

Genuine Lower Bounds for QBF Expansion

Olaf Beyersdorff

School of Computing, University of Leeds, United Kingdom
o.beyersdorff@leeds.ac.uk

Joshua Blinkhorn

School of Computing, University of Leeds, United Kingdom
scjlb@leeds.ac.uk

Abstract

We propose the first general technique for proving genuine lower bounds in expansion-based QBF proof systems. We present the technique in a framework centred on natural properties of winning strategies in the ‘evaluation game’ interpretation of QBF semantics. As applications, we prove an exponential proof-size lower bound for a whole class of formula families, and demonstrate the power of our approach over existing methods by providing alternative short proofs of two known hardness results. We also use our technique to deduce a result with manifest practical import: in the absence of propositional hardness, formulas separating the two major QBF expansion systems must have unbounded quantifier alternations.

2012 ACM Subject Classification Theory of computation → Proof complexity

Keywords and phrases QBF, Proof Complexity, Lower-bound Techniques, Resolution

Digital Object Identifier 10.4230/LIPIcs.STACS.2018.12

Acknowledgements We thank Meena Mahajan and Anil Shukla for helpful discussions on this work. We also thank the anonymous reviewers for useful suggestions, which helped to improve the presentation of this paper. Research was supported by grants from the John Templeton Foundation (grant no. 60842) and the EU (CORCON project).

1 Introduction

The central problem in *proof complexity* is to determine the size of the smallest proof for a given formula in a specified proof system. From its inception the field has borne tight and fruitful connections to open problems in computational complexity (separation of complexity classes [18, 14]) and first-order logic (separation of bounded arithmetic theories [32, 17]).

Proof complexity has since emerged as the natural theoretical counterpart of practical SAT solving, a subfield of automated reasoning that has enjoyed major success in recent years. Indeed, complexity of proofs and efficiency of solving are fundamentally related: the trace of a SAT solver on an unsatisfiable instance can be interpreted as a proof of falsity, whereby the correctness of each SAT solver is underpinned by a proof system. For example, the dominant paradigm in SAT, *conflict-driven clause learning* (CDCL), produces proofs in a system called *resolution* [14]. Lower bounds on resolution proof size therefore correspond to best-case running time for CDCL solvers. Consequently, there has been intense research activity focussed on proof-size lower bounds, and, in particular, *general techniques* for proof-size lower bounds in propositional logic (cf. [40, 14]).

Proof-theoretic techniques are arguably even more valuable in the increasingly challenging settings of modern solving. Consider the logic of *quantified Boolean formulas* (QBF), which extends propositional logic with existential and universal quantification. The succinct



© Olaf Beyersdorff and Joshua Blinkhorn;

licensed under Creative Commons License CC-BY

35th Symposium on Theoretical Aspects of Computer Science (STACS 2018).

Editors: Rolf Niedermeier and Brigitte Vallée; Article No. 12; pp. 12:1–12:15

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



SYMPOSIUM
ON THEORETICAL
ASPECTS
OF COMPUTER
SCIENCE

encodings of problem instances afforded by this PSPACE-complete language [42] foster applications in diverse areas of computer science (automated planning [15, 22, 25, 39], formal verification [4, 35, 36], ontological reasoning [30], and more [20, 41, 28, 33]). Moreover, the more complex setting has spawned two distinct paradigms in solving and associated proof systems.

One approach uses QCDCL [24], the natural extension of the SAT technology, underpinned by the $P+\forall\text{red}$ family of QBF proof systems [6].¹ A second approach, implemented in the solver RAReQs [26], is through expansion of universal variables, embodied by the proof system $\forall\text{Exp}+\text{Res}$ [27]. Research in proof complexity has revealed that these two paradigms are incomparable [27, 7] – that is, their underlying proof systems do not simulate one another.² This observation led to the proposal of the more sophisticated expansion system IR-calc [7], which simulates both approaches.

It is fair to say that there is a distinct lack of general lower-bound techniques for QBF, especially for the expansion-based systems $\forall\text{Exp}+\text{Res}$ and IR-calc. Researchers have of course attempted to lift lower bound techniques from propositional logic, but with mixed success. The celebrated size-width relations for resolution [3] are rendered ineffective in QBF resolution [10]. Prover-delayer games are only applicable to weaker tree-like proofs, both propositionally [38, 12] and in QBF [11]. Feasible interpolation [31] has been successfully transferred to QBF [9], but is tailored towards instances of a rather specific syntactic form.

Moreover, lifting techniques from SAT to QBF can be misleading, since it inevitably entails some degree of *non-genuineness* [16, 13]. The phenomenon of *genuine QBF hardness* – where lower bounds do *not* originate from the propositional level, as formalised in the oracle model of [13] – is a more suitable notion for the comparison of algorithms in quantified logic. Recent work [5] introduced a new technique for genuine QBF lower bounds in the QCDCL systems $P+\forall\text{red}$. In this paper, we show that a semantically-grounded approach can also be employed in expansion-based systems, fostering the general techniques for genuine lower bounds that are currently missing.

Our contributions: framework, technique, and applications

We propose the *first genuine lower-bound technique* for QBF expansion. We introduce a framework built upon two semantically-grounded measures: *strategy size*, the minimum number of responses in a winning strategy; and *weight*, an extension of strategy size for unbounded prenex CNFs. Our technique encompasses three valuable theorems that express proof-size lower bounds for $\forall\text{Exp}+\text{Res}$ and IR-calc solely in terms of these measures:

- Strategy size is an absolute proof-size lower bound in $\forall\text{Exp}+\text{Res}$ (Theorem 7).
- Small strategy size implies short IR-calc proofs for bounded families (Theorem 9).
- Weight is an absolute proof-size lower bound in IR-calc (Theorem 22).

All three theorems are proved by counting *annotations*, a unique feature of expansion systems. Since propositional inferences preserve annotations, corollaries are invariably genuine QBF lower bounds in the formal sense of [13]. Thus, by providing an account of genuine hardness based on semantics, our technique offers valuable insight into the underlying reasons for hardness in expansion systems. Applications of our theorems represent important forward steps on at least three fronts.

¹ More precisely, QCDCL is underpinned by Q-Res [29], the special case of $P+\forall\text{red}$ in which P is resolution.

² Proof system P_1 simulates proof system P_2 whenever P_1 -proofs can be transformed into P_2 -proofs with at most polynomial increase in proof size.

