry of $F \otimes G$ whenever σ, τ are respective symmetries of F, G. For the reverse direction, suppose that $\sigma \cup \tau$ is a symmetry of $F \otimes G$, where the respective domains of σ, τ are the literals of F, G. Let C be a clause in $f \otimes g$, and let C_f, C_g be the respective intersections of C with the literals of F, G. Since each block of $F \otimes G$ is either a block of F or of G, we have

$$\begin{aligned} C_f \in f \text{ and } C_g \in g &\Leftrightarrow C \in f \otimes g \\ &\Leftrightarrow (\sigma \cup \tau)(C) \in f \otimes g \\ &\Leftrightarrow \sigma(C_f) \in f \text{ and } \tau(C_g) \in g\,, \end{aligned}$$

where the last equivalence is due to the fact that both σ, τ are bijections. □

Symmetry blocking. In [10], a technique for obfuscating the independencies of a particular QBF family was introduced. Here, we devise a similar general method for blocking QBF symmetries. The main idea is to add literals in fresh variables, such that only the identity symmetry survives; meanwhile, an assignment to the fresh variables returns the original instance with the symmetries intact.

Definition 11 (blocker). *Given any QBF $F := Q \cdot f$ over variables $Z = \{z_1, \ldots, z_k\}$, the* symmetry blocker *for F is*

$$\mathcal{A}(F) := \exists a b_1 \cdots b_k Q \forall c \cdot \left(\{\{\bar{a}\}\} \otimes f \otimes \{\{c\}\}\right) \cup \{\{a\}\} \cup \{\{b_1 \cdots b_i z_i\} : i \in [k]\}\,,$$

*where $a, b_1, \ldots, b_k, c$ are fresh variables not in Z. For each natural number d, the d-*blocker *of F is*

$$\mathcal{D}_d(F) := \mathcal{A}^d(F) \otimes \cdots \otimes \mathcal{A}^1(F)\,,$$

where $\mathcal{A}^i(F)$ is obtained from $\mathcal{A}(F)$ by adding the superscript i to each occurrence of a variable.

Proposition 12. *For any QBF F and natural number d, the only symmetry of $\mathcal{D}_d(F)$ is the identity.*

Proof. By Proposition 10, every symmetry of $\mathcal{D}_d(F)$ is of the form $\sigma_1 \cup \cdots \cup \sigma_d$, where each σ_i is a symmetry of $\mathcal{A}^i(F)$. Hence, by syntactic equivalence, it suffices to show that the identity is the only symmetry of $\mathcal{A}(F)$.

Let the variables of F be $Z = \{z_1, \ldots, z_k\}$, and let σ be a symmetry of $\mathcal{A}(F)$. The only unit clause in the matrix of $\mathcal{A}(F)$ is $\{a\}$, hence we must have $\sigma(a) = a$. For each $i \in k$, the positive literal b_i occurs only in a clause of size $i+1$, and the negative literal does not occur, hence we must have $\sigma(b_i) = b_i$. For each $i \in k$, the only literal ℓ for which $\{b_1, \ldots, b_i, \ell\}$ is a clause in the matrix of $\mathcal{A}(F)$ is $\ell = z_i$, hence we must have $\sigma(z_i) = z_i$. The remainder of σ is defined by the property $\overline{\sigma(a)} = \sigma(\bar{a})$; hence σ is the identity. □

Lower bound. We use symmetry blocking to prove the following lemma.

Lemma 13. *If a QBF family requires $T(n)$-size* Q-Res *refutations, then its d-blocker requires $T(n)$-size* Q-Res+SL *refutations of degree $d - 1$.*

Our argument uses three low-level propositions. Each of them details a situation in which a refutation-size lower bound can be inferred from that of a simpler QBF. The first is well known, and states the closure of Q-Res under existential restrictions.

