Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 36 (2014)

On Unification of QBF Resolution-Based Calculi

Olaf Beyersdorff[1], Leroy Chew[1], and Mikoláš Janota[2]

[1]School of Computing, University of Leeds, United Kingdom
[2]INESC-ID, Lisbon, Portugal

Abstract. Several calculi for quantified Boolean formulas (QBFs) exist, but relations between them are not yet fully understood. This paper defines a novel calculus, which is resolution-based and enables unification of the principal existing resolution-based QBF calculi, namely Q-resolution, long-distance Q-resolution and the expansion-based calculus ∀Exp+Res. All these calculi play an important role in QBF solving. This paper shows simulation results for the new calculus and some of its variants. Further, we demonstrate how to obtain winning strategies for the universal player from proofs in the calculus. We believe that this new proof system provides an underpinning necessary for formal analysis of modern QBF solvers.

1 Introduction

Traditionally, classifying a problem as NP hard was ultimately understood as evidence for its infeasibility. Sharply contrasting this view, we have today fast algorithms for many important computational tasks with underlying NP-hard problems. One particularly compelling example of tremendous success is the area of SAT solving [26] where fast algorithms are being developed and tested for the classical NP-complete problem of satisfiability of propositional formulas (SAT). Modern SAT-solvers routinely solve industrial instances with even millions of variables. However, from a theoretical perspective, this success of SAT solvers is not well understood. The main theoretical approach to it comes via proof complexity. In particular, resolution and its subsystems have been very successfully analysed in terms of proof complexity and sharp bounds are known on the size and space for many important principles in resolution (cf. [31,7]). This is very important information as the main algorithmic approaches to SAT such as DPLL and CDCL are known to correspond to (tree-like) resolution [3,8,15,27], and therefore bounds on size and space of proofs directly translate into bounds on running time and memory consumption of SAT solvers.

In the last decade, there has been ever increasing interest to transfer the successful approach of SAT-solving to the more expressive case of *quantified propositional formulas (QBF)*. Due to its PSPACE completeness, QBF is far more expressive than SAT and thus applies to further fields such as formal verification or planning [28,5]. As for SAT, proof complexity provides the main theoretical approach towards understanding the performance and limitations of QBF-solving. However, compared to proof complexity of classical propositional logic, QBF proof complexity is at a much earlier stage and also poses additional challenges. Currently, a handful of systems exist, and they correspond to different approaches in QBF-solving. In particular, Kleine Büning et al. [21] define a resolution-like calculus called *Q-resolution*. There are several extensions of Q-resolution; notably *long-distance Q-resolution* [2], which has been shown to be more powerful than plain Q-resolution [11]. Q-resolution and its extensions are important as they model QBF solving based on CDCL [13]. Apart from CDCL, another main approach to QBF-solving is through expansion of quantifiers [6,4,17]. Recently, a proof system ∀*Exp+Res* was introduced with the motivation to trace expansion-based QBF solvers [16].

ISSN 1433-8092

∀Exp+Res also uses resolution, but is conceptually very different from Q-resolution. The precise relation of ∀Exp+Res to Q-resolution is currently open (cf. [18]), but we conjecture that the two systems are incomparable as it has been shown that expansion-based solving can exponentially outperform DPLL-based solving.

In general, it is fair to say that relations between the different types of QBF systems mentioned above are currently not well understood. The objective of the present paper is to unify these approaches. Towards this aim we define a calculus that is able to capture the existing QBF resolution-based calculi and yet remains amenable to machine manipulation. Our main contributions are as follows. (1) We introduce two novel calculi IR-calc and IRM-calc, which are shown to be sound and complete for QBF. (2) IR-calc p-simulates Q-resolution and ∀Exp+Res, i.e., proofs in either Q-resolution or ∀Exp+Res can be efficiently translated into IR-calc. (3) The variant IRM-calc p-simulates long-distance Q-resolution. (4) We show how to extract winning strategies for the universal player from proofs in IR-calc and IRM-calc. Indeed, unified certification of QBF solvers or certification of solvers combining expansion and DPLL is of immense practical importance [14,2,11] and presents one of the main motivations for our research. To the best of our knowledge, constructions of strategies from expansion-based solvers were not known prior to this paper.

The rest of the paper is structured as follows. Section 2 introduces concepts and notation used throughout the paper. Section 3 introduces novel calculi and Section 4 shows how winning strategies for the universal player are constructed; this is used as an argument for soundness. Section 5 shows p-simulation results for the new calculi. Finally, Section 6 concludes the paper and points to directions of future work.

2 Preliminaries

A *literal* is a Boolean variable or its negation; we say that the literal x is *complementary* to the literal $\neg x$ and vice versa. If l is a literal, $\neg l$ denotes the complementary literal, i.e. $\neg\neg x = x$. A *clause* is a disjunction of zero or more literals. The empty clause is denoted by $\bot$, which is semantically equivalent to false. A formula in *conjunctive normal form* (CNF) is a conjunction of clauses. Whenever convenient, a clause is treated as a set of literals and a CNF formula as a set of sets of literals. For a literal $l = x$ or $l = \neg x$, we write $\mathsf{var}(l)$ for x and extend this notation to $\mathsf{var}(C)$ for a clause C and $\mathsf{var}(\psi)$ for a CNF ψ.

A *proof system* (Cook, Reckhow [9]) for a language L over alphabet Γ is a polynomial-time computable partial function $f : \Gamma^* \rightharpoonup \Gamma^*$ with $rng(f) = L$. An *f-proof* of string y is a string x such that $f(x) = y$. In the systems that we consider here, proofs are sequences of clauses; a *refutation* is a proof deriving $\top$.