Intuitive proofs. First, we provide short, semantically-intuitive proofs, supplanting the complicated *ad hoc* arguments that hitherto represented the state-of-the-art in QBF expansion lower bounds. Whereas the authors of [27] needed to invoke Craig’s Interpolation Theorem [19] on the explicit expansion of their unbounded formulas \mathcal{J} ,³ we show their hardness as an immediate consequence of their manifest exponential strategy size. Similarly, the in-depth and lengthy proof of hardness in IR-calc [8] for the unbounded formulas of Kleine Büning et al. [29] is here replaced with a short argument that determines their exponential weight based on game semantics.

New hard formulas. Second, we demonstrate new exponential IR-calc proof-size lower bounds for an entire class of formula families. Using the product constructions of [5, 13], we combine a group of Π_2 CNFs F_i with a minimally unsatisfiable CNF ϕ . Provided the F_i have non-trivial strategy size (a natural stipulation), the strategy size of the product formula grows exponentially with the size of ϕ . We present product formulas with a Σ_3 prefix, but the method easily generalises to arbitrarily many quantifier alternations.

Bounded vs unbounded separations. Third, by applying our second theorem to bounded families in general, we prove that, in the absence of propositional hardness, any separation of the two expansion systems is unbounded. Given that IR-calc simulates Q-Res, this result has a remarkable corollary: any genuine separation of Q-Res from $\forall\text{Exp}+\text{Res}$ is due to an unbounded formula family.

Organisation. We begin with preliminaries in Section 2 followed by the necessary background for QBF expansion in Section 3. We present our lower-bound technique for bounded CNFs and the associated applications in Section 4, and the extension to the unbounded case follows in Section 5. We offer conclusions in Section 6.

2 Preliminaries

Quantified Boolean formulas. A *literal* is a Boolean variable or its negation, a *clause* is a disjunction of literals, and a *CNF* is a conjunction of clauses. Throughout, we refer to a clause as a set of literals and to a CNF as a set of clauses.

A *quantified Boolean formula* (QBF) in *prenex conjunctive normal form* (PCNF) is denoted $F := \mathcal{Q} \cdot \phi$, where (a) $\mathcal{Q} := \mathcal{Q}_1 Z_1 \cdots \mathcal{Q}_n Z_n$ is the *quantifier prefix*, in which the Z_i are pairwise disjoint sets of Boolean variables called *blocks*, $\mathcal{Q}_i \in \{\exists, \forall\}$ for each $i \in [n]$, and $\mathcal{Q}_i \neq \mathcal{Q}_{i+1}$ for each $i \in [n-1]$, and (b) the *matrix* ϕ is a CNF over $\text{vars}(F) := \bigcup_{i=1}^n (E_i \cup U_i)$. A PCNF is *k-bounded* if it has at most k universal blocks.

We denote the existential (resp. universal) variables of F by $\text{vars}_{\exists}(F)$ (resp. $\text{vars}_{\forall}(F)$). For a literal l , we write $\text{var}(l) := z$ if $l = z$ or $l = \neg z$. For a clause C , we write $\text{vars}(C) := \{\text{var}(l) : l \in C\}$, and denote the set of existential (resp. universal) literals in C by C_{\exists} (resp. C_{\forall}). The prefix \mathcal{Q} imposes a linear ordering $<_F$ on the variables of F , such that $z_i <_F z_j$ holds whenever $i < j$, in which case we say that z_j is *right of* z_i and z_i is *left of* z_j . We extend $<_F$ to the blocks of F in the natural way.

A (partial) assignment ρ to the variables of F is represented as a set of literals, typically denoted $\{l_1, \dots, l_k\}$, where literal z (resp. $\neg z$) represents the assignment $z \mapsto 1$ (resp. $z \mapsto 0$).

³ This is our notation; in [27], the formulas are referred to simply as “(2)”.

The CNF $\phi[\rho]$ is obtained from ϕ by removing any clause containing a literal in ρ , and removing the negated literals $\neg l_1, \dots, \neg l_k$ from the remaining clauses. The *restriction of F by ρ* is $F[\rho] := \mathcal{Q}[\rho] \cdot \phi[\rho]$, where $\mathcal{Q}[\rho]$ is obtained from \mathcal{Q} by removing the variables of ρ along with any quantifier whose associated block is rendered empty. For assignments to single variables we may omit the braces; for example, we write $F[l]$ for $F[\{l\}]$.

QBF semantics. Semantics for PCNFs are neatly described by the *two-player evaluation game*. Over the course of a game, the variables of a PCNF are assigned 0/1 values in the order of the prefix, with the \exists -player (\forall -player) choosing the values for the existential (universal) variables. When the game concludes, the players have constructed a total assignment ρ to the variables. The \forall -player wins if and only if ρ falsifies some clause of the matrix.

A \forall -strategy dictates how the \forall -player should respond to every possible move of the \exists -player. A \forall -strategy S for a PCNF F is a mapping from total assignments to $\text{vars}_{\exists}(F)$ into total assignments to $\text{vars}_{\forall}(F)$, such that, for each $i \in [n]$, $S(\alpha)$ and $S(\alpha')$ agree on the first i universal blocks whenever α and α' agree on the first i existential blocks. A strategy S is *winning* if and only if, for each α in the domain of S , $\phi[\alpha \cup S(\alpha)]$ contains the empty clause. We use the terms ‘winning \forall -strategy’ and ‘countermodel’ interchangeably. A PCNF is false if and only if it has a countermodel.

QBF proof systems. A refutational PCNF *proof system* (or *calculus*) P employs a set of axioms and inference rules to prove the falsity of PCNFs. A P derivation of a clause C_m from the *input PCNF* F is a sequence of clauses $\pi := C_1, \dots, C_m$ in which (a) each C_i is either an axiom, or is derivable from previous clauses using an inference rule, and (b) C_m is the unique clause that is not the antecedent of an inference. The subderivation of C_i in π is the subsequence terminating at C_i containing only those clauses used in the derivation of C_i . The size $|\pi|$ of a derivation is the total number of literals appearing in it. A refutation is a derivation of the empty clause.

In this paper, we consider PCNF proof systems based on *resolution*. Resolution is a well-studied refutational proof system for propositional CNF formulas with a single inference rule: the *resolvent* $C_1 \cup C_2$ may be derived from clauses $C_1 \cup \{x\}$ and $C_2 \cup \{\neg x\}$ (variable x is the *pivot*). Resolution is *refutationally* sound and complete: that is, the empty clause can be derived from a CNF iff it is unsatisfiable. There exist a host of resolution-based QBF proof systems – see [8] for a detailed account.

