Dependency Schemes in QBF Calculi: Semantics and Soundness*

Olaf Beyersdorff and Joshua Blinkhorn

School of Computing, University of Leeds, UK
{O.Beyersdorff,scjlb}@leeds.ac.uk

Abstract. We study the parametrisation of QBF resolution calculi by dependency schemes. One of the main problems in this area is to understand for which dependency schemes the resulting calculi are sound. Towards this end we propose a semantic framework for variable independence based on 'exhibition' by QBF models, and use it to express a property of dependency schemes called *full exhibition* that is known to be sufficient for soundness in Q-resolution. Introducing a generalised form of the long-distance resolution rule, we propose a complete parametrisation of classical long-distance Q-resolution, and show that full exhibition remains sufficient for soundness. We demonstrate that our approach applies to the current research frontiers by proving that the reflexive resolution path dependency scheme is fully exhibited.

1 Introduction

The excellent success of SAT solvers in the realm of propositional Boolean formulae has motivated much interest in the corresponding search problem for quantified Boolean formulae (QBF). The greater expressiveness of QBF, afforded by its PSPACE-completeness [23], presents novel challenges in solving, and the array of emerging techniques is motivating a wealth of research in the closely-related field of proof complexity [3, 5–9, 11–14].

There is a natural correspondence between QBF practice and proof theory; when a solver concludes the falsity of an instance, the trace can be interpreted as a formal refutation. Understanding the refutational proof system that underpins a particular solving method, and thereby accounts for its correctness, motivates the proof-theoretic study of specific calculi. Recent work has led to a complete understanding of the relative strength of resolution-based QBF systems [3, 7], including Q-resolution (Q-Res) [15], universal Q-resolution (QU-Res) [13], and long-distance Q-resolution (LD-Q-Res) [1].

Implemented in the state-of-the-art solver DepQBF [16,17], one of the recent and exciting developments in QBF solving has seen the introduction of *dependency schemes*: algorithms that gather information on variable independence by prior appeal to the syntactic form of an instance. The quantifier prefix of a QBF (in prenex normal form) imposes a total order on the variables; due to the

* This is an extended abstract of the paper published in the proceedings of CP 2016 [4]

nesting of quantifier scopes, the value of a Boolean variable z can be dependent upon the variables to its left in the prefix. Naturally, this entails some restrictions on solving methods, and on the rules of the related formal systems. In general, however, z does not necessarily depend on all of the variables to its left. A dependency scheme attempts to replace the linear order of the prefix with a partial order that more accurately reflects the dependency structure of the formula, by identifiying *variable independence*. This approach allows some sets of instances to be solved more effeciently, despite the compuational overhead of computing the dependency scheme [16].

Independence itself is presented as a semantic concept [16, 18]. The truth of a QBF Φ is witnessed by a Skolem-function model, a set of Boolean functions $\{f_x\}$ that produce a propositional tautology when substituted for the existential variables. The arguments to f_x are the universal variables U_x left of x in the quantifier prefix, but it may occur that some circuit computes f_x without using $u \in U_x$ as an input. In this case we say that x *is independent of* u – and a dual notion for false QBFs provides for independence of universals on existentials – even though the Skolem-function model is in general not unique.

This lack of uniqueness has consequences for soundness in QBF calculi. The impact of a dependency scheme in the proof system is to allow some logical steps which previously were prohibited; specifically, the $\forall$-reduction rule of Q-Res receives greater reign. This motivated the proposal of Q($\mathcal{D}$)-Res by Slivovsky and Szeider [22], a parametrization of the classical calculus by dependency schemes. Some schemes which were previously put forward in the literature, such as the triangle [19] and resolution path [24] dependency schemes, have proved too aggressive for soundness in Q($\mathcal{D}$)-Res, admitting refutations of true QBFs. The reflexive resolution path dependency scheme [22] is currently the strongest known scheme for which Q($\mathcal{D}$)-Res is sound, a result which was proved by means of a difficult transformation of a Q($\mathcal{D}$)-Res refutation into a Q-Res refutation [22].