Proposition 14. *Given a* Q-Res *refutation π of a QBF F and an assignment ε to its existentials, $\pi[\varepsilon]$ is a* Q-Res *refutation of $F[\varepsilon]$ whose size is at most $|\pi|$.*

For the second proposition, we say that a QBF $F := Q \cdot f$ *subsumes* another $G := R \cdot g$ if the following two conditions hold:

- each clause in f is a subset of some clause in g;
- $\mathrm{vars}(F) \subseteq \mathrm{vars}(G)$, and, for each $x, y \in \mathrm{vars}(F)$, $x <_Q y \Rightarrow x <_R y$.

It is easy to see that QBFs cannot have smaller Q-Res refutations than those which subsume them.

Proposition 15. *Let F and G be false QBFs. If F subsumes G, then the shortest* Q-Res *refutation of G is no smaller than that of F.*

The third proposition is rather more specific to Q-Res+SL and symmetry blockers; therefore we include a proof.

Proposition 16. *Let F and G be variable-disjoint QBFs satisfying:*

(a) *the rightmost block of F and the leftmost block of G are oppositely quantified;*
(b) *a rightmost variable appears in every clause of F;*
(c) *a leftmost variable appears in every clause of F;*
(d) *the only symmetry of G is the identity.*

Then the shortest degree-d Q-Res+SL refutation of $F \otimes G$ is no smaller than that of G.

Proof. Let $\pi := \pi_0 \circ \rho_1 \circ \cdots \circ \rho_k$ be a degree d Q-Res+SL derivation from $F \otimes G$, and suppose that the identity is the only symmetry of G. Let $\pi' := \pi'_0 \circ \rho'_1 \circ \cdots \circ \rho'_k$ be the sequence obtained from π by removing all literals in variables of F; in particular, let $\pi_0 = C_1, \ldots, C_n$ and $\pi'_0 = C'_1, \ldots, C'_n$. By induction on n and d, we show that π' is a valid degree d Q-Res+SL derivation from G.

The base case $n = 0$ is trivial, since the derivation is empty. For the inductive step, we branch on the inference rule with which C_n was derived.

(a) Suppose that C_n was derived as a standard axiom. Then C'_n belongs to the matrix of G, and can be derived as an axiom.
(b) Suppose that C_n was derived as the resolvent of C_a and C_b over a pivot variable p. If p is in $\mathrm{vars}(G)$, then C'_n can be derived as a resolvent of C'_a and C'_b. On the other hand, if p is in $\mathrm{vars}(F)$, then C'_n can be derived from one of C'_a and C'_b by weakening.
(c) Suppose that C_n was derived by universal reduction from C_a. Then C'_n may be derived by universal reduction from C'_a.
(d) Suppose that C_n was derived by weakening from C_a. Then C'_n may be derived by weakening from C'_a.
(e) Suppose that C_n was derived by application of the symmetry σ to C_a. Then, by Proposition 10, σ is the identity on G, and $C'_n = C'_a$ can be derived by application of the identity symmetry.
(f) Suppose that C_n was introduced as an extra axiom, let α_i be its negation and let ρ_i be the corresponding referenced refutation.

If $\mathrm{vars}(F) \subseteq \mathrm{vars}(C_n)$, then ρ'_i is identical to ρ_i; that is, it is a valid refutation of $F \otimes G[\alpha_i]$ of degree at most $d-1$. Let α_i^F and α_i^G be the subassignments of α_i on the variables of F and G, respectively. It is easy to see that, since $F \otimes G[\alpha_i]$ is false by the soundness of Q-Res+SL, α_i^F falsifies every clause in the matrix of F. It follows that $G[\alpha_i^G] = F \otimes G[\alpha_i]$. Therefore C'_n can be introduced as an extra axiom, referencing ρ'_i.

On the other hand, if $\mathrm{vars}(F) \nsubseteq \mathrm{vars}(C_n)$, then we must have $\mathrm{vars}(C_n) \subset \mathrm{vars}(F)$, by conditions (a), (b) and (c), and the fact that the negation of C_n is a linear assignment to $F \otimes G$. Hence ρ_i is a refutation of $F[\alpha_i] \otimes G$ of degree at most $d-1$. By induction on degree, ρ'_i is a valid degree $d-1$ refutation of G. Thus C'_n, which is the empty clause, may be introduced as an extra axiom, referencing ρ'_i. □

With these three propositions, we are ready to prove Lemma 13.