Quantified Boolean Formulas (QBFs) [20] extend propositional logic with quantifiers with the standard semantics that $\forall x.\Psi$ is satisfied by the same truth assignments as $\Psi[0/x]\wedge\Psi[1/x]$ and $\exists x.\Psi$ as $\Psi[0/x]\vee\Psi[1/x]$. Unless specified otherwise, we assume that QBFs are in *closed prenex* form with a CNF *matrix*, i.e., we consider the form $\mathcal{Q}_1X_1\ldots\mathcal{Q}_kX_k.\ \phi$, where X_i are pairwise disjoint sets of variables; $\mathcal{Q}_i \in \{\exists,\forall\}$ and $\mathcal{Q}_i \neq \mathcal{Q}_{i+1}$. The formula ϕ is in CNF and is defined only on variables $X_1 \cup \ldots \cup X_k$. The propositional part ϕ of a QBF is called the *matrix* and the rest the *prefix*. If a variable x is in the set X_i, we say that x is at *level* i and write $\mathsf{lv}(x) = i$; we write $\mathsf{lv}(l)$ for $\mathsf{lv}(\mathsf{var}(l))$. A closed QBF is *false* (resp. *true*), iff it is semantically equivalent to the constant 0 (resp. 1).

Often it is useful to think of a QBF $\mathcal{Q}_1X_1\ldots\mathcal{Q}_kX_k.\ \phi$ as a *game* between the *universal* and the *existential player*. In the i-th step of the game, the player $\mathcal{Q}_i$ assigns values to the

variables X_i. The existential player wins the game iff the matrix ϕ evaluates to 1 under the assignment constructed in the game. The universal player wins iff the matrix ϕ evaluates to 0. A QBF is false iff there exists a *winning strategy* for the universal player, i.e. if the universal player can win any possible game [1, Sec. 4.2.2].

2.1 Resolution-based Calculi for QBF

This section gives a brief overview of the main existing resolution-based calculi for QBF. *Q-resolution (Q-Res)*, by Kleine Büning et al. [21], is a resolution-like calculus that operates on QBFs in prenex form where the matrix is a CNF. The rules are given in Figure 1.

$$\frac{}{C}\ (\text{Axiom}) \qquad \frac{C_1 \cup \{x\} \qquad C_2 \cup \{\neg x\}}{C_1 \cup C_2}\ (\text{Res})$$

C is a clause in the matrix. Variable x is existential. If $z \in C_1$, then $\neg z \notin C_2$.

$$\frac{C \cup \{u\}}{C}\ (\forall\text{-Red})$$

Variable u is universal. If $x \in C$ is existential, then $\mathsf{lv}(x) < \mathsf{lv}(u)$.

Fig. 1. The rules of Q-Res [21]

Long-distance resolution (LD-Q-Res) appears originally in the work of Zhang and Malik [35] and was formalized into a calculus by Balabanov and Jiang [2]. It merges complementary literals of a universal variable u into the special literal u^*. These special literals prohibit certain resolution steps. In particular, different literals of a universal variable u may be merged only if $\mathsf{lv}(x) < \mathsf{lv}(u)$, where x is the resolution variable. The rules are given in Figure 2. Note that the rules do not prohibit resolving $w^* \vee x \vee C_1$ and $u^* \vee \neg x \vee C_2$ with $\mathsf{lv}(w) \leq \mathsf{lv}(u) < \mathsf{lv}(x)$ as long as $w \neq u$.

$$\frac{}{C}\ (\text{Axiom}) \qquad \frac{D \cup \{u\}}{D}\ (\forall\text{-Red}) \qquad \frac{D \cup \{u^*\}}{D}\ (\forall\text{-Red}^*)$$

C is a clause in the matrix. Literal u is universal and $\mathsf{lv}(u) \geq \mathsf{lv}(l)$ for all $l \in D$.

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

Variable x is existential. If for $l_1 \in C_1, l_2 \in C_2$, $\mathsf{var}(l_1) = \mathsf{var}(l_2) = z$ then $l_1 = l_2 \neq z^*$. U_1, U_2 contain only universal literals with $\mathsf{var}(U_1) = \mathsf{var}(U_2)$. For each $u \in \mathsf{var}(U_1)$ we require $\mathsf{lv}(x) < \mathsf{lv}(u)$. If for $w_1 \in U_1, w_2 \in U_2$, $\mathsf{var}(w_1) = \mathsf{var}(w_2) = u$ then $w_1 = \neg w_2$, $w_1 = u^*$ or $w_2 = u^*$. U is defined as $\{u^* \mid u \in \mathsf{var}(U_1)\}$.

Fig. 2. The rules of LD-Q-Res [2]

A different calculus ∀Exp+Res based on expansions was introduced in [19]. In Figure 3 we present an adapted version of this calculus so that it is congruent with the other resolution-based calculi (semantically it is the same as in [19]). The ∀Exp+Res calculus operates on

$$\frac{}{\{l^{\tau_l} \mid l \in C, l \text{ is existential}\} \cup \{\tau(l) \mid l \in C, l \text{ is universal}\}} \text{ (Axiom)}$$

C is a clause from the matrix and τ is an assignment to all universal variables.
τ_l are partial assignments obtained by restricting τ to variables u with $\mathsf{lv}(u) < \mathsf{lv}(l)$.

$$\frac{C_1 \cup \{x^\tau\} \qquad C_2 \cup \{\neg x^\tau\}}{C_1 \cup C_2} \text{ (Res)}$$

Fig. 3. The rules of ∀Exp+Res (adapted from [19])

clauses that comprise only existential variables from the original QBF; but additionally, each existential variable x is annotated with a substitution to those universal variables that precede x in the quantification order. For instance, the clause $x \vee b^{0/u}$ can be derived from the original clause $x \vee u$ under the prefix $\exists x \forall u \exists b$.

Besides the aforementioned resolution-based calculi, there is a system by Klieber et al. [24,23], which operates on pairs of sets of literals, rather than clauses; this system is in its workings akin to LD-Q-Res. Van Gelder defines an extension of Q-Res, called *QU-resolution*, which additionally supports resolution over universal variables [34]. Another extension of Q-Res are *variable dependencies* [30,32,33] which enable more flexible ∀-reduction than traditional Q-Res. For proofs of true QBFs *term-resolution* was developed [12] or *models* in the form of Boolean functions [22] but those do not provide polynomially-verifiable proof system. Some limitations of term-resolution were shown by Janota et al. [16]. A comparison of sequent calculi [25] and Q-Res was done by Egly [10].

