

Symmetries of Quantified Boolean Formulas^{*}

Manuel Kauers¹ and Martina Seidl²

¹ Institute for Algebra, JKU Linz, Austria
manuel.kauers@jku.at

² Institute for Formal Models and Verification, JKU Linz, Austria
martina.seidl@jku.at

Abstract. While symmetries are well understood for Boolean formulas and successfully exploited in practical SAT solving, less is known about symmetries in quantified Boolean formulas (QBF). There are some works introducing adaptations of propositional symmetry breaking techniques, with a theory covering only very specific parts of QBF symmetries. We present a general framework that gives a concise characterization of symmetries of QBF. Our framework naturally incorporates the duality of universal and existential symmetries resulting in a general basis for QBF symmetry breaking.

1 Introduction

Mathematicians are generally advised [1] to exploit the symmetry in a given problem for solving it. In automated reasoning, however, symmetries are often exploited by destroying them. In this context, to destroy a symmetry means to enrich the given problem by additional constraints which tell the solver that certain parts of the search space are equivalent, so that it investigates only one of them. Such symmetry breaking techniques have been studied for a long time. They are particularly well developed in SAT [2] and CSP [3]. In CSP [4] it has been observed that it is appropriate to distinguish two kinds of symmetries: those of the problem itself and those of the solution set. In the present paper, we apply this idea to Quantified Boolean Formulas (QBF) [5].

Symmetry breaking for QBF has already been studied more than ten years ago [6–9], and it can have a dramatic effect on the performance of QBF solvers. As an extreme example, the instances of the KBKF benchmark set [10] are highly symmetric. For some problem sizes n , we applied the two configurations QRes (standard Q-resolution) and LD (long-distance resolution) of the solver DepQBF [11] to this benchmark set. For LD

n	Solving times (in sec)			
	w/o SB		with SB	
	QRes	LD	QRes	LD
10	0.3	0.5	0.4	0.4
20	160	0.5	0.4	0.4
40	> 3600	0.5	0.4	0.4
80	> 3600	0.7	0.4	0.4
160	> 3600	2.2	0.5	0.4
320	> 3600	12.3	0.6	0.5
640	> 3600	36.8	1.0	0.8
1280	> 3600	241.1	22.6	19.7
2560	> 3600	> 3600	215.7	155.2
5120	> 3600	> 3600	1873.2	1042.6

^{*} Parts of this work were supported by the Austrian Science Fund (FWF) under grant numbers NFN S11408-N23 (RiSE), Y464-N18, and SFB F5004.

it is known that it performs exponentially better than QRes on the KBKF formulas [12]. The table on the previous page shows the runtimes of DepQBF without and with symmetry breaking (SB). In particular, we enriched the formulas with symmetry breaking formulas over the existential variables. While QRes-DepQBF only solves two formulas without symmetry breaking, with symmetry breaking it even outperforms LD-DepQBF. Also for the LD configuration, the symmetry breaking formulas are beneficial. While this is an extreme example, symmetries appear not only in crafted formulas. In fact, we found that about 60% of the benchmarks used in the PCNF track of QBFEval [13] have nontrivial symmetries that could be exploited.

In this paper, we develop an explicit, uniform, and general theory for symmetries of QBFs. The theory is developed from scratch, and we include detailed proofs of all theorems. The pioneering work on QBF symmetries [6–9] largely consisted in translating symmetry breaking techniques well-known from SAT to QBF. This is not trivial, as universal quantifiers require special treatment. Since then, however, research on QBF symmetry breaking almost stagnated. We believe that more work is necessary. For example, we have observed that universal symmetry breakers as introduced in [8] fail to work correctly in modern clause-and-cube-learning QBF solvers when compactly provided as cubes. Although the encoding of the symmetry breaker for universal variables is provably correct in theory, it turns out to be incompatible with pruning techniques like pure literal elimination for which already the compatibility with learning is not obvious [14]. The cubes obtained from symmetry breaking are conceptually different than the learned cubes, because they do not encode a (partial) satisfying assignment of the clauses. As the pruning techniques usually only consider the clausal part of the formula, it can happen that they are wrongly applied in the presence of cubes stemming from a symmetry breaking formula over universal variables, affecting the correctness of the solving result.

We hope that the theory developed in this paper will help to resuscitate the interest in symmetries for QBF, lead to a better understanding of the interplay between symmetry breaking and modern optimization techniques, provide a starting point for translating recent progress made in SAT and CSP to the QBF world, and produce special symmetry breaking formulas that better exploit the unique features of QBF. Potential applications of our framework are the development of novel symmetry breaking formulas based on different orderings than the currently considered lexicographic ordering, the transfer of recent improvements in static symmetry breaking for SAT to QBF, as well as the establishment of dynamic symmetry breaking.

2 Quantified Boolean Formulas

Let $X = \{x_1, \dots, x_n\}$ be a finite set of propositional variables and $\text{BF}(X)$ be a set of *Boolean formulas* over X . The elements of $\text{BF}(X)$ are well-formed formulas built from the variables of X , truth constants \top (true) and \perp (false), as well as logical connectives according to a certain grammar. For most of the paper, we will


not need to be specific about the structure of the elements of $\text{BF}(X)$. We assume a well-defined semantics for the logical connectives, i.e., for every $\phi \in \text{BF}(X)$ and every assignment $\sigma: X \rightarrow \{\top, \perp\}$ there is a designated value $[\phi]_\sigma \in \{\top, \perp\}$ associated to ϕ and σ . In particular, we use \wedge (conjunction), \vee (disjunction), \leftrightarrow (equivalence), \rightarrow (implication), \oplus (xor), and \neg (negation) with their standard semantics for combining and negating formulas. Two formulas $\phi, \psi \in \text{BF}(X)$ are *equivalent* if for every assignment $\sigma: X \rightarrow \{\top, \perp\}$ we have $[\phi]_\sigma = [\psi]_\sigma$. We use lowercase Greek letters for Boolean formulas and assignments.

If $f: \text{BF}(X) \rightarrow \text{BF}(X)$ is a function and $\sigma: X \rightarrow \{\top, \perp\}$ is an assignment, the assignment $f(\sigma): X \rightarrow \{\top, \perp\}$ is defined through $f(\sigma)(x) = [f(x)]_\sigma$ ($x \in X$). A partial assignment is a function $\sigma: Y \rightarrow \{\top, \perp\}$ with $Y \subseteq X$. If σ is such a partial assignment and $\phi \in \text{BF}(X)$, then $[\phi]_\sigma$ shall refer to an element of $\text{BF}(X \setminus Y)$ such that for every assignment $\tau: X \rightarrow \{\top, \perp\}$ with $\tau|_Y = \sigma$ we have $[[\phi]_\sigma]_\tau = [\phi]_\tau$. For example, $[\phi]_\sigma$ could be the formula obtained from ϕ by replacing every variable $y \in Y$ by the truth value $\sigma(y)$ and then simplifying.

We use uppercase Greek letters to denote *quantified Boolean formulas* (QBFs). A QBF has the form $\Phi = P.\phi$ where $\phi \in \text{BF}(X)$ is a Boolean formula and P is a quantifier prefix for X , i.e., $P = Q_1x_1Q_2x_2 \dots Q_nx_n$ for $Q_1, \dots, Q_n \in \{\forall, \exists\}$. We only consider closed formulas, i.e., each element of X appears in the prefix. For a fixed prefix $P = Q_1x_1Q_2x_2 \dots Q_nx_n$, two variables x_i, x_j are said to belong to the same *quantifier block* if $Q_{\min(i,j)} = \dots = Q_{\max(i,j)}$.

Every QBF is either true or false. The truth value is defined recursively as follows: $\forall xP.\phi$ is true iff both $P.[\phi]_{\{x=\top\}}$ and $P.[\phi]_{\{x=\perp\}}$ are true, and $\exists xP.\phi$ is true iff $P.[\phi]_{\{x=\top\}}$ or $P.[\phi]_{\{x=\perp\}}$ is true. For example, $\forall x_1\exists x_2.(x_1 \leftrightarrow x_2)$ is true and $\exists x_1\forall x_2.(x_1 \leftrightarrow x_2)$ is false. The semantics of a QBF $P.\phi$ can also be described as a game for two players [15]: In the i th move, the truth value of x_i is chosen by the existential player if $Q_i = \exists$ and by the universal player if $Q_i = \forall$. The existential player wins if the resulting formula is true and the universal player wins if the resulting formula is false. In this interpretation, a QBF is true if there is a winning strategy for the existential player and it is false if there is a winning strategy for the universal player.

