

Bridging the Gap between Dual Propagation and CNF-based QBF Solving

Alexandra Goultiaeva
Department of Computer Science
University of Toronto
Canada

Martina Seidl, Armin Biere
Institute for Formal Models and Verification,
Johannes Kepler University
Linz, Austria

Abstract—Conjunctive Normal Form (CNF) representation as used by most modern Quantified Boolean Formula (QBF) solvers is simple and powerful when reasoning about conflicts, but is not efficient at dealing with solutions. To overcome this inefficiency a number of specialized non-CNF solvers were created. These solvers were shown to have great advantages. Unfortunately, non-CNF solvers cannot benefit from sophisticated CNF-based techniques developed over the years.

This paper demonstrates how the power of non-CNF structure can be harvested without the need for specialized solvers; in fact, it is easily incorporated into most existing CNF-based QBF solvers using a pre-existing mechanism of cube learning. We demonstrate this using a state-of-the-art QBF solver DepQBF, and experimentally show the effectiveness of our approach.

I. INTRODUCTION

Quantified Boolean Formula satisfiability checking (QBF) is the canonical PSPACE-complete problem. QBF provides a powerful framework for encoding many important verification and reasoning problems like model checking or scheduling [2]. QBF extends propositional logic with existential and universal quantifiers, which allow a more compact representation for problems with adversarial knowledge, incomplete information, or nondeterministic behavior, especially when the propositional form becomes too large to be handled efficiently [11].

Most state-of-the-art QBF solvers adopt techniques of propositional satisfiability checking (SAT), and have inherited *Conjunctive Normal Form* (CNF) representation. CNF encoding, simple but powerful, is widely used in SAT. CNF allows efficient reasoning over conflicts, but conversion to CNF involves a loss of structural information needed to efficiently reason over solutions. A SAT solver only needs a single solution, so in SAT the benefits of CNF seem to far outweigh the losses. However, a QBF solver encounters many solutions for different values of universal variables. The inability to efficiently reason over solutions has been identified as a major obstacle to efficient QBF solving [1].

Several approaches have been developed to take advantage of representations other than CNF [5], [13]. Some of them used the extra structure of non-CNF encodings to efficiently reason over solutions [8], [9], [14]. The core technique in these different approaches, though named differently, is actually the same. In this paper we will refer to it as “dual propagation”.

Dual propagation allows reasoning over conflicts and solutions to be done equally efficiently. It has been shown to yield great performance gains, often exponentially speeding up the search. However, it is not yet common in QBF solving. One reason is that there are too few non-CNF benchmarks currently available. Another reason is that, until now, dual propagation was only employed in specialized non-CNF solvers, which had to sacrifice all specialized techniques, tools, and data structures that were developed over the years of working with CNF.

While the first reason is inherent (efficient QBF reasoning requires more information than available in CNF), we will show that a specialized non-CNF solver is not needed. The mechanism for cube learning, which is already present in most modern QBF solvers, can, if initialized correctly, perform dual propagation exactly as done in non-CNF solvers. The resulting solver combines the best of both worlds: the efficiency of modern CNF-based reasoning, and the power of dual propagation. Our experiments confirm that dual propagation does indeed make a huge impact on performance; also, the resulting solver outperforms the previous non-CNF approaches, demonstrating the benefit of well-engineered CNF-based data structures.

While the idea of dual propagation is not new, this is, to our knowledge, the first time it is recognized that the mechanisms already present in most existing modern QBF solvers are sufficient for dual propagation, and that specialized solvers are not required. We define requirements on the input which retain soundness and are weaker than in previous work. We show that most state-of-the-art CNF-based preprocessing techniques can be soundly applied out-of-the-box. Unfortunately, such preprocessing seems to reduce the impact of dual propagation. An interesting avenue for future work is to extend preprocessing techniques to be not only sound, but also duality preserving.

II. QBF AT A GLANCE

In this paper, we only consider QBF in closed prenex form. A QBF has the structure $Q.\psi$, where Q is a *quantifier prefix*, and ψ is a propositional formula called the *matrix*. Q consists of *quantifier blocks*, which group together all consecutive variables with the same quantifier. The *quantifier level* of a variable x is one plus the number of preceding quantifier blocks. A *literal* is a variable or its negation. A *clause* is a disjunction of literals, a *cube* is a conjunction of literals. A formula is in CNF if it is a conjunction of clauses; it is in

Then $Q_1.\phi \equiv Q.\phi$ and $Q_2.\phi \equiv (\neg Q).\phi$ for any formula ϕ , since Q is simply $Q_1 (\neg Q_2)$ with additional variables inserted.

