

Antichain-Based QBF Solving

Thomas Brihaye¹, Véronique Bruyère², Laurent Doyen³,
Marc Ducobu¹, and Jean-Francois Raskin⁴

¹ Institut de Mathématique – Université de Mons, Belgique

² Institut d'Informatique – Université de Mons, Belgique

³ LSV, ENS Cachan & CNRS, France

⁴ Département d'Informatique – Université Libre de Bruxelles, Belgique

Abstract. We consider the problem of QBF solving viewed as a reachability problem in an exponential And-Or graph. Antichain-based algorithms for reachability analysis in large graphs exploit certain subsumption relations to leverage the inherent structure of the explored graph in order to reduce the effect of state explosion, with high performance in practice.

In this paper, we propose simple notions of subsumption induced by the structural properties of the And-Or graphs for QBF solving. Subsumption is used to reduce the size of the search tree, and to define compact representations of certificates (in the form of antichains) both for positive and negative instances of QBF. We show that efficient exploration of the reduced search tree essentially relies on solving variants of Max-SAT and Min-SAT. Preliminary stand-alone experiments of this algorithm show that the antichain-based approach is promising.

1 Introduction

The problem of evaluating the truth value of a quantified Boolean formula (QBF) is one of the most popular PSPACE-complete problems, like SAT (the satisfiability problem for Boolean formulas) is the typical NP-complete problem. QBF is a simple and elegant formalism in which many problems of practical interest can be encoded, in a large number of areas such as automated planning, artificial intelligence, logic reasoning, and verification [14,6]. For instance, QBF can encode reachability problems more succinctly than SAT with a formula that is logarithmic in the diameter of the system when it can only be done linearly in the diameter with a SAT formula [18]. As another example, SAT and QBF can be integrated for bounded model-checking where the existence of a path is encoded by SAT, and termination is checked with QBF [10].

The simple form of QBF makes the problem appealing and accessible to a large community. However, despite its apparent simplicity, the design of efficient algorithmic solutions remains challenging. Recent progress has been observed in the practical approaches to this problem. In particular, generalizations of heuristics and optimizations used in SAT solving have been applied to QBF with some success [24,27].

Many algorithms have been proposed in the literature to solve QBF and competitive events like QBFEVAL aim at assessing the advances in reasoning about QBF [27,15]. Several leading QBF solvers are search-based. They typically use pruning techniques that extend the DPLL search strategy from SAT to QBF [8]. Common heuristics are unit

$$\psi = \underbrace{(x_1 \vee y_4 \vee y_7)}_{c_1} \wedge \underbrace{(x_2 \vee \bar{y}_4 \vee x_6)}_{c_2} \wedge \underbrace{(x_1 \vee x_3 \vee y_5 \vee \bar{y}_7)}_{c_3} \wedge \underbrace{(\bar{x}_3 \vee \bar{y}_5 \vee \bar{x}_6 \vee y_7)}_{c_4} \wedge \underbrace{(\bar{x}_1 \vee x_2 \vee y_4)}_{c_5} \wedge \underbrace{(\bar{x}_2 \vee \bar{y}_7)}_{c_6} \wedge \underbrace{(x_1 \vee x_2 \vee x_3 \vee y_7)}_{c_7}$$

Fig. 1. The CNF formula ψ for the QBF formula $f = \forall x_1 x_2 x_3 \cdot \exists y_4 y_5 \cdot \forall x_6 \cdot \exists y_7 \cdot \psi$

propagation, conflict learning and back-jumping, which are implemented in tools like QuBE [16,17] and DepQBF [22,23]. Recent works have focused on certifying (rather than just evaluating) QBF formulas, as certificates can help in extracting error traces in QBF-encoded problems. The tool suites ChEQ and sKizzo/ozziKs evaluate and certify QBF formulas [3,17,24].

In verification and automata theory, a typical PSPACE-complete problem is the universality problem for nondeterministic finite automata. Despite its worst-case exponential complexity, dramatic performance improvements have been obtained recently for this problem by *antichain algorithms* [11,12]. One key idea of antichain algorithms is to exploit the underlying structure of automata constructions (classically, powerset-based constructions) to define *subsumption relations*, yielding compact symbolic representations, as well as sound pruning of the search space. Although QBF is also a PSPACE-complete problem, this natural idea has never been used in QBF solving.

In this paper, we identify structural properties of QBF and we define pruning strategies to obtain antichain algorithms for QBF. The purpose is to define and evaluate antichain-based techniques for QBF solving, and to suggest that their integration in other search-based solvers could be valuable. We take the classical view of QBF as a reachability problem in an exponential And-Or graph, where the nodes represent subformulas (the And-nodes correspond to universal quantifications, and the Or-nodes to existential quantifications). We illustrate the main ideas of the algorithm on the following running example. Let $f = \forall x_1 x_2 x_3 \cdot \exists y_4 y_5 \cdot \forall x_6 \cdot \exists y_7 \cdot \psi$ where ψ (shown in Fig. 1) is a CNF formula viewed as the set of clauses $\{c_1, c_2, \dots, c_7\}$. The And-Or graph for f is a DAG where each level corresponds to a block of quantifiers (see a partial expansion in Fig. 2 where each clause c_i is identified with its index i). Nodes of the DAG are subsets $\varphi \subseteq \psi$ of clauses which remain to be satisfied in the evaluation game. The root of the DAG is the set ψ of all clauses. The successors of a node at level i are the sets of clauses obtained by assigning the variables quantified in the i th block of the formula. In the game interpretation, two players P_\forall and P_\exists choose the successor of the nodes (player P_\forall in And-nodes, and P_\exists in Or-nodes) by assigning the variables quantified in the block of the level of the node. The goal of player P_\exists is to reach a node \emptyset where all clauses are satisfied, and to avoid nodes where all literals of a clause are false (denoted by \perp). The QBF formula is true if and only if P_\exists has a winning strategy to reach \emptyset from the initial node ψ in the game. A key observation is that *set inclusion* is a subsumption relation that can be used to substantially reduce the size of the search tree: if player P_\exists has a winning strategy from a node ψ_1 at level i , then player P_\exists also has a winning strategy from all nodes $\psi_2 \subseteq \psi_1$ at level i , because in ψ_2 less clauses remain




Fig. 2. Search tree for the formula of Fig. 1

to be satisfied. This has two implications in the search through the DAG. First, player P_{\exists} should only consider valuations that make true a *maximal subset* of the remaining clauses, while player P_{\forall} should make true a *minimal subset*. Computing such variable assignments reduces to solving variants of Max-SAT and Min-SAT problems [19,21]. Second, the set of winning nodes at level i is downward-closed, and the set of losing nodes at level i is upward-closed. Therefore, antichains of incomparable sets of clauses are the appropriate representation of winning and losing nodes. We exploit this structure when backward propagating the information collected during the exploration of the DAG, and we never explore a node which is smaller than a winning node (or greater than a losing node) at the same level. Finally, the information stored in the antichains at the end of the search is so rich that it immediately provides compact certificates for both positive and negative instances of QBF. Note that compact certificates represented by antichains is a new notion.