Strategies can be described as trees. Let $P = Q_1x_1Q_2x_2 \dots Q_nx_n$ be a prefix. An *existential strategy* for P is a tree of height $n + 1$ where every node at level $k \in \{1, \dots, n\}$ has one child if $Q_k = \exists$ and two children if $Q_k = \forall$. In the case $Q_k = \forall$, the two edges to the children are labeled by \top and \perp , respectively. In the case $Q_k = \exists$, the edge to the only child is labeled by either \top or \perp . *Universal strategies* are defined analogously, the only difference being that the roles of the quantifiers are exchanged, i.e., nodes at level k have two successors if $Q_k = \exists$ (one labeled \perp and one labeled \top) and one successor if $Q_k = \forall$ (labeled either \perp or \top). Here are the four existential strategies and the two universal strategies for the prefix $\forall x_1\exists x_2$:



We write $\mathbb{S}_\exists(P)$ for the set of all existential strategies and $\mathbb{S}_\forall(P)$ for the set of all universal strategies. As shown in the following lemma, the set of paths of a given existential strategy for prefix P is never disjoint from the set of paths of a given universal strategy. Unless otherwise stated, by a path, we mean a complete path starting at the root and ending at a leaf, together with the corresponding truth value labels.

Lemma 1 *If P is a prefix and $s \in \mathbb{S}_\exists(P)$, $t \in \mathbb{S}_\forall(P)$, then s and t have a path in common.*

Proof. A common path can be constructed by induction on the length of the prefix. There is nothing to show for prefixes of length 0. Suppose the claim holds for all prefixes of length n and consider a prefix $P' = P Q_{n+1} x_{n+1}$ of length $n+1$. Let $s \in \mathbb{S}_\exists(P')$, $t \in \mathbb{S}_\forall(P')$ be arbitrary. By chopping off the leafs of s and t , we obtain elements of $\mathbb{S}_\exists(P)$ and $\mathbb{S}_\forall(P)$, respectively, and these share a common path σ_0 by induction hypothesis. If $Q_{n+1} = \exists$, then σ_0 has a unique continuation in s , with an edge labeled either \top or \perp , and σ_0 has two continuations in t , one labeled \top and one labeled \perp , so the continuation of σ_0 in s must also appear in t . If $Q_{n+1} = \forall$, the argumentation is analogous. \square

Every path in a strategy for a prefix P corresponds to an assignment $\sigma: X \rightarrow \{\top, \perp\}$. An existential strategy for QBF $P.\phi$ is a *winning strategy* (for the existential player) if all its paths are assignments for which ϕ is true. A universal strategy is a *winning strategy* (for the universal player) if all its paths are assignments for which ϕ is false. For a QBF $P.\phi$ and an existential strategy $s \in \mathbb{S}_\exists(P)$, we define $[P.\phi]_s = \bigwedge_\sigma [\phi]_\sigma$, where σ ranges over all the assignments corresponding to a path of s . (Recall that our assignments are total unless otherwise stated, and our paths go from the root to a leaf unless otherwise stated.) Then we have $[P.\phi]_s = \top$ if and only if s is an existential winning strategy. For a universal strategy $t \in \mathbb{S}_\forall(P)$, we define $[P.\phi]_t = \bigvee_\tau [\phi]_\tau$, where τ ranges over all the assignments corresponding to a path of t . Then $[P.\phi]_t = \perp$ if and only if t is a universal winning strategy.

The definitions made in the previous paragraph are consistent with the interpretation of QBFs introduced earlier: a QBF is true if and only if there is an existential winning strategy, and it is false if and only if there is a universal winning strategy. Lemma 1 ensures that a QBF is either true or false. As another consequence of Lemma 1, observe that for every QBF $P.\phi$ we have

$$\begin{aligned} & (\exists s \in \mathbb{S}_\exists(P) : [P.\phi]_s = \top) \iff (\forall t \in \mathbb{S}_\forall(P) : [P.\phi]_t = \top) \\ \text{and } & (\forall s \in \mathbb{S}_\exists(P) : [P.\phi]_s = \perp) \iff (\exists t \in \mathbb{S}_\forall(P) : [P.\phi]_t = \perp). \end{aligned}$$

We will also need the following property, the proof of which is straightforward.

Lemma 2 *Let P be a prefix for X , and let $\phi, \psi \in \text{BF}(X)$. Then for all $s \in \mathbb{S}_\exists(P)$ we have $[P.(\phi \wedge \psi)]_s = [P.\phi]_s \wedge [P.\psi]_s$, and for all $t \in \mathbb{S}_\forall(P)$ we have $[P.(\phi \vee \psi)]_t = [P.\phi]_t \vee [P.\psi]_t$.*

3 Groups and Group Actions

Symmetries can be described using groups and group actions [16]. Recall that a group is a set G together with an associative binary operation $G \times G \rightarrow G$, $(g, h) \mapsto gh$. A group has a neutral element and every element has an inverse in G . A typical example for a group is the set \mathbb{Z} of integers together with addition. Another example is the group of permutations. For any fixed $n \in \mathbb{N}$, a permutation is a bijective function $\pi: \{1, \dots, n\} \rightarrow \{1, \dots, n\}$. The set of all such functions together with composition forms a group, called the symmetric group and denoted by S_n .

A (nonempty) subset H of a group G is called a subgroup of G if it is closed under the group operation and taking inverses. For example, the set $2\mathbb{Z}$ of all even integers is a subgroup of \mathbb{Z} , and the set $\{\text{id}, \begin{pmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{pmatrix}\}$ is a subgroup of S_3 . In general, a subset E of G is not a subgroup. However, for every subset E we can consider the intersection of all subgroups of G containing E . This is a subgroup and it is denoted by $\langle E \rangle$. The elements of E are called *generators* of the subgroup. For example, we have $2\mathbb{Z} = \langle 2 \rangle$, but also $2\mathbb{Z} = \langle 4, 6 \rangle$. A set of generators for S_3 is $\{\begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 2 & 3 \\ 2 & 1 & 3 \end{pmatrix}\}$.

If G is a group and S is a set then a *group action* is a map $G \times S \rightarrow S$, $(g, x) \mapsto g(x)$ which is compatible with the group operation, i.e., for all $g, h \in G$ and $x \in S$ we have $(gh)(x) = g(h(x))$ and $e(x) = x$, where e is the neutral element of G . Note that when we have a group action, every element $g \in G$ can be interpreted as a bijective function $g: S \rightarrow S$.

For example, for $G = S_n$ and $S = \{1, \dots, n\}$ we have a group action by the definition of the elements of S_n . Alternatively, we can let S_n act on a set of tuples of length n , say on $S = \{\square, \circ, \triangle\}^n$, via permutation of the indices, i.e., $\pi(x_1, \dots, x_n) = (x_{\pi(1)}, \dots, x_{\pi(n)})$. For example, for $g = \begin{pmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{pmatrix}$ we would have $g(\square, \circ, \square) = (\square, \square, \circ)$, $g(\triangle, \triangle, \square) = g(\triangle, \square, \triangle)$, $g(\circ, \triangle, \triangle) = (\circ, \triangle, \triangle)$, etc. As one more example, we can consider the group $G = S_n \times S_m$ consisting of all pairs of permutations. The operation for this group is defined component-wise, i.e., $(\pi, \sigma)(\pi', \sigma') = (\pi\pi', \sigma\sigma')$. We can let G act on a set of two dimensional arrays with shape $n \times m$, say on $S = \{\square, \circ, \triangle\}^{n \times m}$, by letting the first component of a group element permute the row index and the second component permute the column index. For example, for $g = (\begin{pmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{pmatrix}, \begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{pmatrix})$ we then have

$$g\left(\begin{array}{|c|c|c|} \hline \square & \circ & \triangle \\ \hline \square & \triangle & \square \\ \hline \circ & \circ & \square \\ \hline \end{array}\right) = \begin{array}{|c|c|c|} \hline \circ & \triangle & \square \\ \hline \circ & \square & \circ \\ \hline \triangle & \square & \square \\ \hline \end{array}.$$

If we have a group action $G \times S \rightarrow S$, we can define an equivalence relation on S via $x \sim y \iff \exists g \in G : x = g(y)$. The axioms of groups and group actions ensure that \sim is indeed an equivalence relation. The equivalence classes are called the *orbits* of the group action. For example, for the action of S_3 on $\{\square, \circ, \triangle\}^3$ discussed above, there are some orbits of size 1 (e.g., $\{(\circ, \circ, \circ)\}$), some orbits of size 3 (e.g., $\{(\square, \square, \triangle), (\square, \triangle, \square), (\triangle, \square, \square)\}$), and there is one orbit of size 6 ($\{(\square, \circ, \triangle), (\square, \triangle, \circ), (\circ, \triangle, \square), (\circ, \square, \triangle), (\triangle, \circ, \square), (\triangle, \square, \circ)\}$).