Definition 3. A DCNF representation of a QBF $Q.\phi$ is a pair of CNF formulas $Q_1.\phi_1$ and $Q_2.\phi_2$ such that:

- Q_1 and Q_2 are mergeable
- $Q_1.\phi_1 \equiv Q.\phi$
- If $Q.\phi$ is false, then for $Q' = \text{merge}(Q_1, Q_2)$, at least one counter-model of $Q'.\phi_1$ is also a model of $(\neg Q').\phi_2$.

To get a DCNF for a non-CNF formula $Q.\phi$, we separately convert $Q.\phi$ and $\neg(Q.\phi)$ to CNF using existing transformation tools. From the resulting DCNF $(Q_1.\phi_1, Q_2.\phi_2)$, we use $\text{merge}(Q_1, Q_2).\phi_1$ as the input formula, and add the negations of clauses of ϕ_2 as cubes into the cube database.

Example 4. In our example, $\neg\gamma$ can be represented in CNF as $\forall z\exists xyde.\{\neg d\}, \{\neg e\}, \{d, \neg x, \neg y\}, \{e, \neg z, \neg x\}$. Then, $Q_1 = \{\exists z\forall xy\exists ab\}$ and $Q_2 = \{\forall z\exists xyde\}$, and the merged prefix Q' is $\{\exists z\forall xyde\exists ab\}$. Unit propagation can simplify the formula to $\forall z\exists xy.\{\neg x, \neg y\}, \{\neg z, \neg x\}$. This simplification still satisfies all our desired properties. In that case, $Q' = \{\exists z\forall xy\exists ab\}$.

With the seeded database, as soon as z is selected as decision, the formula gets solved by propagation. Cubes set $\neg x$, which leads to $\neg a, \neg b$, which falsifies the clause $\{a, b\}$. A clause $\{x\}$ is learnt, which is universally reduced to $\{\}$.

We now sketch a proof that given a proper DCNF, a QDPLL solver must produce a correct answer. Soundness of the (stronger) approach mentioned above follows as a corollary.

Proof: A QDPLL solver only returns \perp after deriving an empty clause from $Q'.\phi_1$. Then $Q'.\phi_1$ is \perp , and so is $Q.\phi$.

Suppose that $Q.\phi$ is \perp . Then there is some counter-model m of $Q'.\phi_1$ which is also a model of $(\neg Q').\phi_2$. So, the cube database is consistent. When a cube is created based on a variable assignment, the assignment must satisfy all the clauses, and thus might not appear in m . So, these cubes will not violate consistency, and Q-resolution would never derive an empty cube. Then the solver would never return \top . ■

A. Implementation and Experiments

We have equipped a CNF-based solver DepQBF [10] with dual propagation, yielding Dual DepQBF. Dual DepQBF takes two CNF files as input, and adds the negations of the second CNF's clauses as cubes. These “original” cubes are exempt from deletion, just like original clauses. The merged prefix is computed on the fly, and conflicting variables are renamed. This allows the solver to be used with any converter which does not change the levels or names of original input variables.

The changes broke a few common assumptions: now a variable might occur in the cubes but not the clauses, and the last quantifier is no longer necessarily existential. As a quick work-around for some technical problems encountered, we have turned off pure literal detection, and switched to trivial dependency scheme. These changes do not seem to degrade performance on this dataset. In the experiments below, the original version of DepQBF has both pure literals and the standard dependency scheme; the dual version has neither.

