Noname manuscript No.
(will be inserted by the editor)

Symmetries of Dependency Quantified Boolean Formulas

Clemens Hofstadler · Manuel Kauers · Martina Seidl

Received: date / Accepted: date

Abstract Symmetries have been exploited successfully within the realms of SAT and QBF to improve solver performance in practical applications and to devise more powerful proof systems. As a first step towards extending these advancements to the class of dependency quantified Boolean formulas (DQBFs), which generalize QBF by allowing more nuanced variable dependencies, this work develops a comprehensive theory to characterize symmetries for DQBFs. We also introduce the notion of symmetry breakers of DQBFs, along with a concrete construction, and discuss how to detect DQBF symmetries algorithmically using a graph-based approach.

1 Introduction

Symmetry is an omnipresent phenomenon that we encounter in different forms in all parts of our lives. From the double helix structure of DNA (exhibiting two-fold rotation symmetry) on the microscopic scale to the rotational symmetry of galaxies on the cosmic scale. Symmetries also play a crucial role in automated reasoning, where symmetries of problem instances can be used to simplify the solving process. In practical applications, they can be used to incorporate additional constraints into a problem, which guide a solver away from equivalent parts of the search space, accelerating the search [1,3,15]. On the theoretical side, symmetries can enhance proof systems by introducing new

Parts of this work have been supported by the LIT AI Lab funded by the state of Upper Austria and by the Austrian Science Fund (FWF) [10.55776/COE12].

C. Hofstadler, M. Seidl
Institute for Symbolic Artificial Intelligence, JKU Linz, Austria
{clemens.hofstadler,martina.seidl}@jku.at

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

deduction rules that exploit symmetries, ultimately resulting in exponentially more powerful proof systems [16,21,14]

Such symmetry breaking techniques rely on a solid theoretical foundation for describing and understanding symmetries for different problem classes. This theory has been developed most prominently for the propositional satisfiability problem (SAT) [19] and for constraint satisfaction problems (CSP) [13]. Two of the authors have also developed a theory of symmetries for quantified Boolean formulas (QBFs) [15], extending earlier work on the subject [4,3,5]. In this work, we generalize the theory from [15] from QBFs to dependency quantified Boolean formulas.

Dependency quantified Boolean formulas (DQBFs) [18,6] represent a rich and expressive class of logical formulas that extends QBF by allowing existentially quantified variables to depend on specific subsets of universally quantified variables. In contrast to QBFs, which can only encode linear dependencies between variables, the nuanced quantification of DQBFs allows also for non-linear dependencies. This makes DQBFs a potent framework for encoding a variety of problems in verification, synthesis, and soft-/hardware engineering, see [20,7] and references therein. The extended expressive power, however, comes at the cost of increased computational complexity – the decision problem for DQBFs is NEXPTIME-complete [18]. This necessitates a need for advanced methods for solving DQBF instances efficiently. One promising avenue for mitigating the inherent complexity is the exploitation of symmetries.

In this work, we develop a comprehensive and explicit theory of symmetries for DQBFs, generalizing concepts established for SAT and QBF. In particular, analogous to the case of QBF [15] (and CSP [9]), we distinguish between two kinds of symmetries: those of the problem itself, which we call *syntactic symmetries*, and those of the solutions, which we call *semantic symmetries*. We use the concepts of groups and group actions to formally characterize these symmetries. All required concepts will be recalled, and we provide rigorous proofs of all our results.

One way to exploit symmetries in practice is to extend a given formula with additional constraints that destroy the formula's symmetries and thereby guide a solver away from equivalent areas of the search space. This approach is called *(static) symmetry breaking* and the formula encoding the additional constraints is called a *symmetry breaker*. In this work, we introduce the notion of *(conjunctive) symmetry breakers* for DQBFs and we provide a concrete construction for such symmetry breakers, generalizing ideas from SAT [8] and QBF [15, Sec. 8]. Finally, we also describe how to detect symmetries in DQBFs algorithmically with the help of graph-theoretic methods.

This work extends the symmetry framework for quatified Boolean formulas that was presented at the SAT 2018 conference [15] to the more general case of dependency quantified Boolean formulas.

2 Dependency Quantified Boolean Formulas

Let $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_k\}$ be two finite disjoint sets of propositional variables. For $V \subseteq X \cup Y$, we denote by $\mathrm{BF}(V)$ a set of *(propositional) Boolean formulas* over the variables V. The set $\mathrm{BF}(V)$ contains all well-formed formulas built from the truth constants $\top$ (true) and $\bot$ (false), from the variables in V, and from logical connectives according to some grammar. We note that we make no restrictions on the syntactic structure of the elements in $\mathrm{BF}(V)$ (except for Section 7, where we restrict to formulas in conjunctive normal form). Boolean formulas will be denoted by lowercase Greek letters $\phi, \psi, \dots$.

An *assignment* for a set of variables $V \subseteq X \cup Y$ is a function $\sigma\colon V \to \{\top, \bot\}$. The set of all assignments for V is denoted by $\mathbb{A}(V)$. We assume a well-defined semantics for the logical connectives used to construct the Boolean formulas in $\mathrm{BF}(V)$. In particular, we use the typical operations $\neg$ (negation), $\wedge$ (conjunction), $\vee$ (disjunction), $\leftrightarrow$ (equivalence), $\rightarrow$ (implication), and $\oplus$ (xor) with their standard semantics. Then, every assignment σ extends naturally to a function $[\cdot]_\sigma\colon \mathrm{BF}(V) \to \{\top, \bot\}$, mapping every Boolean formula $\phi \in \mathrm{BF}(V)$ to its *truth value* $[\phi]_\sigma \in \{\top, \bot\}$ under σ.

A *quantified Boolean formula (QBF)* (in prenex form) on a set of variables $V = \{v_1, \dots, v_m\}$ is a formula of the form $Q_1 v_1 Q_2 v_2 \dots Q_m v_m.\phi$, with quantifiers $Q_1, \dots, Q_m \in \{\forall, \exists\}$ and $\phi \in \mathrm{BF}(V)$. In a QBF, if a variable v_i is existentially quantified, i.e., $Q_i = \exists$, then v_i depends semantically on all universally quantified variables v_j with $j < i$. This leads to a linear dependency structure of the variables.

Dependency quantified Boolean formulas (DQBFs) [18] generalize QBFs by allowing *non-linear* dependencies of the variables, see also [6, Ch. 4] for an introduction. These dependencies are specified by explicitly annotating each existential variable with a set of universal variables. This is formalized by considering, for any k subsets $D_1, \dots, D_k \subseteq X$, a *prefix* for X and Y of the form $\forall x_1, \dots, x_n \exists y_1(D_1), \dots, y_k(D_k)$. The set D_i encodes that the existential variable y_i only depends on the universal variables in D_i and is called the *dependency set* of y_i.

Definition 1 Given a prefix $P = \forall x_1, \dots, x_n \exists y_1(D_1), \dots, y_k(D_k)$ for X and Y with dependency sets $D_1, \dots, D_k \subseteq X$ and a Boolean formula $\phi \in \mathrm{BF}(X \cup Y)$, the formula

$$P.\phi = \forall x_1, \dots, x_n \exists y_1(D_1), \dots, y_k(D_k).\phi$$

is called a *dependency quantified Boolean formula (DQBF)*.

We will denote DQBFs by uppercase Greek letters $\Phi, \Psi, \dots$. Note that, by definition, DQBFs are always closed formulas, meaning that each variable in $X \cup Y$ is quantified in the prefix.

Example 1 An example of a DQBF is

$$\forall x_1, x_2 \exists y_1\big(\{x_1\}\big), y_2\big(\{x_2\}\big).\, (\neg x_1 \rightarrow y_1) \wedge (x_2 \vee y_2)\,.$$

Note that this formula cannot be written as a QBF (in prenex form) because the quantifier dependencies cannot be expressed linearly. Conversely, however, every QBF can be expressed as a suitable DQBF. For example, any QBF of the form $\forall x_1 \exists y_1 \forall x_2 \exists y_2 . \phi$, with $\phi \in \mathrm{BF}(\{x_1, x_2, y_1, y_2\})$ can be expressed as $\forall x_1, x_2 \exists y_1\big(\{x_1\}\big), y_2\big(\{x_1, x_2\}\big) . \phi$. Note that the linear dependency structure of the QBF causes the dependency sets of the corresponding DQBF to form an increasing sequence. This is the case for every DQBF that arises from a QBF (in prenex form).

For a prefix $P = \forall x_1, \ldots, x_n \exists y_1(D_1), \ldots, y_k(D_k)$, an *interpretation* for P is a tuple $s = (s_1, \ldots, s_k)$ of functions $s_i \colon \{\top, \bot\}^{|D_i|} \to \{\top, \bot\}$, for $i = 1, \ldots, k$. Each function s_i specifies the truth value of the existential variable y_i in dependence of the truth values of the universal variables in D_i. The functions s_i are called *Skolem functions*. We denote by $\mathbb{S}(P)$ the set of all interpretations for P.

Remark 1 Every Skolem function $s_i \colon \{\top, \bot\}^{|D_i|} \to \{\top, \bot\}$ with dependency set $D_i = \{x_{i_1}, \ldots, x_{i_d}\}$ $(i_1 < \cdots < i_d)$ can be represented by a Boolean formula $\phi_i \in \mathrm{BF}(D_i)$, so that, for every assignment $\sigma \in \mathbb{A}(D_i)$,

$$s_i(\sigma(x_{i_1}), \ldots, \sigma(x_{i_d})) = [\phi_i]_\sigma .$$

Therefore, an interpretation can be represented as a tuple of such Boolean formulas. In the following, we will represent interpretations in this way.

Example 2 Consider the prefix $P = \forall x_1, x_2 \exists y_1\big(\{x_1\}\big), y_2\big(\{x_2\}\big)$. Two possible interpretations of P are $s = (x_1, x_2)$ and $s' = (\top, \neg x_2)$.

An interpretation s for a prefix $P = \forall x_1, \ldots, x_n \exists y_1(D_1), \ldots, y_k(D_k)$ can be visualized as a rooted tree of height $n + k + 1$ with some additional edges to specify the dependencies. The nodes in the first n levels of this tree have two children and the edges to these children are labeled by $\top$ and $\bot$, respectively. These levels represent the universal variables $x_1, \ldots, x_n$ and constitute a complete binary tree. Each path in this complete binary tree corresponds to one assignment of the universal variables.

The nodes in the levels $n + 1, \ldots, n + k$ only have a single child with an edge that is either labelled by $\top$ or by $\bot$. These levels represent the existential variables $y_1, \ldots, y_k$ and each such level consists of 2^n nodes. Each path through these levels corresponds to an assignment of the existential variables.

In order to correctly represent the dependencies of the existential variables on (some of) the universal variables, we introduce additional edges, called *dependency edges*, that connect nodes within one level. We further impose the restriction that, if two nodes are connected by a dependency edge, then the outgoing edge to their respective child has to be labelled equally. The dependency edges are constructed as follows: Given two nodes v and w at level $n + l$ for $l \in \{1, \ldots, k\}$, consider the unique paths from the root node r to v and w, respectively, say

$$r \xrightarrow{a_1} \circ \xrightarrow{a_2} \circ \cdots \circ \xrightarrow{a_{n+l-1}} v \qquad \text{and} \qquad r \xrightarrow{b_1} \circ \xrightarrow{b_2} \circ \cdots \circ \xrightarrow{b_{n+l-1}} w,$$

where $a_i, b_i \in \{\top, \bot\}$ denote the edge labels, for $i = 1, \ldots, n+l-1$. Then, we draw a dependency edge between v and w if and only if $a_i = b_i$ for all i such that $x_i \in D_l$.

Remark 2 Informally, we draw a dependency edge between two nodes v and w at level $n+l$ if and only if the truth values of all universal variables on which y_l depends are equal on the paths to v and w, respectively.