We propose an antichain algorithm which is search-based and reduces the search space using antichains of winning and losing nodes of the And-Or graph. Antichains can be viewed as compact symbolic representations which can be exponentially succinct, thus it also has the flavor of symbolic procedures. In contrast, traditional symbolic QBF solvers rely on binary decision diagrams (BDD) and they are based on quantifier elimination, such as Skolemization-based approaches [2] (with the aim at eliminating existentially quantified variables), or symbolic quantifier elimination by clause resolution or BDD algorithms [25]. The tool sKizzo falls in this category of solvers [5].

We have implemented the ideas presented in this paper in a stand-alone prototype in order to push and evaluate the approach, independently of the established heuristics commonly used in search-based QBF solvers. While some benchmarks are solved more efficiently with our prototype (e.g., see Fig. 5), the results are encouraging beyond the absolute performance. In particular the experiments and comparison with state-of-the-art solvers show that:

- the search trees constructed by our algorithm are generally much smaller (by orders of magnitude) as compared to the entire search space, thanks to subsumption (e.g., see Table 1);
- our algorithm automatically provides certificates with no additional cost, whereas in other approaches, additional computation is required to extract certificates after evaluation of the formula;

- difficult instances (several hundreds of variables, thousands of clauses) are solved by our prototype (e.g., see Table 1), and on several families of formulas, the overall behaviour of our algorithm scales better or similarly as the size of the formulas increases (e.g., see Figs. 5-8).

2 Preliminaries

2.1 Notations and QBF Problem

Let $V = \{x_1, x_2, \dots, x_m\}$ be a set of m Boolean variables, we use X, X_1, X_2, \dots to denote subsets of V . A *literal* ℓ is either a variable $x \in V$ or the negation \bar{x} of a variable $x \in V$, and a *clause* c is a disjunction of literals, or equivalently a set of literals. We use notations such as $\ell \in c, \bar{x} \in c$, etc. A CNF formula is a conjunction of clauses, or equivalently a set of clauses. The empty CNF formula is denoted by 1, and the empty clause by 0. In the figures we use the notations \emptyset and \perp instead of 1 and 0 respectively. Given a set $X \subseteq V$ and a CNF formula ψ over V , we denote by $\pi_X(\psi)$ the projection of ψ over X , with $\pi_X(\psi) = \{\pi_X(c) \mid c \in \psi\}$ and $\pi_X(c) = \{\ell \in c \mid (\ell = x \vee \ell = \bar{x}) \text{ such that } x \in X\}$.

A *quantified Boolean formula* (QBF) is an expression $Q_1 X_1 \cdot Q_2 X_2 \cdots Q_n X_n \cdot \psi$ where each $Q_i \in \{\exists, \forall\}$ for $1 \leq i \leq n$, the sets X_1, \dots, X_n (called *blocks*) form a partition of V , and ψ is a CNF formula over V . We also write $Q_1 x_1 \cdot Q_2 x_2 \cdots Q_n x_n \cdot \psi$ when each block X_i contains one variable ($X_i = \{x_i\}$). Since ψ is in CNF we assume w.l.o.g. that the last block is *existential* (i.e., $Q_n = \exists$). The truth value of a QBF formula is defined as usual. The *QBF evaluation problem* is to decide whether a given QBF formula is true or false. This problem is PSPACE-complete [26].

A *valuation* for $X \subseteq V$ is a function $v : X \rightarrow \{0, 1\}$. The domain of v is $\text{dom}(v) = X$. If $X = \{x_1, \dots, x_k\}$, a valuation $v : X \rightarrow \{0, 1\}$ can be identified with a word $a_1 a_2 \cdots a_k \in \{0, 1\}^{|X|}$ such that $a_l = v(x_l)$ for all $1 \leq l \leq k$. The empty word ϵ corresponds to $\text{dom}(v) = \emptyset$. Given a partition $P = X_1 \cup X_2 \cup \cdots \cup X_n$ of V , let $X_{\leq i}$ be the set of variables $X_1 \cup X_2 \cdots \cup X_i$ (with $X_{\leq 0} = \emptyset$), and let $X_{\geq i} = V \setminus X_{\leq i-1}$. Given the valuations $v : X_{\leq i-1} \rightarrow \{0, 1\}$ and $w : X_i \rightarrow \{0, 1\}$, let vw be the valuation identified with the concatenation of the words representing v and w .

A clause c is *satisfied* by a valuation v (written $v \models c$) if there exists $x \in \text{dom}(v)$ such that either $x \in c$ and $v(x) = 1$, or $\bar{x} \in c$ and $v(x) = 0$. Given a CNF formula ψ , we denote by $\text{sat}_v(\psi)$ the set of clauses $c \in \psi$ such that $v \models c$. We denote by $\psi[v]$ the CNF formula obtained by replacing in ψ each variable $x \in \text{dom}(v)$ by its value $v(x)$. Formula $\psi[v]$ is supposed to be simplified using the laws $c \vee 1 = 1, c \vee 0 = c$ with c being a clause, and $\varphi \wedge 1 = \varphi, \varphi \wedge 0 = 0$ with φ being a CNF formula.

Let ψ be an unsatisfiable CNF formula. An *unsatisfiable core* ψ' of ψ is any subset of clauses of ψ , minimal for the inclusion, such that ψ' is still unsatisfiable.

2.2 QBF Problem as a Game

It is classical to view the QBF evaluation problem as reachability in an And-Or graph, or equivalently as a two-player reachability game [26]. For the formula $f = Q_1 x_1 \cdot$

$Q_2x_2 \cdots Q_mx_m \cdot \psi$ over $V = \{x_1, x_2, \dots, x_m\}$, the game is played in m rounds (numbered $1, \dots, m$) by the existential player P_\exists and the universal player P_\forall . In round i , the truth value of the variable x_i is chosen by player P_{Q_i} . After m rounds, the players have constructed a valuation $v : V \rightarrow \{0, 1\}$, and player P_\exists wins if $\psi[v] = 1$ (all clauses are satisfied by v), otherwise player P_\forall wins. It is easy to see that P_\exists has a winning strategy in this game iff the formula f is true. Note that instead of having one round for each variable, we can also consider a game with one round for each block of variables, such that the blocks correspond to quantifier alternations in f . The players then choose a valuation for all the variables in the block at once, and the number of rounds is equal to the number of quantifier alternations. As the algorithms proposed in this paper are based on this game metaphor, we present the And-Or graph on which the game is played.