4 Syntactic Symmetries

In previous work [9], symmetries are characterized as permutations of literals with certain properties like being closed under negation, taking into account the order of the quantifiers, and, when extended to full formulas, always mapping a QBF to itself. As we will argue in the following, this point of view on QBF symmetries covers only a part of the full theory. We use group actions to describe symmetries of QBFs. Two kinds of group actions are of interest. On the one hand, we consider transformations that map formulas to formulas, i.e., a group action $G \times \text{BF}(X) \rightarrow \text{BF}(X)$. On the other hand, we consider transformations that map strategies to strategies, i.e., a group action $G \times \mathbb{S}_{\exists}(P) \rightarrow \mathbb{S}_{\exists}(P)$ or $G \times \mathbb{S}_{\forall}(P) \rightarrow \mathbb{S}_{\forall}(P)$. In both cases, we consider groups G which preserve the set of winning strategies for a given QBF $P.\phi$.

Let us first consider group actions $G \times \text{BF}(X) \rightarrow \text{BF}(X)$. In this case, we need to impose a technical restriction introduced in the following definition.

Definition 3 *Let P be a prefix for X . A bijective function $f: \text{BF}(X) \rightarrow \text{BF}(X)$ is called admissible (w.r.t. P) if*

1. *for every assignment $\sigma: X \rightarrow \{\top, \perp\}$ and every formula $\phi \in \text{BF}(X)$ we have $[\phi]_{f(\sigma)} = [f(\phi)]_{\sigma}$;*
2. *for every variable $x \in X$ the formula $f(x)$ only contains variables that belong to the same quantifier block of P as x .*

The first condition ensures that an admissible function f preserves propositional satisfiability. In particular, it implies that for any $\phi, \psi \in \text{BF}(X)$, the formulas $f(\neg\phi)$ and $\neg f(\phi)$ are equivalent, as are $f(\phi \circ \psi)$ and $f(\phi) \circ f(\psi)$ for every binary connective \circ . It follows that the inverse of an admissible function is again admissible. It also follows that an admissible function f is essentially determined by its values for the variables. Note that according to Definition 3 variables can be mapped to arbitrary formulas. The second condition may be replaced by a less restricted version, but for simplicity we use the conservative version of above.

Example 4 *Let $X = \{x, y, a, b\}$ and $P = \forall x \forall y \exists a \exists b$. There is an admissible function f with $f(x) = \neg x, f(y) = y, f(a) = b, f(b) = a$. For such a function, we may have $f(x \vee (a \rightarrow y)) = \neg x \vee (b \rightarrow y)$. A function g with $g(x) = b$ cannot be admissible, because of the second condition. By the first condition, a function h with $h(x) = x$ and $h(y) = \neg x$ cannot be admissible.*

Next we show that admissible functions not only preserve satisfiability of Boolean formulas, but also the truth of QBFs.

Theorem 5 *Let P be a prefix for X and $f: \text{BF}(X) \rightarrow \text{BF}(X)$ be admissible for P . For any $\phi \in \text{BF}(X)$ the formula $P.\phi$ is true if and only if $P.f(\phi)$ is true.*

Proof. Since the inverse of an admissible function is admissible, it suffices to show “ \Rightarrow ”. To do so, we proceed by induction on the number of quantifier blocks in P .

There is nothing to show when P is empty. Suppose the claim is true for all prefixes with k quantifier blocks, and consider a prefix $P = Qx_1 Qx_2 \cdots Qx_i P'$

for some $i \in \{1, \dots, n\}$, $Q \in \{\forall, \exists\}$, and a prefix P' for x_{i+1}, \dots, x_n with at most k quantifier blocks whose top quantifier is not Q . By the admissibility, we may view f as a pair of functions $f_1: \text{BF}(\{x_1, \dots, x_i\}) \rightarrow \text{BF}(\{x_1, \dots, x_i\})$ and $f_2: \text{BF}(\{x_{i+1}, \dots, x_n\}) \rightarrow \text{BF}(\{x_{i+1}, \dots, x_n\})$, where f_2 is admissible for P' . Let $s \in \mathbb{S}_\exists(P)$ be a winning strategy for $P.\phi$. We construct a winning strategy $t \in \mathbb{S}_\exists(P)$ for $P.f(\phi)$.

Case 1: $Q = \exists$. In this case, the upper i levels of s and t consist of single paths. Let $\sigma: \{x_1, \dots, x_i\} \rightarrow \{\top, \perp\}$ be the assignment corresponding to the upper i levels of s . The subtree s_σ of s rooted at the end of σ (level $i+1$) is a winning strategy for $P'.[\phi]_\sigma$. By induction hypothesis, $P'.f_2([\phi]_\sigma)$ has a winning strategy. Let t have an initial path corresponding to the assignment $\tau = f_1^{-1}(\sigma)$ followed by a winning strategy of $P'.f_2([\phi]_\sigma)$. (Since f_1 is invertible and independent of x_{i+1}, \dots, x_n , the assignment τ is well-defined.) Then t is a winning strategy of $P.f(\phi)$. To see this, let ρ be an arbitrary path of t . We show that $[f(\phi)]_\rho = \top$. Indeed,

$$\begin{array}{ccccccc}
& & t \text{ starts with } \tau & \text{Def. of } \tau & & \text{Def. of } f_1, f_2 & \\
& & \downarrow & \downarrow & & \downarrow & \\
[f(\phi)]_\rho & \stackrel{=}{=} & [[f(\phi)]_\tau]_\rho & \stackrel{=}{=} & [[f(\phi)]_{f_1^{-1}(\sigma)}]_\rho & \stackrel{=}{=} & [[f_1(f_2(\phi))]_{f_1^{-1}(\sigma)}]_\rho \\
& & \uparrow & & & & \uparrow \\
& & f_1 \text{ admissible} & & & & f_2 \text{ admissible} \quad \text{choice of } t
\end{array}$$

Case 2: $Q = \forall$. In this case, the upper i levels of both s and t form complete binary trees in which every path corresponds to an assignment for the variables x_1, \dots, x_i . Let $\tau: \{x_1, \dots, x_i\} \rightarrow \{\top, \perp\}$ be such an assignment, and let $\sigma = f_1(\tau)$. Let s_σ be the subtree of s rooted at σ . This is a winning strategy for the formula $P'.[\phi]_\sigma$ obtained from $P.\phi$ by instantiating the variables x_1, \dots, x_i according to σ and dropping the corresponding part of the prefix. By induction hypothesis, $P'.f_2([\phi]_\sigma)$ has a winning strategy. Pick one and use it as the subtree of t rooted at τ . The same calculation as in Case 1 shows that t is a winning strategy for $P.f(\phi)$. \square

Example 6 Consider the true QBF $\Phi = P.\phi = \forall x \forall y \exists a \exists b.((x \leftrightarrow a) \wedge (y \leftrightarrow b))$. If f is an admissible function with $f(x) = y$, $f(y) = x$, $f(a) = b$, $f(b) = a$, then obviously, $P.f(\phi)$ is true as well. If g is a non-admissible function with $g(x) = b$, $g(b) = x$, then $P.g(\phi)$ is false.