Proof (of Lemma 13). Let $\{F_n\}_{n\in\mathbb{N}}$ be QBFs requiring $T(n)$-size Q-Res refutations, and let $\{\pi_n\}_{n\in\mathbb{N}}$ be degree $d-1$ Q-Res+SL refutations of $\{\mathcal{D}_d(F_n)\}_{n\in\mathbb{N}}$. We prove that $|\pi_n| \geq T(n)$ by induction on the degree d.

For the base case $d = 1$, recall that the only symmetry of $\mathcal{D}_1(F)$ is the identity, by Proposition 12. It follows that each π_n is a valid Q-Res refutation of $\mathcal{D}_1(F_n)$. Now, let α be the assignment $\{a, \bar{c}\}$. It is easy to see that $\mathcal{D}_1(F_n)[\alpha]$ is syntactically equivalent to F_n. Moreover, by the monotonic existential closure of Q-Res (Proposition 14), restriction of π_n by α yields a valid Q-Res refutation of $\mathcal{D}_1(F_n)[\alpha]$, whose size is no larger than $|\pi_n|$. Thus $|\pi_n| \geq T(n)$.

For the inductive step, let $d \geq 2$. We call an extra axiom A in π_n *short* if $\mathrm{vars}(A) \subseteq \mathrm{vars}(\mathcal{A}^d(F_n))$. We consider two cases.

(a) Suppose that π_n uses a short extra axiom. Observe that a short extra axiom A in π_n references some refutation ρ, whose degree is at most $d-1$, of the QBF

$$\mathcal{A}^d(F_n)[\beta] \otimes \mathcal{A}^{d-1}(F_n) \otimes \cdots \otimes \mathcal{A}^1(F_n)\,,$$

where β is the negation of A. Observe that this QBF is syntactically equivalent to

$$\mathcal{A}^1(F_n)[\beta] \otimes \mathcal{D}_{d-1}(F_n)\,.$$

Thus, by Proposition 16 and the inductive hypothesis, the size of $|\pi_n| \geq |\rho| \geq T(n)$.
(b) On the other hand, suppose that π_n uses no short extra axiom. By definition, π_n is of the form

$$\pi' \circ \rho_1 \circ \cdots \circ \rho_k$$

where π is a Q-Res+S refutation of $\mathcal{D}_d(F_n)$ using $k \geq 0$ extra axioms $A_1, \ldots, A_k$, referencing the refutations $\rho_1, \ldots, \rho_k$. By Proposition 12, the only symmetry of $\mathcal{D}_d(F_n)$ is the identity; therefore π' is in fact a Q-Res refutation of $\mathcal{D}_d(F_n)$ with the same extra axioms. We claim that every extra axiom in π_n is subsumed by some clause in $\mathcal{A}^d(F_n)$, in which case π' is a Q-Res refutation of a QBF that is subsumed by $\mathcal{A}^d(F_n)$. Since $\mathcal{A}^d(F_n)$ and $\mathcal{D}_1(F_n)$ are syntactically equivalent, and the latter requires $T(n)$-size Q-Res refutations, we have $|\pi_n| \geq T(n)$ by Proposition 15.