Example 3 Consider the prefix $P = \forall x_1, x_2 \exists y_1\big(\{x_1\}\big), y_2\big(\{x_2\}\big)$. The two interpretations for P given in Example 2 can be visualized as follows:

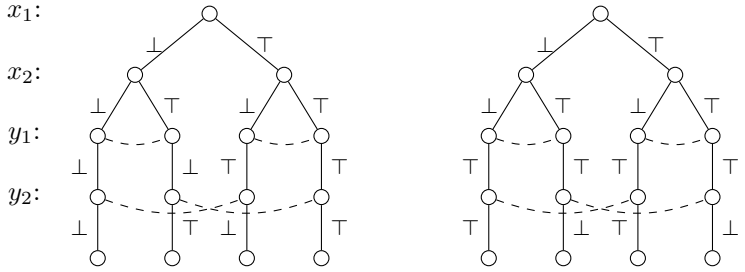

The dependency edges are depicted by dashed edges. Note that whenever two nodes are connected by a dependency edge, the outgoing edge to their respective child is labelled equally. The converse, however, need not hold. Nodes that are not connected by a dependency edge can still have outgoing edges with the same label (as witnessed by the tree on the right).

Each path from the root through the first $n+1$ layers of such a tree representing an interpretation $s \in \mathbb{S}(P)$ corresponds to one assignment $\sigma \in \mathbb{A}(X)$ of the universal variables $x_1, \ldots, x_n$. Extending this path to a leaf (ignoring dependency edges) yields an assignment $\sigma_s \in \mathbb{A}(X \cup Y)$ of all variables, called the *induced assignment* of σ and s. Formally, it is defined by

$$\begin{aligned} \sigma_s(x_i) &= \sigma(x_i) && \text{for } i \in \{1, \ldots, n\}, \\ \sigma_s(y_i) &= s_i(\sigma(x_{i_1}), \ldots, \sigma(x_{i_d})) && \text{for } i \in \{1, \ldots, k\}, \end{aligned}$$

where $D_i = \{x_{i_1}, \ldots, x_{i_d}\}$ is the dependency set of y_i and $i_1 < \cdots < i_d$.

The truth value of a DQBF under an interpretation s can then be obtained by considering all possible induced assignments of s, i.e., all complete paths from the root node to a leaf (ignoring dependency edges) in the tree representing s. Formally, we arrive at the following definition.

Definition 2 Let $\Phi = P.\phi$ be a DQBF and let $s \in \mathbb{S}(P)$ be an interpretation for the prefix P. The *truth value* of Φ under s is

$$[\Phi]_s = \bigwedge_{\sigma \in \mathbb{A}(X)} [\phi]_{\sigma_s}.$$

The DQBF Φ is *true* if there exists $s \in \mathbb{S}(P)$ with $[\Phi]_s = \top$ and it is *false* otherwise. If Φ is true, then any interpretation s with $[\Phi]_s = \top$ is called a *model* for Φ.

Example 4 Consider the prefix $P = \forall x_1, x_2 \exists y_1(\{x_1\}), y_2(\{x_2\})$ and let $s = (x_1, x_2)$ and $s' = (\top, \neg x_2)$ be the interpretations for P from Example 2. Furthermore, let $\Phi = P.\phi$ be the DQBF with

$$\phi = (\neg x_1 \to y_1) \wedge (x_2 \vee y_2)\,.$$

The truth value of Φ under s is $[\Phi]_s = \bot$ because the induced assignment σ_s that maps all variables to $\bot$ yields $[\phi]_{\sigma_s} = \bot$. Note that σ_s corresponds to the leftmost path in the left tree in Example 3. The truth value of Φ under s' is $[\Phi]_{s'} = \top$. There are four induced assignments $\sigma_{s'}$ of s' (one corresponding to each complete path in the right tree in Example 3) and one can check that $[\phi]_{\sigma_{s'}} = \top$ for all of them. Therefore, we can conclude that the DQBF Φ is true and that s' is a model for Φ.

The following lemma, which shall prove useful later, follows easily from the definitions above.

Lemma 1 *Let P be a prefix for X and Y, and let $\phi, \psi \in \mathrm{BF}(X \cup Y)$. Then, we have $[P.(\phi \wedge \psi)]_s = [P.\phi]_s \wedge [P.\psi]_s$ for all $s \in \mathbb{S}(P)$.*

Proof By Definition 2 and by the semantics of conjunction, we obtain

$$\begin{aligned}[P.(\phi \wedge \psi)]_s &= \bigwedge_{\sigma \in \mathbb{A}(X)} [\phi \wedge \psi]_{\sigma_s} = \bigwedge_{\sigma \in \mathbb{A}(X)} ([\phi]_{\sigma_s} \wedge [\psi]_{\sigma_s}) \\ &= \bigwedge_{\sigma \in \mathbb{A}(X)} [\phi]_{\sigma_s} \wedge \bigwedge_{\sigma \in \mathbb{A}(X)} [\psi]_{\sigma_s} = [P.\phi]_s \wedge [P.\psi]_s. \blacksquare\end{aligned}$$

3 Groups and Group Actions

Symmetries of an object can be described formally using *groups* and *group actions* [2]. We recall these concepts in this section. A group is a set G equipped with a binary associative operation $*\colon G \times G \to G$, such that G contains a neutral element and such that every element in G also has an inverse in G. A prototypical example of a group is the set of integers $\mathbb{Z}$ together with addition as the binary operation.

Another important example of a group, one particularly relevant for describing symmetries, is the *symmetric group* S_n. For any fixed $n \in \mathbb{N}$, the symmetric group S_n is the set of all bijective functions $\pi\colon \{1, \ldots, n\} \to \{1, \ldots, n\}$ together with function composition as the binary operation. The elements in S_n are called *permutations*. A permutation $\pi \in S_n$ can be conveniently denoted as a two dimensional array with two rows and n columns $\pi = \left(\begin{smallmatrix} 1 & 2 & \cdots & n \\ \pi(1) & \pi(2) & \cdots & \pi(n) \end{smallmatrix}\right)$. Permutations lend themselves nicely to describing symmetries of (geometric) objects.

Example 5 Consider the following square with vertices labelled by the symbols ♣, ♢, ♠, ♡.

If we assign to every symbol a number, say $\clubsuit \leftrightarrow 1, \diamondsuit \leftrightarrow 2, \spadesuit \leftrightarrow 3, \heartsuit \leftrightarrow 4$, then we can use permutations $\pi \in S_4$ to shuffle around the symbols, moving each symbol from vertex v to $\pi(v)$. For example, the permutation $\pi = \left(\begin{smallmatrix}1&2&3&4\\2&3&4&1\end{smallmatrix}\right)$ rotates the square by 90 degrees clockwise, leaving the relative order of the symbols unchanged (see Figure 1). The permutation $\sigma = \left(\begin{smallmatrix}1&2&3&4\\1&2&4&3\end{smallmatrix}\right)$, on the other hand, changes the relative order of the symbols. In fact, π describes a symmetry of the square and σ does not.

Fig. 1: Transformation of a square by two permutations $\pi = \left(\begin{smallmatrix}1&2&3&4\\2&3&4&1\end{smallmatrix}\right)$ and $\sigma = \left(\begin{smallmatrix}1&2&3&4\\1&2&4&3\end{smallmatrix}\right)$.

To formally describe that a group element (e.g., a permutation) can be used to transform an object (e.g., a square), we consider the notion of a *group action*. If G is a group with binary operation $*$ and S is any set, then a *group action* of G on S is a map $G \times S \to S$, $(g, s) \mapsto g(s)$, which is compatible with the group operation. This means that, for all $g, h \in G$ and $s \in S$, we have $(g * h)(s) = g(h(s))$ as well as $e(s) = s$, where e is the neutral element of G. If we have such a group action, we also say that G *acts* on the set S.

Example 6 The symmetric group S_n yields a group action on the set $S = \{1, \ldots, n\}$ by mapping every pair $(\pi, s) \in S_n \times S$ to $\pi(s)$. More generally, if S is any nonempty set and G is a group consisting of bijective functions $g\colon S \to S$ (with function decomposition as the binary operation of G), then a group action of G on S is given by mapping each pair $(g, s) \in G \times S$ to the element $g(s) \in S$.

As another example, consider the group action implicitly described in Example 5. It can be made explicit by letting the symmetric group S_4 act on the set $S = \{\clubsuit, \diamondsuit, \spadesuit, \heartsuit\}^4$ of 4-tuples by permuting indices, i.e., $\pi((x_1, \ldots, x_4)) = (x_{\pi(1)}, \ldots, x_{\pi(4)})$. For example, if we consider the original square as the tuple $s = (\clubsuit, \diamondsuit, \spadesuit, \heartsuit)$ and let $\pi = \left(\begin{smallmatrix}1&2&3&4\\2&3&4&1\end{smallmatrix}\right)$ be as in Example 5, then $\pi(s) = (\heartsuit, \clubsuit, \diamondsuit, \spadesuit)$. Analogously, for $\sigma = \left(\begin{smallmatrix}1&2&3&4\\1&2&4&3\end{smallmatrix}\right)$, we obtain $\sigma(s) = (\clubsuit, \diamondsuit, \heartsuit, \spadesuit)$.

As yet another example, let $V = \{v_1, \ldots, v_n\}$ be a set of propositional variables. A group action of the symmetric group S_n on the set of Boolean formulas $\mathrm{BF}(V)$ is given by permuting the variables, i.e., $\pi(\phi) = \phi'$, where the formula $\phi' \in \mathrm{BF}(V)$ is obtained from $\phi \in \mathrm{BF}(V)$ by replacing each variable v_i in ϕ by $v_{\pi(i)}$. For instance, for $\pi = \left(\begin{smallmatrix}1&2&3\\3&1&2\end{smallmatrix}\right)$ and $\phi = v_2 \oplus (v_1 \to \neg v_3)$, we get $\pi(\phi) = v_1 \oplus (v_3 \to \neg v_2)$.

Remark 3 In Example 6, we have seen that a group G consisting of bijective functions on a set S naturally induces a group action of G on S. Conversely, if we have a group action $G \times S \to S$, where G is now an arbitrary group, then we can associate to each group element $g \in G$ a unique bijective function $S \to S$, namely the one given by $s \mapsto g(s)$. The defining properties of a group action imply that this map is indeed bijective for each $g \in G$, with its inverse given by the map $s \mapsto g^{-1}(s)$. Therefore, in the following, when working with a group action, we may identify the group elements with their corresponding bijective functions on S.

Often, not all elements of a group are relevant in a particular context. For instance, in Example 5, we have seen that some elements of the symmetric group S_4 describe symmetries of a square, while others do not. Therefore, we recall the concept of a *subgroup*. A nonempty subset $H \subseteq G$ of a group G is a *subgroup* if it is closed under the group operation and under taking inverses. For any subset $E \subseteq G$, we can consider the smallest subgroup of G that contains E. This unique subgroup is denoted by $\langle E \rangle$ and the elements of E are called *generators* of $\langle E \rangle$.

Example 7 The set $42\mathbb{Z} = \{\dots, -84, -42, 0, 42, 84, \dots\}$ of integer multiples of 42 is a subgroup of $\mathbb{Z}$. It is generated by 42, i.e., $42\mathbb{Z} = \langle 42 \rangle$. A subgroup of the symmetric group S_4 is the eight element set

$$\{\mathrm{id}, \left(\begin{smallmatrix}1&2&3&4\\2&3&4&1\end{smallmatrix}\right), \left(\begin{smallmatrix}1&2&3&4\\3&4&1&2\end{smallmatrix}\right), \left(\begin{smallmatrix}1&2&3&4\\4&1&2&3\end{smallmatrix}\right), \left(\begin{smallmatrix}1&2&3&4\\2&1&4&3\end{smallmatrix}\right), \left(\begin{smallmatrix}1&2&3&4\\4&3&2&1\end{smallmatrix}\right), \left(\begin{smallmatrix}1&2&3&4\\1&4&3&2\end{smallmatrix}\right), \left(\begin{smallmatrix}1&2&3&4\\3&2&1&4\end{smallmatrix}\right)\}.$$

This subgroup describes all symmetries of a square. A set of generators is given by $\{\left(\begin{smallmatrix}1&2&3&4\\2&3&4&1\end{smallmatrix}\right), \left(\begin{smallmatrix}1&2&3&4\\2&1&4&3\end{smallmatrix}\right)\}$.