For two PCNF proof systems P_1 and P_2 , a PCNF family \mathcal{F} *separates* P_1 from P_2 if \mathcal{F} has polynomial-size refutations in P_1 but not in P_2 . P_1 *p-simulates* P_2 if each P_2 -proof can be transformed in polynomial time into a P_1 -proof of the same formula [18].

3 Fundamentals of expansion-based calculi

In this section, we recall the definitions of $\forall\text{Exp}+\text{Res}$ [27] and IR-calc [7] and discuss the underlying concepts of the calculi, including their use of annotations. We also cover proof restrictions and strategy extraction, both of which are central to the following discourse.

Intuition and definition. To explain the concept of expansion, we consider the example PCNF $\exists x \forall u \exists t. \phi(x, u, t)$. The formula is semantically equivalent to $\exists x \exists t^0 \exists t^1. \phi(x, 0, t^0) \wedge \phi(x, 1, t^1)$, in which the universal variable u has been ‘expanded out’, yielding a fully existentially quantified formula. Note that variable x , which is left of u , remains unchanged, while we have to create two duplicate copies t^0 and t^1 for the variable t , which is right of u .

To keep track of why we created these copies of t , we annotate them with the reason for their creation, i.e., we write $t^{\neg u}$ instead of t^0 (where $\neg u$ corresponds to the assignment $u \mapsto 0$) and likewise t^u instead of t^1 . Syntactically, $t^{\neg u}$ and t^u are just new, distinct existential variables.

Since a single expansion doubles the formula size in the worst case, the complete expansion of a PCNF can blow up exponentially. In the worst case, an existential in the scope of n universals will require 2^n duplicate copies. Keeping track of all these duplicates requires annotations that are assignments to *sets* of the preceding universal variables.

In the basic theoretical model $\forall\text{Exp}+\text{Res}$ [27], each axiom clause is immediately annotated with a *fixed, complete* assignment to the universal variables. The proof then proceeds exactly as a propositional resolution proof, with clauses in fully annotated variables. In short, $\forall\text{Exp}+\text{Res}$ is propositional resolution on the conjuncts of a PCNF's complete expansion.

IR-calc, defined in [7], improves on this approach by working instead with *partial* assignments. In addition to resolution, the calculus is equipped with an *instantiation* rule by which partial annotations are grown throughout the course of the proof. To facilitate instantiation, the \circ operator describes how partial assignments are combined. Formally, for each PCNF F , we define $\text{ann}(F)$ to be the set of partial assignments to $\text{vars}_{\forall}(F)$. Then for each $\tau, \sigma \in \text{ann}(F)$, we define $\tau \circ \sigma := \tau \cup \{l \in \sigma \mid \neg l \notin \tau\}$.

The rules of both systems are given in Figure 1. Note that we write annotations as literal strings (e.g. $u_1\neg u_3\neg u_6u_7$) rather than as sets.

Restrictions. This paper makes frequent use of restrictions of PCNFs and IR-calc refutations, operations that derive from their counterparts in propositional logic. Let π be an IR-calc refutation of a PCNF F .

As we will see, the purpose of restricting π by an assignment ρ is to obtain a refutation of the restricted formula $F[\rho]$. Naturally, one applies the assignment to the refutation and simplifies the result, eliminating all satisfied clauses in the process. The procedure differs depending on the quantification of the assigned variable.

For an *existential literal* l , the restricted refutation $\pi[l]$ is obtained as follows. First, remove all clauses containing a literal of the form l^τ for some $\tau \in \text{ann}(F)$, and from the remaining clauses remove all literals of the form $\neg l^\tau$ for some $\tau \in \text{ann}(F)$. Then $\pi[l]$ is the subderivation of the first occurrence of the empty clause in the resulting sequence.⁴

For a *universal literal* l that is *unopposed* in π (meaning that $\neg l$ does not appear in the annotations), the restricted derivation $\pi[l]$ is obtained from π simply by removing l from the annotations. We need only define restriction for unopposed universal literals.

Finally, for restriction by a partial assignment $\rho := \{l_1, \dots, l_n\}$ with $\text{var}(l_i)$ left of $\text{var}(l_{i+1})$ for each $i \in [n-1]$, we define $\pi[\rho] := \pi_n$, where $\pi_0 := \pi$ and $\pi_i := \pi_{i-1}[l_i]$ for each $i \in [n]$, provided that each intermediate restriction is defined.

Restrictions of IR-calc refutations are central to strategy extraction, which rests upon the following two propositions. The first implies that first block universal literals are always unopposed. The second states that IR-calc refutations are closed under restrictions.

► **Proposition 1.** *Let π be an IR-calc derivation from a PCNF F whose leftmost block U is universal. There exists a function f such that, for each clause C in π , (a) for each annotation τ in C , the projection of τ to U is $f(C)$, and (b) $f(C') \subseteq f(C)$ for each C' in the subderivation of C .*

⁴ That such a clause and its subderivation remain is proved as part of Proposition 2. We note that this subderivation may include weakening steps – the addition of arbitrary literals to a clause – but such steps are easily erased from a refutation.

$\frac{}{\{l^{\tau(l)} \mid l \in C_{\exists}\}} \text{ [axiom}(C, \tau)]$	<ul style="list-style-type: none"> ■ C is a clause in the matrix of F ■ τ is a total assignment to $\text{vars}_{\forall}(F)$ falsifying C_{\forall} ■ $\tau(l)$ is the projection of τ to the universal variables left of $\text{var}(l)$
$\frac{C_1 \cup \{x^{\tau}\} \quad C_2 \cup \{\neg x^{\tau}\}}{C_1 \cup C_2} \text{ [res}(C_1, C_2, x^{\tau})]$	
$\frac{}{\{l^{\tau(l)} \mid l \in C_{\exists}\}} \text{ [axiom}(C)]$	<ul style="list-style-type: none"> ■ C is a clause in the matrix of F. ■ τ is the smallest assignment falsifying C_{\forall} ■ $\tau(l)$ is the projection of τ to the universal variables left of $\text{var}(l)$
$\frac{C}{\{l^{\sigma \circ \tau(l)} \mid l \in C\}} \text{ [inst}(C, \tau)]$	<ul style="list-style-type: none"> ■ τ is a partial assignment to $\text{vars}_{\forall}(F)$. ■ $\tau(l)$ is the projection of τ to the universal variables left of $\text{var}(l)$.
$\frac{C_1 \cup \{x^{\tau}\} \quad C_2 \cup \{\neg x^{\tau}\}}{C_1 \cup C_2} \text{ [res}(C_1, C_2, x^{\tau})]$	

■ **Figure 1** The rules of $\forall\text{Exp+Res}$ [27] (top) and IR-calc [7] (bottom). Note that $F = \mathcal{Q} \cdot \phi$ is the input PCNF.