It remains to show that every extra axiom in π' is subsumed by some clause in $\mathcal{A}^d(F_n)$. To that end, let A_i be an extra axiom in π_n, and let α_i be its negation. Now, a rightmost variable of $\mathcal{A}^d(F_n)$ appears in every clause, namely c^d; the same is true of a leftmost variable of $\mathcal{A}^{d-1}(F_n)$, namely a^{d-1}. Since c^d and a^{d-1} are oppositely quantified, and α_i is a linear assignment to $\mathcal{D}_d(F_n)$, variable c^d must be absent from the variables of $\mathcal{D}_d(F_n)[\alpha_i]$. On the one hand, this implies that $\mathcal{A}^d(F_n)[\alpha_i]$ has no variables. On the other hand, $\mathcal{A}^d(F_n)[\alpha_i]$ cannot be true, for otherwise $\mathcal{D}_d(F_n)[\alpha_i]$, which is refuted by ρ_i, would also be true, contradicting the soundness of Q-Res+SL. It follows that the matrix of $\mathcal{A}^d(F_n)[\alpha_i]$ contains the empty clause. Thus some clause of $\mathcal{A}^d(F_n)$ subsumes the negation of α_i, and the claim follows. □

Upper bound. For the corresponding upper bound, we construct short degree-d refutations of the d-blocker for the equality family. A general construction, matching the upper-bound argument, is not possible here; to prove the correctness of our construction, we need to use the specifics of the formulas inside the d-blocker.

We give a brief description of the construction. The central idea is that restriction of $\mathcal{A}(\mathrm{EQ}_n)$ by the assignment $a \mapsto 1$ yields $\mathrm{EQ}_n \otimes \forall c \cdot \{\{c\}\}$, whose symmetries are (essentially) those of equality itself. This establishes a short refutation for $d = 1$.

For larger d, the construction is iterated. Restriction of $\mathcal{A}^1(\mathrm{EQ}_n)$ by $a^1 \mapsto 1$ unlocks new symmetries (again, essentially those of EQ_n), which are used in conjunction with extra axioms that reference short proofs in a repeated, nested fashion.

Lemma 17. *For each d in $\mathbb{N}$, the d-blocker of the equality family has $\mathcal{O}(n)$-size* Q-Res+SL *refutations of degree d.*

Proof. Let d be a fixed natural number. For each i in $[d]$, we let EQ_n^i be the QBF obtained from EQ_n by adding the superscript i to each variable occurrence, and let eq_n^i be its matrix; further, we define a set of clauses

$$f_n^i := \{\{x_1^i \cdots x_n^i u_1^i \cdots u_n^i\}\} \otimes \{\{\bar{z}_1^i\}\} \cdots \{\{\bar{z}_n^i\}\{z_1^i \cdots z_n^i\}\} \otimes \{\{c^i\}\}\,,$$

and a sequence π_n^i, read as two columns:

$$\begin{array}{ll}
\{x_1^i \cdots x_n^i u_1^i \cdots u_n^i z_1^i \cdots z_n^i\} & \{x_1^i \cdots x_n^i\} \\
\{x_1^i \cdots x_n^i u_1^i \cdots u_n^i \bar{z}_n^i\} & \{x_1^i \cdots \bar{x}_n^i\} \\
\{x_1^i \cdots x_n^i u_1^i \cdots u_n^i z_1^i \cdots z_{n-1}^i\} & \{x_1^i \cdots x_{n-1}^i\} \\
\{x_1^i \cdots x_n^i u_1^i \cdots u_n^i \bar{z}_{n-1}^i\} & \{x_1^i \cdots \bar{x}_{n-1}^i\} \\
\vdots & \vdots \\
\{x_1^i \cdots x_n^i u_1^i \cdots u_n^i z_1^1\} & \{x_1^i\} \\
\{x_1^i \cdots x_n^i u_1^i \cdots u_n^i \bar{z}_1^1\} & \{\bar{x}_1^i\} \\
\{x_1^i \cdots x_n^i u_1^i \cdots u_n^i\} & \square
\end{array}$$

Based on the proof of Theorem 4, it is easy to see that $\mathrm{seq}(f_n^i) \circ \pi_n^i$ is a Q-Res+S refutation of $\mathrm{EQ}_n^i \otimes F_c^i$ using extra axioms f_n^i, where $F_c^i := \forall c^i \cdot \{\{c^i\}\}$ and seq() denotes the clauses of a CNF written in an arbitrary fixed sequence.