The action of a group G on a set S allows to define an equivalence relation on S via $s \sim t \iff \exists g \in G : t = g(s)$. It is straightforward to verify that the properties of a group action ensure that $\sim$ is indeed an equivalence relation. The equivalence classes are called the *orbits* of the group action. So, the orbit of $s \in S$ is the set $\{t \in S \mid s \sim t\} = \{g(s) \mid g \in G\}$.

Example 8 We reconsider the group action of S_4 on the set $S = \{\clubsuit, \diamondsuit, \spadesuit, \heartsuit\}^4$ discussed in Example 6. In the example, we have seen that $(\clubsuit, \diamondsuit, \spadesuit, \heartsuit) \sim (\heartsuit, \clubsuit, \diamondsuit, \spadesuit)$. In fact, the orbit of $(\clubsuit, \diamondsuit, \spadesuit, \heartsuit)$ consists of 24 elements (all the possible permutations of the four symbols $\clubsuit, \diamondsuit, \spadesuit, \heartsuit$). The orbit of $s = (\heartsuit, \heartsuit, \heartsuit, \heartsuit)$ only consists of a single element, namely s itself.

From a group G with binary operation $*$, we can construct another group, called the *opposite group* and denoted by G^{op}. This group has the same underlying set as G, i.e., $G^{\mathrm{op}} = G$, and its group operation $*^{\mathrm{op}} \colon G^{\mathrm{op}} \times G^{\mathrm{op}} \to G^{\mathrm{op}}$ is defined as $g *^{\mathrm{op}} g' \coloneqq g' * g$. Thus, the operation in G^{op} is the operation from the original group G but with the order of the arguments reversed.

Example 9 Consider the group $\mathbb{Z}$ of integers together with addition. In this case, the opposite group $\mathbb{Z}^{\mathrm{op}}$ is simply $\mathbb{Z}$ itself. This follows from the fact that

integer addition is commutative, i.e., $a + b = b + a$ for all $a, b \in \mathbb{Z}$. More generally, for any commutative group G, the opposite group is simply G itself.

For the noncommutative group S_4, the opposite group S_4^{op} actually has a different structure. For example, for $\pi = \left(\begin{smallmatrix}1&2&3&4\\2&3&4&1\end{smallmatrix}\right)$ and $\sigma = \left(\begin{smallmatrix}1&2&3&4\\2&1&4&3\end{smallmatrix}\right)$ we have

$$\pi \circ^{\mathrm{op}} \sigma = \sigma \circ \pi = \left(\begin{smallmatrix}1&2&3&4\\1&4&3&2\end{smallmatrix}\right) \neq \left(\begin{smallmatrix}1&2&3&4\\3&2&1&4\end{smallmatrix}\right) = \pi \circ \sigma.$$

4 Symmetries of DQBFs

We can use group actions to describe and study symmetries of DQBFs. Like in the case of QBFs [15], we distinguish between *syntactic symmetries* and *semantic symmetries*. The former concern transformations of the syntactic structure of a formula and arise from group actions of the form $G \times \mathrm{BF}(X \cup Y) \to \mathrm{BF}(X \cup Y)$, which transform formulas into other formulas. Semantic symmetries, on the other hand, concern the semantics of a formula and arise from group actions of the form $G \times \mathbb{S}(P) \to \mathbb{S}(P)$, which transform interpretations of a prefix P into other interpretations. In both cases, we consider groups G that preserve models of a given DQBF. We first discuss syntactic symmetries.

4.1 Syntactic Symmetries

In this section, we study symmetries of the syntactic structure of DQBFs. To this end, we consider group actions of the form $G \times \mathrm{BF}(X \cup Y) \to \mathrm{BF}(X \cup Y)$, for suitable groups G. We cannot allow arbitrary transformations of Boolean formulas. As a technicality, we have to require that a group action respects the semantics of propositional satisfiability and that it also respects the dependency structure of DQBFs. We will formalize these compatibility requirements in Definition 4 and 5 below. An analogous restriction is also required in the case of QBFs, cf. [15, Def. 3].

For the following, it is convenient to introduce the following auxiliary notion. For a set of variables $V \subseteq X \cup Y$, a function $g\colon \mathrm{BF}(V) \to \mathrm{BF}(V)$, and an assignment $\sigma \in \mathbb{A}(V)$, we denote by $g(\sigma)$ the assignment $g(\sigma) \in \mathbb{A}(V)$ defined by $g(\sigma)(v) = [g(v)]_\sigma$ for all $v \in V$.

Definition 3 A function $g\colon \mathrm{BF}(V) \to \mathrm{BF}(V)$ *preserves propositional satisfiability* if $[g(\phi)]_\sigma = [\phi]_{g(\sigma)}$ for every assignment $\sigma \in \mathbb{A}(V)$ and every formula $\phi \in \mathrm{BF}(V)$.

It follows from the definition that a function g that preserves propositional satisfiability is compatible with the logical connectives in the sense that $g(\neg\phi)$ and $\neg g(\phi)$ are logically equivalent, as are $g(\phi \circ \psi)$ and $g(\phi) \circ g(\psi)$ for all $\phi, \psi \in \mathrm{BF}(V)$ and every binary connective $\circ$. Therefore, such a function is essentially determined by its values on the variables.

Example 10 Let $V = \{x, y, z\}$. There is a function $g\colon \mathrm{BF}(V) \to \mathrm{BF}(V)$ that preserves propositional satisfiability and satisfies $g(x) = \neg x$, $g(y) = z$, $g(z) = y$. For such a function, we have, for example, $g((x \oplus y) \wedge \neg z) = (\neg x \oplus z) \wedge \neg y$. If a function $h\colon \mathrm{BF}(V) \to \mathrm{BF}(V)$ satisfies $h(x \wedge y) = x$, $h(x) = x$, $h(y) = y$, then it cannot preserve propositional satisfiability, because the formulas $h(x \wedge y) = x$ and $h(x) \wedge h(y) = x \wedge y$ are not logically equivalent.

To formally specify that syntactic symmetries have to respect the dependency structure of DQBFs, we fix a prefix $P = \forall x_1, \ldots, x_n \exists y_1(D_1), \ldots, y_k(D_k)$. In the following, we say that a Boolean formula $\phi \in \mathrm{BF}(Y)$ *depends on* a variable x_i if ϕ contains a variable y_j such that $x_i \in D_j$.

Definition 4 Let $P = \forall x_1, \ldots, x_n \exists y_1(D_1), \ldots, y_k(D_k)$ be a prefix for X and Y. A bijective function $g\colon \mathrm{BF}(X \cup Y) \to \mathrm{BF}(X \cup Y)$ is *admissible* (w.r.t. P) if the following conditions are satisfied for all $i \in \{1, \ldots, n\}$ and $j \in \{1, \ldots, k\}$:

1. g preserves propositional satisfiability;
2. $g(x_i) \in \mathrm{BF}(X)$ and $g(y_j) \in \mathrm{BF}(Y)$;
3. if $g(y_j)$ depends on x_i, then $g^{-1}(x_i) \in \mathrm{BF}(D_j)$;

The first condition is the same as the first condition in [15, Def. 3] for the QBF case. The other two conditions in Definition 4 generalize the second condition in [15, Def. 3]. They ensure that admissible functions transform existential (resp. universal) variables into existential (resp. universal) formulas and that admissible functions preserve the dependencies of the prefix P.

Example 11 Consider the prefix $P = \forall x_1, x_2 \exists y_1\big(\{x_1\}\big), y_2\big(\{x_2\}\big)$. There is an admissible function $g\colon \mathrm{BF}(X \cup Y) \to \mathrm{BF}(X \cup Y)$ satisfying $g(x_1) = x_2$, $g(x_2) = x_1$, $g(y_1) = y_2$, $g(y_2) = y_1$. A function $h\colon \mathrm{BF}(X \cup Y) \to \mathrm{BF}(X \cup Y)$ with $h(x_1) = y_1$ cannot be admissible because of the second condition. Neither can be a function h which leaves x_1 and x_2 fixed but exchanges y_1 and y_2. This follows from the third condition, as then $h(y_1) = y_2$ depends on x_2, but $h^{-1}(x_2) = x_2 \notin \mathrm{BF}(\{x_1\})$.

Like in the case of QBFs, cf. [15, Thm. 5], admissible functions preserve the truth value of DQBFs. More precisely, if g is an admissible function w.r.t. a prefix P, then a DQBF $P.\phi$ is true if and only if $P.g(\phi)$ is true. More precisely, there is an explicit correspondence between models of $P.\phi$ and those of $P.g(\phi)$. We defer the formalization of this statement and its proof to Proposition 1 in the next section, since they rely on constructions introduced later.

By Remark 3, we can consider the elements of a group G which acts on $\mathrm{BF}(X \cup Y)$ as bijective functions on $\mathrm{BF}(X \cup Y)$. Therefore, we extend Definition 4 above to the elements of such a group and say that $g \in G$ is *admissible* (w.r.t. P) if the corresponding bijective function is admissible (w.r.t. P).

Definition 5 Let P be a prefix for X and Y and let G be a group. A group action $G \times \mathrm{BF}(X \cup Y) \to \mathrm{BF}(X \cup Y)$ is *admissible* (w.r.t. P) if all elements $g \in G$ are admissible w.r.t. P.

If we have an admissible group action $G \times \mathrm{BF}(X \cup Y) \to \mathrm{BF}(X \cup Y)$, we may also say that G acts *admissibly* on $\mathrm{BF}(X \cup Y)$. Using admissible group actions, we can now introduce the concept of *syntactic symmetry groups*.

Definition 6 Let $P.\phi$ be a DQBF and let G be a group acting admissibly on $\mathrm{BF}(X \cup Y)$ w.r.t. P. We call G a *syntactic symmetry group* for $P.\phi$ if $[P.\phi]_s = [P.g(\phi)]_s$ for all $g \in G$ and all interpretations $s \in \mathbb{S}(P)$.

Like in the case of QBFs, cf. the discussion after [15, Def. 6], being a symmetry group is strictly speaking not a property of the group G itself but rather of the action of G on the Boolean formulas. We call the elements $g \in G$ of a syntactic symmetry group G *(syntactic) symmetries* of the DQBF $P.\phi$. Moreover, for any syntactic symmetry g of $P.\phi$, the definition above implies that $P.\phi$ and $P.g(\phi)$ have the same models.

We note that this definition of syntactic symmetry groups generalizes the corresponding notion for QBFs introduced in [15, Def. 6] in two ways: Not only does Definition 6 apply to DBQFs while [15, Def. 6] only applies to QBFs, but also, more interestingly, the latter requires logical equivalence of the Boolean formulas ϕ and $g(\phi)$, whereas the former only requires equivalence of the quantified formulas $P.\phi$ and $P.g(\phi)$. This more general definition now allows to deal with syntactic symmetries of (D)QBFs that were not covered by the previous definition.

Example 12 Consider the DQBF $\Phi = P.(x \wedge y)$ with prefix $P = \forall x \exists y(\{x\})$. A syntactic symmetry group for Φ is $G = \{\mathrm{id}, g\}$, where g is an admissible function with $g(x) = \neg x$ and $g(y) = y$. It is easy to see that

$$[\Phi]_s = [P.(x \wedge y)]_s = \bot = [P.(\neg x \wedge y)]_s = [P.g(x \wedge y)]_s$$

for all $s \in \mathbb{S}(P)$. This shows, on the one hand, that g is indeed a symmetry of Φ, and, on the other hand, that Φ is false. We note that, according to [15, Def. 6], the function g would not be a symmetry of the QBF $\forall x \exists y.(x \wedge y)$ because the Boolean formulas $x \wedge y$ and $g(x \wedge y) = \neg x \wedge y$ are not equivalent.