3 Instantiation-based Calculi IR-calc and IRM-calc

We begin by setting up a framework allowing us to define our new calculi. The framework hinges on the concept of annotated clauses. An *extended assignment* is a partial mapping from the boolean variables to $\{0, 1, *\}$. Two assignments τ and μ are called *contradictory* if there exists a variable $x \in \mathsf{dom}(\tau) \cap \mathsf{dom}(\mu)$ with $\tau(x) \neq \mu(x)$. An *annotated clause* is a clause where each literal is annotated by an extended assignment to universal variables. For an extended assignment σ to universal variables we write $l^{[\sigma]}$ to denote an annotated literal where $[\sigma] = \{c/u \in \sigma \mid \mathsf{lv}(u) < \mathsf{lv}(l)\}$.

Further we define operations that let us modify annotations of a clause by *instantiation*. For (extended) assignments τ and μ, we write $\tau \veebar \mu$ for the assignment σ defined as follows: $\sigma(x) = \tau(x)$ if $x \in \mathsf{dom}(\tau)$, otherwise $\sigma(x) = \mu(x)$ if $x \in \mathsf{dom}(\mu)$. The operation $\tau \veebar \mu$ is referred to as *completion* because μ provides values for variables that are not defined in τ. The operation is associative and therefore we can omit parentheses. In contrast, it is *not* commutative.

Lemma 1. *The following equalities hold.*

1. $(\mu \veebar \tau) \veebar \sigma = \mu \veebar (\tau \veebar \sigma)$
2. *For non-contradictory* μ *and* τ*, it holds that* $\mu \veebar \tau = \tau \veebar \mu$.
3. $\tau \veebar \tau = \tau$.
4. *For non-contradictory* μ *and* τ*, it holds that* $\mu \veebar \tau = \mu \cup \tau$.

We consider an auxiliary function $\mathsf{inst}(\tau, C)$, which for an extended assignment τ and an annotated clause C returns the annotated clause $\left\{l^{[\sigma \trianglelefteq \tau]} \mid l^\sigma \in C\right\}$.

Our first new system IR-calc operates on clauses annotated with usual assignments with range $\{0, 1\}$. The calculus introduces clauses from the matrix and allows to **i**nstantiate and **r**esolve clauses; hence the name IR-calc. It comprises the rules in Figure 4.

$$\frac{}{\left\{x^{[\tau]} \mid x \in C, x \text{ is existential}\right\}} \text{ (Axiom)}$$

C is a non-tautological clause from the matrix. $\tau = \{0/u \mid u \text{ is universal in } C\}$, where the notation $0/u$ for literals u is shorthand for $0/x$ if $u = x$ and $1/x$ if $u = \neg x$.

$$\frac{x^\tau \vee C_1 \qquad \neg x^\tau \vee C_2}{C_1 \cup C_2} \text{ (Resolution)} \qquad \frac{C}{\mathsf{inst}(\tau, C)} \text{ (Instantiation)}$$

τ is an assignment to universal variables.

Fig. 4. The rules of IR-calc.

Our second system IRM-calc is an extension of IR-calc where we allow extended assignments with range $\{0, 1, *\}$. To introduce $*$ we include a new rule called *merging*. IRM-calc is defined in Figure 5. The resolution rule has been adapted to deal with $*$, but when σ, ξ are empty we have exactly the resolution rule from Figure 4.

$$\frac{}{\left\{x^{[\tau]} \mid x \in C, x \text{ is existential}\right\}} \text{ (Axiom)}$$

C is a non-tautological clause from the matrix. $\tau = \{0/u \mid u \text{ is universal in } C\}$

$$\frac{x^{\tau \cup \xi} \vee C_1 \qquad \neg x^{\tau \cup \sigma} \vee C_2}{\mathsf{inst}(\sigma, C_1) \cup \mathsf{inst}(\xi, C_2)} \text{ (Resolution)}$$

$\mathsf{dom}(\tau)$, $\mathsf{dom}(\xi)$ and $\mathsf{dom}(\sigma)$ are mutually disjoint. $\mathsf{rng}(\tau) = \{0, 1\}$

$$\frac{C \vee b^\mu \vee b^\sigma}{C \vee b^\xi} \text{ (Merging)}$$

$\mathsf{dom}(\mu) = \mathsf{dom}(\sigma)$. $\xi = \{c/u \mid c/u \in \mu, c/u \in \sigma\} \cup \{*/u \mid c/u \in \mu, d/u \in \sigma, c \neq d\}$

$$\frac{C}{\mathsf{inst}(\tau, C)} \text{ (Instantiation)} \qquad \tau \text{ is an assignment to universal variables with } \mathsf{rng}(\tau) \subseteq \{0, 1\}.$$

Fig. 5. The rules of IRM-calc.

Example 2. Consider the (true) QBF $\exists x \forall u w \exists b.\, (x \vee u \vee b) \wedge (\neg x \vee \neg u \vee b) \wedge (u \vee w \vee \neg b)$. In both calculi axioms yield $x \vee b^{0/u}$, $\neg x \vee b^{1/u}$, and $\neg b^{0/w, 0/u}$. In IR-calc we resolve to get $b^{0/u} \vee b^{1/u}$. IRM-calc further derives $b^{*/u}$ by merging. Intuitively, $b^{0/u} \vee b^{1/u}$ means that the existential player must play so that for any assignment to w either $b = 1$ if $u = 0$, or $b = 0$ if $u = 1$. So for instance, the player might choose to play $b = 1$ if $w = 0$ and $u = 1$, and if

$w = 1$ and $u = 0$. The clause $b^{*/u}$ means that b must be 1 but the existential player has the freedom to chose whether $b = 1$ when $u = 0$ or when $u = 1$, based on the value of w, i.e. it is a more compact representation of the previous clause. Note that it would be *unsound* to derive that $b = 1$ for any move of the universal player as b needs to be 0 when $u = w = 0$ due to the third axiom.