Now we build short degree-d refutations ζ_n^d of $\mathcal{D}_d(\mathrm{EQ}_n)$. We define

$$\zeta_n^1 := \{a^1\}, \{\bar{a}^1\}, \square \circ \rho_n^1 \quad \text{and} \quad \rho_n^1 := \mathrm{seq}(\mathrm{eq}_n^1) \circ \pi_n^1\,;$$

further, for each $2 \leq i \leq d$, we define $\zeta_n^i := \{a^i\}, \{\bar{a}^i\}, \square \circ \zeta_n^{i-1} \circ \rho_n^i$, where

$$\rho_n^i := \mathrm{seq}(f_n^i \otimes \{\{a^{i-1}\}\}) \circ \mathrm{seq}(f_n^i \otimes \{\{\bar{a}^{i-1}\}\}) \circ \pi_n^i \circ \zeta_n^{i-2} \circ \rho_n^{i-1}\,.$$

Here, we take $\zeta_n^0 := \emptyset$.

We observe that $|\zeta_n^d| = \mathcal{O}(n)$. To see this, observe that there exists constant c_1, c_2 such that $|\zeta_n^1| \leq c_1 \cdot n$ and, for each i in $[d]$, $|\zeta_n^i| \leq c_2|\zeta_n^{i-1}|$. Hence $|\zeta_n^d| \leq c_1 c_2^{d-1} \cdot n$.

To finish the proof, we show two invariants by induction on d:

(1) ρ_n^{d-1} is a refutation of $\mathcal{D}_d(\mathrm{EQ}_n)[\alpha_d]$ of degree $d-1$;
(2) ζ_n^d is a refutation of $\mathcal{D}_d(\mathrm{EQ}_n)$ of degree d.

We make use of $\alpha_d : a^d \mapsto 1$ and $\bar{\alpha}_d : a^d \mapsto 0$, which are both linear assignments to $\mathcal{D}_d(\mathrm{EQ}_n)$.

For the base case $d = 1$, observe that ζ_n^1 is a degree-0 refutation of $\mathcal{D}_1(\mathrm{EQ}_n)$ using a single extra axiom $\{\bar{a}^d\}$ whose negation is α_1. Since $\mathrm{eq}_n^1 \otimes \{\{c^1\}\}$ subsumes f_n^1, ρ_n^1 is a degree 0 refutation of $\mathrm{EQ}_n^1 \otimes F_c^1$. This establishes invariant (1), since $\mathcal{D}_1(\mathrm{EQ}_n)[\alpha_1] = \mathrm{EQ}_n^1 \otimes F_c^1$. It follows that ζ_n^1 is indeed a degree-1 refutation of $\mathcal{D}_1(\mathrm{EQ}_n)$, establishing invariant (2).

For the inductive step, let $d \geq 2$. Every symmetry of $\mathrm{EQ}_n^d \otimes F_c^d$, when extended by the identity on the variables of $\mathcal{A}^{d-1}(\mathrm{EQ}_n) \otimes \cdots \otimes \mathcal{A}^1(\mathrm{EQ}_n)$, is a symmetry of

$$\mathcal{D}_d(\mathrm{EQ}_n)[\alpha_d] = \mathrm{EQ}_n^d \otimes F_c^d \otimes \mathcal{A}^{d-1}(\mathrm{EQ}_n) \otimes \cdots \otimes \mathcal{A}^1(\mathrm{EQ}_n)$$

by Proposition 10. It follows that

$$\mathrm{seq}(f_n^d \otimes \{\{a^{d-1}\}\}) \circ \mathrm{seq}(f_n^d \otimes \{\{\bar{a}^{d-1}\}\}) \circ \pi_n^d$$

is a degree-0 refutation of $\mathcal{D}_d(\mathrm{EQ}_n)[\alpha_d]$, given the first two terms as extra axioms. Let β_1 and β_2 be the negations of any clause in $\mathrm{seq}(f_n^d \otimes \{\{a^{d-1}\}\})$ and $\mathrm{seq}(f_n^d \otimes \{\{\bar{a}^{d-1}\}\})$, respectively. It is readily verified that both β_1 and β_2 are a linear assignments to $(\mathcal{D}_d(\mathrm{EQ}_n) \otimes F_c^d)[\alpha_d]$. By the inductive hypothesis, ζ_n^{d-2} is a refutation of