Let $P = X_1 \cup X_2 \cup \cdots \cup X_n$ be a partition of $V = \{x_1, x_2, \dots, x_m\}$, and let $f = Q_1X_1 \cdot Q_2X_2 \cdots Q_nX_n \cdot \psi$ be a QBF formula over V . We define the *And-Or graph* $G_f = (S, S_\exists, S_\forall, s_0, E, F)$ where:

- $S = \{\psi[v] \mid \text{dom}(v) = X_{\leq i-1}, \text{ for } i, 1 \leq i \leq n+1\}$;
- $S_\exists = \{\psi[v] \mid \text{dom}(v) = X_{\leq i-1} \wedge Q_i = \exists, \text{ for } i, 1 \leq i \leq n\}$ is the set of P_\exists nodes;
- $S_\forall = \{\psi[v] \mid \text{dom}(v) = X_{\leq i-1} \wedge Q_i = \forall, \text{ for } i, 1 \leq i \leq n\}$ is the set of P_\forall nodes;
- $s_0 = \psi$ is the initial node;
- $E = \{(\psi[v], \psi[vw]) \mid \text{dom}(v) = X_{\leq i-1} \wedge \text{dom}(w) = X_i, \text{ for } i, 1 \leq i \leq n\}$ is the set of edges;
- $F = \{\psi[v] \in S \mid \psi[v] = 1\}$ is the set of final nodes.

The set S is naturally partitioned into *levels* as follows: $S = \text{Level}_1 \cup \text{Level}_2 \cup \cdots \cup \text{Level}_{n+1}$ where $\text{Level}_i = \{\psi[v] \mid \text{dom}(v) = X_{\leq i-1}\}$ for each $1 \leq i \leq n+1$. The objective of player P_\exists is to reach the set F of nodes $\psi[v]$ such that all clauses of ψ are satisfied by v . The game starts in node s_0 and player P_Q ($Q \in \{\exists, \forall\}$) chooses the successor of node s if $s \in S_Q$. Thus if $s = \psi[v] \in S_\exists$ and $\text{dom}(v) = X_{\leq i-1}$, then player P_\exists chooses one of the $2^{|X_i|}$ possible successors of s in E , corresponding to a valuation $w : X_i \rightarrow \{0, 1\}$. A node s is *winning* for player P_\exists if he has a strategy to force reaching a node in F from s , no matter the choices of P_\forall ; otherwise it is *losing*. We denote by W the set of winning nodes for player P_\exists , and by $L = S \setminus W$ the set of losing nodes for P_\exists . We say that P_\exists is *winning the game* if $s_0 \in W$. In the sequel, we use the notations W_i (resp. L_i) to denote $W \cap \text{Level}_i$ (resp. $L \cap \text{Level}_i$).

Proposition 1. *A QBF formula f is true iff player P_\exists is winning the game G_f .*

Note that in the graph G_f , each node $\psi[v]$ with $\text{dom}(v) = X_{\leq i-1}$ can be associated with the formula $\text{Formula}(\psi[v]) \equiv Q_iX_i \cdots Q_nX_n \cdot \psi[v]$, and we can strengthen the previous proposition as follows.

Proposition 2. *Given a QBF formula f , the set of winning nodes in the graph G_f is $W = \{\psi[v] \in S \mid \text{Formula}(\psi[v]) \text{ is true}\}$, and the set of losing nodes is $L = \{\psi[v] \in S \mid \text{Formula}(\psi[v]) \text{ is false}\}$.*

2.3 Structure in the And-Or Graph and Antichains

We present in the next section an algorithm to solve the game played on G_f which exploits the following *subsumption* relation on QBF formulas. We write $f_1 \sqsubseteq f_2$ if




Fig. 3. Level₁ and Level₂ of G_f and the 5 minimal valuations of P_\forall

$f_1 = Q_i X_i \cdots Q_n X_n \cdot \psi_1$ and $f_2 = Q_i X_i \cdots Q_n X_n \cdot \psi_2$ are two QBF formulas with the same quantifier prefix, and $\psi_1 \subseteq \psi_2$. Intuitively, f_1 is more promising than f_2 for player P_\exists because all strategies that are winning from ψ_2 are also winning from ψ_1 .

Proposition 3. *Suppose that $f_1 \sqsubseteq f_2$. If f_2 is true, then f_1 is true; and if f_1 is false, then f_2 is false.*

As a direct consequence of Propositions 2 and 3, we obtain the next corollary.

Corollary 1. *In the graph G_f , for all nodes $s_1, s_2 \in \text{Level}_i$ such that $s_1 \subseteq s_2$, i.e. $\text{Formula}(s_1) \sqsubseteq \text{Formula}(s_2)$, if $s_2 \in W_i$, then $s_1 \in W_i$; and if $s_1 \in L_i$, then $s_2 \in L_i$.*

Hence, W_i is \subseteq -downward closed and L_i is \subseteq -upward closed. The set of \subseteq -maximal elements of W_i , noted $\lceil W_i \rceil$, is an *antichain* for the partial order \subseteq (i.e., a set of pairwise incomparable elements) that canonically and compactly represents W_i . Similarly, the set of \subseteq -minimal elements of L_i , noted $\lfloor L_i \rfloor$, is an antichain that canonically and compactly represents L_i . Elements of these antichains are denoted α, β .

3 Algorithms

In Section 3.1, we discuss the computation of optimal valuations to explore only the most promising nodes, and in Section 3.2, we propose an antichain-based algorithm for solving the QBF evaluation game.

3.1 Maximal and Minimal Valuations

According to Corollary 1, when it is the turn for P_\exists to play in node $s = \varphi$ in Level_i , he can restrict his choices among valuations $w : X_i \rightarrow \{0, 1\}$ that *maximize* the set of clauses of φ that are satisfied. Symmetrically, player P_\forall can restrict his choices among valuations $w : X_i \rightarrow \{0, 1\}$ that *minimize* the set of clauses of φ that are satisfied.