4.2 Semantic Symmetries

In the following, we study symmetries of the semantic structure of DQBFs. To this end, we fix a prefix P and consider transformations of interpretations for P, i.e., we look at group actions of the form $G \times \mathbb{S}(P) \to \mathbb{S}(P)$, for suitable groups G. Given a DQBF Φ, we are interested in group actions that transform models of Φ into other models. In contrast to the previous section, we now have to impose no technical restrictions on the considered group actions.

Definition 7 Let $P.\phi$ be a DQBF and let $G \times \mathbb{S}(P) \to \mathbb{S}(P)$ be a group action on $\mathbb{S}(P)$. We call the group G a *semantic symmetry group* for $P.\phi$ if $[P.\phi]_s = [P.\phi]_{g(s)}$ for all $g \in G$ and all interpretations $s \in \mathbb{S}(P)$.

Analogous to syntactic symmetry groups, we call the elements of a semantic symmetry group *(semantic) symmetries*. It was observed for QBFs that many semantic symmetries arise from syntactic symmetries, cf. [15, Sec. 5]. This observation also generalizes to DQBFs. To see this, we first note that syntactic transformations of Boolean formulas naturally lead to semantic transformations of assignments. In particular, for any set of variables $V \subseteq X \cup Y$, a group action on $\mathrm{BF}(V)$ naturally induces a group action of the opposite group on the set $\mathbb{A}(V)$ of assignments for V. Recall that, for a function $g\colon \mathrm{BF}(V) \to \mathrm{BF}(V)$ and an assignment $\sigma \in \mathbb{A}(V)$, the assignment $g(\sigma)$ is given by $g(\sigma)(v) = [g(v)]_\sigma$ for all $v \in V$.

Lemma 2 *Let $V \subseteq X \cup Y$ be a set of variables. If a group G acts on $\mathrm{BF}(V)$, then the opposite group G^{op} acts on $\mathbb{A}(V)$ via the map $(g, \sigma) \mapsto g(\sigma)$.*

Proof We have to verify that the map $G^{\mathrm{op}} \times \mathbb{A}(V) \to \mathbb{A}(V), (g, \sigma) \mapsto g(\sigma)$ is compatible with the group operation $*^{\mathrm{op}}$ in G^{op}, i.e., that for all $g, h \in G^{\mathrm{op}}$ and $\sigma \in \mathbb{A}(V)$, we have $(g *^{\mathrm{op}} h)(\sigma) = g(h(\sigma))$ as well as $e(\sigma) = \sigma$, where e is the neutral element of G^{op}.

For the first property, observe that, for all $v \in V$, we have

$$\begin{aligned}(g *^{\mathrm{op}} h)(\sigma)(v) &= [(g *^{\mathrm{op}} h)(v)]_\sigma = [(h * g)(v)]_\sigma \\ &= [h(g(v))]_\sigma = [g(v)]_{h(\sigma)} = [v]_{g(h(\sigma))},\end{aligned}$$

where the third equality follows from the fact that G acts on $\mathrm{BF}(V)$. Similarly, the second property follows, for all $v \in V$, from

$$e(\sigma)(v) = [e(v)]_\sigma = [v]_\sigma = \sigma(v),$$

where the second equality uses the fact G acts on $\mathrm{BF}(V)$. ∎

The following example shows how a syntactic symmetry of a DQBF naturally gives rise to a semantic symmetry.

Example 13 Consider the DQBF

$$P.\phi = \forall x_1, x_2 \exists y_1\big(\{x_1\}\big), y_2\big(\{x_2\}\big).\, (x_1 \vee y_1) \wedge (x_2 \vee y_2)\,.$$

A syntactic symmetry of $P.\phi$ is given by an admissible function g, which exchanges x_1 with x_2 and y_1 with y_2. We describe how to translate this syntactic symmetry into a semantic one.

Each interpretation $s = (s_1, s_2) \in \mathbb{S}(P)$ of P consists of two Skolem functions, which, by Remark 1, can be represented by Boolean formulas $s_i \in \mathrm{BF}(\{x_i\})$ $(i = 1, 2)$. Now, exchanging x_1 with x_2 in the original formula $P.\phi$ corresponds semantically to exchanging the roles of x_1 and x_2 in the Skolem functions s_1 and s_2, respectively, i.e., it corresponds to replacing s_i by $g(s_i)$ for $i = 1, 2$. Further, exchanging y_1 with y_2 in $P.\phi$ corresponds semantically to exchanging the order of the Skolem functions in the interpretation s, i.e, it

corresponds to replacing $s = (s_1, s_2)$ by (s_2, s_1). Combining these two steps, we consider the function $f\colon \mathbb{S}(P) \to \mathbb{S}(P)$ defined by

$$(s_1, s_2) \mapsto (g(s_2), g(s_1)).$$

This function f satisfies

$$[P.g(\phi)]_s = [P.\phi]_{f(s)},$$

for all $s \in \mathbb{S}(P)$. Moreover, since g is a syntactic symmetry of $P.\phi$, this implies

$$[P.\phi]_s = [P.\phi]_{f(s)},$$

showing that f is a semantic symmetry of $P.\phi$.

For example, for $s = (\neg x_1, x_2) \in \mathbb{S}(P)$, we have $f(s) = (x_1, \neg x_2)$ and $[P.\phi]_s = \bot = [P.\phi]_{f(s)}$. Analogously, for $s' = (\neg x_1, \top)$, we get $f(s') = (\top, \neg x_2)$ and $[P.\phi]_{s'} = \top = [P.\phi]_{f(s')}$.

The construction from the previous example is formalized in the definition below. It allows to construct, starting from a syntactic symmetry, a semantic one, mimicking the same behaviour.

For what follows, we have to generalize one definition slightly. So far, for $V \subseteq X \cup Y$, the assignment $g(\sigma) \in \mathbb{A}(V)$ is defined for a function $g\colon \mathrm{BF}(V) \to \mathrm{BF}(V)$ and an assignment $\sigma \in \mathbb{A}(V)$ on the same set of variables. In the following, however, we need to consider cases where $g\colon \mathrm{BF}(X \cup Y) \to \mathrm{BF}(X \cup Y)$, but $\sigma \in \mathbb{A}(X)$. For an arbitrary function g, considering $g(\sigma)$ in this setting would not make sense, because, for $x \in X$, the formula $g(x)$ could contain variables from Y and thus $g(\sigma)(x) = [g(x)]_\sigma$ would not be well-defined. However, if g is an admissible function, then, by definition, $g(x) \in \mathrm{BF}(X)$ for all $x \in X$, and thus, in this case, we can define $g(\sigma) \in \mathbb{A}(X)$ as the assignment defined by $g(\sigma)(x) = [g(x)]_\sigma$ for all $x \in X$.

Definition 8 Let G be a group acting admissibly on $\mathrm{BF}(X \cup Y)$ w.r.t. a prefix P. For $g \in G$ and $s \in \mathbb{S}(P)$, we define $g(s) \in \mathbb{S}(P)$ as the interpretation $t \in \mathbb{S}(P)$ with the property that $\sigma_t = g(g^{-1}(\sigma)_s)$ for all assignments $\sigma \in \mathbb{A}(X)$.

In order to justify this definition, we have to show that the expression $t = g(s)$ is well-defined. To see this, let $j \in \{1, \ldots, k\}$ be arbitrary, and let D_j be the dependency set of y_j. We have to show that, for any two assignments $\sigma, \sigma' \in \mathbb{A}(X)$ with $\sigma(x) = \sigma'(x)$ for all $x \in D_j$, we have

$$[y_j]_{g(g^{-1}(\sigma)_s)} = [y_j]_{g(g^{-1}(\sigma')_s)}.$$

Suppose otherwise. Then, the admissibility of g implies

$$[g(y_j)]_{g^{-1}(\sigma)_s} \neq [g(y_j)]_{g^{-1}(\sigma')_s},$$

which means that $g(y_j)$ contains a variable $y \in Y$ such that

$$[y]_{g^{-1}(\sigma)_s} \neq [y]_{g^{-1}(\sigma')_s}.$$

This implies that there must be a variable x in the dependency set of y with

$$[x]_{g^{-1}(\sigma)_s} \neq [x]_{g^{-1}(\sigma')_s},$$

i.e., $[g^{-1}(x)]_\sigma \neq [g^{-1}(x)]_{\sigma'}$. Then, $g^{-1}(x)$ contains some variable $x_i \in X$ such that $\sigma(x_i) \neq \sigma'(x_i)$. However, by the admissibility of the group action, x_i must belong to D_j, which gives a contradiction to the choice of σ, σ'. Thus, t is well-defined.

We collect some properties of the interpretation $g(s)$.

Lemma 3 *Let G_{syn} be a group acting admissibly on $\text{BF}(X \cup Y)$ w.r.t. a prefix P. For every $g \in G_{\text{syn}}$, the bijective function $\mathbb{S}(P) \to \mathbb{S}(P), s \mapsto g(s)$ satisfies $g(\sigma)_{g(s)} = g(\sigma_s)$ for all $s \in \mathbb{S}(P)$ and $\sigma \in \mathbb{A}(X)$.*

Proof The map $s \mapsto g(s)$ is clearly a bijective function with inverse given by $s \mapsto g^{-1}(s)$. Furthermore, for any $s \in \mathbb{S}(P)$ and $\sigma \in \mathbb{A}(X)$, we have

$$g(\sigma)_{g(s)} = g(g^{-1}(g(\sigma))_s) = g(\sigma_s). \blacksquare$$

The following result formalizes the statement that we can transform syntactic symmetries into semantic ones.

Proposition 1 *Let $P.\phi$ be a DQBF and let G_{syn} be a group acting admissibly on $\text{BF}(X \cup Y)$ w.r.t. P. Then, for every $g \in G_{\text{syn}}$ and all $s \in \mathbb{S}(P)$, we have*

$$[P.g(\phi)]_s = [P.\phi]_{g(s)}.$$

In particular, $P.\phi$ is true if and only if $P.g(\phi)$ is true, and s is a model of $P.g(\phi)$ if and only if $g(s)$ is a model of $P.\phi$.

Proof We show that $[P.g(\phi)]_s = \top$ if and only if $[P.\phi]_{g(s)} = \top$. In fact, since G is a group, it suffices to show only one direction, say "$\Leftarrow$". To this end, assume that $[P.\phi]_{g(s)} = \top$, i.e., $[\phi]_{\sigma_{g(s)}} = \top$ for all $\sigma \in \mathbb{A}(X)$. Note that this implies that also $[\phi]_{g(\sigma)_{g(s)}} = \top$ for all $\sigma \in \mathbb{A}(X)$. Then, Lemma 3 yields

$$[g(\phi)]_{\sigma_s} = [\phi]_{g(\sigma_s)} = [\phi]_{g(\sigma)_{g(s)}} = \top,$$

for all $\sigma \in \mathbb{A}(X)$. $\blacksquare$

Starting from a syntactic symmetry group G_{syn}, we can collect all bijective functions that satisfy a similar condition like the ones constructed in Lemma 3. This yields a semantic symmetry group, which we call the *associated group* of G_{syn}. Note that the definition below is slightly more general than the construction in Lemma 3, in the sense that, in Lemma 3, the element $g \in G_{\text{syn}}$ is fixed, while, in the definition below, g may depend on s and σ.

Definition 9 Let G_{syn} be a group acting admissibly on $\text{BF}(X \cup Y)$ w.r.t. a prefix P. Furthermore, let G_{sem} be the set of all bijective functions $f\colon \mathbb{S}(P) \to \mathbb{S}(P)$ such that for every $s \in \mathbb{S}(P)$ and every assignment $\sigma \in \mathbb{A}(X)$ there exists $g \in G_{\text{syn}}$ with $g(\sigma)_{f(s)} = g(\sigma_s)$. Then G_{sem} is called the *associated group* of G_{syn}.

We record the following result for later use. It follows immediately from Lemma 3 and from the definition of the associated group.