► **Proposition 2** ([7]). *Let π be an IR-calc refutation of a PCNF F and let l be a literal with $\text{var}(l) \in \text{vars}(F)$. Then $\pi[l]$ is an IR-calc refutation of $F[l]$ if (a) l is existential, or (b) l is universal and unopposed in π .*

Strategy extraction. Strategy extraction is a prevalent paradigm in QBF proof complexity [23, 6, 1, 37], and has already been studied for IR-calc [7]. In summary, there exists an algorithm that takes a refutation and returns a countermodel (the *extracted strategy*).

Starting with an IR-calc refutation π of a PCNF $F := \exists X_1 \forall U_1 \cdots \exists X_n \forall U_n \exists X_{n+1} \cdot \phi$, we build a winning \forall -strategy S , viewing F as a game of n rounds. In round one, the \exists -player chooses some total assignment α_1 to X_1 , and we collect the \forall -player's response β_1 simply by negating the U_1 literals appearing in the annotations of $\pi[\alpha_1]$. By Proposition 1, all such literals are unopposed, so β_1 is indeed an assignment. Any absent variables are assigned to 0, extending β_1 to a total assignment to U_1 . By Proposition 2, $\pi[\alpha_1 \cup \beta_1]$ is a refutation of $\exists X_2 \forall U_2 \cdots \exists X_n \forall U_n \exists X_{n+1} \cdot \phi[\alpha_1 \cup \beta_1]$, i.e. of $F[\alpha_1 \cup \beta_1]$, so we repeat the process to obtain the \forall -player's response for the next round.

In this way, one obtains a full response $S(\alpha)$ to each total assignment α to the existentials, such that $\alpha \cup S(\alpha)$ falsifies ϕ . Moreover, $S(\alpha)$ and $S(\alpha')$ must agree up to block U_i if α and α' agree up to block X_i . This serves as a proof sketch for the following proposition.

► **Proposition 3** ([7]). *If π is an IR-calc refutation of a PCNF F , then the extracted strategy for π is a winning \forall -strategy for F .*

4 A technique for bounded formula families

In this section, we present our results for bounded PCNF families, culminating in a theorem with obvious practical relevance: in the absence of propositional hardness, separation of IR-calc from $\forall\text{Exp}+\text{Res}$ is due to an unbounded formula family. We employ the following (unbounded) formulas from [27] (equation (2) in Section 6) as a running example.

► **Definition 4** ([27]). Let \mathcal{J} be the PCNF family defined by $\mathcal{J}(n) := \mathcal{Q}_{\mathcal{J}}(n) \cdot \phi_{\mathcal{J}}(n)$, where

$$\begin{aligned} \mathcal{Q}_{\mathcal{J}}(n) &:= \mathcal{Q}_1 \cdots \mathcal{Q}_n, \quad \text{where } \mathcal{Q}_i := \exists x_i \forall u_i \exists t_{2i-1} t_{2i} \text{ for each } i \in [n], \\ \phi_{\mathcal{J}}(n) &:= \{(\neg t_1, \dots, \neg t_{2n})\} \cup \bigcup_{i=1}^n \{(\neg x_i, t_{2i-1}), (\neg u_i, t_{2i-1}), (x_i, t_{2i}), (u_i, t_{2i})\}. \end{aligned}$$

The authors of [27] showed that this PCNF family separates IR-calc from $\forall\text{Exp}+\text{Res}$.⁵ In light of our results, the fact that \mathcal{J} is an unbounded family is not coincidental; indeed, we show that coercing \mathcal{J} into a bounded family by variable reordering yields a PCNF family that is hard even for IR-calc.

4.1 IR-calc lower bounds by strategy size

Our principal insight for bounded families is that proof-size lower-bounds can be obtained by appealing to a natural and semantically-grounded measure we call *strategy size*. The strategy size of a false PCNF is the minimum number of responses in a winning strategy for the \forall -player. We recall that a winning strategy, or countermodel, is represented formally as a function (cf. Section 2) whose range is the set of responses. Strategy size is therefore defined as the minimum cardinality of the range of a countermodel.

► **Definition 5** (strategy size). The *strategy size* of a false QBF F is the minimum cardinality of the range of a countermodel for F . The strategy size of a PCNF family \mathcal{F} is the function $\nabla_{\mathcal{F}} : \mathbb{N} \rightarrow \mathbb{N}$ mapping n to the strategy size of $\mathcal{F}(n)$.

► **Example 6.** For each $n \in \mathbb{N}$, the strategy size of $\mathcal{J}(n)$ is 2^n , so the strategy size of \mathcal{J} is $\nabla_{\mathcal{J}}(n) = 2^n$. To see this, observe that the only way for the \forall -player to win the evaluation game by force is to set u_i not equal to x_i for each $i \in [n]$. This necessitates at least 2^n distinct responses. On the other hand, the range of a countermodel for $\mathcal{J}(n)$ is at most 2^n , since there are exactly n universal variables.

Now, recall that $\forall\text{Exp}+\text{Res}$ works by applying propositional resolution to the clauses in the complete universal expansion of a PCNF. In fact, the conjuncts of the full expansion are exactly the allowable axiom clauses. An interesting question arises: how many such clauses must be introduced as axioms? It is perhaps not too difficult to see that the smallest unsatisfiable subset of the allowable axioms has cardinality not less than strategy size – this holds because the initial instantiations, one per axiom, encompass a complete set of responses for a winning strategy. Hence strategy size is an absolute proof-size lower-bound in $\forall\text{Exp}+\text{Res}$.

⁵ In fact, the authors separated Q-Res from $\forall\text{Exp}+\text{Res}$; since IR-calc p -simulates Q-Res [7], the result stated in the text is an immediate corollary.