We define the notion of maximal and minimal valuations as follows. Let φ be a CNF formula over $X_{\geq i}$. A valuation $w : X_i \rightarrow \{0, 1\}$ is φ -*maximal* if for all $w' : X_i \rightarrow \{0, 1\}$, $\text{sat}_w(\varphi) \subseteq \text{sat}_{w'}(\varphi)$ implies $\text{sat}_w(\varphi) = \text{sat}_{w'}(\varphi)$. Symmetrically, w is φ -*minimal* if for all $w' : X_i \rightarrow \{0, 1\}$, $\text{sat}_{w'}(\varphi) \subseteq \text{sat}_w(\varphi)$ implies $\text{sat}_w(\varphi) = \text{sat}_{w'}(\varphi)$.

Example 1. Consider the CNF formula ψ of Fig. 1, viewed as the set of clauses $\{c_1, c_2, \dots, c_7\}$. In Level_1 , we have $X_1 = \{x_1, x_2, x_3\}$ which are universal variables. Among the $2^3 = 8$ valuations, 5 are ψ -minimal (shaded in Fig. 3, where each clause c_i is identified with i). Remember that the nodes in the And-Or graph are the clauses that *remain* to be satisfied, thus maximal such sets correspond to minimal valuations. Note also that we may need to compute *all* maximal (or minimal) valuations in a node.

Maximal and minimal valuations can be computed by multiple calls to a SAT solver. Let us give the intuition for maximal valuations. Let φ be a set of clauses over $X_{\geq i}$. First notice that a valuation $w : X_i \rightarrow \{0, 1\}$ is φ -maximal if and only if it is $\pi_{X_i}(\varphi)$ -maximal. Thus we can assume w.l.o.g. that φ is a set of clauses over X_i (instead of $X_{\geq i}$). Using a set of new variables $Y = \{y_c \mid c \in \varphi\}$, called selectors, we transform the set of clauses φ into a set of clauses φ' over $X_i \cup Y$ with the following property : if a valuation $w : X_i \cup Y \rightarrow \{0, 1\}$ satisfying φ' is such that $w(y_c) = 1$, then w satisfies c . By a first call to a SAT solver on φ' , we get a valuation w and a subset C of clauses of φ that are satisfied by w . Then we modify φ' into φ'' by imposing the constraint that at least one of the variables y_c such that $w(y_c) = 0$ is true in φ , so that a second call to a SAT solver provides a subset of satisfied clauses of φ that strictly contains C . Iterating this procedure, we finally obtain a valuation that satisfies a maximal set of clauses in φ .

Computing maximal and minimal valuations can also be achieved thanks to solvers for variants of the *Maximum Satisfiability* (Max-SAT) and *Minimum Satisfiability* (Min-SAT) problems [19,21]. Given a CNF formula φ , the Max-Sat problem asks to compute a valuation that maximizes the *number* of satisfied clauses in φ (Min-Sat is defined symmetrically). Note that such a valuation is φ -maximal but the converse is not necessarily true. Given a CNF formula $\varphi = \varphi_h \wedge \varphi_s$ where φ_h represents the *hard* clauses and φ_s represents the *soft* clauses, the *partial Max-SAT* problem consists in finding a valuation such that all hard clauses are satisfied and the number of satisfied soft clauses is maximized. This variant of Max-SAT can be used to generate *all* φ -maximal valuations as follows. The first φ -maximal valuation is computed by a call to a Max-SAT solver. The next ones are computed thanks to a partial Max-SAT solver, such that hard clauses with selectors impose that for each already computed φ -maximal valuation w , at least one new clause $c \notin \text{sat}_w(\varphi)$ is satisfied.

3.2 Antichain-Based Algorithm

In this section we present an antichain-based algorithm to evaluate a QBF formula $f = Q_1 X_1 \cdots Q_n X_n \cdot \psi$. It is a search-based algorithm of the And-Or graph G_f with backward propagation of the information collected during the exploration. Such a forward-backward exploration was also used with success in timed games [9].

Our algorithm consists of two recursive procedures named $\text{ATCSearch}\exists(\varphi, i)$ and $\text{ATCSearch}\forall(\varphi, i)$ where φ is a node of G_f and i is the recursion level (see Algorithms 1 and 2). Initially, we make a call to $\text{ATCSearch}\exists(\psi, 1)$ if $Q_1 = \exists$, and to $\text{ATCSearch}\forall(\psi, 1)$ if $Q_1 = \forall$. These procedures determine whether a node φ is winning or losing for P_{\exists} , i.e. whether $\varphi \in W_i$ or $\varphi \in L_i$. The sets W_i and L_i are updated as global variables and compactly stored by antichains $\lceil W_i \rceil$ and $\lfloor L_i \rfloor$ respectively. They are used to prune the search by the *subsumption* checks (see lines 8, 11 in Algorithm 1).

The details of $\text{ATCSearch}\exists(\varphi, i)$ are as follows. If φ is not even satisfiable, then it is a losing node; if φ is satisfiable and $i = n$, then it is winning since φ belongs to S_{\exists} ; otherwise, the procedure enumerates the φ -maximal valuations w (line 7) and checks if $\varphi[w]$ is winning at level $i + 1$. For player P_{\exists} , maximal valuations are sufficient because the set of winning nodes is downward-closed (see Corollary 1). The recursive call to $\text{ATCSearch}\forall(\varphi[w], i + 1)$ can be avoided if $\varphi[w]$ is in the downward-closure of $\lceil W_{i+1} \rceil$ (line 8), or if $\varphi[w]$ is in the upward-closure of $\lfloor L_{i+1} \rfloor$ (line 11). Finally, if all φ -maximal valuations have been explored, then φ is losing (line 17).

Algorithm 1. $\text{ATCSearch}\exists(\varphi, i)$

Require: node $\varphi \in S_{\exists} \cap \text{Level}_i, i \leq n$.
Ensure: Win if $\varphi \in W_i$, Lose if $\varphi \in L_i$.

- 1: **if** $\neg \text{IsSat}(\varphi)$ **then**
- 2: Add($\varphi, \lfloor L_i \rfloor$)
- 3: **return** Lose
- 4: **if** $i = n$ **then**
- 5: Add($\varphi, \lceil W_i \rceil$)
- 6: **return** Win
- 7: **for** each φ -maximal valuation $w : X_i \rightarrow \{0, 1\}$ **do**
- 8: **if** $\exists \alpha \in \lceil W_{i+1} \rceil$ s.t. $\varphi[w] \subseteq \alpha$ **then**
- 9: Add($\varphi, \lceil W_i \rceil$)
- 10: **return** Win
- 11: **if** $\neg(\exists \alpha \in \lfloor L_{i+1} \rfloor$ s.t. $\alpha \subseteq \varphi[w])$ **then**
- 12: $R \leftarrow \text{ATCSearch}\forall(\varphi[w], i + 1)$
- 13: **if** $R = \text{Win}$ **then**
- 14: Add($\varphi, \lceil W_i \rceil$)
- 15: **return** Win
- 16: Add($\varphi, \lfloor L_i \rfloor$)
- 17: **return** Lose

Algorithm 2. $\text{ATCSearch}\forall(\varphi, i)$

Require: node $\varphi \in S_{\forall} \cap \text{Level}_i, i < n$.
Ensure: Win if $\varphi \in W_i$, Lose if $\varphi \in L_i$.