If the third clause of the formula is changed to $\neg b$, the formula becomes false, which is shown by instantiating $\neg b$ to $\neg b^{0/u}$ and to $\neg b^{1/u}$, using those to obtain x and $\neg x$ by resolution and deriving the empty clause.

Example 3. Consider the QBF $\exists x \forall u \exists bc.\, (x \vee u \vee b) \wedge (\neg x \vee \neg u \vee c) \wedge (\neg b \vee c) \wedge (\neg c)$. The following derivation is possible in IR-calc. Resolving $x \vee b^{0/u}$ and $\neg x \vee c^{1/u}$ yields $b^{0/u} \vee c^{1/u}$. Instantiating $\neg b \vee c$ by $0/u$ gives $\neg b^{0/u} \vee c^{0/u}$, resolving this with the previous resolvent yields $c^{0/u} \vee c^{1/u}$. Refutation can be obtained by instantiation of $\neg c$ once by $0/u$ and once by $1/u$ and subseqent two resolution steps. In IRM-calc it is possible to obtain $c^{*/u}$ by merging and resolve that directly with $\neg c$, which yields $\bot$.

4 Soundness and Extraction of Winning Strategies

The purpose of this section is twofold: show how to obtain a *winning strategy* for the universal player given an IRM-calc proof, and, to show that IRM-calc is *sound* (and therefore also IR-calc). First we show how to obtain a winning strategy for the universal player from a proof. From this, the soundness of the calculus follows because a QBF is false if and only if such strategy exists.

The approach we follow is similar to the one used for Q-Res [14] or LD-Q-Res [11]. Consider a QBF $\Gamma = \exists E \forall U.\, \Phi$, where E and U are sets of variables and Φ is a QBF (potentially with further quantification). Let π be an IRM-calc refutation of Γ, and let ϵ be a total assignment to E. The assignment ϵ represents a move of the existential player. Reduce π to a refutation π_ϵ of $\forall U.\, \Phi|_\epsilon$. To obtain a response of the universal player, we construct an assignment μ to the variables U such that reducing π_ϵ by μ gives a refutation of $\Phi|_{\epsilon \cup \mu}$.

Let $\pi_{\epsilon,\mu}$ be the proof resulting from reducing π_ϵ by μ. The game continues with $\phi|_{\epsilon\cup\mu}$ and $\pi_{\epsilon,\mu}$. In each of these steps, two quantifier levels are removed from the given QBF and a refutation for each of the intermediate formulas is produced. This guarantees a winning strategy for the universal player because in the end the existential player will be faced with an unsatisfiable formula without universal variables. We follow this notation for the rest of the section.

To reduce a refutation π by the existential assignment ϵ, we reduce the leaves of π by ϵ and repeat the steps of π with certain modifications. Instantiation steps are repeated with no discrimination. Merging is repeated in the reduced proof unless either of the merged literals is not in the reduced clause and then the clause is left as it is. Whenever a resolution step is possible, just repeat it in the reduced proof. If it is not possible, the resolvent in the reduced proof is obtained from the antecedent that is not $\top$ and that does *not* contain the pivot literal. If such does not exist, the resolvent is marked as $\top$ (effectively removing it from the proof). When producing a resolvent from a single antecedent, additional instantiation is required. This instantiation is the same one as done by the original resolution step but any $*$ is replaced by 0 (indeed, we can choose the constant arbitrarily). Like so, domains of annotations are preserved. In the end, any clauses marked as $\top$ are removed. An algorithmic description of this transformation is given in Algorithm 1.

Algorithm 1: Reduction of an IRM-calc proof by existential assignment ϵ. The literal $x^{\tau'}$ denotes literal that resulted in reduction from x^{τ}; such τ' may contain 0 or 1 instead of some $*$. Likewise for $x^{\sigma'}$.

Function Reduce (C, ϵ)
begin
  if *C is an axiom derived from C'* **then**
    return D derived from $C'|_\epsilon$ by the axiom rule
  if *C is derived by resolution of $x^\tau \vee C_1$ and $\neg x^\sigma \vee C_2$* **then**
    $D_1 \leftarrow$ **Reduce**(C_1, ϵ)
    $D_2 \leftarrow$ **Reduce**(C_2, ϵ)
    if $D_1 \neq \top$ *and* $x^{\tau'} \notin D_1$ **then**
      return $\mathsf{inst}(\{c/u \mid u \notin \mathsf{dom}(\tau),$ either $c/u \in \sigma, c \in \{0,1\}$ or $c = 0, */u \in \sigma\}, D_1)$
    else if $D_2 \neq \top$ *and* $\neg x^{\tau'} \notin D_2$ **then**
      return $\mathsf{inst}(\{c/u \mid u \notin \mathsf{dom}(\sigma),$ either $c/u \in \tau, c \in \{0,1\}$ or $c = 0, */u \in \tau\}, D_2)$
    else if $D_1 \neq \top$ *and* $D_2 \neq \top$ *and* $x^{\tau'} \in D_1$ *and* $\neg x^{\sigma'} \in D_2$ **then**
      return resolvent of D_1 and D_2
    else return $\top$
  if *C is obtained from C' by merging literals l^τ and l^σ* **then**
    $D' \leftarrow$ **Reduce**(C', ϵ)
    if $l^{\tau'} \in D' \wedge l^{\sigma'} \in D'$ **then**
      return D' where $l^{\tau'}$ and $l^{\sigma'}$ are merged
    else return D'
  if $C = \mathit{inst}(\tau, C')$ **then**
    return $\mathsf{inst}(\tau,$ **Reduce**$(C', \epsilon))$

To obtain an assignment to the variables U, collect all the assignments μ to U appearing in annotations in π_ϵ; any variables not appearing in π_ϵ are given an arbitrary value. To obtain $\pi_{\epsilon,\mu}$, remove occurrences of U-variables from the annotation in the of proof π_ϵ. This will leave us with a valid refutation because we will show in Lemma 5 that only a single value constant annotation can appear in the entire proof π_ϵ for each variable in U.

To show that this procedure is correct, we need to argue that the reduction returns a valid IRM-calc refutation π_ϵ, and that π_ϵ does not contain annotations giving contradictory values to variables in U. We start with the first claim.