► **Theorem 7.** *A PCNF family \mathcal{F} requires $\forall\text{Exp}+\text{Res}$ refutations of size $\nabla_{\mathcal{F}}(n)$.*

The hardness of \mathcal{J} in $\forall\text{Exp}+\text{Res}$ is an immediate corollary to Theorem 7. Moreover, the fact that \mathcal{J} has short IR-calc refutations implies that Theorem 7 does not lift to IR-calc. As we will see, the crux of this counterexample is that \mathcal{J} is unbounded. We can in fact use strategy size as the basis for a lower-bound technique in IR-calc if we restrict our attention to bounded families. We introduce a technique based on counting annotations in the refutation. (Refutation size is clearly greater than the number of distinct annotations.) Of particular importance is the *final annotation*, the annotation to the final pivot.

► **Definition 8.** Let π be an IR-calc refutation, and let x^τ be the unique Boolean variable for which the empty clause is derived in π by resolution over the pivot variable x^τ . Then τ is the *final annotation* of π .

Now, if we dig into the details of the strategy extraction paradigm, we unearth a useful corollary to Proposition 1 from Section 3: Given a refutation of a PCNF whose first block U is universal, *all* the U -literals appearing in the annotations of π occur in the final annotation. This fact is crucial in the proof of the following theorem.

► **Theorem 9.** *A k -bounded PCNF family \mathcal{F} requires IR-calc refutations of size $\sqrt[k]{\nabla_{\mathcal{F}}(n)}$.*

Proof sketch. Let \mathcal{F} be a k -bounded PCNF family. We apply the pigeonhole principle multiplicatively to deduce the following: for any countermodel S for $\mathcal{F}(n)$, there exists some $i \in [k]$ for which the assignments to the i^{th} universal block U_i number at least $\lfloor \sqrt[k]{\nabla_{\mathcal{F}}(n)} \rfloor$. By Proposition 1 and the definition of strategy extraction, each such partial response appears as the projection to U_i of the final annotation of $\pi[\alpha]$, extended by zeros to a total assignment to U_i , for some existential assignment α . By the definition of restriction, each such final annotation is in fact the projection to U_i of an annotation in the original refutation. It follows that π contains at least $\sqrt[k]{\nabla_{\mathcal{F}}(n)}$ distinct annotations. ◀

We illustrate the effectiveness of Theorem 9 by proving that natural Σ_3 versions of \mathcal{J} are hard even in IR-calc. We transform \mathcal{J} into a bounded family \mathcal{J}' by reordering the quantifier prefix, while preserving the strategy size.

► **Definition 10.** Let \mathcal{J}' be the PCNF family defined by $\mathcal{J}'(n) := \mathcal{Q}_{\mathcal{J}'}(n) \cdot \phi_{\mathcal{J}}(n)$, where $\mathcal{Q}_{\mathcal{J}'}(n) := \exists x_1 \cdots x_n \forall u_1 \cdots u_n \exists t_1 \cdots t_{2n}$.

To see that exponential strategy size is preserved in \mathcal{J}' , observe that the *unique* winning strategy for the \forall -player is to play u_i not equal to x_i for each $i \in [n]$. Since \mathcal{J}' is a 1-bounded PCNF family with $\nabla_{\mathcal{J}'}(n) = 2^n$, Theorem 9 yields an exponential proof-size lower bound.

► **Theorem 11.** *The PCNF family \mathcal{J}' requires exponential-size IR-calc refutations.*

4.2 A new class of bounded hard families

Applying Theorem 7, we present a blueprint for a PCNF family with large strategy size, yielding a whole class of bounded families that are hard for IR-calc. For any CNF ϕ and clause C , let us write $\phi \otimes C := \{C' \cup C : C' \in \phi\}$ for the CNF obtained by augmenting each clause in ϕ with the literals of C . Consider the following construction, which is inspired by the random QBFs in [5].

► **Definition 12.** Let $k : \mathbb{N} \rightarrow \mathbb{N}$ be a function satisfying $k(n) = n^{\Omega(1)}$. Further, for each $n \in \mathbb{N}$, let $\{C_1^n, \dots, C_{k(n)}^n\}$ be a minimally unsatisfiable CNF over variables T^n , and, for

$i \in [k(n)]$, let $\exists X_i^n \forall U_i^n \cdot \phi_i^n$ be variable-disjoint false PCNFs with strategy size greater than 1. Then the PCNF family \mathcal{P} defined by $\mathcal{P}(n) := \mathcal{Q}_{\mathcal{P}}(n) \cdot \phi_{\mathcal{P}}(n)$ is a *linear product*, where

$$\mathcal{Q}_{\mathcal{P}}(n) := \exists X_1^n \cdots X_{k(n)}^n \forall U_1^n \cdots U_{k(n)}^n \exists T^n, \quad \text{and} \quad \phi_{\mathcal{P}}(n) := \bigcup_{i=1}^{k(n)} (\phi_i^n \otimes C_i^n).$$

The intuition behind the construction of a linear product is this: to win the evaluation game on $\mathcal{P}(n)$, the \forall -player must win each ‘subgame’ $\exists X_i^n \forall U_i^n \cdot \phi_i^n$ in order to leave each clause C_i^n on the board. The non-trivial strategy size of the subgames causes the overall strategy size to blow up exponentially.

► **Lemma 13.** *Let \mathcal{P} be a linear product. Then $\nabla_{\mathcal{P}}(n) = \exp(n^{\Omega(1)})$.*

Proof sketch. Let \mathcal{P} be defined as in Definition 12. The only winning approach for the \forall -player – to reduce each CNF $\phi_i^n \otimes C_i^n$ to the clause C_i^n – encompasses winning strategies for each PCNF $F_i^n := \exists X_i^n \forall U_i^n \cdot \phi_i^n$. Since the F_i^n are pairwise variable disjoint, and therefore semantically independent from one another, one may deduce that the strategy size of $\mathcal{P}(n)$ is at least the product of the individual strategy sizes of the F_i^n . Hence the strategy size of $\mathcal{P}(n)$ is at least $2^{k(n)}$. It follows that $\nabla_{\mathcal{P}}(n) = \exp(n^{\Omega(1)})$. ◀

Since a linear product is 1-bounded, applying Theorem 9 yields an IR-calc lower bound.

► **Theorem 14.** *Any linear product requires superpolynomial-size IR-calc refutations.*

4.3 Separations and propositional hardness

As a further application of Theorem 9, we prove an interesting theorem with clear relevance to QBF solving.

First, consider a PCNF $F := \mathcal{Q} \cdot \phi$ that has a countermodel S . The elements of the range of S are all total assignments to the universal variables of F , and it should be clear that instantiating each clause in ϕ by each element of $\text{rng}(S)$ gives rise to an unsatisfiable set of clauses in annotated variables. Let us denote this set $\psi := \text{inst}(\phi, \text{rng}(S))$, and say that F *expands* to ψ . Further, let us say that a PCNF family \mathcal{F} *expands to a CNF family* f if and only if $\mathcal{F}(n)$ expands to $f(n)$, for each natural number n .