Next we introduce the concept of a *syntactic symmetry group*. The attribute ‘syntactic’ shall emphasize that this group acts on formulas, in contrast to the ‘semantic’ symmetry group introduced later, which acts on strategies. Our distinction between syntactic and semantic symmetries corresponds to the distinction between the problem and solution symmetries made in CSP [4].

Definition 7 Let $P.\phi$ be a QBF and let $G \times \text{BF}(X) \rightarrow \text{BF}(X)$ be a group action such that every $g \in G$ is admissible w.r.t. P . We call G a syntactic symmetry group for $P.\phi$ if ϕ and $g(\phi)$ are equivalent (i.e. $\phi \leftrightarrow g(\phi)$ is a tautology) for all $g \in G$.

It should be noticed that being a ‘symmetry group’ is strictly speaking not a property of the group itself but rather a property of the action of G on $\text{BF}(X)$. The elements of a symmetry group are called symmetries. In general, we call a group action admissible if every $g \in G$ is admissible. Definition 7 implies that when G is a syntactic symmetry group for $P.\phi$, then for every element $g \in G$ the QBF $P.g(\phi)$ has the same set of winning strategies as $P.\phi$. Note that this is not already a consequence of Thm. 5, which only said that $P.g(\phi)$ is true if and only if $P.\phi$ is true, which does not imply that they have the same winning strategies.

Example 8 Consider the QBF $\Phi = P.\phi = \forall x \forall y \exists a \exists b.((x \leftrightarrow a) \wedge (y \leftrightarrow b))$. A syntactic symmetry group for Φ is $G = \{\text{id}, f\}$, where f is an admissible function with $f(x) = y$, $f(y) = x$, $f(a) = b$, $f(b) = a$.

Symmetries are often restricted to functions which map variables to literals [9]. But this restriction is not necessary. Also the admissible function g defined by $g(x) = x$, $g(y) = x \oplus y$, $g(a) = a$, $g(b) = a \oplus b$ is a syntactic symmetry for Φ .


5 Semantic Symmetries

In SAT, considering syntactic symmetries is enough, because the solutions of Boolean formulas are variable assignments. As introduced in Section 2, the solutions of QBFs are tree-shaped strategies. In order to be able to permute certain subtrees of a strategy while keeping others untouched, we introduce semantic symmetry groups. For the definition of semantic symmetry groups, no technical requirement like the admissibility is needed. Every permutation of strategies that maps winning strategies to winning strategies is fine.

Definition 9 Let $\Phi = P.\phi$ be a QBF and let G be a group acting on $\mathbb{S}_{\exists}(P)$ (or on $\mathbb{S}_{\forall}(P)$). We call G a semantic symmetry group for Φ if for all $g \in G$ and all $s \in \mathbb{S}_{\exists}(P)$ (or all $s \in \mathbb{S}_{\forall}(P)$) we have $[\Phi]_s = [\Phi]_{g(s)}$.

A single syntactic symmetry can give rise to several distinct semantic symmetries, as shown in the following example.

Example 10 Consider again $\Phi = P.\phi = \forall x \forall y \exists a \exists b.((x \leftrightarrow a) \wedge (y \leftrightarrow b))$. The function f of the previous example, which exchanges x with y and a with b in the formula, can be translated to a semantic symmetry \tilde{f} :



This symmetry exchanges the labels of level 3 and level 4 and swaps the existential parts of the two paths in the middle. Regardless of the choice of $\alpha, \dots, \eta \in$

$\{\perp, \top\}$, the strategy on the left is winning if and only if the strategy on the right is winning, so \tilde{f} maps winning strategies to winning strategies.

Some further semantic symmetries can be constructed from f . For example, in order to be a winning strategy, it is necessary that $\alpha = \beta = \perp$. So we can take a function that just flips α and β but does not touch the rest of the tree. For the same reason, also a function that just flips η and ϑ but does not affect the rest of the tree is a semantic symmetry. The composition of these two functions and the function \tilde{f} described before (in an arbitrary order) yields a symmetry that exchanges γ with ζ and δ with ϵ but keeps $\alpha, \beta, \eta, \vartheta$ fixed. Also this function is a semantic symmetry.

The construction described in the example above works in general. Recall that for an assignment $\sigma: X \rightarrow \{\top, \perp\}$ and a function $f: \text{BF}(X) \rightarrow \text{BF}(X)$, the assignment $f(\sigma): X \rightarrow \{\top, \perp\}$ is defined by $f(\sigma)(x) = [f(x)]_\sigma$ for $x \in X$.

Lemma 11 *Let P be a prefix for X and g be an element of a group acting admissibly on $\text{BF}(X)$. Then there is a function $f: \mathbb{S}_\exists(P) \rightarrow \mathbb{S}_\exists(P)$ such that for all $s \in \mathbb{S}_\exists(P)$ we have that σ is a path of $f(s)$ if and only if $g(\sigma)$ is a path of s .*

Proof. Since g is an admissible function, it acts independently on variables belonging to different quantifier blocks. Therefore it suffices to consider the case where P consists of a single quantifier block. If all quantifiers are existential, then s consists of a single path, so the claim is obvious. If there are only universal quantifiers, then s consists of a complete binary tree containing all possible paths, so the claim is obvious as well. \square

Starting from a syntactic symmetry group G_{syn} , we can consider all the semantic symmetries that can be obtained from it like in the example above. All these semantic symmetries form a semantic symmetry group, which we call the semantic symmetry group associated to G_{syn} .

Definition 12 *Let P be a prefix for X and let $G_{\text{syn}} \times \text{BF}(X) \rightarrow \text{BF}(X)$ be an admissible group action. Let G_{sem} be the set of all bijective functions $f: \mathbb{S}_\exists(P) \rightarrow \mathbb{S}_\exists(P)$ such that for all $s \in \mathbb{S}_\exists(P)$ and every path σ of $f(s)$ there exists a $g \in G_{\text{syn}}$ such that $g(\sigma)$ is a path of s . This G_{sem} is called the associated group of G_{syn} .*

Again, it would be formally more accurate but less convenient to say that the action of G_{sem} on $\mathbb{S}_\exists(P)$ is associated to the action of G_{syn} on $\text{BF}(X)$.

Theorem 13 *If G_{syn} is a syntactic symmetry group for a QBF Φ , then the associated group G_{sem} of G_{syn} is a semantic symmetry group for Φ .*


Proof. Let $\Phi = P.\phi$. Obviously, G_{sem} is a group. To show that it is a symmetry group, let $s \in \mathbb{S}_\exists(P)$ be a winning strategy for Φ , and let $g_{\text{sem}} \in G_{\text{sem}}$. We show that $g_{\text{sem}}(s)$ is again a winning strategy. Let σ be a path of $g_{\text{sem}}(s)$. By Definition 12, there exists a $g_{\text{syn}} \in G_{\text{syn}}$ such that $g_{\text{syn}}(\sigma)$ is a path of s . Since s is a winning strategy, $[\phi]_{g_{\text{syn}}(\sigma)} = \top$, and since g_{syn} is admissible, $[\phi]_{g_{\text{syn}}(\sigma)} = [g_{\text{syn}}(\phi)]_\sigma$. Since g_{syn} is a symmetry, $[g_{\text{syn}}(\phi)]_\sigma = [\phi]_\sigma$, so reading backwards we have $[\phi]_\sigma = [g_{\text{syn}}(\phi)]_\sigma = [g_{\text{syn}}(\phi)]_\sigma = [\phi]_{g_{\text{syn}}(\sigma)} = \top$. Hence every path of $g_{\text{sem}}(s)$ is a satisfying assignment for ϕ , so $g_{\text{sem}}(s)$ is a winning strategy. \square

The distinction between a syntactic and a semantic symmetry group is immaterial when the prefix consists of a single quantifier block. In particular, SAT problems can be viewed as QBFs in which all quantifiers are \exists . For such formulas, each tree in $\mathbb{S}_{\exists}(P)$ consists of a single path, so in this case the requirement $\forall s \in \mathbb{S}_{\exists}(P) : [\mathcal{F}]_s = [\mathcal{F}]_{g(s)}$ from Definition 9 boils down to the requirement that $[\phi]_{\sigma} = [\phi]_{f(\sigma)}$ should hold for all assignments $\sigma : X \rightarrow \{\top, \perp\}$. This reflects the condition of Definition 7 that ϕ and $f(\phi)$ are equivalent.