- 1: **if** $\neg \text{IsSat}(\varphi)$ **then**
- 2: Add($\varphi, \lfloor L_i \rfloor$)
- 3: **return** Lose
- 4: **for** each φ -minimal valuation $w : X_i \rightarrow \{0, 1\}$ **do**
- 5: **if** $\exists \alpha \in \lfloor L_{i+1} \rfloor$ s.t. $\alpha \subseteq \varphi[w]$ **then**
- 6: Add($\varphi, \lfloor L_i \rfloor$)
- 7: **return** Lose
- 8: **if** $\neg(\exists \alpha \in \lceil W_{i+1} \rceil$ s.t. $\varphi[w] \subseteq \alpha)$ **then**
- 9: $R \leftarrow \text{ATCSearch}\exists(\varphi[w], i + 1)$
- 10: **if** $R = \text{Lose}$ **then**
- 11: Add($\varphi, \lfloor L_i \rfloor$)
- 12: **return** Lose
- 13: Add($\varphi, \lceil W_i \rceil$)
- 14: **return** Win

The procedure $\text{ATCSearch}\forall(\varphi, i)$ for nodes $\varphi \in S_{\forall}$ is dual. Note that the case $i = n$ is not relevant since $Q_n = \exists$. By a symmetrical argument, P_{\forall} needs to consider only the φ -minimal valuations.

In these two procedures, the φ -maximal and φ -minimal valuations are computed as explained in Section 3.1, by either using a SAT solver or a partial Max-SAT solver. Procedure $\text{IsSat}(\varphi)$ tests whether formula φ is satisfiable by a call to a SAT solver. The antichains $\lceil W_i \rceil$ and $\lfloor L_i \rfloor$ (for $1 \leq i \leq n$) are initially empty. The antichain structure is maintained by the procedure **Add** which computes $\lceil \{\varphi\} \cup W_i \rceil$ and $\lfloor \{\varphi\} \cup L_i \rfloor$.

Example 2. Consider the CNF formula ψ of Fig. 1. Since $Q_1 = \forall$, the algorithm starts with $\text{ATCSearch}\forall(\psi, 1)$ which needs to explore the 5 minimal valuations of Fig. 3. Assume that the first valuation is $(x_1 \mapsto 0, x_2 \mapsto 0, x_3 \mapsto 0)$, denoted 000, which satisfies clauses 4, 5, 6. Then, the game proceeds to the node $\psi[w] = \psi[000] = \{1, 2, 3, 7\}$ of remaining clauses where the turn is to player P_{\exists} . The subgraph of G_f rooted at $\psi[000]$ is shown in the first tree of Fig. 4. Among the 4 possible valuations for player P_{\exists} , only 2 are maximal, namely 01 and 11. For valuation 01, only clauses 1 and 7 remain. At this point, all variables are instantiated except x_6 and y_7 , and player P_{\exists} wins by choosing $y_7 \mapsto 1$ which satisfies ψ no matter the value of x_6 chosen by player P_{\forall} .

Thus nodes $\{1, 7\}$ at Level_3 , and $\{1, 2, 3, 7\}$ at Level_2 are winning, and the related antichains are updated as follows: $\lceil W_3 \rceil = \{\{1, 7\}\}$ and $\lceil W_2 \rceil = \{\{1, 2, 3, 7\}\}$.

At the root node of G_f , the valuation 000 is not a good choice for player P_{\forall} , and no conclusion can be drawn yet for this node (Fig. 3). We need to explore another choice for player P_{\forall} . The second tree of Fig. 4 shows the subgame rooted at node $\psi[001]$. In this case, with minimal valuation 00, player P_{\exists} reaches node $\{1\}$ in Level_3 . Since




Fig. 4. Subgames rooted at $\psi[000]$, $\psi[001]$, and $\psi[101]$

$\{1\} \in W_3$ (indeed $\{1\} \subseteq \{1, 7\} \in \lceil W_3 \rceil$, and W_3 is the downward-closure of $\lceil W_3 \rceil$), he knows immediately that he is winning without further exploring the graph. This situation illustrates the power of the subsumption which allows to prune the search for nodes smaller than previously visited winning ones. The antichain $\lceil W_2 \rceil$ is then updated to $\{\{1, 2, 3, 7\}, \{1, 2, 4\}\}$. Valuation 001 is again a bad choice for P_\forall .

One can check that valuations 010 and 011 are bad choices for P_\forall and their exploration leads to the following update of the antichains: $\lceil W_2 \rceil = \{\{1, 2, 3, 7\}, \{1, 2, 4\}, \{1, 3, 6\}, \{1, 4, 6\}\}$ and $\lceil W_3 \rceil = \{\{1, 7\}, \{6\}\}$. The last minimal valuation is 101 and for all choices of player P_\exists , there is a choice of player P_\forall to falsify the formula (see the last tree in Fig. 4). Therefore P_\exists is losing the game and the formula f is false.

The correctness of this algorithm is established using the notion of certificate presented in the next section.

Theorem 1. *Let f be a QBF formula. Applying Algorithms 1 and 2 on f returns Win if and only if f is true.*

4 Certificates

In the previous section we have described a search-based algorithm to evaluate a QBF formula f . This algorithm computes the sets of winning nodes and losing nodes for each level of the graph G_f , and these sets are compactly represented by antichains.

Our algorithm gathers enough information in these antichains to easily build *compact certificates* for both true and false QBF formulas. Intuitively, if f is true, that is, player P_\exists is winning the game G_f , then a certificate is given by the antichains $\lceil W_i \rceil$, $1 \leq i \leq n$, and for each $\alpha \in W_i$ by the maximal valuation computed by Algorithm 1 when α has been declared winning. If f is false, a certificate is defined similarly from the antichains $\lfloor L_i \rfloor$. To the best of our knowledge these certificates are different from the ones considered in the literature [3,24].