Lemma 4 *Let G_{syn} be a group acting admissibly on $\mathrm{BF}(X \cup Y)$ w.r.t. a prefix P. For any $g \in G_{\text{syn}}$, the function $\mathbb{S}(P) \to \mathbb{S}(P), s \mapsto g(s)$ lies in the associated group of G_{syn}.*

If G_{syn} is a syntactic group, then the associated group G_{sem} is a semantic symmetry group.

Lemma 5 *If G_{syn} is a syntactic symmetry group for a DQBF Φ, then the associated group G_{sem} of G_{syn} is a semantic symmetry group for Φ.*

Proof First, we show that G_{sem} is indeed a group. To this end, note that it contains the identity function. To see that G_{sem} is closed under the binary operation of function decomposition, let $f, f' \in G_{\text{sem}}$, and let $s \in \mathbb{S}(P)$ and $\sigma \in \mathbb{A}(X)$ be arbitrary. We have to show that there exists a $g \in G_{\text{syn}}$ such that $g(\sigma)_{(f' \circ f)(s)} = g(\sigma_s)$. By assumption on f and f', we know that there exist $h, h' \in G_{\text{syn}}$ such that $h(\sigma)_{f(s)} = h(\sigma_s)$ and $h'(h(\sigma))_{f'(f(s))} = h'(h(\sigma)_{f(s)})$. Now, with $g = h' *^{\text{op}} h = h * h'$, where $*$ is the group operation in G_{syn}, we obtain,

$$\begin{aligned} g(\sigma)_{(f' \circ f)(s)} &= (h' *^{\text{op}} h)(\sigma)_{(f' \circ f)(s)} = h'(h(\sigma))_{f'(f(s))} \\ &= h'(h(\sigma)_{f(s)}) = h'(h(\sigma_s)) = (h' *^{\text{op}} h)(\sigma_s) = g(\sigma_s), \end{aligned}$$

where the second and fifth equality follow from Lemma 2. Finally, G_{sem} is also closed under taking inverses. To see this, note that $g(\sigma)_{f(s)} = g(\sigma_s)$ implies $g^{-1}(g(\sigma)_{f(s)}) = \sigma_s$. But every s can be written as $s = f^{-1}(s')$ for some $s' \in \mathbb{S}(P)$ and every σ can be written as $\sigma = g^{-1}(\sigma')$ for some $\sigma' \in \mathbb{A}(X)$. This yields $g^{-1}(\sigma'_{s'}) = g^{-1}(\sigma')_{f^{-1}(s')}$ for all $s' \in \mathbb{S}(P)$ and $\sigma' \in \mathbb{A}(X)$, and hence, $f^{-1} \in G_{\text{sem}}$.

Next, we show that G_{sem} is a semantic symmetry group. To this end, let $f \in G_{\text{sem}}$ and $s \in \mathbb{S}(P)$. We have to show that $[P.\phi]_s = [P.\phi]_{f(s)}$. It suffices to show that $[P.\phi]_s = \bot \iff [P.\phi]_{f(s)} = \bot$. In fact, since G_{sem} is a group, it even suffices to only show one direction, say "$\Rightarrow$". Recall that $[P.\phi]_{f(s)} = \bot$ if and only if there exists an assignment $\tau \in \mathbb{A}(X)$ such that $[\phi]_{\tau_{f(s)}} = \bot$. By assumption, there exists an assignment $\sigma \in \mathbb{A}(X)$ such that $[\phi]_{\sigma_s} = \bot$. Fix such a σ and note that, since G_{sem} is the associated group of G_{syn}, there exists a $g \in G_{\text{syn}}$ such that $g(\sigma)_{f(s)} = g(\sigma_s)$. Then,

$$[\phi]_{g(\sigma)_{f(s)}} \overset{\text{choice of } g}{=} [\phi]_{g(\sigma_s)} \overset{g \text{ admissible}}{=} [g(\phi)]_{\sigma_s} \overset{g \text{ symmetry}}{=} [\phi]_{\sigma_s} \overset{\text{choice of } \sigma}{=} \bot.$$

This shows that, for $\tau = g(\sigma) \in \mathbb{A}(X)$, we have $[\phi]_{\tau_{f(s)}} = \bot$, implying that $[P.\phi]_{f(s)} = \bot$ as claimed. ∎

The associated semantic group is very versatile and typically contains a lot more symmetries than the corresponding syntactic symmetry group. In particular, if two interpretations are related via one semantic symmetry, then the associated group G_{sem} also contains elements that allow to exchange and combine these interpretations. For example, for any two interpretations $s, s' \in \mathbb{S}(P)$ that are related via some $f \in G_{\text{sem}}$ via $s' = f(s)$, there exists another symmetry $h \in G_{\text{sem}}$ with $h(s) = s'$, $h(s') = s$, and $h(t) = t$ for all other $t \in \mathbb{S}(P) \setminus \{s, s'\}$. More generally, the associated group contains elements that allow to exchange subtrees of interpretations. To formalize this statement, we introduce the following notion of a *section* of an interpretation. Recall from Remark 1 that every Skolem function s_i with dependency set D_i can be represented by a Boolean formula in $\text{BF}(D_i)$.

Definition 10 Let $P = \forall x_1, \ldots, x_n \exists y_1(D_1), \ldots, y_k(D_k)$ be a prefix for X and Y and let $X_0 \subseteq X$. Furthermore, let $s = (s_1, \ldots, s_k) \in \mathbb{S}(P)$ be an interpretation and let $\sigma \in \mathbb{A}(X_0)$ be an assignment. For $i = 1, \ldots, k$, let $\phi_i \in \text{BF}(D_i)$ denote a Boolean formula representing the Skolem function s_i and let $\phi_i|_\sigma \in \text{BF}(D_i \setminus X_0)$ denote the formula obtained from ϕ_i by assigning all the variables in X_0 as specified by σ.

The *section* of the Skolem function s_i with respect to σ is the Skolem function $s_i|_\sigma \colon \{\top, \bot\}^{|D_i \setminus X_0|} \to \{\top, \bot\}$ represented by the formula $\phi_i|_\sigma$. The *section* of the interpretation s with respect to σ is the interpretation $s|_\sigma = (s_1|_\sigma, \ldots, s_k|_\sigma) \in \mathbb{S}(P_0)$, where P_0 denotes the prefix obtained from P by discarding all variables in X_0.

In other words, the Skolem function $s_i|_\sigma$ is obtained from s_i by setting those inputs of s_i that are in X_0 to the values specified by the assignment σ. The following example illustrates how this is done.

Example 14 Consider a Skolem function $s_i \colon \{\top, \bot\}^3 \to \{\top, \bot\}$ for an existential variable y_i with dependency set $D_i = \{x_1, x_3, x_5\}$. Say that s_i can be represented by the Boolean formula

$$\phi_i = (\neg x_1 \vee x_3) \wedge (\neg x_1 \vee x_5) \wedge (x_1 \vee \neg x_3 \vee \neg x_5).$$

Then, for $X_0 = \{x_2, x_4, x_5\}$, the section of s_i w.r.t. an assignment $\sigma \in \mathbb{A}(X_0)$ such that $\sigma(x_5) = \bot$ is the Skolem function $s_i|_\sigma \colon \{\top, \bot\}^2 \to \{\top, \bot\}$ represented by the Boolean formula

$$\phi_i|_\sigma = (\neg x_1 \vee x_3) \wedge (\neg x_1 \vee \bot) \wedge (x_1 \vee \neg x_3 \vee \top) \equiv \neg x_1.$$

In other words, $s_i|_\sigma$ is the binary function which returns $\top$ if and only if its first input is $\bot$. Equivalently, $s_i|_\sigma$ can also be described as $s_i|_\sigma(\xi_1, \xi_2) = s_i(\xi_1, \xi_2, \bot)$. Note that, although one input would suffice to describe $s_i|_\sigma$, we still consider $s_i|_\sigma$ as a binary function.

For $\tau \in \mathbb{A}(X_0)$ with $\tau(x_5) = \top$, we obtain

$$\phi_i|_\tau = (\neg x_1 \vee x_3) \wedge (\neg x_1 \vee \top) \wedge (x_1 \vee \neg x_3 \vee \bot) \equiv x_1 \leftrightarrow x_3.$$

Therefore, $s_i|_\tau$ is the binary function which returns $\top$ if and only if both its inputs are equal.

The section $s|_\sigma$ of an interpretation s with respect to σ can also be visualized nicely as a subtree of the tree representing the interpretation s. We assume that X_0 contains the first l universal variables, i.e., $X_0 = \{x_1, \ldots, x_l\}$ and $X \setminus X_0 = \{x_{l+1}, \ldots, x_k\}$ for some $l \in \{1, \ldots, k\}$. We note that this can always be achieved by renaming the variables.

The following tree represents an interpretation s. The assignment σ is visualized as a path starting at the root. The shaded area is the subtree representing the section $s|_\sigma$.

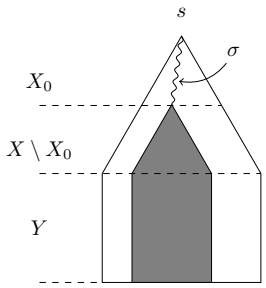

Using the notion of sections, we can now better describe the structure of the associated semantic group G_{sem}. In particular, the associated group contains elements that allow to exchange sections of an interpretation $s \in \mathbb{S}(P)$ with that of $f(s)$ for any $f \in G_{\text{sem}}$. In the simplest case, where we assume that X_0 contains the first l universal variables and where we ignore some technicalities, this fact can be visualized as follows:

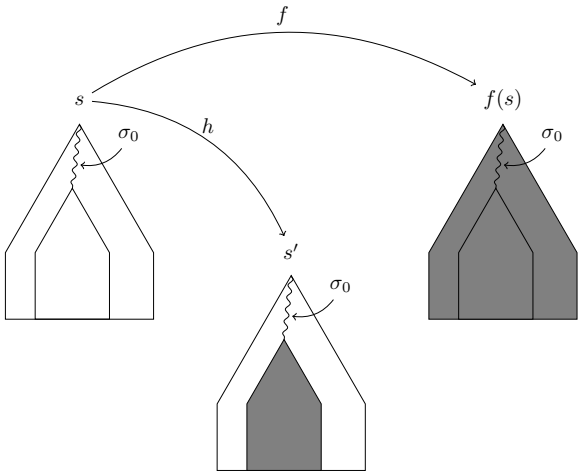

In the first row, we see an interpretation s and its image $f(s)$ under some semantic symmetry $f \in G_{\text{sem}}$. In both interpretations, the section with respect to some assignment $\sigma_0 \in \mathbb{A}(X_0)$ is highlighted. The associated semantic group now contains an element $h \in G_{\text{sem}}$ that allows to transform s into the interpretation s' depicted in the second row, which coincides with s except for the fact that the section $s|_{\sigma_0}$ has been replaced by $f(s)|_{\sigma_0}$.

The following lemma formalizes this fact and generalizes it to arbitrary subsets $X_0 \subseteq X$.

Lemma 6 *Let P be a prefix for X and Y, let G_{syn} be a group acting admissibly on $\text{BF}(X \cup Y)$, and let G_{sem} be the associated semantic group. Let $X_0 \subseteq X$ and $\sigma_0 \in \mathbb{A}(X_0)$ be such that $\rho|_{X_0} = \sigma_0$ implies $g(\rho)|_{X_0} = \sigma_0$ for all $g \in G_{\text{syn}}$ and all $\rho \in \mathbb{A}(X)$. Then, for any $s \in \mathbb{S}(P)$ and $f \in G_{\text{sem}}$, there exists an interpretation $s' \in \mathbb{S}(P)$ such that*

$$s'|_\tau = \begin{cases} f(s)|_{\sigma_0} & \text{if } \tau = \sigma_0, \\ s|_\tau & \text{if } \tau \neq \sigma_0, \end{cases}$$

for all $\tau \in \mathbb{A}(X_0)$. Furthermore, there exists $h \in G_{\text{sem}}$ with $h(s) = s'$.