An immediate corollary to Theorem 9 is that any bounded PCNF family with polynomial-size IR-calc refutations must have polynomial strategy size; hence any such family expands to a CNF family of polynomial-size. This observation leads to the following theorem.

► **Theorem 15.** *Let \mathcal{F} be a bounded PCNF family separating IR-calc from $\forall\text{Exp}+\text{Res}$. Then \mathcal{F} expands to a polynomial-size CNF family requiring superpolynomial-size resolution refutations.*

Proof. Let $\mathcal{F}(n) := \mathcal{Q}_{\mathcal{F}}(n) \cdot \phi_{\mathcal{F}}(n)$. Since \mathcal{F} has polynomial-size IR-calc refutations, $\nabla_{\mathcal{F}}$ is polynomially bounded, by Theorem 9. Hence, there exist countermodels $S(n)$ for $\mathcal{F}(n)$ for which $|\text{rng}(S(n))|$ is polynomially bounded, and the number of literals in the CNF $f(n) := \text{inst}(\phi_{\mathcal{F}}(n), \text{rng}(S(n)))$ is polynomially bounded. Therefore, the function $f : n \mapsto f(n)$ is a CNF family. Observe that every clause in $f(n)$ may be downloaded as an axiom in a $\forall\text{Exp}+\text{Res}$ derivation from $\mathcal{F}(n)$, and that \mathcal{F} requires superpolynomial-size $\forall\text{Exp}+\text{Res}$ refutations. It follows that polynomial-size resolution refutations of f do not exist. ◀

The import of Theorem 15. The essence of the result can perhaps be captured as follows: if the lower bound is not derived from propositional hardness, a separation of IR-calc from $\forall\text{Exp}+\text{Res}$ must be due to an unbounded family of PCNFs. In the spirit of [13], it is natural to label this kind of separation as *genuine*, since the $\forall\text{Exp}+\text{Res}$ lower bound is due to a large expansion, rather than a large number of resolution steps.

Moreover, since IR-calc simulates the well-studied QBF proof system Q-resolution (Q-Res [29]), Theorem 15 holds when IR-calc is replaced by Q-Res. Thus, a ‘genuine separation’ of Q-Res from $\forall\text{Exp}+\text{Res}$ requires an unbounded PCNF family.

As the theoretical models of $\forall\text{Exp}+\text{Res}$ and Q-Res underpin the two major paradigms in QBF practice – expansion-based solving [27] and QCDCL [24] – Theorem 15 has a clear practical import. A typical QBF expansion solver will use a SAT solver as an oracle, assuming that SAT calls are inexpensive. According to Theorem 15, if bounded formulas separating Q-Res from $\forall\text{Exp}+\text{Res}$ exist, they may still be easy for an expansion-based algorithm given access to a SAT oracle, and hence offer no insight into how to improve the algorithm.

5 The Weight Theorem: conquering unbounded families

In this section, we extend the lower-bound technique to cover unbounded PCNF families. Since the technical details are quite demanding, the proof of the main theorem is preceded by a brief overview of the technique. We conclude with an application: a very short proof of hardness for what is arguably the most famous PCNF family.

5.1 Outline of the technique

We invite the reader to consider once again the example PCNF family \mathcal{J} (Definition 4) from the previous section. That family has exponential strategy size and linear-size IR-calc refutations. This illustrates that the responses from the extracted strategy do not always appear as annotations in an IR-calc refutation. However, with careful analysis, we can show that *certain portions of the responses always will*.


Our method makes use of a particular class of assignments: assignments to all existentials except those in the final block. We call such assignments *restrictors*.

► **Definition 16.** Let Z be the rightmost block of a PCNF F . Any total assignment to the variables $\text{vars}_{\exists}(F) \setminus Z$ is a *restrictor* of F .

Now, take a refutation π of a PCNF F and select a restrictor α . First, apply strategy extraction to π , and consider the response $S(\alpha)$ in the extracted strategy S . Then compare this response with the final annotation τ_{α} of the restricted refutation $\pi[\alpha]$. On the one hand, the definition of strategy extraction ensures that the literals in τ_{α} are a subset of the response $S(\alpha)$. We combine this with a proof that certain *critical variables* must occur in τ_{α} . As a result, we obtain a subset of the response to α , called the *critical response*, that must be contained in the annotation τ_{α} . This is the central observation of our method, depicted in Figure 2. Note that τ_{α} occurs also as an annotation in the original refutation.

Proof of the Weight Theorem. The *critical variables* of a PCNF are those universals that appear in every subset of the matrix that is false under the quantifier prefix. The projection of a restrictor’s response to its critical variables is termed the *critical response*.

► **Definition 17.** Let S be a countermodel for a false PCNF $F := \mathcal{Q} \cdot \phi$, and let α be a restrictor of F . The *critical variables* of F are the universal variables appearing in every CNF ϕ' for which (a) $\phi' \subseteq \phi$ and (b) $\mathcal{Q} \cdot \phi'$ is false. The *critical response* to α with respect to S and F is the projection of $S(\alpha)$ to the critical variables of $F[\alpha]$.



■ **Figure 2** Depiction of the central observation of our lower-bound technique. The final statement is proved in Lemma 18.

The key notion in our argument is the following relationship between the critical response to a restrictor and the final annotation of the restricted refutation.

► **Lemma 18.** *Let S be the extracted strategy for a IR-calc refutation π of a PCNF F . Then, for each restrictor α of F , the final annotation of $\pi[\alpha]$ contains the critical response to α with respect to S and F .*

Proof sketch. The lemma is vacuously true if F contains no universal variables, so we assume otherwise. Let α be a restrictor of F , and let $\tau[\alpha]$ be the final annotation of $\pi[\alpha]$. In combination with Proposition 1, the fact that $F[\alpha]$ has a Π_2 prefix is enough to deduce that $\text{vars}(\tau_\alpha)$ contains the critical variables of $F[\alpha]$. Hence, the lemma follows from the claim that $\tau_\alpha \subseteq S(\alpha)$, a fairly straightforward consequence of the definition of strategy extraction, and Propositions 1 and 2. ◀

Since the final annotation of $\pi[\alpha]$ appears also in π , any k mutually inconsistent critical responses give rise to k distinct annotations in π . For that reason, given a winning \forall -strategy S , we define the *critical response graph* that has a vertex for each critical response and an edge between each inconsistent pair. Hence, as we prove subsequently, the number of distinct annotations in a refutation is lower bounded by the clique number of the critical response graph for the extracted strategy. The clique number of a graph G is denoted $\omega(G)$.