We first define positive certificates as a witness for true QBF formulas. A *positive certificate* for a formula $f \equiv Q_1 X_1 \cdots Q_n X_n \cdot \psi$ is a pair $\langle (C_i^+)_{1 \leq i \leq n}, w \rangle$ such that:

- each C_i^+ is a set of nodes at the i th level of the graph G_f , that is, $C_i^+ \subseteq \text{Level}_i$;
- for each i such that $\text{Level}_i \subseteq S_\exists$, w is a function that assigns a valuation $w(\alpha) : X_i \rightarrow \{0, 1\}$ to each $\alpha \in C_i^+$;
- and the following properties are verified:

1. $C_1^+ = \{\psi\}$.
2. for each $i < n$ such that $\text{Level}_i \subseteq S_{\exists}$, for all $\alpha \in C_i^+$, there exists $\beta \in C_{i+1}^+$ such that $\alpha[w(\alpha)] \subseteq \beta$.
3. for each $i < n$ such that $\text{Level}_i \subseteq S_{\forall}$, for all $\alpha \in C_i^+$, for all $w : X_i \rightarrow \{0, 1\}$, there exists $\beta \in C_{i+1}^+$ such that $\alpha[w] \subseteq \beta$.
4. for $i = n$, for all $\alpha \in C_i^+$, $\alpha[w(\alpha)] = 1$.

Clearly, there exists a nondeterministic polynomial time algorithm to recognize pairs $\langle (C_i^+)_{1 \leq i \leq n}, w \rangle$ that are not positive certificate. All the verification related to Conditions 1, 2 and 4 can be done in deterministic polynomial time while Condition 3 requires nondeterminism. Therefore, verifying the validity of a positive certificate is a problem in coNP.

Lemma 1. *If $\langle (C_i^+)_{1 \leq i \leq n}, w \rangle$ is a positive certificate for a QBF formula f , then f is true.*

Lemma 2. *Let $\lceil W_i \rceil$, $1 \leq i \leq n$, be the antichains built by the execution of Algorithms 1 and 2 on formula f . For each i such that $\text{Level}_i \subseteq S_{\exists}$, for all $\alpha \in \lceil W_i \rceil$, let $w(\alpha)$ be the valuation used by Algorithms 1 when α has been declared winning. If f is true, then $\langle (\lceil W_i \rceil)_{1 \leq i \leq n}, w \rangle$ is a positive certificate for f .*

The next theorem directly follows from the two previous lemmas.

Theorem 2. *Let f be a QBF formula. Then f is true if and only if there exists a positive certificate for f .*

Negative certificates are defined in a way similar to positive certificates; they are a witness for false formulas f .

5 Optimizations

5.1 Guiding the Search to Promising Valuations

We recall that in Algorithm 1, when φ is a satisfiable formula, and $i \leq n - 1$, player P_{\exists} traverses all the maximal valuations $w : X_i \rightarrow \{0, 1\}$ in the hope to find w such that $\varphi[w]$ is winning. The first optimization that we consider tries to guide the search to *promising* maximal valuations.

Improving ATCSearch \exists . When considering all the maximal valuations $w : X_i \rightarrow \{0, 1\}$, we observe that player P_{\exists} is winning as soon as he can find a valuation w such that $\varphi[w] \subseteq \alpha$ for some $\alpha \in \lceil W_{i+1} \rceil$ (see Corollary 1). So, it is better for P_{\exists} to first try to find such a valuation w . Suppose now that player P_{\exists} cannot win by exploiting the elements of $\lceil W_{i+1} \rceil$. Then he should avoid considering maximal valuations w such that $\alpha \subseteq \varphi[w]$ for some $\alpha \in \lfloor L_{i+1} \rfloor$ as $\varphi[w]$ is losing (see Corollary 1). These two observations are exploited by the new algorithm.

Improving ATCSearch \forall . We can improve the choice of minimal valuations for player P_{\forall} in a symmetric manner. Indeed, P_{\forall} should first look for a valuation w such that there exists $\alpha \in \lfloor L_{i+1} \rfloor$ with $\alpha \subseteq \varphi[w]$. If such a valuation does not exist, then he should only consider minimal valuations that avoid the set $\lceil W_{i+1} \rceil$ of winning nodes.

5.2 Improving the Information about Losing and Winning Nodes

Our algorithms can be seen as mixing a forward exploration of the And-Or graph G_f with a backward propagation of the information about the winning and losing nodes. We present below several ways of improving the propagation phase.

At initialization. Recall that the antichain $[L_i]$ is initially empty in Algorithms 1 and 2. We can add some useful information to $[L_i]$ in the following case. Consider the initial node ψ of G_f , and let $i, 1 \leq i \leq n - 1$. Suppose that the projection $d = \pi_{X_{\geq i}}(c)$ of a clause c in ψ on $X_{\geq i}$ has all its literals *universally quantified*. Then the clause d must be satisfied by a valuation generated before reaching Level $_i$ since otherwise player P_{\forall} can falsify it. So, at initialization we can add $\{d\}$ to $[L_i]$ for any such d (all sets of clauses that contain d are losing at Level $_i$).

Some other information can be initially stored in the antichains $[W_i]$ and $[L_i]$ to better guide the search. Consider the initial node ψ of G_f , and its projection $\varphi = \pi_{X_{\geq i}}(\psi)$. If φ is unsatisfiable, then any unsatisfiable core of φ can be extracted and added to $[L_i]$. On the other hand, if φ is satisfiable and we further restrict φ to the existentially quantified variables, then we can compute a φ -maximal valuation w to get a maximal subset of satisfied clauses $\varphi' = \text{sat}_w(\varphi)$. The set φ' can be added to $[W_i]$ because player P_{\exists} has a winning strategy at Level $_i$ if the set of clauses not in φ' are satisfied when the game enters this level.

When updating $[W_i]$ and $[L_i]$. We now give an optimization that can be applied when updating the sets of winning and losing nodes during the search. We consider two scenarios. First, assume that node φ is declared winning by the algorithm $\text{ATCSearch}_{\exists}$ in Level $_i$. This means that there exists a valuation $w : X_i \rightarrow \{0, 1\}$ such that either $\varphi[w] \subseteq \alpha$ for some $\alpha \in [W_{i+1}]$, or $\varphi[w]$ is declared winning by the recursive call to $\text{ATCSearch}_{\forall}$. Notice that φ is a subset of the set of clauses in the root ψ (projected on variables $X_{\geq i}$). It may happen that other clauses $c \in \psi$ are also satisfied by the valuation w . Thus, in a way to have bigger elements in $[W_i]$, it is preferable to add the set $\varphi' = \varphi \cup \{c \mid c \in \text{sat}_w(\psi)\}$ instead of φ to $[W_i]$.

Second, assume that φ is declared losing by Algorithm $\text{ATCSearch}_{\exists}$ in Level $_i$ with $i = n$. This means that φ is unsatisfiable. Instead of adding φ to $[L_i]$, we can add any unsatisfiable core $\varphi' \subseteq \varphi$ instead.