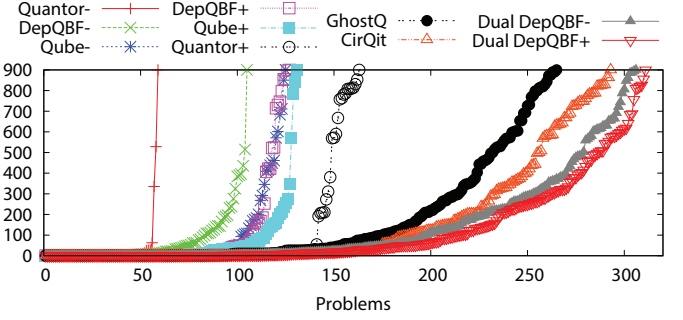


Fig. 2: Number of problems solved within a time (in seconds)




Fig. 3: Comparison of cubes and clauses learnt after negating the problem. DepQBF (left) and Dual DepQBF (right)

Fig. 2 shows the comparison of Dual DepQBF with state-of-the-art CNF and non-CNF solvers. We used nonprenex non-CNF benchmark set from QBFEval’10 [6]. Non-CNF solvers were given the formulas in qpro format. We obtained CNF and DCNF using Plaisted-Greenbaum transformation [12].


We evaluated the following solvers: Qube, a state-of-the-art DPLL-based CNF solver (version 7.2) [7]; quantor, an expansion-based CNF solver [3]; and non-CNF solvers GhostQ [9] and CirQit [8]. The experiments were obtained on a cluster with Intel Core 2 Duo Quad Q9550 2.8-GHz processors, 8-GB main memory, running Ubuntu Linux. We used the timeout of 900 seconds and memory limit of 7-GB.

For CNF solvers, we separately report solving time on positive and negative polarities of problems (suffixing the names with “+” and “-” respectively). Dual DepQBF took both representations. The entry “Dual DepQBF+” used the positive versions as the problem and negative CNF as the cube database, and vice versa for “Dual DepQBF-”.

Fig. 2 clearly shows that dual propagation is effective: note the substantial gap between the solvers that use dual propagation and those that do not. Also, the effectiveness of fast CNF-based reasoning is reflected by the fact that Dual DepQBF is more efficient than the current structural solvers.

The effect of Dual Propagation is visualized in Fig. 3. For the same problem, we compare the number of cubes learned while solving the problem to the number of clauses learned while solving its negation. If the problem was not solved, we take the number of cubes/clauses learnt within the timeout.

Theoretically, we would expect these numbers to be similar. After all, any clause in a problem is a cube of its negation. However, this is not the case for the original DepQBF. Because of the bias introduced by the formula representation, it learns substantially more cubes than clauses. Note that the plot is



(a) Valid DCNF from Ex. 4 (b) Not a valid DCNF, Ex. 5

Fig. 4: Truncated decision trees, (counter)models highlighted

logarithmic, so the gap is many orders of magnitude. For the dual version the plot is as one would expect. Our approach has removed the bias that weakened cube learning.

V. PREPROCESSING

Our approach allows both input formulas to be preprocessed, as long as it is ensured that the properties of Def. 3 are retained. The nontrivial property is the third one. To show that a technique preserves the third property, it suffices to show that it never destroys models or countermodels of the formula. We found that all the techniques employed by the preprocessor *bloqger* [4] preserve both models and countermodels. **Blocked clause elimination** might change a value on a node of a decision tree from \perp to \top . But, by properties of blocked clauses, it can be shown that the node must have an existential ancestor whose other child is \top . This means that the changed node must not be a part of any countermodel. Similar and dual arguments can be used to show that **pure** and **unit literal elimination** do not destroy any models or countermodels. **Equivalence replacement** is simply propositional transformations followed by blocked clause removal. **Variable expansion** can be seen as simply setting both universal node’s children to \perp if at least one is \perp . Obviously, if the node changed value, it could not have been a part of a model or a countermodel. So, all the techniques employed in *bloqger* can be used with dual propagation. However, that is not true in general. For example, it is unsound to apply blocked clause insertion.

Example 5. Fig 4a shows the decision tree for the DCNF from Ex. 4. Each node is marked with two values, the left for $Q'.\phi_1$ and the right for $(\neg Q').\phi_2$. The model for $Q'.\phi_1$, which is also a countermodel for $(\neg Q').\phi_2$, is highlighted. By blocked clause insertion and elimination, $(\neg Q').\phi_2$ can be transformed into $\phi'_2 = \{x, y\}, \{x, \neg y\}, \{\neg x, y\}$. The updated decision tree is shown in Fig. 4b, which highlights the new model for $\neg Q'.\phi'_2$. The pair is no longer a DCNF, and QDPLL might produce the wrong result \top by analyzing solution $\{a, b, z, x, y\}$, getting the dual clause $\{\neg z, \neg x, \neg y\}$; it can be resolved with existing dual clauses to obtain an empty cube.