$$\mathcal{D}_d(\mathrm{EQ}_n)[\alpha_d][\beta_1] = \mathcal{D}_{d-2}(\mathrm{EQ}_n)$$

of degree $d-2$, where $\mathcal{D}_0(\mathrm{EQ}_n)$ is the QBF on the empty set of variables whose matrix contains only the empty clause. Also by the inductive hypothesis, ρ_n^{d-1} is a refutation of

$$\mathcal{D}_d(\mathrm{EQ}_n)[\alpha_d][\beta_2] = \mathcal{D}_{d-1}(\mathrm{EQ}_n)[\alpha_{d-1}]$$

of degree $d-1$. This establishes invariant (1).

Now, observe that ζ_n^d is a degree-0 refutation of $\mathcal{D}_d(\mathrm{EQ}_n)$ using two extra axioms $\{a^d\}$ and $\{\bar{a}^d\}$, whose respective negations are $\bar{\alpha}_d$ and α_d. By the inductive hypothesis, ζ_n^{d-1} is a refutation of $\mathcal{D}_{d-1}(\mathrm{EQ}_n)$ of degree $d-1$, and that QBF is equal to $\mathcal{D}_d(\mathrm{EQ}_n)[\bar{\alpha}_d]$. Hence, by invariant (1), ζ_n^d is indeed a degree-d refutation of $\mathcal{D}_d(\mathrm{EQ}_n)$, which establishes invariant (2). □

Our main result is an immediate consequence of Lemmata 13 and 17, and the fact that the equality family requires 2^n-size Q-Res refutations [7].

Theorem 18. *For each d in $\mathbb{N}$, there exists a QBF family that has $\mathcal{O}(n)$-size* Q-Res+SL *refutations of degree d and requires $2^{\Omega(n)}$-size refutations of degree $d-1$.*

6 Constraint Propagation

From a practical point of view, the following observation could be made: linear assignments do not cover constraint propagation (i.e. unit propagation and pure literal elimination, cf. [20]), whereas our lower bound formulas contain both unit clauses and pure literals.

However, as we show below, any lower bounds are easily adapted to hold in the presence of constraint propagation; in fact, one can easily modify a formula in such a way that propagation is rendered ineffective. Moreover, the current setup shows tighter upper bounds, since the system without constraint propagation is certainly no stronger. Thus, as far as proof complexity is concerned, including propagation in the system merely introduces unnecessary complications.

Indeed, with a simple addition, one can block all propagation, while (essentially) preserving the symmetry group. Given a QBF $F := Q \cdot f$ whose last block is universal, add a fresh block of existential variables p_1, p_2, q_1, q_2 to the end of the prefix Q, and replace the matrix f with

$$(f \otimes \{\{\bar{p}_1, \bar{p}_2\}, \{\bar{p}_1, p_2\}, \{p_1, \bar{p}_2\}, \{p_1, p_2\}\}) \cup (g \otimes \{\{q_1, \bar{q}_2\}, \{\bar{q}_1, q_2\}\}),$$

where g consists of the full set of unit clauses for F:

$$g := \{\{\bar{z}\} : z \in \mathrm{vars}(F)\} \cup \{\{z\} : z \in \mathrm{vars}(F)\}\,.$$

Let us call the result of this modification F'.

It is easy to see that no unit propagation or pure literal elimination can take place until all the variables of F are assigned, leaving only the fresh variables in the final block. Hence, propagation cannot enlarge the set of linear assignments. Notice also that $\{\{q_1, \bar{q}_2\}, \{\bar{q}_1, q_2\}\}$ is a satisfiable CNF. As a result, the clauses in $g \otimes \{\{q_1, \bar{q}_2\}, \{\bar{q}_1, q_2\}\}$ never contribute positively to a refutation, and one can assume without loss of generality that they never appear. Thus, Q-Res+SL hardness for F lifts straightforwardly to F', even if propagation is built into the proof system (viz. [20]).