Lemma 4. *The above reduction yields a valid IRM-calc refutation π_ϵ of $\forall U. \Phi|_\epsilon$.*

Proof. By induction on the derivation depth. The induction hypothesis states that any derived clause C' in the reduced proof has a valid derivation $\pi_{C'}$. Further, if there is a literal $l^{\sigma'} \in C'$, then the original clause C corresponding to the clause C' contains a literal l^σ, where σ satisfies the following: for every $c/u \in \sigma'$ there is exactly one $d/u \in \sigma$ where $d = c$ or $d = *$. Further $\mathsf{dom}(\sigma') = \mathsf{dom}(\sigma)$, i.e. a $*$ may become a constant in the reduction.

Base Case. Since ϵ assigns only existential variables, the clause in the reduced proof is either unchanged, or some literals are removed from the axiom clause, or the whole axiom is set to $\top$.

Instantiation. If an instantiation step was present in the original proof, it is also present after the reduction. Having already replaced $*$ by constants does not affect the instantiation step as the value of that annotation does not change in either the reduced or original proof.

Resolution. If the actual resolution step is performed, consider the resolution of $x^{\tau \cup \xi} \vee C_1$ and $\neg x^{\tau \cup \sigma} \vee C_2$ with the following conditions: $\mathsf{dom}(\tau)$, $\mathsf{dom}(\xi)$ and $\mathsf{dom}(\sigma)$ are mutually disjoint; $\mathsf{rng}(\tau) = \{0, 1\}$.

By induction hypothesis this is reduced to $x^{\tau \cup \xi'} \vee C_1'$ and $\neg x^{\tau \cup \sigma'} \vee C_2'$ with $\mathsf{dom}(\xi') = \mathsf{dom}(\xi)$ and $\mathsf{dom}(\sigma') = \mathsf{dom}(\sigma)$. We can perform a resolution step as τ, ξ', σ' have disjoint domains. Now the instantiation steps are performed: if $*/u$ was introduced as annotation in a literal in the original proof, the literal (if it exists) will be annotated with c/u in the new proof for $c \in \{0, 1, *\}$. Otherwise, we perform instantiation steps, which by the disjointness of ξ and σ lead to non-contradictory annotations and preserve the induction hypothesis.

If the resolution step is not performed and instead one of its antecedents is used, then the induction hypothesis is preserved because in the reduced proof we only modify instantiations from $*/u$ to $0/u$.

Merging. Consider a merge of l^τ and l^σ in the original proof. If in the reduced clause these literals became $l^{\tau'}$ and $l^{\sigma'}$, some $*$ in those annotations might be 0 or 1 due to the induction hypothesis but the domains of all τ, τ', σ, and σ' are equal. The reduced proof will contain a merge of $l^{\tau'}$ and $l^{\sigma'}$, which preserves the induction hypothesis. If the reduced clause contains only one of the literals, say $l^{\tau'}$, the merge step is not performed in π_ϵ. This preserves the induction hypothesis as the literal resulting from merge of l^τ and l^σ has an annotation with domain equal to τ (and therefore to τ'), and, with some constants 0 or 1 changed to $*$ compared to τ (and therefore to τ'). □

Lemma 5. *Let π be an IRM-calc refutation of a QBF formula starting with a block of universally quantified variables U. Consider the set of annotations μ on variables U that appear anywhere in π. Then μ is non-contradictory and does not contain instances of $*$.*

Proof. The proof proceeds by induction on the derivation depth. Let μ_C denote the set of annotations to variables in U appearing anywhere in the derivation of C (i.e., we only consider the connected component of the proof dag with sink C). The induction hypothesis states:

(i) The set μ_C is non-contradictory.
(ii) For every literal $l^\sigma \in C$, it holds that $\mu_C \subseteq \sigma$.
(iii) $*/u \notin \mu_C$, for any $u \in U$.

Base Case. Condition (i) is satisfied by the axioms because we are assuming there are no complementary literals in clauses in the matrix. Condition (ii) is satisfied because all existential literals are at a higher level than the variables of U. Condition (iii) holds because we do not instantiate by $*/u$ in the axiom rule.

Instantiation. Let $u \in U$ and $C = \mathsf{inst}(c/u, C')$ in the proof π. By induction hypothesis, u either appears in the annotations of all the literals l^ξ in C' or it does not appear in any of them. In the first case, the instantiation step is ineffective. In the second case, c/u is added to all literals in C. By induction hypothesis u does not appear in any annotation of any clause in the sub-proof deriving C', and hence C is the first clause containing u.

Resolution. Let C be derived by resolving $x^{\tau \cup \xi} \vee C_1$ and $\neg x^{\tau \cup \sigma} \vee C_2$. Let $u \in U$, consider the following cases.

Case 1. For some $c \in \{0, 1\}$, $c/u \in \sigma$ and $u \notin \mathsf{dom}(\xi)$. By induction hypothesis, u does not appear in the annotations of C_1. Hence $\mathsf{inst}(\sigma, C_1)$ adds c/u to all the annotations in C_1.

Case 2. $c/u \in \tau$. By induction hypothesis, c/u appears in all annotations of C_1, C_2 and hence in all annotations of the resolvent.

Case 3. $u \notin \mathsf{dom}(\tau) \cup \mathsf{dom}(\sigma) \cup \mathsf{dom}(\xi)$. Then u does not appear as annotation anywhere in the derivation of either of the antecedents and neither it will appear in the resolvent.

Merging. Because of (i) we do not obtain $*$ for variables in U. □

Theorem 6. *The construction above yields a winning strategy for the universal player.*

Proof. For any QBF $\Gamma = \exists E \forall U.\, \Phi$, and ϵ, the construction provides an IRM-calc refutation $\pi_{\epsilon,\mu}$ of $\Phi|_{\epsilon \cup \mu}$. This process is iterated until no universal variables are left in the formula. Hence we get an IRM-calc refutation of whatever was left from the matrix of Γ. Since an IRM-calc refutation on a formula with no universal variables is in fact a classical propositional resolution refutation, we are left with an unsatisfiable formula, i.e. a formula with no winning move for the existential player. Hence, all the considered assignments correspond to a game won by the universal player. Since this process works for any assignment made by the existential player, it yields a winning strategy for the universal player. □

The soundness of IRM-calc follows directly from Theorem 6.