Proof The existence of s' is easy to see. The ith component of s' is the function which evaluates to the value of the ith component of $f(s)$ for all assignments in $\mathbb{A}(X)$ whose restriction to X_0 is σ_0 and which evaluates to the value of the ith component of s for all remaining assignments.

Define $h\colon \mathbb{S}(P) \to \mathbb{S}(P)$ by $h(s) = s'$, $h(s') = s$, and $h(t) = t$ for all $t \in \mathbb{S}(P) \setminus \{s, s'\}$. Obviously, h is a bijective function. To show that h belongs to G_{sem}, we must show that for every $t \in \mathbb{S}(P)$ and every assignment $\rho \in \mathbb{A}(X)$ there exists $g \in G_{\text{syn}}$ such that $g(\rho)_{h(t)} = g(\rho_t)$. For $t \in \mathbb{S}(P) \setminus \{s, s'\}$ we have $h(t) = t$, so g can be chosen as the neutral element of G_{syn}. For the other cases $t \in \{s, s'\}$, let $\rho \in \mathbb{A}(X)$ be an assignment. If $\rho|_{X_0} \neq \sigma_0$, then $\rho_{s'} = \rho_s$ by definition of s' and we can again choose g as the neutral element of G_{syn}. If $\rho|_{X_0} = \sigma_0$, then $\rho_{s'} = \rho_{f(s)}$, again by definition of s'. By definition of the associated group, there exists $g \in G_{\text{syn}}$ with

$$g(\rho)_{f(s)} = g(\rho_s).$$

Now, by assumption $g(\rho)|_{X_0} = \sigma_0$, and thus, $g(\rho)_{s'} = g(\rho)_{f(s)}$. This yields

$$g(\rho)_{h(s)} = g(\rho)_{s'} = g(\rho)_{f(s)} = g(\rho_s).$$

Analogously, by definition of G_{sem}, there exists $g' \in G_{\text{syn}}$ such that

$$g'(\rho)_{f^{-1}(s')} = g'(\rho_{s'}).$$

Furthermore, by assumption $g'(\rho)|_{X_0} = \sigma_0$, which implies $g'(\rho)_{s'} = g'(\rho)_{f(s)}$. Applying f^{-1} to both interpretations in this identity yields $g'(\rho)_{f^{-1}(s')} = g'(\rho)_s$. Thus, ultimately we obtain

$$g'(\rho)_{h(s')} = g'(\rho)_s = g'(\rho)_{f^{-1}(s')} = g'(\rho_{s'}).$$

This covers all cases. ∎

5 Conjunctive Symmetry Breakers

We note that the following discussion is completely analogous to the case of QBFs, cf. the beginning of [15, Sec. 6].

The action of a syntactic symmetry group for a DQBF $P.\phi$ splits the set of Boolean formulas $\mathrm{BF}(X \cup Y)$ into orbits. By definition, for all formulas ψ in the orbit of ϕ, the original formula $P.\phi$ and $P.\psi$ share the same models:

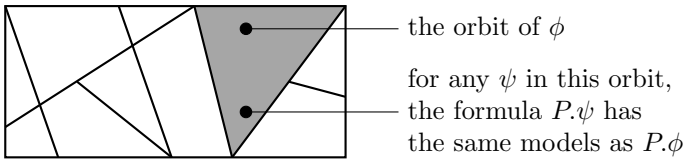

Therefore, for finding a model for $P.\phi$, we can replace ϕ by any formula in its orbit.

A semantic symmetry group for $P.\phi$, on the other hand, splits the set of interpretations $\mathbb{S}(P)$ into orbits so that each orbit either contains no models at all for $P.\phi$ or only models for $P.\phi$:

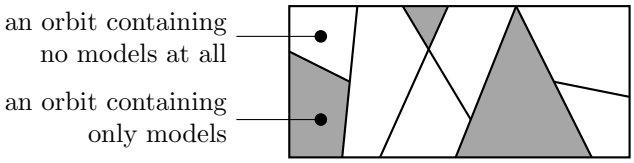

Therefore, for finding a model for $P.\phi$, it suffices to check only one interpretation per orbit.

To goal of *symmetry breaking* is to exploit this fact and to construct a Boolean formula $\psi \in \mathrm{BF}(X \cup Y)$, called a *(conjunctive) symmetry breaker*, in a such way that $P.\psi$ has at least one model in every orbit. Then, instead of solving $P.\phi$, we can solve $P.(\phi \wedge \psi)$. By Lemma 1, every model for the latter is also a model for the former. Moreover, if $P.\phi$ has a model, then there exists a whole orbit consisting only of models. Thus, by construction, this orbit also contains a model of $P.(\phi \wedge \psi)$. Ideally, we want to construct ψ in such a way that $P.\psi$ contains precisely one model per orbit. In this way, we have to inspect only one element per orbit when solving $P.(\phi \wedge \psi)$, the one model for $P.\psi$.

When constructing symmetry breakers, we can also consider the effect of a syntactic symmetry group for $P.\phi$. Such a symmetry group allows us to exchange a symmetry breaker ψ by $g(\psi)$ for any syntactic symmetry g. Thus, ultimately we arrive at the following definition.

Definition 11 Let P be a prefix for X and Y, let G_{syn} be a group acting admissibly on $\mathrm{BF}(X \cup Y)$, and let G_{sem} be a group acting on $\mathbb{S}(P)$. A formula $\psi \in \mathrm{BF}(X \cup Y)$ is called a *conjunctive symmetry breaker* for G_{syn} and G_{sem} if for every $s \in \mathbb{S}(P)$ there exist $g_{\mathrm{syn}} \in G_{\mathrm{syn}}$ and $g_{\mathrm{sem}} \in G_{\mathrm{sem}}$ such that $[P.g_{\mathrm{syn}}(\psi)]_{g_{\mathrm{sem}}(s)} = \top$.

Example 15 Consider the DQBF

$$P.\phi = \forall x_1, x_2 \exists y_1\big(\{x_1\}\big), y_2\big(\{x_2\}\big).\, (x_1 \vee y_1) \wedge (x_2 \vee y_2)\,.$$

As noted in Example 13, a syntactic symmetry for $P.\phi$ is given by an admissible function g, which exchanges x_1 with x_2 and y_1 with y_2. So, we can take $G_{\text{syn}} = \{\text{id}, g\}$ as a syntactic symmetry group. Concerning semantic symmetries, we note that every interpretation in $\mathbb{S}(P)$ is of the form:

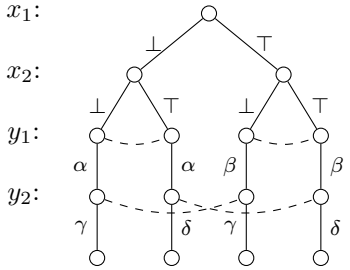

Semantic symmetries for $P.\phi$ are given by the function $f_\beta \colon \mathbb{S}(P) \to \mathbb{S}(P)$, which replaces β by $\neg\beta$ and leaves everything else unchanged, and by $f_\delta \colon \mathbb{S}(P) \to \mathbb{S}(P)$, which replaces δ by $\neg\delta$ and leaves everything else unchanged. Thus, the group $G_{\text{sem}} = \langle f_\beta, f_\delta \rangle$ is a semantic symmetry group for $P.\phi$.

We claim that $\psi = y_1 \to y_2$ is a conjunctive symmetry breaker for G_{syn} and G_{sem}. To prove this, let $s \in \mathbb{S}(P)$ be an arbitrary interpretation. Note that $[P.\psi]_s = \top$ if and only if the propositional formula

$$\alpha \to \gamma \;\wedge\; \alpha \to \delta \;\wedge\; \beta \to \gamma \;\wedge\; \beta \to \delta, \tag{1}$$

with $\alpha, \beta, \gamma, \delta$ as specified by s, holds.

Using the syntactic symmetry g and replacing ψ by $g(\psi) = y_2 \to y_1$ if necessary, we can always assume that $\alpha \to \gamma$ holds for the interpretation s. Then, using the semantic symmetries f_β and f_δ, we can replace s by $g_{\text{sem}}(s)$ so that $\beta = \bot$ and $\delta = \top$. Under this interpretation $g_{\text{sem}}(s)$, the formula (1) evaluates to true. This shows that, for all $s \in \mathbb{S}(P)$, there exist $g_{\text{syn}} \in G_{\text{syn}}$ and $g_{\text{sem}} \in G_{\text{sem}}$ such that $[P.g_{\text{syn}}(\psi)]_{g_{\text{sem}}(s)} = \top$. Thus, ψ is a conjunctive symmetry breaker for G_{syn} and G_{sem} as claimed.

The following theorem is the main property of conjunctive symmetry breakers. It generalizes the analogous result [15, Thm. 16] for QBFs.

Theorem 1 *Let $P.\phi$ be a DQBF. Furthermore, let G_{syn} and G_{sem} be a syntactic and a semantic symmetry group, respectively, for $P.\phi$. If $\psi \in \text{BF}(X \cup Y)$ is a conjunctive symmetry breaker for G_{syn} and G_{sem}, then $P.\phi$ is true if and only if $P.(\phi \wedge \psi)$ is true.*

Proof The implication "$\Leftarrow$" follows immediately from Lemma 1. For the other implication "$\Rightarrow$", suppose that $P.\phi$ is true. Then there exists $s \in \mathbb{S}(P)$ such that $[P.\phi]_s = \top$. We have to show that there exists $t \in \mathbb{S}(P)$ such that also

$[P.(\phi \wedge \psi)]_t = \top$. Since ψ is a conjunctive symmetry breaker for G_{syn} and G_{sem}, there exist $g_{\mathrm{syn}} \in G_{\mathrm{syn}}$ and $g_{\mathrm{sem}} \in G_{\mathrm{sem}}$ such that $[P.g_{\mathrm{syn}}(\psi)]_{g_{\mathrm{sem}}(s)} = \top$. Using Lemma 1 and the fact that G_{syn} and G_{sem} are symmetry groups for Φ, we get

$$\begin{aligned}[P.g_{\mathrm{syn}}(\phi \wedge \psi)]_{g_{\mathrm{sem}}(s)} &= [P.(g_{\mathrm{syn}}(\phi) \wedge g_{\mathrm{syn}}(\psi))]_{g_{\mathrm{sem}}(s)} \\ &= [P.g_{\mathrm{syn}}(\phi)]_{g_{\mathrm{sem}}(s)} \wedge [P.g_{\mathrm{syn}}(\psi)]_{g_{\mathrm{sem}}(s)} = \top \wedge \top = \top,\end{aligned}$$

showing that $P.g_{\mathrm{syn}}(\phi \wedge \psi)$ is true. By Proposition 1, it follows that $P.(\phi \wedge \psi)$ is true. ■

Example 16 Reconsider the DQBF

$$P.\phi = \forall x_1, x_2 \exists y_1\big(\{x_1\}\big), y_2\big(\{x_2\}\big).\, (x_1 \vee y_1) \wedge (x_2 \vee y_2)\,.$$

and the conjunctive symmetry breaker $\psi = y_1 \to y_2$ from Example 15. Clearly, $P.\phi$ is true and a model is, for example, $s = (\neg x_1, \neg x_2) \in \mathbb{S}(P)$. Moreover, also $P.(\phi \wedge \psi)$ is true, with model $s' = (\neg x_1, \top) \in \mathbb{S}(P)$. Note that s is not a model for $P.(\phi \wedge \psi)$, but s' can be obtained from s by applying the semantic symmetry f_δ from Example 15.

6 Construction of Symmetry Breakers

In the following, we discuss the construction of a conjunctive symmetry breaker for a given DQBF $P.\phi$. What is worth noting here is that such a symmetry breaker can be constructed without the explicit knowledge of a semantic symmetry group. It suffices to know a syntactic symmetry group for $P.\phi$; the associated semantic group will act as the corresponding semantic symmetry group.