► **Definition 19.** Let S be a countermodel for a PCNF F . The *critical response graph* of S with respect to F is the undirected graph $G(S, F)$ defined as follows: (a) For each restrictor α of F , $G(S, F)$ has a vertex labelled with the critical response to α with respect to S and F ; (b) $G(S, F)$ has an edge between two vertices if and only if their labels are inconsistent.

► **Lemma 20.** *Let S be the strategy extracted from an IR-calc refutation π of a PCNF F . Then there are at least $\omega(G(S, F))$ distinct annotations in π .*

Proof. Let $k := \omega(G(S, F))$, and let $\alpha_1, \dots, \alpha_k$ be restrictors of F whose critical responses (with respect to S and F) are pairwise inconsistent. For each $i \in [k]$, the final annotation τ_{α_i}

12:12 Genuine Lower Bounds for QBF Expansion

of $\pi[\alpha_i]$ contains the critical response to α_i , by Lemma 18, and τ_{α_i} appears as an annotation in π (an existential restriction of π preserves any annotation that is not deleted). Hence, for each $i, j \in [k]$ with $i \neq j$, τ_{α_i} and τ_{α_j} are distinct annotations appearing in π . ◀

An IR-calc proof is at least as large as the number of distinct annotations; hence, the minimal clique number of a critical response graph for a countermodel yields a refutation-size lower bound. This motivates the following definition, in which we define the *weight* of a PCNF F , denoted $\mu(F)$, to be equal to this minimal clique number.

► **Definition 21.** The *weight* $\mu(F)$ of a false PCNF F is the minimum value of $\omega(G(S, F))$ over the countermodels S of F .

The main result of this section, the Weight Theorem, is almost immediate from Lemma 20.

► **Theorem 22 (Weight Theorem).** *The size of any IR-calc refutation of a PCNF F is at least the weight of F .*

Proof. Let S be the strategy extracted from a refutation π of F . Since S is a winning \forall -strategy by Proposition 3, the weight of F is at most $\omega(G(S, F))$. By Lemma 20, at least $\omega(G(S, F))$ distinct annotations, and at least as many distinct literals, appear in π . ◀

5.2 Application to the formulas of Kleine Büning et al.

The final application of our framework is to the familiar QBFs introduced in [29] which occupy a central place in the QBF proof complexity literature (e.g. [21, 8, 2, 34]; the original formulas from [29] are called Φ_t and appear there in the proof of Theorem 3.2). We state the formulas and then prove that they have exponential weight. The IR-calc lower bound follows immediately, by the Weight Theorem (Theorem 22).

► **Definition 23 ([29]).** Let \mathcal{K} be the PCNF family defined by $\mathcal{K}(n) := \mathcal{Q}_{\mathcal{K}}(n) \cdot \phi_{\mathcal{K}}(n)$, where

$$\begin{aligned} \mathcal{Q}_{\mathcal{K}}(n) &:= \exists x_1 y_1 \forall u_1 \cdots \exists x_n y_n \forall u_n \exists t_1 \cdots t_n, \\ \phi_{\mathcal{K}}(n) &:= \{(\neg x_1, \neg y_1), (x_n, u_n, \neg t_1, \dots, \neg t_n), (y_n, \neg u_n, \neg t_1, \dots, \neg t_n)\} \\ &\quad \bigcup_{i=1}^{n-1} \{(x_i, u_i, \neg x_{i+1}, \neg y_{i+1}), (y_i, \neg u_i, \neg x_{i+1}, \neg y_{i+1})\} \\ &\quad \bigcup_{i=1}^n \{(u_i, t_i), (\neg u_i, t_i)\}. \end{aligned}$$

► **Lemma 24.** *For each $n \in \mathbb{N}$, the weight of $\mathcal{K}(n)$ is at least 2^n .*

Proof sketch. Consider the set A of restrictors of $\mathcal{K}(n)$ that contain exactly one of $\neg x_i$ and $\neg y_i$ for each $i \in [n]$, and let $\alpha \in A$. For any countermodel S of $\mathcal{K}(n)$, the gameplay implies that $\neg u_i \in S(\alpha) \Leftrightarrow \neg x_i \in \alpha$ and $u_i \in S(\alpha) \Leftrightarrow \neg y_i \in \alpha$, for each $i \in [n]$. Moreover, it can be verified that $\text{vars}_{\forall}(\mathcal{K}(n))$ are all critical in $\mathcal{K}(n)[\alpha]$. It follows that every total assignment to the universals is the critical response to some restrictor in A . Hence, the critical response graph $G(S, \mathcal{K}(n))$ has a 2^n -clique. ◀

Applying the Weight Theorem concludes a very short proof of this historic QBF result.

► **Theorem 25 ([29, 8]).** *The family $\mathcal{K}(n)$ requires exponential-size IR-calc refutations.*

6 Conclusions

We introduced the first technique for genuine QBF lower bounds in expansion systems. As applications, we proved exponential IR-calc lower bounds for a new class of formula families, and produced greatly simplified proofs of two known hardness results. Whereas our work on unbounded families was based on restrictions up to the penultimate existential block, the technique could be explored in greater generality by considering restrictions up to the i^{th} block. We also applied the technique to prove that any bounded separation of IR-calc from $\forall\text{Exp}+\text{Res}$ is due to a non-genuine lower bound. It remains an open problem whether such a bounded separation exists.