What is currently absent in the literature is a deeper understanding of soundness based on classification of dependency schemes; moreover, the lack of general methods may frustrate future developments. It is natural to propose the parametrization by dependency schemes of stronger QBF calculi, of the other CDCL-based QBF resolution systems and QBF Frege [5], whereupon methods for proving soundness based on *properties* of dependency schemes will carry over. In this paper we demonstrate that semantic notions of independence are indeed equipped for this; our contributions are summarized below.

1. New QBF calculi parametrized by dependency schemes. We extend the parametrisation by dependency schemes to all the CDCL-based resolution calculi for QBF: with the new long-distance calculus LD-Q($\mathcal{D}$)-Res, with universal resolution QU($\mathcal{D}$)-Res, and with their combination LQU($\mathcal{D}$)-Res. Our new long-distance calculus presents the greatest challenge. Of the two inference rules employed classically, parametrization of $\forall$-reduction can be lifted straight from Q($\mathcal{D}$)-Res; here we investigate the additional effects of parametrizing the long-distance resolution rule as well, by relaxing the conditions under which so-called

'merged literals' can be introduced. Progressing from Q-resolution, we demonstrate that variable independence and merging have a more subtle interaction; in LD-Q($\mathcal{D}$)-Res, we must supplant merged literals with *annotated literals*, which record existential pivots to prevent unsound $\forall$-reduction steps.

2. A semantic framework for independence and soundness. We unify some existing approaches in the literature towards a more fruitful understanding of the interplay between Q-resolution and dependency schemes. Building on the work of Samer [18] and Lonsing [16] we propose a semantic framework for variable independence. Central to the framework is a property of dependency schemes called *full exhibition*, which was shown to be sufficient for soundness in Q($\mathcal{D}$)-Res by Slivovsky [21]. We further the potential of this approach to show that full exhibition is sufficient for soundness in all the dependency calculi we introduce. To that end, we handle the semantic obstacles of long-distance resolution by incorporating techniques from strategy extraction due to Balabanov et al. [2].

3. Demonstrating full exhibition. We conclude by proving Slivovsky's conjecture [21, p. 37] that the reflexive resolution path dependency scheme $\mathcal{D}^{\mathrm{rrs}}$ is fully exhibited. Currently, $\mathcal{D}^{\mathrm{rrs}}$ is arguably the most important dependency scheme, capable of revealing more cases of independence than any other tractable scheme known to be sound for Q($\mathcal{D}$)-Res. As such, we show that everything currently known about soundness in this setting can be explained by full exhibition. On the technical level, the result is obtained by an algorithmic transformation of an arbitrary model for a true QBF Φ into a model that exhibits all the required independencies. We therefore reveal the possibility for QBF solving to implement long-distance techniques fully parametrized by $\mathcal{D}^{\mathrm{rrs}}$, or any other fully exhibited scheme.

2 Preliminaries

Quantified Boolean Formulas. A *Quantified Boolean Formula* (QBF) Φ over a set $V = \{z_1, \ldots, z_n\}$ of n variables is a formula in quantified Boolean logic with variables ranging over $\{0, 1\}$. We consider only formulas in *prenex conjunctive normal form* (PCNF), denoted $\Phi = \mathcal{Q}\,.\,\phi$, in which all variables are quantified either existentially or universally in the *quantifier prefix* $\mathcal{Q} = \mathcal{Q}_1 z_1 \cdots \mathcal{Q}_n z_n$, $\mathcal{Q}_i \in \{\exists, \forall\}$ for $i \in [n]$, and ϕ is a propositional conjunctive normal form (CNF) formula called the *matrix*. A CNF matrix is a conjunction of clauses, each clause is a disjunction of literals, and a literal is a variable or its negation. Whenever convenient, we refer to a clause as a set of literals and to a matrix as a set of clauses. We typically write x for existential variables, u and v for universals, and z for either. We denote the sets of existentially and universally quantified variables of Φ by $V_\exists = \{z_i \in V \mid \mathcal{Q}_i = \exists\}$ and $V_\forall = \{z_i \in V \mid \mathcal{Q}_i = \forall\}$ respectively. The prefix $\mathcal{Q}$ imposes a linear ordering $<_\Phi$ on the variables of Φ, such that $z_i <_\Phi z_j$ holds whenever $i < j$, in which case we say that z_j *is right of*

z_i, or that z_i *is left of* z_j. The sets of variables right and left of z are denoted $R_\Phi(z) = \{z' \in V \mid z <_\Phi z'\}$ and $L_\Phi(z) = \{z' \in V \mid z' <_\Phi z\}$.