As we have seen in Example 10, there is more diversity for prefixes with several quantifier blocks. In such cases, a single element of a syntactic symmetry group can give rise to a lot of elements of the associated semantic symmetry group. In fact, the associated semantic symmetry group is very versatile. For example, when there are two strategies $s, s' \in \mathbb{S}_{\exists}(P)$ and some element f of an associated semantic symmetry group G_{sem} such that $f(s) = s'$, then there is also an element $h \in G_{\text{sem}}$ with $h(s) = s'$, $h(s') = s$ and $h(r) = r$ for all $r \in \mathbb{S}_{\exists}(P) \setminus \{s, s'\}$. The next lemma is a generalization of this observation which indicates that G_{sem} contains elements that exchange subtrees across strategies.

Lemma 14 *Let $P = Q_1x_1 \dots Q_nx_n$ be a prefix and $G_{\text{syn}} \times \text{BF}(X) \rightarrow \text{BF}(X)$ be an admissible group action. Let G_{sem} be the associated group of G_{syn} . Let $s \in \mathbb{S}_{\exists}(P)$ and let σ be a path of s . Let $i \in \{1, \dots, n\}$ be such that $[x_j]_{\sigma} = [g(x_j)]_{\sigma}$ for all $g \in G_{\text{syn}}$ and all $j < i$.*

Further, let $f \in G_{\text{sem}}$ and $s' = f(s)$. Let σ' be a path of s' such that the first $i - 1$ edges of σ' agree with the first $i - 1$ edges of σ . By the choice of i such a σ' exists. Let $t, t' \in \mathbb{S}_{\exists}(Q_ix_i \dots Q_nx_n)$ be the subtrees of s, s' rooted at the i th node of σ, σ' , respectively, and let $s'' \in \mathbb{S}_{\exists}(P)$ be the strategy obtained from s by replacing t by t' , as illustrated in the picture below. Then there exists $h \in G_{\text{sem}}$ with $h(s) = s''$.



Proof. Define $h : \mathbb{S}_{\exists}(P) \rightarrow \mathbb{S}_{\exists}(P)$ by $h(s) = s''$, $h(s'') = s$, and $h(r) = r$ for all $r \in \mathbb{S}_{\exists}(P) \setminus \{s, s''\}$. Obviously, h is a bijective function from $\mathbb{S}_{\exists}(P)$ to $\mathbb{S}_{\exists}(P)$. To show that h belongs to G_{sem} , we must show that for every $r \in \mathbb{S}_{\exists}(P)$ and every path ρ of $h(r)$ there exists $g \in G_{\text{syn}}$ such that $g(\rho)$ is a path of r . For $r \in \mathbb{S}_{\exists}(P) \setminus \{s, s''\}$ we have $h(r) = r$, so there is nothing to show.


Consider the case $r = s$. Let ρ be a path of $h(r) = s''$. If ρ does not end in the subtree t' , then the same path ρ also appears in r and we can take $g = \text{id}$. Now suppose that ρ does end in the subtree t' . Then ρ is also a path of $s' = f(s)$, because all paths of s and s' ending in t or t' agree above the i th node. Since $f \in G_{\text{sem}}$, there exists $g \in G_{\text{syn}}$ such that $g(\rho)$ is a path of s .

Finally, consider the case $r = s''$. Let ρ be a path of $h(r) = s$. If ρ does not end in the subtree t , then the same path ρ also appears in r and we can take $g = \text{id}$. Now suppose that ρ does end in the subtree t . Then the first $i - 1$ edges of ρ agree with those of σ . Since $s = f^{-1}(s')$, there exists $g \in G_{\text{syn}}$ such that $g(\rho)$ is a path of s' . By assumption on G_{syn} , the element g fixes first $i - 1$ edges of ρ , so $g(\rho)$ ends in t' and is therefore a path of s'' , as required. \square

6 Existential Symmetry Breakers

The action of a syntactic symmetry group of a QBF $P.\phi$ splits $\text{BF}(X)$ into orbits. For all the formulas ψ in the orbit of ϕ , the QBF $P.\psi$ has exactly the same winning strategies as $P.\phi$. For finding a winning strategy, we therefore have the freedom of exchanging ϕ with any other formula in its orbit.

The action of a semantic symmetry group on $\mathbb{S}_{\exists}(P)$ splits $\mathbb{S}_{\exists}(P)$ into orbits. In this case, every orbit either contains only winning strategies for $P.\phi$ or no winning strategies for $P.\phi$ at all:




Instead of checking all elements of $\mathbb{S}_{\exists}(P)$, it is sufficient to check one element per orbit. If a winning strategy exists, then any such sample contains one.

To avoid inspecting strategies that belong to the same orbit symmetry breaking introduces a formula $\psi \in \text{BF}(X)$ which is such that $P.\psi$ has at least one winning strategy in every orbit. Such a formula is called a *symmetry breaker*. The key observation is that instead of solving $P.\phi$, we can solve $P.(\phi \wedge \psi)$. Every winning strategy for the latter will be a winning strategy for the former, and if the former has at least one winning strategy, then so does the latter. By furthermore allowing transformations of ϕ via a syntactic symmetry group, we get the following definition.

Definition 15 Let P be a prefix for X , let G_{syn} be a group acting admissibly on $\text{BF}(X)$ and let G_{sem} be a group action on $\mathbb{S}_{\exists}(P)$. A formula $\psi \in \text{BF}(X)$ is called an *existential symmetry breaker* for P (w.r.t. the actions of G_{syn} and G_{sem}) if for every $s \in \mathbb{S}_{\exists}(P)$ there exist $g_{\text{syn}} \in G_{\text{syn}}$ and $g_{\text{sem}} \in G_{\text{sem}}$ such that $[P.g_{\text{syn}}(\psi)]_{g_{\text{sem}}(s)} = \top$.

Example 16 Consider the formula $\Phi = P.\phi = \forall x \exists y \exists z.(y \leftrightarrow z)$. All the elements of $\mathbb{S}_{\exists}(P)$ have the form depicted on the right. As syntactic symmetries, we have the admissible functions $f, g: \text{BF}(X) \rightarrow \text{BF}(X)$ defined by $f(x) = x$, $f(y) = z$, $f(z) = y$, and $g(x) = x$, $g(y) = \neg y$, $g(z) = \neg z$, respectively, so we can take $G_{\text{syn}} = \langle f, g \rangle$ as a syntactic symmetry group.



According to [8, 9] the formula $\neg y$ is a symmetry breaker for $P.\phi$. When considering G_{syn} together with $G_{\text{sem}} = \{\text{id}\}$ (what would be sufficient for SAT), the complications for QBF become obvious. The orbit of $\neg y$ is $O = \{y, z, \neg y, \neg z\}$. Now consider the strategy with $\alpha = \top, \beta = \perp, \gamma = \perp, \delta = \top$. For any $\psi \in O$, this strategy does not satisfy $P.\psi$, because ψ is true on one branch, but false on the other. Using semantic symmetries can overcome this problem.

Semantic symmetries can act differently on different paths. Let $f_1: \mathbb{S}_{\exists}(P) \rightarrow \mathbb{S}_{\exists}(P)$ be the function which exchanges α, β and leaves γ, δ fixed, let $g_1: \mathbb{S}_{\exists}(P) \rightarrow \mathbb{S}_{\exists}(P)$ be the function which replaces α, β by $\neg\alpha, \neg\beta$ and leaves γ, δ fixed, and let $f_2, g_2: \mathbb{S}_{\exists}(P) \rightarrow \mathbb{S}_{\exists}(P)$ be defined like f_1, g_1 but with the roles of α, β and γ, δ exchanged. The group $G_{\text{sem}} = \langle f_1, g_1, f_2, g_2 \rangle$ is a semantic symmetry group for Φ . This group splits $\mathbb{S}_{\exists}(P)$ into four orbits: one orbit consists of all strategies with $\alpha = \beta, \gamma = \delta$, one consists of those with $\alpha = \beta, \gamma \neq \delta$, one consists of those with $\alpha \neq \beta, \gamma = \delta$, and one consists of those with $\alpha \neq \beta, \gamma \neq \delta$.