Corollary 7. *The calculi IR-calc and IRM-calc are sound.*

5 Simulations of Known QBF Proof Systems

In this section we prove that our calculi simulate the main existing resolution-based QBF proof systems. As a by-product, this also shows completeness of our proof systems IR-calc and IRM-calc. We start by simulating Q-resolution, which is even possible with our simpler calculus IR-calc.

Theorem 8. *IR-calc p-simulates Q-Res.*

Proof. Let $C_1, \ldots, C_k$ be a Q-Res proof. We translate the clauses into $D_1, \ldots, D_k$, which will form the skeleton of a proof in IR-calc.

- For an axiom C_i in Q-Res we introduce the same clause D_i by the axiom rule of IR-calc, i.e., we remove all universal variables and add annotations.
- If C_i is obtained via $\forall$-reduction from C_j, then $D_i = D_j$.
- Consider now the case that C_i is derived by resolving C_j and C_k with pivot variable x. Then $D_j = x^\tau \vee K_j$ and $D_k = x^\sigma \vee K_k$. We instantiate to get $D'_j = \mathsf{inst}(\sigma, D_j)$ and $D'_k = \mathsf{inst}(\tau, D_k)$. Define D'_i as the resolvent of D'_j and D'_k. In order to obtain D_i we must ensure that there are no identical literals with different annotations. For this consider the set $\zeta = \{c/u \mid c/u \in t, l^t \in D'_i\}$ and define $D_i = \mathsf{inst}(\zeta, D'_i)$. This guarantees that we will always have fewer literals in D_i than in C_i, and we get a refutation.

We will prove inductively that the resolution steps are valid, by showing that τ and σ are not contradictory and ζ does not contain contradictory annotations.

Induction Hypothesis. For all existential literals l we have $l \in C_i$ iff $l^t \in D_i$ for some annotation t. Additionally, if $0/u \in t$ for a literal u, then $u \in C_i$ (where for a variable x, we equivalently denote the annotation $1/x$ by $0/\neg x$).

Before proving the induction we argue that this yields the claim above. Assume for a contradiction that τ contradicts σ. This means that for some universal variable u, both u and $\neg u$ appear in C_i, which is not allowed; similarly if ζ contains contradictory annotations.

We now show the inductive claim by induction on the proof length.

Base case: Axiom. $l^t \in D_i$ iff $l \in C_u$ by definition. As annotations falsify all universal literals in the original clause, $0/u \in t$ for literal u implies $u \in C_i$.

Inductive step: $\forall$-Reduction. Suppose C_i is obtained via universal reduction from C_j. We have $l^t \in D_i$ iff $l^t \in D_j$, iff $l \in C_j$. Since l is existential it is not reduced and $l \in C_i$. Assume now $l^t \in D_j$ and $0/u \in t$. By inductive hypothesis we have $u \in C_j$. Further, u cannot be reduced in this step because it is blocked by l; hence $u \in C_i$.

Resolution. Suppose that C_i is derived by resolving C_j and C_k over variable x, and $D_j = x^\tau \vee K_j$ and $D_k = \neg x^\sigma \vee K_k$. Then $l^t \in D_i$ iff $l \in C_j \cup C_k \backslash \{x, \neg x\} = C_i$. Without loss of generality, if $0/u \in t$ then there is some literal $p^{t'} \trianglelefteq \sigma \in D'_i$ (with $p^{t'} \in D_j$) such that $0/u \in t' \trianglelefteq \sigma$. If $0/u \in t'$ then $u \in C_j$ by inductive hypothesis, and if $0/u \in \sigma$ then $u \in C_k$, again by inductive hypothesis; hence $u \in C_i$. □

Despite its simplicity, IR-calc is powerful enough to also simulate the expansion based proof system $\forall$Exp+Res from [19].

Theorem 9. *IR-calc p-simulates $\forall$Exp+Res.*

Proof. Consider an $\forall$Exp+Res proof $C_1, \ldots, C_k$. We use that to form the skeleton of a proof $D_1, \ldots, D_k$ in IR-calc.

- If C_i is an axiom from clause C and assignment τ we construct D_i by taking the axiom in IR-calc of C and then instantiating with $\mathsf{inst}(\tau, C)$.
- If C_i is derived by resolving C_j, C_k over variable x^τ, then D_i is derived by resolving D_j, D_k over variable x^τ.

This yields a valid IR-calc proof because $l^t \in D_i$ iff $l^t \in C_i$, which is preserved under applications of both rules. □

We now come to the simulation of a more powerful system than Q-resolution, namely LD-Q-Res from [2]. We show that this system is simulated by IRM-calc. The proof uses a similar, but more involved technique as in Theorem 8.

Theorem 10. *IRM-calc p-simulates LD-Q-Res.*

Proof. Consider an LD-Q-Res refutation $C_1, \ldots, C_n$. We construct clauses $D_1, \ldots, D_n$, which will form the skeleton of the IRM-calc proof. The construction proceeds as follows. If C_i is an axiom, D_i is constructed by the axiom rule from the same clause. If C_i is a $\forall$-reduction of C_j with $j < i$, then we set D_i equal to D_j. If C_i is obtained by a resolution step from C_j and C_k with $j < k < i$, the clause D_i is obtained by a resolution step from D_j and D_k, yielding clause K, and by performing some additional steps on K. Firstly, we let $\theta = \{c/u \mid c \in \{0,1\}, c/u \in t, l^t \in K\} \cup \{0/u \mid */u \in t, l^t \in K\}$ and we perform instantiation on K by θ (in any order) to derive K'. Since θ has the same domain as the entire set of annotations in K' we apply merging on K' for all literals to derive D_i.