Assignment Trees and Models. Assignment trees for PCNF were first introduced in [20]. We represent an assignment tree formally as a set of paths. Let Φ be a PCNF over variables $V = \{z_1, \ldots, z_n\}$ and let $V_\forall = \{u_1, \ldots, u_k\}$. A path is a set of literals $P = \{l_1, \ldots, l_n\}$ with $\mathrm{var}(l_i) = z_i$ for all $i \in [n]$, and we write $P[z_i] = l_i$. A set of paths T is well-formed for Φ iff (1) for all $u \in V_\forall$ and for all $P, Q \in T$, if $P[v] = Q[v]$ for all $v \in L_\Phi(u) \cap V_\forall$, then $P[x] = Q[x]$ for each $x \in L_\Phi(u) \cap V_\exists$, and (2) there is a unique path $P \in T$ with $U \subseteq P$ for each set of literals $U = \{l_1, \ldots, l_k\}$ such that $\mathrm{var}(l_i) = u_i$ for $i \in [k]$. A set of paths that is well-formed for Φ is an *assignment tree* for Φ. We also use P to denote the total assignment $P : V \to \{\top, \bot\}$ given by $P(z_i) = \bot$ if $l_i = \neg z_i$ and $P(z_i) = \top$ if $l_i = z_i$, and extend this notation to literals with $P(\neg z_i) = \neg P(z_i)$, where $\top = \neg\bot$ and vice versa. An assignment tree for Φ is a *model* for Φ, typically denoted M, iff $P(C) = \top$ for all paths $P \in T$ and all clauses $C \in \phi$, where $P(C) = \top$ iff $P(l) = \top$ for some $l \in C$. A PCNF which has a model is *true*, otherwise it is *false*. An assignment tree is depicted as a tree with root r.

Dependency Schemes. The trivial dependency scheme $\mathcal{D}^{\mathrm{trv}}$ is a mapping which associates each PCNF $\Phi = \mathcal{Q}_1 z_1 \cdots \mathcal{Q}_n z_n \,.\, \phi$ over variables V to the trivial dependency relation $\mathcal{D}^{\mathrm{trv}}_\Phi = \{(z_i, z_j) \mid i < j \text{ and } \mathcal{Q}_i \neq Q_j\}$. A *proto-dependency scheme*[1] $\mathcal{D}$ is a function that maps each PCNF Φ to a binary relation $\mathcal{D}_\Phi \subseteq \mathcal{D}^{\mathrm{trv}}_\Phi$ called the *dependency relation*. If $(z_i, z_j) \in \mathcal{D}_\Phi$, then (z_i, z_j) is a $\mathcal{D}$-dependency and z_j is a $\mathcal{D}$-dependent of z_i, otherwise z_j is $\mathcal{D}$-independent of z_i. A proto-dependency scheme $\mathcal{D}'$ is said to be *at least as general* as another $\mathcal{D}$ if $\mathcal{D}'_\Phi \subseteq \mathcal{D}_\Phi$ for all PCNFs Φ, and is *strictly more general* if the inclusion is strict for some formula. For a PCNF Φ over variables V and $u \in V_\forall$, we write $\bar{\mathcal{D}}_\Phi(u) = \{(u, x) \mid x \in V_\exists \text{ and } (u, x) \notin \mathcal{D}_\Phi\}$.

QBF Resolution Calculi. We give a brief overview of four resolution-based CDCL QBF calculi – see [7] for a more detailed survey. A refutational QBF calculus is *sound* iff the empty clause cannot be derived from any true formula.

Q-resolution (Q-Res) introduced in [15] is the standard refutational calculus for PCNF. In addition to resolution over existential pivots with non-tautologous resolvents, the calculus has a universal reduction rule which allows a clause C to be derived from $C \cup \{u\}$, where u is a universal literal and all existential literals in C are left of u. *QU-resolution* (QU-Res) [13] is a natural extension of Q-Res that allows universal resolution pivots.