On the other hand, the modified formulas are certainly no harder to refute; the symmetries of F are preserved when extended by the identity to the fresh variables, and every clause of the original matrix can be recovered in three resolution steps.

If the final block of F is existential, the same effect is achieved by inserting a fresh universal variable v directly before p_1, and taking a further matrix product with $\{\{v\}, \{\bar{v}\}\}$.

7 Conclusions

We introduced a theoretical model of QBF symmetry learning, which forms a hierarchy of proof systems of strictly increasing strength when the degree is bounded by a constant. Bounding the degree of a refutation corresponds to bounding the number of recomputations allowed on any single search branch of runs of QCDCL algorithms.

Since the number of paths explored correlates approximately with total running time of solvers, a sensible bound on degree – some fraction of the number of variables, for example – limits the total number of recomputations in terms of the length of the search. Our strict hierarchy shows the best possible separations for bounds of this type.

Our investigation of QBF here also applies to SAT as a special case. Considering purely existentially quantified formulas, the system Q-Res+S coincides with resolution with the symmetry rule as introduced by Krishnamurthy [19]. Similarly, when only allowing existential formulas, we obtain a propositional version of Q-Res+SL for symmetry learning in SAT. The methods we employed for QBF are sufficient to show that propositional symmetry learning is exponentially stronger than Krishnamurthy's system. We need only apply our symmetry blocking technique to formulas separating the latter from resolution (e.g. the pigeonhole formulas [25]).

Acknowledgments. Research was supported by grants from the John Templeton Foundation (grant no. 60842) and the Carl-Zeiss Foundation.

References

1. Aloul, F.A., Ramani, A., Markov, I.L., Sakallah, K.A.: Dynamic symmetry-breaking for boolean satisfiability. Ann. Math. Artif. Intell. 57(1), 59–73 (2009)
2. Audemard, G., Jabbour, S., Sais, L.: Symmetry breaking in quantified boolean formulae. In: Veloso, M.M. (ed.) IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence. pp. 2262–2267 (2007)
3. Audemard, G., Mazure, B., Sais, L.: Dealing with symmetries in quantified boolean formulas. In: SAT 2004 - The Seventh International Conference on Theory and Applications of Satisfiability Testing (2004)
4. Balabanov, V., Widl, M., Jiang, J.R.: QBF resolution systems and their proof complexities. In: Sinz, C., Egly, U. (eds.) International Conference on Theory and Practice of Satisfiability Testing (SAT). Lecture Notes in Computer Science, vol. 8561, pp. 154–169. Springer (2014)
5. Beame, P., Kautz, H.A., Sabharwal, A.: Towards understanding and harnessing the potential of clause learning. J. Artif. Intell. Res. (JAIR) 22, 319–351 (2004)
6. Benedetti, M., Mangassarian, H.: QBF-based formal verification: Experience and perspectives. Journal of Satisfiability, Boolean Modeling and Computation 5(1-4), 133–191 (2008)