References

- 1 Valeriy Balabanov, Jie-Hong Roland Jiang, Mikoláš Janota, and Magdalena Widl. Efficient extraction of QBF (counter)models from long-distance resolution proofs. In *Conference on Artificial Intelligence (AAAI)*, pages 3694–3701, 2015.
- 2 Valeriy Balabanov, Magdalena Widl, and Jie-Hong R. Jiang. QBF resolution systems and their proof complexities. In *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 154–169, 2014.
- 3 Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow - resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001.
- 4 Marco Benedetti and Hratch Mangassarian. QBF-based formal verification: Experience and perspectives. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)*, 5(1-4):133–191, 2008.
- 5 Olaf Beyersdorff, Joshua Blinkhorn, and Luke Hinde. Size, cost, and capacity: A semantic technique for hard random QBFs. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:35, 2017.
- 6 Olaf Beyersdorff, Ilario Bonacina, and Leroy Chew. Lower bounds: From circuits to QBF proof systems. In *ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 249–260, 2016.
- 7 Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. On unification of QBF resolution-based calculi. In *International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 81–93, 2014.
- 8 Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. Proof complexity of resolution-based QBF calculi. In *International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 30, pages 76–89, 2015.
- 9 Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Feasible interpolation for QBF resolution calculi. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 180–192, 2015.
- 10 Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Are short proofs narrow? QBF resolution is not simple. In *Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 15:1–15:14, 2016.
- 11 Olaf Beyersdorff, Leroy Chew, and Karteek Sreenivasaiah. A game characterisation of tree-like Q-resolution size. *Journal of Computer and System Sciences (in press)*, 2017.
- 12 Olaf Beyersdorff, Nicola Galesi, and Massimo Lauria. A characterization of tree-like resolution size. *Information Processing Letters*, 113(18):666–671, 2013.
- 13 Olaf Beyersdorff, Luke Hinde, and Ján Pich. Reasons for hardness in QBF proof systems. In *Conference on Foundations of Software Technology and Theoretical Computer Science (FCTTCS)*, 2017. Preprint available at ECCC, TR17-044.

- 14 Samuel R. Buss. Towards NP-P via proof complexity and search. *Ann. Pure Appl. Logic*, 163(7):906–917, 2012.
- 15 Michael Cashmore, Maria Fox, and Enrico Giunchiglia. Partially grounded planning as quantified Boolean formula. In *International Conference on Automated Planning and Scheduling (ICAPS)*, 2013.
- 16 Hubie Chen. Proof complexity modulo the polynomial hierarchy: Understanding alternation as a source of hardness. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 94:1–94:14, 2016.
- 17 Stephen A. Cook and Phuong Nguyen. *Logical Foundations of Proof Complexity*. Cambridge University Press, Cambridge, 2010.
- 18 Stephen A. Cook and Robert A. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, 1979.
- 19 William Craig. Linear reasoning. A new form of the Herbrand-Gentzen Theorem. *J. Symb. Log.*, 22(3):250–268, 1957.
- 20 Nachum Dershowitz, Ziyad Hanna, and Jacob Katz. Space-efficient bounded model checking. In *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 502–518. Springer, 2005.
- 21 Uwe Egly. On stronger calculi for QBFs. In *Theory and Applications of Satisfiability Testing (SAT)*, pages 419–434, 2016.
- 22 Uwe Egly, Martin Kronegger, Florian Lonsing, and Andreas Pfandler. Conformant planning as a case study of incremental QBF solving. *Annals of Mathematics and Artificial Intelligence*, 80(1):21–45, 2017.
- 23 Uwe Egly, Florian Lonsing, and Magdalena Widl. 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)*, pages 291–308, 2013.
- 24 Enrico Giunchiglia, Paolo Marin, and Massimo Narizzano. Reasoning with quantified Boolean formulas. In *Handbook of Satisfiability*, pages 761–780. IOS Press, 2009.
- 25 Enrico Giunchiglia, Massimo Narizzano, and Armando Tacchella. QBF reasoning on real-world instances. In *International Conference on Theory and Applications of Satisfiability Testing (SAT), online proceedings*, 2004.
- 26 Mikoláš Janota, William Klieber, Joao Marques-Silva, and Edmund M. Clarke. Solving QBF with counterexample guided refinement. *Journal of Artificial Intelligence*, 234:1–25, 2016.
- 27 Mikoláš Janota and João Marques-Silva. Expansion-based QBF solving versus Q-resolution. *Theoretical Computer Science*, 577:25–42, 2015.
- 28 Charles Jordan and Lukasz Kaiser. Experiments with reduction finding. In *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 192–207, 2013.
- 29 Hans Kleine Büning, Marek Karpinski, and Andreas Flögel. Resolution for quantified Boolean formulas. *Information and Computation*, 117(1):12–18, 1995.
- 30 Roman Kontchakov, Luca Pulina, Ulrike Sattler, Thomas Schneider, Petra Selmer, Frank Wolter, and Michael Zakharyashev. Minimal module extraction from DL-lite ontologies using QBF solvers. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 836–841. AAAI Press, 2009.
- 31 Jan Krajíček. Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. *Journal of Symbolic Logic*, 62(2):457–486, 1997.
- 32 Jan Krajíček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*, volume 60 of *Encyclopedia of Mathematics and Its Applications*. Cambridge University Press, Cambridge, 1995.

- 33 Andrew C. Ling, Deshanand P. Singh, and Stephen Dean Brown. FPGA logic synthesis using quantified boolean satisfiability. In *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 444–450, 2005.
- 34 Florian Lonsing, Uwe Egly, and Martina Seidl. Q-resolution with generalized axioms. In *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 435–452, 2016.
- 35 Hratch Mangassarian, Andreas G. Veneris, and Marco Benedetti. Robust QBF encodings for sequential circuits with applications to verification, debug, and test. *IEEE Transactions on Computers*, 59(7):981–994, 2010.
- 36 Hratch Mangassarian, Andreas G. Veneris, Sean Safarpour, Marco Benedetti, and Duncan Exon Smith. A performance-driven QBF-based iterative logic array representation with applications to verification, debug and test. In *International Conference on Computer-Aided Design (ICCAD)*, pages 240–245, 2007.
- 37 Tomáš Peitl, Friedrich Slivovsky, and Stefan Szeider. Long distance Q-resolution with dependency schemes. In *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 500–518, 2016.
- 38 Pavel Pudlák and Russell Impagliazzo. A lower bound for DLL algorithms for sat (preliminary version). In *Symposium on Discrete Algorithms*, pages 128–136, 2000.
- 39 Jussi Rintanen. Asymptotically optimal encodings of conformant planning in QBF. In *National Conference on Artificial Intelligence (AAAI)*, pages 1045–1050. AAAI Press, 2007.
- 40 Nathan Segerlind. The complexity of propositional proofs. *Bulletin of Symbolic Logic*, 13(4):417–481, 2007.
- 41 Stefan Staber and Roderick Bloem. Fault localization and correction with QBF. In *International Conference on Theory and Applications of Satisfiability Testing (SAT) 2007*, pages 355–368, 2007.
- 42 Larry J. Stockmeyer and Albert R. Meyer. Word problems requiring exponential time: Preliminary report. In *Annual Symposium on Theory of Computing*, pages 1–9. ACM, 1973.