Long-distance resolution, which was introduced in [25] and formalised as the calculus LD-Q-Res [1], allows tautologous resolvents under certain conditions,

[1] The term 'dependency scheme' was first introduced to denote a subset of proto-dependency schemes with a more technical definition [19]; for consistency with the literature we will use 'proto-dependency scheme' in technical portions of this paper.

using the special merged literal u^* to represent the tautology $\{u, \neg u\}$. The resulting system is exponentially stronger than Q-Res [12]. Finally, the calculus LQU-Res [3] combines naturally the features of QU-Res and LD-Q-Res, allowing merged literals and resolution over universal pivots.

3 Our Contributions

In this section, we give a brief survey of the contributions contained in the full-length version of the paper [4], to which we refer the reader for proofs and further detailed commentry.

3.1 A Semantic Framework and New QBF Calculi

We reformulate the definition of independence in terms of assignment trees from [16, 18]; we feel our notation is better suited to the aims of the current work. We introduce the new idea of *complementary paths* in an assignment tree, and define a property of dependency schemes called *full exhibition*.

Definition 1 (Complementary path). *Let Φ be a QBF over variables V, let U be a non-tautologous set of literals such that* $\mathrm{var}(U) = V_\forall$*, let T be an assignment tree for Φ and let $P \in T$ be the unique path such that $U \subseteq P$. Then, for any $u \in V_\forall$, $P_u \in T$ is the unique path such that $U' \subseteq P_u$, where $U' = (U \setminus \{l\}) \cup \{\neg l\}$, $l \in U$ and* $\mathrm{var}(l) = u$.

Definition 2 (Independence of existentials from universals [16,18]). *Let Φ be a true QBF over variables V and let $u \in V_\forall$, $x \in V_\exists$. We say that x is independent of u in Φ if there exists a model M for Φ in which $P(x) = P_u(x)$ for all paths $P \in M$. For such a model M we write $M \prec (u, x)$, and we say that M* exhibits the independence *of x from u in Φ.*

Definition 3 (Fully exhibited dependency scheme). *Let $\mathcal{D}$ be a proto-dependency scheme. We say that $\mathcal{D}$ is* fully exhibited *iff for each true PCNF Φ there is a model M for Φ such that $M \prec (u, x)$ for each pair $(u, x) \notin D_\Phi$, with $u \in V_\forall$ and $x \in V_\exists$.*

In [21], it was proved that Q($\mathcal{D}$)-Res is sound for fully exhibited[2] $\mathcal{D}$, and this was combined with the fact that the standard dependency scheme $\mathcal{D}^{\mathrm{std}}$ is fully exhibited (attributed to [10]). We show that this approach scales up to the dependency versions of stronger QBF calculi. To do this, we introduce the new long-distance calculi LD-Q($\mathcal{D}$)-Res and LQU($\mathcal{D}$)-Res (Fig. 1), the respective dependency versions of LD-Q-Res and LQU-Res. Parametrising long-distance resolution calls for the introduction of *annotations* that prevent unsound $\forall$-reduction steps.

[2] Full exhibition is treated equivalently, as a property of models.

$$\frac{}{C}\ (\text{Axiom}) \qquad C \text{ is a clause in the matrix of } \Phi.$$

$$\frac{D \cup \{u^X\}}{D}\ (\forall\text{-Red}) \qquad \text{Variable } u \text{ is universal. If } l \in D \text{ and } \mathrm{var}(l) = z, \text{ then } (u, z) \notin \mathcal{D}_\Phi, \text{ and if } l = z^{X'} \text{ then } (u, x) \notin \mathcal{D}_\Phi \text{ for all } x \in X'. \text{ If } X = \emptyset \text{ then literal } u^X \text{ is either } u \text{ or } \neg u.$$

$$\frac{C_1 \cup U_1 \cup \{x\} \qquad C_2 \cup U_2 \cup \{\neg x\}}{C_1 \cup C_2 \cup U}\ (\text{Res})$$

If for $l_1 \in C_1, l_2 \in C_2, \mathrm{var}(l_1) = \mathrm{var}(l_2)$, then $l_1 = l_2$ is not annotated. $\mathrm{var}(U_1) = \mathrm{var}(U_2) \subseteq V_\forall$, and $(x, u) \notin \mathcal{D}_\Phi$ for each $u \in \mathrm{var}(U_1)$. If for $u_1 \in U_1, u_2 \in U_2, \mathrm{var}(u_1) = \mathrm{var}(u_2) = u$, then $u_1 = \neg u_2$, or at least one of u_1, u_2 is annotated. U is defined as $\{u^X \mid u \in \mathrm{var}(U_1)\}$, where X is the union of $\{x\}$ with any annotations on u in $U_1 \cup U_2$. In LD-Q($\mathcal{D}$)-Res $\mathrm{var}(x)$ is existential. In LQU($\mathcal{D}$)-Res, $\mathrm{var}(x)$ is existential or universal.