We have to show that this construction yields a valid IRM-calc refutation. For this we claim that the construction preserves the following invariants for $i = 1, \ldots, n$:

(1) For an existential literal l, it holds that $l \in C_i$ iff $l^t \in D_i$ for some t.
(2) The clause D_i has no literals l^{t_1} and l^{t_2} such that $t_1 \neq t_2$.

(3) If $l^t \in D_i$ with $0/u \in t$, then $u \in C_i$ or $u^* \in C_i$, likewise if $l^t \in D_i$ with $1/u \in t$, then $\neg u \in C_i$ or $u^* \in C_i$.
(4) If $l^t \in D_i$ with $*/u \in t$, then $u^* \in C_i$.

We show these invariants by induction on i.

Base case (axiom). Because we do not remove or add any existential literals in the axiom case, condition (1) holds. Likewise we do not create duplicates, so (2) holds. As we do not obtain any $*$ annotations from axioms, (4) holds. Any 0/1 annotation corresponds exactly to the opposite literal appearing in the clause, by definition of the axiom rule, hence (3) holds.

Inductive step ($\forall$-reduction). Consider a $\forall$-reduction step from C_j to C_i on universal variable u. Because we do not alter the existential literals in a $\forall$-reduction and the corresponding clause D_i in the IRM-calc proof remains unchanged, conditions (1) and (2) are satisfied by induction hypothesis. For conditions (4) and (3) we note that D_j cannot contain any annotations involving u. This holds because u would only appear as annotation on existential literals with level higher than u. These cannot exist as they would be blocking the reduction by (1).

Resolution step. Consider C_j, C_k being resolved in LD-QRes to obtain C_i. As only the resolved variable is removed, which is removed completely due to condition (2), D_i fulfills (1). By induction hypothesis we know that there can be at most two copies of each variable when we derive K. Their annotations have the same domain in K', because instantiation by θ applies the entire domain of all annotations in the clause to all its literals. It then follows that all copies of identical literals are merged into one literal in D_i. Therefore (2) holds for D_i.

To prove (3) consider the case where $l^t \in D_i$ with $0/u \in t$. The case with $1/u \in t$ is analogous. We know that $0/u$ appearing in D_i means that $0/u$ must appear in K' as merging cannot produce a new annotation $0/u$. Existence of $0/u$ in K' means that either $*/u$ appears in K or $0/u$ appears in K (the case with $1/u$ is simpler as it guarantees that $1/u$ appears in K). No new annotations are created in a resolution step, so either $*/u$ or $0/u$ must appear in one or more of D_j, D_k. By induction hypothesis this means that u or u^* appears in $C_j \cup C_k$, hence also in C_i.

To show condition (4), let $l^t \in D_i$ with $*/u \in t$. Then either $*/u$ is present in K', or $0/u$ and $1/u$ are present in K' and will be merged. In the first case it is clear that some $*/u$ annotation appears in K and thus in D_j or in D_k, in which case from (4) of the induction hypothesis u^* must appear in C_i. In the second case it is possible that $0/u$ in K' was obtained from $*/u$ in K. Thus as already argued, u^* must appear in C_i. If instead $1/u$, $0/u$ are both present in K then they must come from the original clauses D_j, D_k. If they both appear in the same clause D_j, then by condition (3) it must be the case that u^* appears in C_j and thus in C_i. If, however, they appear in different clauses, then by (3) either of the clauses C_j, C_k contains u^* or they contain literals over u of opposite polarity. Both situations merge the literals to $u^* \in C_i$.

We now show that these invariants imply that we indeed obtain a valid IRM-calc proof. We only need to consider the resolution steps. Suppose $x^{t_1} \in D_j$ and $\neg x^{t_2} \in D_k$ where C_j and C_k are resolved on x to get C_i in the LD-Q-Res proof. To perform the resolution step between D_j and D_k we need to ensure that we do not have $c/u \in t_1$, $d/u \in t_2$ where $c \neq d$ or $c = d = *$. Assume on the contrary that $*/u \in t_1$ and $c/u \in t_2$. By (4) we have $u^* \in C_j$, and by (3) some literal of u is in C_k. But as $\mathsf{lv}(u) < \mathsf{lv}(x)$ the LD-resolution of C_j and C_k on

variable x is forbidden, giving a contradiction. Similarly, if there is $0/u \in t_1$ and $1/u \in t_2$, then either we get the same situation or we have two opposite literals of u in the different clauses C_j, C_k. In either case the resolution of C_j, C_k is forbidden. Hence the IRM-calc proof is correct.

Further, the IRM-calc proof is indeed a refutation. Namely, if we derive $\bot$ in an LD-Q-Res proof, by (1) we derive a clause with no existential literals in the IRM-calc proof. That clause also contains no universal variables as these are all instantiated at the axioms. Therefore, we derive the empty clause $\bot$.

Finally, we observe that all steps of the construction can be performed in polynomial time, thus we obtain a p-simulation. □

6 Conclusion

This paper introduces two novel calculi for quantified Boolean formulas. Both of these calculi are anchored in a common framework of *annotated clauses*. The first calculus, IR-calc, provides the rules of resolution and instantiation of clauses. Instantiation resembles *specialization* in first-order logic, i.e., an annotated literal $b^{0/u}$ specifies that b is true whenever $u = 0$. Resolution in IR-calc can be seen as a simplified version of *Robinson's resolution* [29]. The second calculus, IRM-calc, additionally enables *merging* contradictory annotations. By merging, a disjunction $b^{1/u} \vee b^{0/u}$ is shortened to a single literal $b^{*/u}$. The paper demonstrates that the simple calculus IR-calc already p-simulates Q-resolution and the expansion-based system ∀Exp+Res. The extended version IRM-calc additionally p-simulates long-distance Q-resolution. The paper further demonstrates that refutations in the introduced calculi enable generation of winning strategies of the universal player—a favorable property from a practical perspective [2].

The contribution of the paper is both practical and theoretical. From a practical perspective, a calculus unifying the existing calculi for QBF enables a uniform certification of off-the-shelf QBF solvers. From a theoretical perspective, a unifying calculus provides an underpinning necessary for complexity characterizations of existing solvers as well as for furthering our understanding of the strengths of the underlying proof systems.