Taking $G_{\text{syn}} = \{\text{id}\}$ together with this group G_{sem} , the formula $\neg y$ is a symmetry breaker, because each orbit contains one element with $\alpha = \gamma = \perp$.

The following theorem is the main property of symmetry breakers.

Theorem 17 *Let $\Phi = P.\phi$ be a QBF. Let G_{syn} be a syntactic symmetry group and G_{sem} be a semantic symmetry group acting on $\mathbb{S}_{\exists}(P)$. Let ψ be an existential symmetry breaker for G_{syn} and G_{sem} . Then $P.\phi$ is true iff $P.(\phi \wedge \psi)$ is true.*

Proof. The direction “ \Leftarrow ” is obvious (by Lemma 2). We show “ \Rightarrow ”. Let $s \in \mathbb{S}_{\exists}(P)$ be such that $[\Phi]_s = \top$. Since Φ is true, such an s exists. Let $g_{\text{syn}} \in G_{\text{syn}}$ and $g_{\text{sem}} \in G_{\text{sem}}$ be such that $[P.g_{\text{syn}}(\psi)]_{g_{\text{sem}}(s)} = \top$. Since ψ is an existential symmetry breaker, such elements exist. Since G_{syn} and G_{sem} are symmetry groups, $[P.g_{\text{syn}}(\phi)]_{g_{\text{sem}}(s)} = [P.\phi]_s = \top$. Lemma 2 implies $[P.(g_{\text{syn}}(\phi) \wedge g_{\text{syn}}(\psi))]_{g_{\text{sem}}(s)} = \top$. By the compatibility with logical operations (admissibility),

$$[P.g_{\text{syn}}(\phi \wedge \psi)]_{g_{\text{sem}}(s)} = [P.(g_{\text{syn}}(\phi) \wedge g_{\text{syn}}(\psi))]_{g_{\text{sem}}(s)} = \top.$$

Now by Thm. 5 applied with g_{syn}^{-1} to $P.g_{\text{syn}}(\phi \wedge \psi)$, it follows that there exists s' such that $[P.(\phi \wedge \psi)]_{s'} = \top$, as claimed. \square

As a corollary, we may remark that for an existential symmetry breaker ψ for the prefix P the formula $P.\psi$ is always true. To see this, choose $\phi = \top$ and observe that any groups G_{syn} and G_{sem} are symmetry groups for ϕ . By the theorem, $P.(\phi \wedge \psi)$ is true, so $P.\psi$ is true.

7 Universal Symmetry Breakers

An inherent property of reasoning about QBFs is the duality between “existential” and “universal” reasoning [17], i.e., the duality between proving and refuting a QBF. For showing that a QBF is true, an existential strategy has to be found that is an existential winning strategy. An existential symmetry breaker

tightens the pool of existential strategies among which the existential winning strategy can be found (in case there is one).

If the given QBF is false, then a universal strategy has to be found that is a universal winning strategy. In this case, an existential symmetry breaker is not useful. Recall that a universal winning strategy is a tree in which all paths are falsifying assignments. Using an existential symmetry breaker as in Thm. 17 tends to increase the number of such paths and thus increases the number of potential candidates. To aid the search for a universal winning strategy, it would be better to increase the number of paths corresponding to satisfying assignments, because this reduces the search space for universal winning strategies. For getting symmetry breakers serving this purpose, we can use a theory that is analogous to the theory of the previous section.

Definition 18 *Let P be a prefix for X , let G_{syn} be a group acting admissibly on $\text{BF}(X)$ and let G_{sem} be a group action on $\mathbb{S}_{\forall}(P)$. A formula $\psi \in \text{BF}(X)$ is called a universal symmetry breaker for P (w.r.t. the actions of G_{syn} and G_{sem}) if for every $t \in \mathbb{S}_{\forall}(P)$ there exist $g_{\text{syn}} \in G_{\text{syn}}$ and $g_{\text{sem}} \in G_{\text{sem}}$ such that $[P.g_{\text{syn}}(\psi)]_{g_{\text{sem}}(t)} = \perp$.*

No change is needed for the definition of syntactic symmetry groups. A semantic symmetry group for $\Phi = P.\phi$ is now a group acting on $\mathbb{S}_{\forall}(P)$ in such a way that $[P.\phi]_t = [P.\phi]_{g(t)}$ for all $g \in G$ and all $t \in \mathbb{S}_{\forall}(P)$. With these adaptations, we have the following analog of Thm. 17.

Theorem 19 *Let $\Phi = P.\phi$ be a QBF. Let G_{syn} be a syntactic symmetry group and G_{sem} be a semantic symmetry group acting on $\mathbb{S}_{\forall}(P)$. Let ψ be a universal symmetry breaker for G_{syn} and G_{sem} . Then $P.\phi$ is false iff $P.(\phi \vee \psi)$ is false.*

The proof is obtained from the proof of Thm. 17 by replacing $\mathbb{S}_{\exists}(P)$ by $\mathbb{S}_{\forall}(P)$, every \wedge by \vee , every \top by \perp , and “existential” by “universal”.

We have seen before that for an existential symmetry breaker ψ_{\exists} the QBF $P.\psi_{\exists}$ is necessarily true. Likewise, for a universal symmetry breaker ψ_{\forall} , the QBF $P.\psi_{\forall}$ is necessarily false. This has the important consequence that existential and universal symmetry breakers can be used in combination, even if they are not defined with respect to the same group actions.

Theorem 20 *Let $\Phi = P.\phi$ be a QBF. Let G_{syn}^{\exists} and G_{syn}^{\forall} be syntactic symmetry groups of Φ , let G_{sem}^{\exists} be a semantic symmetry group of Φ acting on $\mathbb{S}_{\exists}(P)$ and let G_{sem}^{\forall} be a semantic symmetry group of Φ acting on $\mathbb{S}_{\forall}(P)$. Let ψ_{\exists} be an existential symmetry breaker for G_{syn}^{\exists} and G_{sem}^{\exists} , and let ψ_{\forall} be a universal symmetry breaker for G_{syn}^{\forall} and G_{sem}^{\forall} . Then $P.\phi$ is true iff $P.((\phi \vee \psi_{\forall}) \wedge \psi_{\exists})$ is true iff $P.((\phi \wedge \psi_{\exists}) \vee \psi_{\forall})$ is true.*

Proof. For the first equivalence, we have

$$\begin{aligned}
P.\phi \text{ is true} & \stackrel{\text{Thm. 19}}{\iff} P.(\phi \vee \psi_{\forall}) \text{ is true} \\
& \stackrel{\text{Def.}}{\iff} \exists s \in \mathbb{S}_{\exists}(P) : [P.(\phi \vee \psi_{\forall})]_s = \top \\
& \iff \exists s \in \mathbb{S}_{\exists}(P) : [P.(\phi \vee \psi_{\forall})]_s \wedge \underbrace{[P.\psi_{\exists}]_s}_{=\top} = \top \\
& \stackrel{\text{Lem. 2}}{\iff} \exists s \in \mathbb{S}_{\exists}(P) : [P.((\phi \vee \psi_{\forall}) \wedge \psi_{\exists})]_s = \top \\
& \stackrel{\text{Def.}}{\iff} P.((\phi \vee \psi_{\forall}) \wedge \psi_{\exists}) \text{ is true.}
\end{aligned}$$