6 Experimental Results

Setting. We have implemented the algorithms of Section 3 with the optimizations of Section 5 in a stand-alone prototype in order to evaluate the impact of the antichain approach. Thus none of the classical heuristics used in search-based QBF solvers, like backjumping, unit propagation, and monotone literals elimination [8] has been integrated. The code is written partly in C for the low level operations on the data structures, and partly in Python to implement high level operations like the exploration of the And-Or graph, and for the construction of CNF formulas submitted to the SAT or partial Max-SAT solvers. We use Python to facilitate the fast evaluation of different ideas even if some price has to be paid at the performance level. We use the SAT solver MiniSat [13] and the partial Max-SAT solver Akmaxsat [20] to compute the maximal

and minimal valuations as explained in Section 3.1, and PicoSAT [7] to compute unsatisfiable cores as described in Section 5.

In a preprocessing step, the formula is simplified with PreQuel [28] and then presented in a tree-like structure [4]. This is standard practice in QBF solving. The experiments were run on a PC equipped with a Intel i7 2.8GHz processor, 6 GB of RAM and running Linux Ubuntu 2.6.

Instances. We tested our algorithm on several instances proposed during the seventh QBF solvers evaluation (QBFEVAL'10) [27] and compared the results with the ones obtained with three state-of-the-art QBF solvers: QuBE-7.0, sKizzo-v0.8.2-beta and DepQBF-0.1 (winner of QBFEVAL'10). QuBE and DepQBF are search-based solvers, whereas sKizzo uses symbolic Skolemization.

The families k^* correspond to the encoding of the satisfiability problem for modal K formulas into QBF, and they are known to give difficult instances for search-based solvers. The families k^* have a deep level of quantifier alternations while the families Toilet* and aim-* have three quantifier alternations.

Results. In Table 1, **Var** (resp. **Cl**, **Blocs**) gives the number of variables (resp. clauses, blocs) of the instance and **Value** is its truth value. **Nodes** is the number of nodes of the And-Or graph that have been explored. The columns **ATC**, **QuBE**, **DepQBF**, and **sKizzo** present the execution times in seconds (with a timeout of 600 seconds) of our antichain-based solver and the three other solvers. The experimental results obtained with our solver are encouraging.

- According to pure performance, our prototype already performs better than the three state-of-the-art solvers for the families $k\text{-path-n}$ and $k\text{-path-p}$.

For other families in k^* , QuBE and DepQBF solvers (which are search-based) have often an execution time beyond 600 seconds.

See Table 1, Fig. 5 and 6.

- The antichain approach leads to search trees that are *amazingly small* as compared to the size of the entire search space (see the **Nodes** entry in Table 1 as compared to the size of the entire search space in $O(2^{\text{Var}})$); the antichains are thus also very small as they are composed of nodes of the search tree.
- For several families from the QBFLIB library [15] (including Toilet* and aim-*), the execution times of our solver follow the same shape of curve (at logarithmic scale, with a timeout of 600 seconds) as for the other three state-of-the-art solvers, see for instance the family aim-100 in Fig. 7 (true instances) and Fig. 8 (false instances).

Table 1. Family $k\text{-path-n}$

Instance	Var	Cl	Blocs	Value	Nodes	ATC	QuBE	DepQBF	sKizzo
k-path-n-01	108	275	7	1	14	0.22	0.01	0.005	0.01
k-path-n-02	180	481	9	1	43	0.93	0.01	0.02	0.05
k-path-n-05	384	1051	15	1	120	4.17	2.53	39.97	0.08
k-path-n-06	456	1257	17	1	142	7.54	13.74	/	0.26
k-path-n-10	732	2033	25	1	247	16.33	/	/	43.87
k-path-n-11	804	2234	27	1	269	27.36	/	/	445.71
k-path-n-20	1428	3992	45	1	503	95.19	/	/	/
k-path-n-21	1488	4155	47	1	452	88.81	/	/	/




Fig. 5. Family k-path-n




Fig. 6. Family k-path-p




Fig. 7. Family aim-100 (true - log scale)




Fig. 8. Family aim-100 (false - log scale)

As explained in Section 3.1, the maximal and minimal valuations can be computed either with a SAT solver or with a partial Max-SAT solver. In Section 5.2, we described several ways to improve the information about losing and winning nodes in the antichains. These different approaches have been tested, and the experiments show the best approach can vary with the family of formulas that is tested (the table and figures always present the results for the best approach). For instance, depending on the family, the size of the search space can be either increased or decreased when using a partial MaxSAT solver instead of a SAT solver.

Along with deciding if a given QBF formula is true or false, our solver provides compact certificates in the form of antichains both for positive and negative instances of QBF, and *without any additional computation* (see Section 4). DepQBF solver does not construct certificates. For sKizzo and QuBE solvers, additional computation is required to extract certificates [3,17,24]. The log produced by sKizzo is evaluated by ozziKs to construct certificates (only for true instances); QuBE-cert (in the suite ChEQ) is an extension of QuBE that adds to QuBE the instrumentation required to generate certificates (for both true and false instances). In [3], experiments have been done with the family `adder*` with “*the surprising phenomenon that the time taken to reconstruct a model may overcome the time needed to solve the instance*”.

7 Conclusion and Perspectives

The approach presented in this paper for QBF solving is inspired by previous works on effective antichain algorithms for PSPACE-complete problems in automata theory which are based on simple subsumption relations [11,12]. While the And-Or graph of QBF formulas enjoys such a subsumption relation, this idea has not been exploited in search-based QBF solvers. In a prototypical implementation, we have evaluated the feasibility of antichain-based algorithms for QBF. Experimental results show that on several benchmarks the size of the search tree is drastically reduced, and that some instances are solved more efficiently than by the leading QBF solvers. This shows that the antichain approach is promising and it provides a new research direction in the area. Its integration in standard search-based QBF solvers is worth investigating. On the other hand, our algorithm provides automatically compact certificates represented by antichains with no additional cost, and for both true and false instances.

We now plan to improve our algorithm with respect to the computation of the valuations. Indeed, we observed on the experiments that the execution time is partly spent when computing maximal and minimal valuations. Our current solver computes the *best* valuations since they are restricted to maximal/minimal ones, and exploit as much as possible the information stored in the antichains. We could instead compute *approximate* valuations in a way to decrease the execution time while keeping the advantages of the antichains [1].