Unfortunately, out-of-the-box preprocessing seems to harm the benefits of dual propagation. Dual DepQBF is only able to solve two more formulas than DepQBF on the preprocessed instances. We conjecture that separate preprocessing pushes the formulas further apart, so that they no longer define the exact dual problem. We compared the cube and clause sizes,

and found that the duality is once again broken: the transformations done by the preprocessor (especially the introduction of new variables) limits the duality of the resulting formulas.

Lastly, we note that the combination of DepQBF and bloqger is able to solve 469 of out 478 formulas (471 with Dual Propagation). Similar trend occurs on all other non-CNF formulas available to us. While problems that are hard for bloqger do exist, their non-CNF versions are not available.

VI. CONCLUSIONS

We have shown that dual propagation does not require a specialized solver, but can be combined with existing CNF-based datastructures and techniques. We verified its effectiveness, both at improving runtime and at removing the bias which forced CNF-based solvers to learn excessively many of cubes. We experimentally verified that the result noticeably outperforms existing solvers, both CNF and non-CNF.

Our approach decouples the encoding of the problem from dual propagation. We can use most existing CNF encoding methods out of the box, whereas specialized solvers such as IQTest are limited to a single built-in method. Our relaxed constraints guarantee soundness while applying sophisticated preprocessing techniques to the formulas, which has, to our knowledge, never been applied with dual propagation.

An interesting direction for further research is to develop duality-aware preprocessing tools, which would preserve duality and perhaps yield stronger preprocessing techniques. Another avenue to explore is the application to SAT, where a dual CNF can be used to produce partial models.

REFERENCES

- [1] Carlos Ansótegui, Carla P. Gomes, and Bart Selman. The achilles’ heel of QBF. In *Proc. AAAI*, pages 275–281, 2005.
- [2] M. Benedetti and H. Mangassarian. QBF-based formal verification: Experience and perspectives. *JSAT*, 5(1-4):133–191, 2008.
- [3] Armin Biere. Resolve and expand. In *Proc. SAT*, pages 238–246. Springer, 2004.
- [4] Armin Biere, Florian Lonsing, and Martina Seidl. Blocked clause elimination for QBF. In *CADE*, pages 101–115, 2011.
- [5] Uwe Egly, Martina Seidl, and Stefan Woltran. A solver for QBFs in negation normal form. *Constraints*, 14(1):38–79, 2009.
- [6] E. Giunchiglia, M. Narizzano, and A. Tacchella. Quantified Boolean Formulas satisfiability library (QBFLIB), 2001. www.qbflib.org.
- [7] Enrico Giunchiglia, Massimo Narizzano, and Armando Tacchella. QUBE: A system for deciding Quantified Boolean Formulas satisfiability. In *Proc. IJCAR*, pages 364–369, 2001.
- [8] Alexandra Goultiaeva and Fahiem Bacchus. Exploiting QBF duality on a circuit representation. In *AAAI*, 2010.
- [9] William Klieber, Samir Sapra, Sicun Gao, and Edmund M. Clarke. A non-prenex, non-clausal QBF solver with game-state learning. In *Proc. SAT*, pages 128–142, 2010.
- [10] Florian Lonsing and Armin Biere. DepQBF: A dependency-aware QBF solver. *JSAT*, 7(2-3):71–76, 2010.
- [11] Hratch Mangassarian, Andreas G. Veneris, Sean Safarpour, Marco Benedetti, and Duncan Smith. A performance-driven QBF-based iterative logic array representation with applications to verification, debug and test. In *Proc. ICCAD*, pages 240–245, 2007.
- [12] David A. Plaisted and Steven Greenbaum. A structure-preserving clause form translation. *Journal of Symbolic Computation*, 2(3):293–304, 1986.
- [13] Ashish Sabharwal, Carlos Ansótegui, Carla P. Gomes, Justin W. Hart, and Bart Selman. QBF modeling: Exploiting player symmetry for simplicity and efficiency. In *Proc. SAT*, pages 382–395, 2006.
- [14] Lintao Zhang. Solving QBF with combined conjunctive and disjunctive normal form. In *Proc. AAAI*, 2006.