Fig. 1. The rules of LD-Q($\mathcal{D}$)-Res and LQU($\mathcal{D}$)-Res

3.2 Results

We first prove that full exhibition is a sufficient condition for soundness in the new long-distance QBF calculi.

Theorem 4. *Let $\mathcal{D}$ be a fully exhibited proto-dependency scheme. Then* LD-Q($\mathcal{D}$)-Res *is sound.*

Since the (omitted) proof of Theorem 4 makes no use of the fact that the pivot is existential, it also shows the soundness of LQU($\mathcal{D}$)-Res, the 'dependency version' of LQU-Res, for any fully exhibited $\mathcal{D}$.

Theorem 5. *Let $\mathcal{D}$ be a fully exhibited proto-dependency scheme. Then* LQU($\mathcal{D}$)-Res *is sound.*

Also, since LQU($\mathcal{D}$)-Res clearly simulates QU($\mathcal{D}$)-Res simply by disallowing long-distance resolution steps, we obtain the same result for QU($\mathcal{D}$)-Res.

Theorem 6. *Let $\mathcal{D}$ be a fully exhibited proto-dependency scheme. Then* QU($\mathcal{D}$)-Res *is sound.*

Theorems 4, 5 and 6 together constitute the generalisation to all the CDCL QBF calculi of Slivovsky's result [21] that Q($\mathcal{D}$)-Res is sound for fully exhibited $\mathcal{D}$. Whereas full exhibition is a sufficient condition for each calculus, it is not a necessary condition for any of them.

Proposition 7. *There exists a proto-dependency scheme $\mathcal{D}$ that is not fully-exhibited for which* LQU($\mathcal{D}$)-Res *is sound.*

For proof of concept, we demonstrate that the reflexive resolution path dependency scheme $\mathcal{D}^{\mathrm{rrs}}$ [22] is fully exhibited, thereby proving the conjecture of Slivovsky [21, p.37]. This result provides a better understanding of soundness in Q-resolution with dependency schemes; since $\mathcal{D}^{\mathrm{rrs}}$ is the most general scheme known to be sound in $\mathsf{Q}(\mathcal{D})$-Res, what is already known about soundness for that calculus can subsequently be explained entirely by full exhibition.

Theorem 8. *$\mathcal{D}^{rrs}$ is fully exhibited.*

Our concluding result now follows immediately from Theorems 4, 5 and 6.

Corollary 9. $\mathsf{QU}(\mathcal{D}^{\mathrm{rrs}})$-$\mathsf{Res}$, LD-$\mathsf{Q}(\mathcal{D}^{\mathrm{rrs}})$-$\mathsf{Res}$ *and* $\mathsf{LQU}(\mathcal{D}^{\mathrm{rrs}})$-$\mathsf{Res}$ *are sound proof systems.*

4 Conclusions

As we have shown, the parametrization by dependency schemes can be extended to all four CDCL QBF calculi, and the property of full exhibition – which is possessed by the reflexive resolution path dependency scheme – is sufficient for soundness in each case. Showing by counterexample that full-exhibition is not a necessary condition, our work leads naturally to the open problem of finding a characterization for soundness in this setting.

Acknowledgments. This research was supported by grant no. 48138 from the John Templeton Foundation and EPSRC grant EP/L024233/1.