Our work also provides a new application of the Max-SAT problem. We intend to submit instances coming from our experiments to the *Evaluation of Max-SAT Solvers* that is yearly organized as an affiliated event of the International Conference on Theory and Applications of Satisfiability Testing (SAT)¹. We believe that the difficult instances we produce could be of interest to the Max-SAT community and that antichain-based QBF solving would benefit from their improvements.

Acknowledgements. We thank Marco Benedetti, the author of sKizzo, for his great help, Florian Lonsing for explanations about DepQBF, the QuBE's team for answers about their solver, and Nicolas Maquet for his guidance in the implementation.

References

1. Asano, T., Williamson, D.P.: Improved approximation algorithms for max sat. *J. Algorithms* 42(1), 173–202 (2002)
2. Benedetti, M.: Evaluating QBFs via Symbolic Skolemization. In: Baader, F., Voronkov, A. (eds.) LPAR 2004. LNCS (LNAI), vol. 3452, pp. 285–300. Springer, Heidelberg (2005)
3. Benedetti, M.: Extracting Certificates from Quantified Boolean Formulas. In: Kaelbling, L.P., Saffioti, A. (eds.) IJCAI, pp. 47–53. Professional Book Center (2005)
4. Benedetti, M.: Quantifier Trees for QBFs. In: Bacchus, F., Walsh, T. (eds.) SAT 2005. LNCS, vol. 3569, pp. 378–385. Springer, Heidelberg (2005)
5. Benedetti, M.: sKizzo: A Suite to Evaluate and Certify QBFs. In: Nieuwenhuis, R. (ed.) CADE 2005. LNCS (LNAI), vol. 3632, pp. 369–376. Springer, Heidelberg (2005)

¹ See <http://www.maxsat.udl.cat/>

6. Benedetti, M., Mangassarian, H.: Qbf-based formal verification: Experience and perspectives. *JSAT* 5(1-4), 133–191 (2008)
7. Biere, A.: PicoSAT Essentials. *JSAT* 4(2-4), 75–97 (2008)
8. Cadoli, M., Giovanardi, A., Schaerf, M.: An algorithm to evaluate quantified Boolean formulae. In: Proc. of AAAI 1998/IAAI 1998, pp. 262–267. MIT Press, Cambridge (1998)
9. Cassez, F., David, A., Fleury, E., Larsen, K.G., Lime, D.: Efficient on-the-fly algorithms for the analysis of timed games. In: Abadi, M., de Alfaro, L. (eds.) CONCUR 2005. LNCS, vol. 3653, pp. 66–80. Springer, Heidelberg (2005)
10. Cook, B., Kroening, D., Sharygina, N.: Verification of boolean programs with unbounded thread creation. *Theor. Comput. Sci.* 388(1-3), 227–242 (2007)
11. De Wulf, M., Doyen, L., Henzinger, T.A., Raskin, J.-F.: Antichains: A New Algorithm for Checking Universality of Finite Automata. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 17–30. Springer, Heidelberg (2006)
12. Doyen, L., Raskin, J.-F.: Antichain Algorithms for Finite Automata. In: Esparza, J., Majumdar, R. (eds.) TACAS 2010. LNCS, vol. 6015, pp. 2–22. Springer, Heidelberg (2010)
13. Eén, N., Sörensson, N.: An extensible SAT-solver. In: Giunchiglia, E., Tacchella, A. (eds.) SAT 2003. LNCS, vol. 2919, pp. 502–518. Springer, Heidelberg (2004)
14. Egly, U., Eiter, T., Tompits, H., Woltran, S.: Solving Advanced Reasoning Tasks Using Quantified Boolean Formulas. In: Proc. of IAAI, pp. 417–422. AAAI Press, Menlo Park (2000)
15. Giunchiglia, E., Narizzano, M., Tacchella, A.: Quantified Boolean Formulas satisfiability library (QBFLIB) (2001), www.qbflib.org
16. Giunchiglia, E., Narizzano, M., Tacchella, A.: QuBE++: An Efficient QBF Solver. In: Hu, A.J., Martin, A.K. (eds.) FMCAD 2004. LNCS, vol. 3312, pp. 201–213. Springer, Heidelberg (2004)
17. Giunchiglia, E., Narizzano, M., Tacchella, A.: Clause/Term Resolution and Learning in the Evaluation of Quantified Boolean Formulas. *J. Artif. Intell. Res.* 26, 371–416 (2006)
18. Jussila, T., Biere, A.: Compressing BMC Encodings with QBF. *Electron. Notes Theor. Comput. Sci.* 174, 45–56 (2007)
19. Kohli, R., Krishnamurti, R., Mirchandani, P.: The Minimum Satisfiability Problem. *SIAM J. Discrete Math.* 7(2), 275–283 (1994)
20. Kügel, A.: Improved Exact Solver for the Weighted Max-SAT problem. Accepted at the workshop Pragmatics of SAT. To appear in easychair electronic proceedings (2011)
21. Li, C.M., Manyà, F.: MaxSAT, Hard and Soft Constraints. In: Handbook of Satisfiability. Frontiers in Artificial Intelligence and Applications, vol. 185, pp. 613–631. IOS Press, Amsterdam (2009)
22. Lonsing, F., Biere, A.: DepQBF: A dependency-aware QBF solver (System Description). *Journal on Satisfiability, Boolean Modeling and Computation* 7, 71–76 (2010)
23. Lonsing, F., Biere, A.: Integrating Dependency Schemes in Search-Based QBF Solvers. In: Strichman, O., Szeider, S. (eds.) SAT 2010. LNCS, vol. 6175, pp. 158–171. Springer, Heidelberg (2010)
24. Narizzano, M., Peschiera, C., Pulina, L., Tacchella, A.: Evaluating and certifying QBFs: A comparison of state-of-the-art tools. *AI Commun.* 22, 191–210 (2009)
25. Pan, G., Vardi, M.Y.: Symbolic Decision Procedures for QBF. In: Wallace, M. (ed.) CP 2004. LNCS, vol. 3258, pp. 453–467. Springer, Heidelberg (2004)
26. Papadimitriou, C.H.: Computational complexity. Addison-Wesley Publishing Company, Reading (1994)
27. Peschiera, C., Pulina, L., Tacchella, A., Bubeck, U., Kullmann, O., Lynce, I.: The Seventh QBF Solvers Evaluation (QBFEVAL 2010). In: Strichman, O., Szeider, S. (eds.) SAT 2010. LNCS, vol. 6175, pp. 237–250. Springer, Heidelberg (2010)
28. Samulowitz, H., Davies, J., Bacchus, F.: Preprocessing QBF. In: Benhamou, F. (ed.) CP 2006. LNCS, vol. 4204, pp. 514–529. Springer, Heidelberg (2006)