The general idea to construct a conjunctive symmetry breaker for $P.\phi$ is the same as for QBF [15, Sec. 8] and similar to the approach for SAT introduced in [8], see also [19]. First, we impose an order on the set of interpretations $\mathbb{S}(P)$. Then, using the information provided by a syntactic symmetry group, we construct a formula $\psi \in \mathrm{BF}(X \cup Y)$ so that $P.\psi$ has (at least) the minimal element in each orbit (of the associated semantic symmetry group) as a model. Any such formula is, by construction, a conjunctive symmetry breaker for $P.\phi$. The following theorem provides one way of constructing such a symmetry breaker. It is a direct generalization of the symmetry breaker construction for QBF introduced in [15, Thm. 21].

Theorem 2 *Let $P = \forall x_1, \ldots, x_n \exists y_1(D_1), \ldots, y_k(D_k)$ be a prefix for X and Y. Furthermore, let G_{syn} be a group acting admissibly on $\mathrm{BF}(X \cup Y)$ and let G_{sem} be the associated group of G_{syn}. Then*

$$\psi = \bigwedge_{g \in G_{\mathrm{syn}}} \bigwedge_{i=1}^{k} \Big(\Big(\bigwedge_{x \in D_1 \cup \cdots \cup D_i} (x \leftrightarrow g(x)) \wedge \bigwedge_{j<i} (y_j \leftrightarrow g(y_j)) \Big) \to (y_i \to g(y_i)) \Big)$$

is a conjunctive symmetry breaker for G_{syn} and G_{sem}.

Proof Fix an arbitrary order $<$ on the set of assignments $\mathbb{A}(X)$. On the set of interpretations $\mathbb{S}(P)$, define an order $s < s'$ for $s = (s_1, \ldots, s_k)$ and $s' = (s'_1, \ldots, s'_k)$, if $s \neq s'$ and for the smallest index $i \in \{1, \ldots, k\}$ with $s_i \neq s'_i$ and the smallest assignment $\sigma \in \mathbb{A}(X)$ with $[y_i]_{\sigma_s} \neq [y_i]_{\sigma_{s'}}$, we have $[y_i]_{\sigma_s} = \bot$ and $[y_i]_{\sigma_{s'}} = \top$.

Let $s_0 \in \mathbb{S}(P)$. We need to show that there are $g_{\text{syn}} \in G_{\text{syn}}$ and $g_{\text{sem}} \in G_{\text{sem}}$ such that $[g_{\text{syn}}(\psi)]_{g_{\text{sem}}(s_0)} = \top$. Let $g_{\text{syn}} = \text{id}$ and let g_{sem} be such that $s \coloneqq g_{\text{sem}}(s_0)$ is as small as possible in the order defined above. Note that such a choice of s is always possible since the set of interpretations $\mathbb{S}(P)$ is finite.

We show that $[\psi]_s = \top$. Assume, for contradiction, that $[\psi]_s = \bot$. Then there exists an assignment $\sigma \in \mathbb{A}(X)$ such that $[\psi]_{\sigma_s} = \bot$. In particular, there exist $g \in G_{\text{syn}}$ and $i \in \{1, \ldots, k\}$ satisfying the following properties:

1. $[x]_\sigma = [g(x)]_\sigma$ for all $x \in X_0 \coloneqq D_1 \cup \cdots \cup D_i$;
2. $[y_j]_{\sigma_s} = [g(y_j)]_{\sigma_s}$ for all $j < i$;
3. $[y_i]_{\sigma_s} = \top \neq \bot = [g(y_i)]_{\sigma_s}$.

Fix such an i and such an assignment σ. We may assume that the chosen σ is minimal with respect to the order fixed at the beginning (among all σ's that qualify for the chosen i).

We will now construct another interpretation $s' = g_{\text{sem}}(s_0)$ for a suitable $g_{\text{sem}} \in G_{\text{sem}}$ satisfying $s' < s$, which will contradict the minimality of s.

By Lemma 3 and 4, the element $g \in G_{\text{syn}}$ can be translated into an element $f \in G_{\text{sem}}$ such that $g(\sigma)_{f(s)} = g(\sigma_s)$. In particular, the three conditions above imply:

1. $[y_j]_{\sigma_s} = [y_j]_{\sigma_{f(s)}}$ for all $j < i$;
2. $[y_i]_{\sigma_{f(s)}} = \bot$.

By Lemma 6, applied to $\sigma_0 \coloneqq \sigma|_{X_0}$ and the subgroup $G_0 \subseteq G_{\text{syn}}$ consisting of all $g_0 \in G_{\text{syn}}$ with $[x]_\sigma = [g_0(x)]_\sigma$ for all $x \in X_0$, there exists an interpretation s' such that

$$s'|_\tau = \begin{cases} f(s)|_{\sigma_0} & \text{if } \tau = \sigma_0, \\ s|_\tau & \text{if } \tau \neq \sigma_0, \end{cases}$$

for all $\tau \in \mathbb{A}(X_0)$. Furthermore, there is $h \in G_{\text{sem}}$ with $h(s) = s'$. By construction, we have $[y_j]_{\tau_s} = [y_j]_{\tau_{s'}}$ for all $\tau \in \mathbb{A}(X)$ and all $j < i$, i.e., the functions in the jth components of s and s' agree for all $j < i$. Furthermore, at the ith component, we have $[y_i]_{\tau_s} = [y_i]_{\tau_{s'}}$ for all $\tau < \sigma$ by the minimality of σ and the choice of s'. Finally, we have $[y_i]_{\sigma_s} = \top \neq \bot = [y_i]_{\sigma_{s'}}$. Therefore, $s' < s$, in contradiction to the minimality of s. ∎

Note that, if a formula $\psi_1 \wedge \psi_2$ is a conjunctive symmetry breaker, then so are ψ_1 and ψ_2. Therefore, when constructing the symmetry breaker from Theorem 2, we are free to limit the outermost conjunction to a subset of the elements from G_{syn}. This can be beneficial in situations where the syntactic symmetry group contains a lot of elements, as it often happens in practice. In such cases, picking a set E of generators for G_{syn} and using only (some of)

the elements from E to construct the symmetry breaker can help maintain a manageable formula size.

Like in the case of SAT or QBF, also DQBF solvers typically expect their input to be in conjunctive normal form (CNF). Recall that a DQBF $P.\phi$ is in CNF if ϕ is a disjunction of clauses, where a clause is a conjunction of literals and a literal is either a variable or its negation. While the symmetry breaker from Theorem 2 as presented is not in CNF, it can be readily encoded in this form. To this end, we generalize the encoding from [15, Sec. 8] for QBFs, which, in turn, is based on the propositional case [11,19].

Fix a prefix $P = \forall x_1, \ldots, x_n \exists y_1(D_1), \ldots, y_k(D_k)$ for X and Y. To simplify the following discussion, we consider the following order of the propositional variables $X \cup Y$:

$$\boxed{D_1}, y_1, \boxed{D_2 \setminus D_1}, y_2, \ldots, \boxed{D_i \setminus (\bigcup_{j<i} D_j)}, y_i, \ldots, \boxed{D_k \setminus (\bigcup_{j<k} D_j)}, y_k$$

Within each block $\boxed{D_i \setminus (\bigcup_{j<i} D_j)}$ of universal variables, we assume an arbitrary but fixed order. We denote by v_j the jth variable in this sequence, for $j \in \{1, \ldots, n+k\}$. Using this order, we can, for $g \in G_{\text{syn}}$ and $i \in \{1, \ldots, k\}$, write the subformula

$$\Big(\bigwedge_{x \in D_1 \cup \cdots \cup D_i} (x \leftrightarrow g(x)) \wedge \bigwedge_{j<i} (y_j \leftrightarrow g(y_j)) \Big) \rightarrow (y_i \rightarrow g(y_i))$$

of ψ from Theorem 2 as

$$\Big(\bigwedge_{j=1}^{d_i - 1} (v_j \leftrightarrow g(v_j)) \Big) \rightarrow (v_{d_i} \rightarrow g(v_{d_i})),$$

where $d_i = |D_1 \cup \cdots \cup D_i| + i$.

Now, with a set of new variables $\{z_0^g, \ldots, z_{n+k-1}^g\}$, we recursively encode the antecedent of the outer implication above by setting

$$z_j^g \leftrightarrow (z_{j-1}^g \wedge (v_j \leftrightarrow g(v_j))$$

for $j \in \{1, \ldots, n+k-1\}$ and assuming the base case z_0^g to be true. Thus, the variable z_j^g encodes that v_r and $g(v_r)$ are equivalent for all $1 \leq r \leq j$. With this, ψ is equivalent to the formula

$$z_0^g \wedge \tag{2}$$

$$\bigwedge_{j=1}^{n+k-1} \Big(z_j^g \leftrightarrow \big(z_{j-1}^g \wedge (v_j \leftrightarrow g(v_j)) \big) \Big) \wedge \tag{3}$$

$$\bigwedge_{i=1}^{k} \Big(z_{d_i - 1}^g \rightarrow (v_{d_i} \rightarrow g(v_{d_i})) \Big), \tag{4}$$

where again $d_i = |D_1 \cup \cdots \cup D_i| + i$.

Before translating this formula into CNF, we note that the subformula (4) can be used to simplify the conjunction (3). In particular, for each j, the outer

equivalence in (3) and be replaced by an implication $\leftarrow$, and if v_j appears in (4), then also the inner equivalence can be replaced by an implication $\leftarrow$. For further details, see the proof of [11, Thm. 1]. Note that v_j appears in (4) if and only if $v_j \in Y$. With this, the CNF encoding of the symmetry breaker from Theorem 2 is given by the conjunction of the following formula for all desired $g \in G_{\text{syn}}$:

$$
\begin{aligned}
& z_0^g \wedge \\
& \bigwedge_{\substack{j=1 \\ v_j \in X}}^{n+k-1} \Big(\big(z_j^g \vee \neg z_{j-1}^g \vee v_j \vee g(v_j)\big) \wedge \big(z_j^g \vee \neg z_{j-1}^g \vee \neg v_j \vee \neg g(v_j)\big) \Big) \wedge \\
& \bigwedge_{\substack{j=1 \\ v_j \in Y}}^{n+k-1} \Big(\big(z_j^g \vee \neg z_{j-1}^g \vee \neg v_j\big) \wedge \big(z_j^g \vee \neg z_{j-1}^g \vee g(v_j)\big) \Big) \wedge \\
& \bigwedge_{i=1}^{k} \big(\neg z_{d_i-1}^g \vee \neg v_{d_i} \vee g(v_{d_i})\big).
\end{aligned}
$$

When using this encoding, the prefix P has to be extended with the existential variables $z_0^g, \ldots, z_{n+k-1}^g$ for all used $g \in G_{\text{syn}}$. The dependency set of z_i^g is given by $\{v_j \in X \mid j \leq i\}$.

7 Detection of Symmetries

To detect symmetries of DQBFs in conjunctive normal form, we introduce a representation of DQBFs as undirected, colored graphs. Based on these graphs we can employ tools like Saucy[1] to detect the symmetries. This is also the standard approach for detecting symmetries in SAT [19]. In this encoding, also the different types of quantifiers as well as the dependencies have to be taken into account. The (D)QBF is translated to a colored graph as follows.

Definition 12 Let

$$\Phi = \forall x_1, \ldots, x_n \exists y_1(D_1), \ldots y_k(D_k). \bigwedge_{i=1}^{d} C_i$$

be a DQBF in CNF, that is, $C_1, \ldots, C_d$ are clauses, with universal variables $X = \{x_1, \ldots, x_n\}$ and existential variables $Y = \{y_1, \ldots, y_k\}$. The *DQBF graph* (V, E, f) of Φ is a directed colored graph with vertices V, edges E, and coloring $f\colon V \to \{1, 2, 3\}$. The set of vertices $V = X \cup Y \cup L \cup C$ is composed of the disjoint sets

1. variables nodes $X \cup Y$,
2. literal nodes nodes $L = \bigcup_{v \in X \cup Y} \{+v, -v\}$,

[1] `http://vlsicad.eecs.umich.edu/BK/SAUCY/`

3. clause nodes $C = \{C_1, \ldots, C_d\}$.