Acknowledgments

The second author was supported by a Doctoral Training Grant from EPSRC. This work was partially supported by FCT grants ATTEST (CMU-PT-/ELE/0009/2009), POLARIS (PTDC/EIA-CCO/123051/2010), INESC-ID's multiannual PIDDAC funding PEst-OE/EEI/LA0021/2011, and a grant from the John Templeton Foundation.

References

1. Arora, S., Barak, B.: Computational Complexity - A Modern Approach. Cambridge University Press (2009)
2. Balabanov, V., Jiang, J.H.R.: Unified QBF certification and its applications. Formal Methods in System Design 41(1), 45–65 (2012)
3. Beame, P., Kautz, H.A., Sabharwal, A.: Towards understanding and harnessing the potential of clause learning. J. Artif. Intell. Res. (JAIR) 22, 319–351 (2004)
4. Benedetti, M.: Evaluating QBFs via symbolic Skolemization. In: LPAR (2004)

5. Benedetti, M., Mangassarian, H.: QBF-based formal verification: Experience and perspectives. JSAT 5(1-4), 133–191 (2008)
6. Biere, A.: Resolve and expand. In: SAT. pp. 238–246 (2004)
7. Buss, S.R.: Towards NP-P via proof complexity and search. Ann. Pure Appl. Logic 163(7), 906–917 (2012)
8. Buss, S.R., Hoffmann, J., Johannsen, J.: Resolution trees with lemmas: Resolution refinements that characterize DLL algorithms with clause learning. Logical Methods in Computer Science 4(4) (2008)
9. Cook, S.A., Reckhow, R.A.: The relative efficiency of propositional proof systems. J. Symb. Log. 44(1), 36–50 (1979)
10. Egly, U.: On sequent systems and resolution for QBFs. In: SAT. pp. 100–113 (2012)
11. Egly, U., Lonsing, F., Widl, M.: Long-distance resolution: Proof generation and strategy extraction in search-based QBF solving. In: LPAR (2013)
12. Giunchiglia, E., Narizzano, M., Tacchella, A.: Clause/term resolution and learning in the evaluation of quantified Boolean formulas. JAIR 26(1), 371–416 (2006)
13. Giunchiglia, E., Marin, P., Narizzano, M.: Reasoning with quantified boolean formulas. In: Handbook of Satisfiability, pp. 761–780. IOS Press (2009)
14. Goultiaeva, A., Van Gelder, A., Bacchus, F.: A uniform approach for generating proofs and strategies for both true and false QBF formulas. In: IJCAI (2011)
15. Hertel, P., Bacchus, F., Pitassi, T., Van Gelder, A.: Clause learning can effectively p-simulate general propositional resolution. In: AAAI (2008)
16. Janota, M., Grigore, R., Marques-Silva, J.: On QBF proofs and preprocessing. In: LPAR. pp. 473–489 (2013)
17. Janota, M., Klieber, W., Marques-Silva, J., Clarke, E.M.: Solving QBF with counterexample guided refinement. In: SAT. pp. 114–128 (2012)
18. Janota, M., Marques-Silva, J.: ∀Exp+Res does not P-Simulate Q-resolution. International Workshop on Quantified Boolean Formulas (2013)
19. Janota, M., Marques-Silva, J.: On propositional QBF expansions and Q-resolution. In: SAT. pp. 67–82 (2013)
20. Kleine Büning, H., Bubeck, U.: Theory of quantified boolean formulas. In: Handbook of Satisfiability, pp. 735–760. IOS Press (2009)
21. Kleine Büning, H., Karpinski, M., Flögel, A.: Resolution for quantified Boolean formulas. Inf. Comput. 117(1), 12–18 (1995)
22. Kleine Büning, H., Subramani, K., Zhao, X.: Boolean functions as models for quantified boolean formulas. J. Autom. Reasoning 39(1), 49–75 (2007)
23. Klieber, W., Janota, M., Marques-Silva, J., Clarke, E.M.: Solving QBF with free variables. In: Schulte, C. (ed.) CP. vol. 8124, pp. 415–431. Springer (2013)
24. Klieber, W., Sapra, S., Gao, S., Clarke, E.M.: A non-prenex, non-clausal QBF solver with game-state learning. In: SAT (2010)
25. Krajíček, J., Pudlák, P.: Quantified propositional calculi and fragments of bounded arithmetic. Mathematical Logic Quarterly 36(1), 29–46 (1990)
26. Marques Silva, J.P., Lynce, I., Malik, S.: Conflict-driven clause learning SAT solvers. In: Handbook of Satisfiability. IOS Press (2009)
27. Pipatsrisawat, K., Darwiche, A.: On the power of clause-learning SAT solvers as resolution engines. Artif. Intell. 175(2), 512–525 (2011)
28. Rintanen, J.: Asymptotically optimal encodings of conformant planning in QBF. In: AAAI. pp. 1045–1050. AAAI Press (2007)
29. Robinson, J.A.: A machine-oriented logic based on the resolution principle. J. ACM 12(1), 23–41 (1965)
30. Samer, M., Szeider, S.: Backdoor sets of quantified Boolean formulas. J. Autom. Reasoning 42(1), 77–97 (2009)
31. Segerlind, N.: The complexity of propositional proofs. Bulletin of Symbolic Logic 13(4), 417–481 (2007)
32. Slivovsky, F., Szeider, S.: Variable dependencies and Q-Resolution. International Workshop on Quantified Boolean Formulas (2013)
33. Van Gelder, A.: Variable independence and resolution paths for quantified Boolean formulas. In: Lee, J.H.M. (ed.) CP. vol. 6876, pp. 789–803. Springer (2011)
34. Van Gelder, A.: Contributions to the theory of practical quantified Boolean formula solving. In: Milano, M. (ed.) CP. vol. 7514, pp. 647–663. Springer (2012)
35. Zhang, L., Malik, S.: Conflict driven learning in a quantified Boolean satisfiability solver. In: ICCAD. pp. 442–449 (2002)

ECCC ISSN 1433-8092

http://eccc.hpi-web.de