References

1. Balabanov, V., Jiang, J.R.: Unified QBF certification and its applications. Formal Methods in System Design 41(1), 45–65 (2012)
2. Balabanov, V., Jiang, J.R., Janota, M., Widl, M.: Efficient extraction of QBF (counter)models from long-distance resolution proofs. In: Conference on Artificial Intelligence (AAAI). pp. 3694–3701 (2015)
3. Balabanov, V., Widl, M., Jiang, J.R.: QBF resolution systems and their proof complexities. In: International Conference on Theory and Applications of Satisfiability Testing (SAT). pp. 154–169 (2014)
4. Beyersdorff, O., Blinkhorn, J.: Dependency schemes in QBF calculi: Semantics and soundness. In: Principles and Practice of Constraint Programming (CP). pp. 96–112 (2016)
5. Beyersdorff, O., Bonacina, I., Chew, L.: Lower bounds: From circuits to QBF proof systems. In: ACM Conference on Innovations in Theoretical Computer Science (ITCS). pp. 249–260 (2016)
6. Beyersdorff, O., Chew, L., Janota, M.: On unification of QBF resolution-based calculi. In: International Symposium on Mathematical Foundations of Computer Science (MFCS). pp. 81–93 (2014)

7. Beyersdorff, O., Chew, L., Janota, M.: Proof complexity of resolution-based QBF calculi. In: International Symposium on Theoretical Aspects of Computer Science (STACS). Leibniz International Proceedings in Informatics (LIPIcs), vol. 30, pp. 76–89 (2015)
8. Beyersdorff, O., Chew, L., Mahajan, M., Shukla, A.: Feasible interpolation for QBF resolution calculi. In: International Colloquium on Automata, Languages, and Programming (ICALP). pp. 180–192 (2015)
9. Beyersdorff, O., Chew, L., Mahajan, M., Shukla, A.: Are short proofs narrow? QBF resolution is not simple. In: Symposium on Theoretical Aspects of Computer Science (STACS). pp. 15:1–15:14 (2016)
10. Bubeck, U.: Model-based transformations for quantified boolean formulas (2010)
11. Egly, U.: On sequent systems and resolution for QBFs. In: International Conference on Theory and Applications of Satisfiability Testing (SAT). pp. 100–113 (2012)
12. Egly, U., Lonsing, F., Widl, M.: Long-distance resolution: Proof generation and strategy extraction in search-based QBF solving. In: International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR). pp. 291–308 (2013)
13. Gelder, A.V.: Contributions to the theory of practical quantified boolean formula solving. In: International Conference on Principles and Practice of Constraint Programming (CP). pp. 647–663 (2012)
14. Janota, M., Marques-Silva, J.: Expansion-based QBF solving versus Q-resolution. Theorical Compututer Science 577, 25–42 (2015)
15. Kleine Büning, H., Karpinski, M., Flögel, A.: Resolution for quantified boolean formulas. Information and Computation 117(1), 12–18 (1995)
16. Lonsing, F.: Dependency Schemes and Search-Based QBF Solving: Theory and Practice. Ph.D. thesis, Johannes Kepler University (2012)
17. Lonsing, F., Egly, U.: Incrementally computing minimal unsatisfiable cores of QBFs via a clause group solver API. In: International Conference on Theory and Applications of Satisfiability Testing (SAT). pp. 191–198 (2015)
18. Samer, M.: Variable dependencies of quantified csps. In: International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR). pp. 512–527 (2008)
19. Samer, M., Szeider, S.: Backdoor sets of quantified boolean formulas. Journal of Automated Reasoning 42(1), 77–97 (2009)
20. Samulowitz, H., Bacchus, F.: Using SAT in QBF. In: International Conference on Principles and Practice of Constraint Programming (CP). pp. 578–592 (2005)
21. Slivovsky, F.: Structure in #SAT and QBF. Ph.D. thesis, Vienna University of Technology (2015)
22. Slivovsky, F., Szeider, S.: Soundness of Q-resolution with dependency schemes. TCS 612, 83–101 (2016)
23. Stockmeyer, L.J., Meyer, A.R.: Word problems requiring exponential time: Preliminary report. In: Annual Symposium on Theory of Computing. pp. 1–9. ACM (1973)
24. Van Gelder, A.: Variable independence and resolution paths for quantified boolean formulas. In: International Conference on Principles and Practice of Constraint Programming (CP). pp. 789–803. Springer (2011)
25. Zhang, L., Malik, S.: Conflict driven learning in a quantified boolean satisfiability solver. In: International Conference on Computer-aided Design (ICCAD). pp. 442–449 (2002)