The coloring $f\colon V \to \{1, 2, 3\}$ is defined as follows:

$$f(v) = \begin{cases} 1 & \text{if } v \in \bigcup_{x \in X} \{x, +x, -x\} \\ 2 & \text{if } v \in \bigcup_{y \in Y} \{y, +y, -y\} \\ 3 & \text{if } v \in C \end{cases}$$

Finally, the set of edges $E = E_v \cup E_d \cup E_c$ is defined by

1. variable edges $E_v = \bigcup_{v \in X \cup Y} \{(v, +v), (v, -v), (+v, -v), (-v, +v)\}$,
2. dependency edges $E_d = \{(y_i, x) \mid x \in D_i, i = 1, \ldots, k\}$,
3. clause occurrence edges $E_c = \{(c, l) \mid c \in C, l \in L, l \text{ appears in } c\}$.

In the graph, we distinguish between variable nodes $X \cup Y$, literal nodes L that represent the positive and negative literal of a variable, and clause nodes C that represent the different clauses of the formula. With the coloring, we partition the nodes in universal variables and literals (color 1), existential variables and literals (color 2), and clause nodes (color 3). This coloring ensures that only nodes of the right type are matched by the symmetry detection algorithm. Note that a variable and its two corresponding literal nodes are colored in the same color which indicates the type of quantification. The existential variables are connected to the universal variables on which they depend.

Example 17 The DQBF

$$\forall x_1, x_2 \exists y_1\big(\{x_1, x_2\}\big).\, (x_1 \vee x_2 \vee y_1) \wedge (\neg x_1 \vee \neg x_2 \vee y_1) \wedge (x_1 \vee x_2 \vee \neg y_1)$$

has the following DQBF graph.

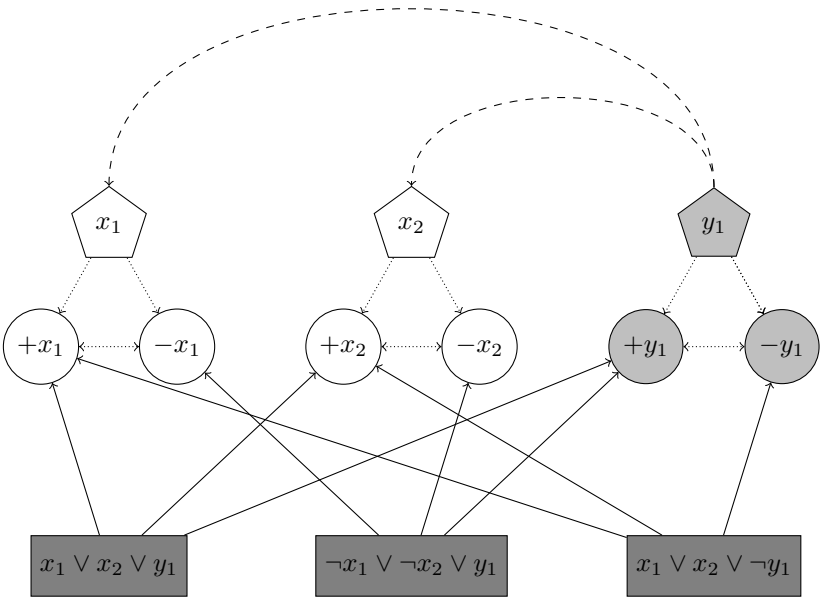

In the illustration, we distinguish between the different node types by using various shapes: variable nodes are represented as pentagons, literal nodes as circles, and clause nodes as rectangles. Similarly, the different edge types are differentiated by distinct line styles: variable edges are shown with dotted arrows, dependency edges with dashed arrows, and clause occurrence edges with solid arrows. Finally, the coloring is illustrated by different shades of the nodes: universal nodes are white, existential nodes are in a lighter gray, and clause nodes are in a darker gray.

We have implemented the translation in a tool called dqsym that can process formulas in the DQDIMACS format. This format is a more general version of the QDIMACS format and it allows for the explicit specification of quantifier dependencies. Our tool is able to process both QBFs and DQBFs in prenex conjunctive normal form (PCNF), which is also the supported format of most state-of-the-art (D)QBF solvers.

We have applied symmetry detection to the QBFs from the PCNF track and to the formulas of the DQBF track used in the QBFGallery 2023, the most recent QBF competition event.[2] The QBF set contains 377 formulas and the DQBF set contains 354 formulas.

For each DQBF, generating the graph encoding and detecting the symmetries took less than a second. The sizes of the symmetry groups are shown on the left of Figure 2. In particular, the figure presents a histogram showing the number of instances with group size of at most $10^0, 10^1, 10^2$, and 10^3, respectively, as well as those with group size greater than 10^3.

More than half of the formulas (190) do not have any symmetries. The group size of 116 formulas is between 2 and 10, indicating that a few variables can be exchanged safely. There are, however, 14 formulas with huge group sizes, the largest having a size of $7.622\,442 \times 10^{30}$.

For 350 of the 377 QBFs, the generation of the graph and the symmetry detection took less than 10 seconds. For one formula, the symmetries could not be detected within a time limit of 15 seconds, and for three formulas the graph became too large to be processed. One of these graphs had almost two billion edges, which can be explained as follows. In the DQBF graph, the dependencies between variables are represented by edges between the variable nodes. In this case, there were many universally quantified variables occurring to the left of the huge last quantifier block. Therefore, it was necessary to include an edge between each of these existential and universal variables.

In order to get a more compact encoding, the different quantifier blocks could be colored in different colors. This approach would, however, only work for QBFs, but not for DQBFs.

The right side of Figure 2 shows some statistics on the group sizes of the QBFs. Here, almost half of the instances have a lot of symmetries. It remains to be explored to what extent these symmetries can be exploited in practice.

[2] `https://qbf23.pages.sai.jku.at/gallery/`

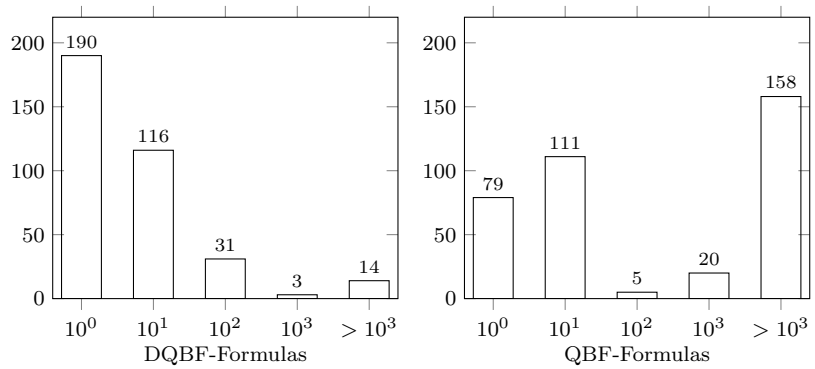

Fig. 2: Histograms of symmetry group sizes for different (D)QBFs

8 Conclusion

With this work, we lay a solid theoretical foundation for the study of symmetries of DQBFs, which hopefully sparks further exploration and innovation in both QBF/DQBF theory and solver development. Based on the concise definition of symmetry breakers given in this paper, there are many promising directions for future work. For example, one could investigate different constructions for symmetry breakers [17] or try to lift more recent improvements in symmetry breaking from SAT [11] to DQBF. Further, in this work, we focus solely on *static* symmetry breaking, where a formula is extended by a symmetry breaker as a preprocessing step. It could be beneficial to investigate also *dynamic* symmetry breaking techniques [12,10], which interfere directly in the solving process. Another promising direction of future work could be to extend existing DQBF proof systems with symmetry rules, analogous to [16,21,14], and investigate their properties.

References

1. F. A. Aloul, K. A. Sakallah, and I. L. Markov. *Efficient Symmetry Breaking for Boolean Satisfiability*. IEEE Transactions on Computers 55(5), p. 549–558 (2006)
2. M. Artin. *Algebra*. Pearson Prentice Hall (2011)
3. G. Audemard, S. Jabbour, and L. Saïs. *Symmetry Breaking in Quantified Boolean Formulae*. In: Proceedings of International Joint Conference on Artificial Intelligence, p. 2262–2267. (2007)
4. G. Audemard, B. Mazure, and L. Saïs. *Dealing with Symmetries in Quantified Boolean Formulas*. In: Proceedings of Theory and Applications of Satisfiability Testing. Online Proceedings (2004)
5. G. Audemard, S. Jabbour, and L. Saïs. *Efficient symmetry breaking predicates for Quantified Boolean Formulae*. In: Proceedings of Workshop on Symmetry and Constraint Satisfaction Problems, 7 pages (2007)
6. U. Bubeck. *Model-Based Transformations for Quantified Boolean Formulas*. PhD thesis, University of Paderborn (2009)
7. F. H. Chen, S.-C. Huang, Y.-C. Lu, and T. Tan *Reducing NEXP-complete problems to DQBF*. In: Proceedings of FMCAD, p. 199–204. TU Wien Academic Press (2022)

8. J. M. Crawford, M. L. Ginsberg, E. M. Luks, A. Roy. *Symmetry-Breaking Predicates for Search Problems*. In: Proceedings of the 5th International Conference on Principles of Knowledge Representation and Reasoning, p. 148–159. Morgan Kaufmann (1996)
9. D. Cohen, P. Jeavons, C. Jefferson, K. E. Petrie, and B. M. Smith. *Constraint Symmetry and Solution Symmetry*. In: Proceedings of the National Conference on Artificial Intelligence, p. 1589–1592. AAAI Press (2006)
10. J. Devriendt, B. Bogaerts, B., and M. Bruynooghe. *Symmetric explanation learning: Effective dynamic symmetry handling for SAT*. In: Proceedings of Theory and Applications of Satisfiability Testing, p. 83–100. Springer (2017)
11. J. Devriendt, B. Bogaerts, M. Bruynooghe, and M. Denecker. *Improved static symmetry breaking for SAT*. In: Proceedings of Theory and Applications of Satisfiability Testing, p. 104–122. Springer (2016)
12. J. Devriendt, B. Bogaerts, B. De Cat, M. Denecker, C. Mears. *Symmetry Propagation: Improved Dynamic Symmetry Breaking in SAT*. In: IEEE 24th International Conference on Tools with Artificial Intelligence, p. 49–56. IEEE (2012)
13. I. P. Gent, K. E. Petrie, and J. Puget. *Symmetry in Constraint Programming*. In: Handbook of Constraint Programming, p. 329–376. Elsevier (2006)
14. M. Kauers and M. Seidl. *Short proofs for some symmetric quantified Boolean formulas*. Information Processing Letters 140, p. 4–7 (2018)
15. M. Kauers and M. Seidl. *Symmetries of Quantified Boolean Formulas*. In: International Conference on Theory and Applications of Satisfiability Testing, p. 199–216. Springer (2018)
16. B. Krishnamurthy. *Short Proofs for Tricky Formulas*. Acta informatica 22, p. 253–275 (1985)
17. N. Narodytska and T. Walsh. *Breaking Symmetry with Different Orderings*. In: International Conference on Principles and Practice of Constraint Programming, p. 545–561. Springer (2013)
18. G. Peterson, J. Reif, and S. Azhar. *Lower Bounds for Multiplayer Non-Cooperative Games of Incomplete Information*. Computers and Mathematics with Applications 41(7–8), p. 957–992 (2001)
19. Karem A. Sakallah. *Symmetry and Satisfiability*. In: Handbook of Satisfiability, 2nd edition, p. 289–338. IOS Press (2021)
20. C. Scholl and R. Wimmer. *Dependency Quantified Boolean Formulas: An Overview of Solution Methods and Applications* In: International Conference on Theory and Applications of Satisfiability Testing, p. 3–16. Springer (2018)
21. A. Urquhart. *The symmetry rule in propositional logic*. Discrete Applied Mathematics 96, p. 177–193 (1999)