The proof of the second equivalence is analogous. \square

Next we relate existential symmetry breakers to universal symmetry breakers. Observe that when P is a prefix and \tilde{P} is the prefix obtained from P by changing all quantifiers, i.e., replacing each \exists by \forall and each \forall by \exists , then $\mathbb{S}_{\exists}(P) = \mathbb{S}_{\forall}(\tilde{P})$. For any formula $\phi \in \text{BF}(X)$ and any $s \in \mathbb{S}_{\exists}(P) = \mathbb{S}_{\forall}(\tilde{P})$ we have $\neg[P.\phi]_s = [\tilde{P}.\neg\phi]_s$. Therefore, if G_{syn} is a group acting admissibly on $\text{BF}(X)$ and G_{sem} is a group acting on $\mathbb{S}_{\exists}(P) = \mathbb{S}_{\forall}(\tilde{P})$, we have

$$\begin{aligned}
& \psi \text{ is an existential symmetry breaker for } G_{\text{syn}} \text{ and } G_{\text{sem}} \\
& \iff \forall s \in \mathbb{S}_{\exists}(P) \exists g_{\text{syn}} \in G_{\text{syn}}, g_{\text{sem}} \in G_{\text{sem}} : [P.g_{\text{syn}}(\psi)]_{g_{\text{sem}}(s)} = \top \\
& \iff \forall s \in \mathbb{S}_{\forall}(\tilde{P}) \exists g_{\text{syn}} \in G_{\text{syn}}, g_{\text{sem}} \in G_{\text{sem}} : [\tilde{P}.\neg g_{\text{syn}}(\psi)]_{g_{\text{sem}}(s)} = \perp \\
& \iff \forall s \in \mathbb{S}_{\forall}(\tilde{P}) \exists g_{\text{syn}} \in G_{\text{syn}}, g_{\text{sem}} \in G_{\text{sem}} : [\tilde{P}.g_{\text{syn}}(\neg\psi)]_{g_{\text{sem}}(s)} = \perp \\
& \iff \neg\psi \text{ is a universal symmetry breaker for } G_{\text{syn}} \text{ and } G_{\text{sem}},
\end{aligned}$$

where admissibility of g_{syn} is used in the third step. We have thus proven the following theorem, which captures Property 2 of the symmetry breaker introduced in [8] by relating existential and universal symmetry breakers.

Theorem 21 *Let P be a prefix for X and let \tilde{P} be the prefix obtained from P by flipping all the quantifiers. Let G_{syn} be a group acting admissibly on $\text{BF}(X)$ and let G_{sem} be a group acting on $\mathbb{S}_{\exists}(P) = \mathbb{S}_{\forall}(\tilde{P})$. Then $\psi \in \text{BF}(X)$ is an existential symmetry breaker for G_{syn} and G_{sem} if and only if $\neg\psi$ is a universal symmetry breaker for G_{syn} and G_{sem} .*

8 Construction of Symmetry Breakers

Because of Thm. 21, it suffices to discuss the construction of existential symmetry breakers. A universal symmetry breaker is obtained in a dual manner. Given a symmetry group, the basic idea is similar as for SAT (see also the French thesis of Jabbour [9] for a detailed discussion on lifting SAT symmetry breaking techniques to QBF). First an order on $\mathbb{S}_{\exists}(P)$ is imposed, so that every orbit contains an element which is minimal with respect to the order. Then we construct

a formula ψ_{\exists} for which (at least) the minimal elements of the orbits are winning strategies. Any such formula is an existential symmetry breaker. One way of constructing an existential symmetry breaker is given in the following theorem, which generalizes the symmetry breaking technique by Crawford et al. [18]. We give a formal proof that we obtain indeed a QBF symmetry breaker and conclude with lifting a CNF encoding used in recent SAT solving technology [19] to QBF.

Theorem 22 *Let $P = Q_1x_1 \dots Q_nx_n$ be a prefix for X , let G_{syn} be a group acting admissibly on $\text{BF}(X)$, and let G_{sem} be the associated group of G_{syn} . Then*

$$\psi = \bigwedge_{\substack{i=1 \\ Q_i = \exists}}^n \bigwedge_{g \in G_{\text{syn}}} \left(\left(\bigwedge_{j < i} (x_j \leftrightarrow g(x_j)) \right) \rightarrow (x_i \rightarrow g(x_i)) \right)$$

is an existential symmetry breaker for G_{syn} and G_{sem} .

Proof. All elements of $\mathbb{S}_{\exists}(P)$ are trees with the same shape. Fix a numbering of the edge positions in these trees which is such that whenever two edges are connected by a path, the edge closer to the root has the smaller index. (One possibility is breadth first search order.) For any two distinct strategies $s_1, s_2 \in \mathbb{S}_{\exists}(P)$, there is then a minimal k such that the labels of the k th edges of s_1, s_2 differ. Define $s_1 < s_2$ if the label is \perp for s_1 and \top for s_2 , and $s_1 > s_2$ otherwise.

Let $s \in \mathbb{S}_{\exists}(P)$. We need to show that there are $g_{\text{syn}} \in G_{\text{syn}}$ and $g_{\text{sem}} \in G_{\text{sem}}$ such that $[g_{\text{syn}}(\psi)]_{g_{\text{sem}}(s)} = \top$. Let $g_{\text{syn}} = \text{id}$ and let g_{sem} be such that $\tilde{s} := g_{\text{sem}}(s)$ is as small as possible in the order defined above. We show that $[\psi]_{\tilde{s}} = \top$. Assume otherwise. Then there exists $i \in \{1, \dots, n\}$ with $Q_i = \exists$ and $g \in G_{\text{syn}}$ and a path σ in \tilde{s} with $[x_j]_{\sigma} = [g(x_j)]_{\sigma}$ for all $j < i$ and $[x_i]_{\sigma} = \top$ and $[g(x_i)]_{\sigma} = \perp$. By Lemma 11, the element $g \in G_{\text{syn}}$ can be translated into an element $f \in G_{\text{sem}}$ which maps \tilde{s} to a strategy $f(\tilde{s})$ which contains a path that agrees with σ on the upper $i-1$ edges but not on the i th. By Lemma 14, applied to the subgroup $H \subseteq G_{\text{syn}}$ consisting of all $h \in G_{\text{syn}}$ with $[x_j]_{\sigma} = [h(x_j)]_{\sigma}$ for all $j < i$, we may assume that $f(\tilde{s})$ and \tilde{s} only differ in edges that belong to the subtree rooted at the i th node of σ . As all these edges have higher indices, we have $\tilde{s} < s$, in contradiction to the minimality assumption on s . \square

Note that we do not need to know the group G_{sem} explicitly. It is only used implicitly in the proof.

In nontrivial applications, G_{syn} will have a lot of elements. It is not necessary (and not advisable) to use them all, although Thm. 22 would allow us to do so. In general, if a formula $\psi_1 \wedge \psi_2$ is an existential symmetry breaker, then so are ψ_1 and ψ_2 , so we are free to use only parts of the large conjunctions. A reasonable choice is to pick a set E of generators for G_{syn} and let the inner conjunction run over (some of) the elements of E .

The formula ψ of Thm. 22 can be efficiently encoded as conjunctive normal form (CNF), adopting the propositional encoding of [19, 2]: let $g \in G_{\text{syn}}$ and let

$\{y_0^g, \dots, y_{n-1}^g\}$ be a set of fresh variables. First, we define a set I^g of clauses that represent all implications $x_i \rightarrow g(x_i)$ of ψ from Thm. 17,

$$I^g = \{(\neg y_{i-1}^g \vee \neg x_i \vee g(x_i)) \mid 1 \leq i \leq n, Q_i = \exists\}.$$

When x_i is existentially quantified, by using Tseitin variables y_{i-1}^g we can recycle the implications $x_i \rightarrow g(x_i)$ in the encoding of the equivalences $x_j \leftrightarrow g(x_j)$ that appear in the outer implication:

$$E^g = \{(y_j^g \vee \neg y_{j-1}^g \vee \neg x_j) \wedge (y_j \vee \neg y_{j-1} \vee g(x_j)) \mid 1 \leq j < n, Q_j = \exists\}.$$

If variable x_j is universally quantified, the recycling is not possible, so we use

$$U^g = \{(y_j^g \vee \neg y_{j-1}^g \vee \neg x_j \vee \neg g(x_j)) \wedge (y_j^g \vee \neg y_{j-1}^g \vee x_j \vee g(x_j)) \mid 1 \leq j < n, Q_j = \forall\}$$

instead. The CNF encoding of ψ is then the conjunction of y_0^g and all the clauses in I^g , E^g , and U^g , for all desired $g \in G_{\text{syn}}$. The prefix P has to be extended by additional quantifiers which bind the Tseitin variables y_i^g . As explained in [20], the position of such a new variable in the prefix has to be behind the quantifiers of the variables occurring in its definition. The encoding of universal symmetry breakers works similarly and results in a formula in disjunctive normal form (DNF), i.e., a disjunction of cubes (where a cube is a conjunction of literals). In this case the auxiliary variables are universally quantified. The obtained cubes could be used by solvers that simultaneously reason on the CNF and DNF representation of a formula (e.g., [21, 22]) or by solvers that operate on formulas of arbitrary structure (e.g., [23, 24, 22]). The practical evaluation of this approach is a separate topic which we leave to future work.