7. Beyersdorff, O., Blinkhorn, J., Hinde, L.: Size, cost and capacity: A semantic technique for hard random QBFs. In: Karlin, A.R. (ed.) ACM Conference on Innovations in Theoretical Computer Science (ITCS). Leibniz International Proceedings in Informatics (LIPIcs), vol. 94, pp. 9:1–9:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018)
8. Beyersdorff, O., Chew, L., Janota, M.: Proof complexity of resolution-based QBF calculi. In: Mayr, E.W., Ollinger, N. (eds.) International Symposium on Theoretical Aspects of Computer Science (STACS). Leibniz International Proceedings in Informatics (LIPIcs), vol. 30, pp. 76–89. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2015)
9. Beyersdorff, O., Hinde, L., Pich, J.: Reasons for hardness in QBF proof systems. In: Lokam, S.V., Ramanujam, R. (eds.) Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS). LIPIcs, vol. 93, pp. 14:1–14:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017)
10. Blinkhorn, J., Beyersdorff, O.: Shortening QBF proofs with dependency schemes. In: Gaspers, S., Walsh, T. (eds.) International Conference on Theory and Practice of Satisfiability Testing (SAT). Lecture Notes in Computer Science, vol. 10491, pp. 263–280. Springer (2017)
11. Cook, S.A., Reckhow, R.A.: The relative efficiency of propositional proof systems. Journal of Symbolic Logic 44(1), 36–50 (1979)
12. Devriendt, J., Bogaerts, B., Bruynooghe, M., Denecker, M.: Improved static symmetry breaking for SAT. In: Creignou, N., Berre, D.L. (eds.) Theory and Applications of Satisfiability Testing - SAT 2016 - 19th International Conference. Lecture Notes in Computer Science, vol. 9710, pp. 104–122. Springer (2016)
13. Devriendt, J., Bogaerts, B., Cat, B.D., Denecker, M., Mears, C.: Symmetry propagation: Improved dynamic symmetry breaking in SAT. In: IEEE 24th International Conference on Tools with Artificial Intelligence, ICTAI 2012. pp. 49–56. IEEE Computer Society (2012)
14. Egly, U., Kronegger, M., Lonsing, F., Pfandler, A.: Conformant planning as a case study of incremental QBF solving. Annals of Mathematics and Artificial Intelligence 80(1), 21–45 (2017)
15. Faymonville, P., Finkbeiner, B., Rabe, M.N., Tentrup, L.: Encodings of bounded synthesis. In: Legay, A., Margaria, T. (eds.) International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS). Lecture Notes in Computer Science, vol. 10205, pp. 354–370. Springer (2017)
16. Kauers, M., Seidl, M.: Short proofs for some symmetric quantified boolean formulas. Inf. Process. Lett. 140, 4–7 (2018)
17. Kauers, M., Seidl, M.: Symmetries of quantified boolean formulas. In: Beyersdorff, O., Wintersteiger, C.M. (eds.) Theory and Applications of Satisfiability Testing - SAT 2018 - 21st International Conference, SAT 2018. Lecture Notes in Computer Science, vol. 10929, pp. 199–216. Springer (2018)
18. Kleine Büning, H., Karpinski, M., Flögel, A.: Resolution for quantified Boolean formulas. Information and Computation 117(1), 12–18 (1995)
19. Krishnamurthy, B.: Short proofs for tricky formulas. Acta Inf. 22(3), 253–275 (1985)
20. Lonsing, F., Egly, U., Seidl, M.: Q-resolution with generalized axioms. In: Creignou, N., Berre, D.L. (eds.) International Conference on Theory and Practice of Satisfiability Testing (SAT). Lecture Notes in Computer Science, vol. 9710, pp. 435–452. Springer (2016)
21. Pipatsrisawat, K., Darwiche, A.: On the power of clause-learning SAT solvers as resolution engines. Artificial Intelligence 175(2), 512–525 (2011)
22. Sakallah, K.A.: Symmetry and satisfiability. In: Biere, A., Heule, M., van Maaren, H., Walsh, T. (eds.) Handbook of Satisfiability, Frontiers in Artificial Intelligence and Applications, vol. 185, pp. 289–338. IOS Press (2009)
23. Schaafsma, B., Heule, M., van Maaren, H.: Dynamic symmetry breaking by simulating Zykov contraction. In: Kullmann, O. (ed.) Theory and Applications of Satisfiability Testing - SAT 2009, 12th International Conference. Lecture Notes in Computer Science, vol. 5584, pp. 223–236. Springer (2009)
24. Szeider, S.: The complexity of resolution with generalized symmetry rules. Theory Comput. Syst. 38(2), 171–188 (2005)
25. Urquhart, A.: The symmetry rule in propositional logic. Discrete Applied Mathematics 96-97, 177–193 (1999)

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il