Besides the practical evaluation of the discussed symmetry breakers in connection with recent QBF solving technologies there are many more promising directions for future work. Also different orderings than the lexicographic order applied in Thm. 22 could be used [25] for the construction of novel symmetry breakers. Recent improvements of static symmetry breaking [19] for SAT could be lifted to QBF and applied in combination with recent preprocessing techniques. Also dynamic symmetry breaking during the solving could be beneficial, for example in the form of symmetric explanation learning [26].

An other interesting direction would be the relaxation of the quantifier ordering. Our symmetry framework assumes a fixed quantifier prefix with a strict ordering. In recent works it has been shown that relaxing this order by the means of dependency schemes is beneficial for QBF solving both in theory and in practice [27, 28]. In a similar way as proof systems have been parameterized with dependency schemes, our symmetry framework can also be parameterized with dependency schemes. It can be expected that a more relaxed notion of quantifier dependencies induces more symmetries resulting in more powerful symmetry breakers.

References

1. Polya, G.: How to solve it: A new aspect of mathematical method. Princeton university press (1945)

2. Sakallah, K.A.: Symmetry and satisfiability. In: Handbook of Satisfiability. Volume 185 of Frontiers in Artificial Intelligence and Applications. IOS Press (2009) 289–338
3. Gent, I.P., Petrie, K.E., Puget, J.: Symmetry in constraint programming. In: Handbook of Constraint Programming. Volume 2 of Foundations of Artificial Intelligence. Elsevier (2006) 329–376
4. Cohen, D.A., Jeavons, P., Jefferson, C., Petrie, K.E., Smith, B.M.: Constraint symmetry and solution symmetry. In: Proc. of the 21st Nat. Conf. on Artificial Intelligence and the 18th Innovative Applications of Artificial Intelligence Conf. (AAAI/IAAI'06), AAAI Press (2006) 1589–1592
5. Kleine Büning, H., Bubeck, U.: Theory of quantified boolean formulas. In: Handbook of Satisfiability. Volume 185 of Frontiers in Artificial Intelligence and Applications. IOS Press (2009) 735–760
6. Audemard, G., Mazure, B., Sais, L.: Dealing with Symmetries in Quantified Boolean Formulas. In: Proc. of the 7th Int. Conf. on Theory and Applications of Satisfiability Testing (SAT'04), Online Proceedings. (2004)
7. Audemard, G., Jabbour, S., Sais, L.: Symmetry Breaking in Quantified Boolean Formulae. In: Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI'07). (2007) 2262–2267
8. Audemard, G., Jabbour, S., Sais, L.: Efficient symmetry breaking predicates for Quantified Boolean Formulae. In: Proc. of Workshop on Symmetry and Constraint Satisfaction Problems (SymCon'07). (2007) 7 pages
9. Jabbour, S.: De la satisfiabilité propositionnelle aux formules booléennes quantifiées. PhD thesis, CRIL, Lens, France, 2008 (2008)
10. Kleine Büning, H., Karpinski, M., Flögel, A.: Resolution for quantified boolean formulas. *Inf. Comput.* **117**(1) (1995) 12–18
11. Lonsing, F., Egly, U.: Depqbf 6.0: A search-based QBF solver beyond traditional QCDCL. In: Proc. of the 26th Int. Conf. on Automated Deduction (CADE'17). Volume 10395 of LNCS., Springer (2017) 371–384
12. Egly, U., Lonsing, F., Widl, M.: Long-distance resolution: Proof generation and strategy extraction in search-based QBF solving. In: Proc. of the 19th Int. Conf. on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR'13). Volume 8312 of LNCS., Springer (2013) 291–308
13. Pulina, L., Seidl, M.: The QBF Eval 2017. <http://www.qbflib.org/qbfeval17>
14. Giunchiglia, E., Narizzano, M., Tacchella, A.: Monotone literals and learning in QBF reasoning. In: Proc. of the 10th Int. Conf. on Principles and Practice of Constraint Programming (CP'04). Volume 3258 of LNCS., Springer (2004) 260–273
15. Papadimitriou, C.H.: Computational complexity. Addison-Wesley (1994)
16. Artin, M.: Algebra. Pearson Prentice Hall (2011)
17. Sabharwal, A., Ansótegui, C., Gomes, C.P., Hart, J.W., Selman, B.: QBF modeling: Exploiting player symmetry for simplicity and efficiency. In: Proc. of the 9th Int. Conf on Theory and Applications of Satisfiability Testing (SAT'06). Volume 4121 of LNCS., Springer (2006) 382–395
18. Crawford, J.M., Ginsberg, M.L., Luks, E.M., Roy, A.: Symmetry-breaking predicates for search problems. In: Proc. of the 5th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'96), Morgan Kaufmann (1996) 148–159
19. Devriendt, J., Bogaerts, B., Bruynooghe, M., Denecker, M.: Improved static symmetry breaking for SAT. In: Proc. of the 19th Int. Conf. on Theory and Applications of Satisfiability Testing (SAT'16). Volume 9710 of LNCS., Springer (2016) 104–122

20. Egly, U., Seidl, M., Tompits, H., Woltran, S., Zolda, M.: Comparing different prenexing strategies for quantified boolean formulas. In: Proc. of the 6th Int. Conf. on Theory and Applications of Satisfiability Testing (SAT'03). Volume 2919 of LNCS., Springer (2003) 214–228
21. Goultiaeva, A., Seidl, M., Biere, A.: Bridging the gap between dual propagation and CNF-based QBF solving. In: Proc. of the Int. Conf. on Design, Automation and Test in Europe (DATE'13), EDA Consortium San Jose, CA, USA / ACM DL (2013) 811–814
22. Janota, M., Klieber, W., Marques-Silva, J., Clarke, E.M.: Solving QBF with counterexample guided refinement. *Artif. Intell.* **234** (2016) 1–25
23. Janota, M.: QFUN: towards machine learning in QBF. *CoRR* **abs/1710.02198** (2017)
24. Tentrup, L.: Non-prenex QBF solving using abstraction. In: Proc. of the 19th Int. Conf. on Theory and Applications of Satisfiability Testing (SAT'16). Volume 9710 of LNCS., Springer (2016) 393–401
25. Narodytska, N., Walsh, T.: Breaking symmetry with different orderings. In: Proc. of the 19th Int. Conf. on Principles and Practice of Constraint Programming (CP'13). Volume 8124 of LNCS., Springer (2013) 545–561
26. Devriendt, J., Bogaerts, B., Bruynooghe, M.: Symmetric explanation learning: Effective dynamic symmetry handling for SAT. In: Proc. of the 20th Int. Conf. on Theory and Applications of Satisfiability Testing (SAT'17). Volume 10491 of LNCS., Springer (2017) 83–100
27. Blinkhorn, J., Beyersdorff, O.: Shortening QBF proofs with dependency schemes. In: Proc. of the 20th Int. Conf. on Theory and Applications of Satisfiability Testing (SAT'17). Volume 10491 of LNCS., Springer (2017) 263–280
28. Peitl, T., Slivovsky, F., Szeider, S.: Dependency learning for QBF. In: Proc. of the 20th Int. Conf on Theory and Applications of Satisfiability Testing (SAT'17). Volume 10491 of LNCS., Springer (2017) 298–313