

FEASIBLE INTERPOLATION FOR QBF RESOLUTION CALCULI*

OLAF BEYERSDORFF^a, LEROY CHEW^a, MEENA MAHAJAN^b, AND ANIL SHUKLA^b

^a School of Computing, University of Leeds, United Kingdom

^b The Institute of Mathematical Sciences (HBNI), Chennai, India

ABSTRACT. In sharp contrast to classical proof complexity we are currently short of lower bound techniques for QBF proof systems. In this paper we establish the feasible interpolation technique for all resolution-based QBF systems, whether modelling CDCL or expansion-based solving. This both provides the first general lower bound method for QBF proof systems as well as largely extends the scope of classical feasible interpolation. We apply our technique to obtain new exponential lower bounds to all resolution-based QBF systems for a new class of QBF formulas based on the clique problem. Finally, we show how feasible interpolation relates to the recently established lower bound method based on strategy extraction by Beyersdorff et al. [BCJ15, BBC16].

1. INTRODUCTION

The main aim in proof complexity is to understand the complexity of theorem proving. Arguably, what is even more important is to establish techniques for lower bounds, and the recent history of computational complexity speaks volumes on how difficult it is to develop general lower bound techniques. Understanding the size of proofs is important for at least two reasons. The first is its tight relation to the separation of complexity classes: NP vs. coNP for propositional proofs, and NP vs. PSPACE in the case of proof systems for quantified boolean formulas (QBF). New superpolynomial lower bounds for specific proof systems rule out specific classes of non-deterministic poly-time algorithms for problems in co-NP or PSPACE, thereby providing an orthogonal approach to the predominantly machine-oriented view of computational complexity.

The second reason to study lower bounds for proofs is the analysis of SAT and QBF solvers: powerful algorithms that efficiently solve the classically hard problems of SAT and QBF for large classes of practically relevant formulas. Modern SAT solvers routinely solve industrial instances in even millions of variables for various applications. Even though QBF

2012 ACM CCS: [Theory of computation]: Proof complexity.

Key words and phrases: Feasible interpolation, Proof complexity, QBF, Resolution.

* This work was supported by the EU Marie Curie IRSES grant CORCON, grant no. 48138 from the John Templeton Foundation, EPSRC grant EP/L024233/1, and a Doctoral Training Grant from EPSRC (2nd author).

A preliminary version of this article appeared in the proceedings of the conference ICALP'15 [BCMS15].

solving is at a much earlier state, due to its power to express problems more succinctly, QBF even applies to further fields such as formal verification or planning [Rin07, BM08, EKL14]. Each successful run of a solver on an unsatisfiable instance can be interpreted as a proof of unsatisfiability; and many modern SAT solvers based on conflict-driven clause learning (CDCL) are known to implicitly generate resolution proofs. Thus, understanding the complexity of resolution proofs helps obtain worst-case bounds for the performance of CDCL-based SAT solvers.

The picture is more complex for QBF solving, as there exist two main, yet conceptually very different paradigms: CDCL-based and expansion-based solving. A variety of QBF resolution systems have been designed to capture the power of QBF solvers based on these paradigms. The core system of these is Q-Resolution (**Q-Res**), introduced by Kleine Büning et al. [KKF95]. This has been augmented to capture ideas from CDCL solving, leading to long-distance resolution (**LD-Q-Res**) [BJ12], universal resolution (**QU-Res**) [VG12], or its combinations like **LQU⁺-Res** [BWJ14].

Powerful proof systems for expansion-based solving were developed in the form of **∃Exp+Res** [JM15], and the stronger **IR-calc** and **IRM-calc** [BCJ14]. Recent findings show that CDCL and expansion are indeed orthogonal paradigms as the underlying proof systems from the two categories are incomparable with respect to simulations [BCJ15].

Understanding which general techniques can be used to show lower bounds for proof systems is of paramount importance in proof complexity. For propositional proof systems we have a number of very effective techniques, most notably the size-width technique of Ben-Sasson and Wigderson [BSW01], deriving size from width bounds, game characterisations (e.g. [Pud00, BK14]), the approach via proof-complexity generators (cf. [Kra11]), and feasible interpolation. Feasible interpolation, first introduced by Krajíček [Kra97], is a particularly successful paradigm that transfers circuit lower bounds to size of proof lower bounds. The technique has been shown to be effective for resolution [Kra97], cutting planes [Pud97] and even strong Frege systems for modal and intuitionistic logics [Hru09]. However, feasible interpolation fails for strong propositional systems as Frege systems under plausible cryptographic and number-theoretic assumptions [KP98, BPR00, BDG⁺04].

The situation is drastically different for QBF proof systems, where we currently possess a very limited bag of techniques. In particular, the classical size-width technique of Ben-Sasson and Wigderson [BSW01] by which most resolution lower bounds are obtained drastically fails in **Q-Res** [BCMS16a]. At present we only have the recent strategy extraction technique of Beyersdorff et al. [BCJ15], which works for **Q-Res** as well as for stronger QBF Frege systems [BBC16, BP16], a game characterisation of the very weak tree-like **Q-Res** [BCS15], and ad-hoc lower bound arguments for various systems [BCJ15, KKF95]. In addition, Balabanov et al. [BWJ14] develop methods to lift some previous lower bounds from **Q-Res** to stronger systems.

We now proceed to explain the main contributions of the article.

1. A general lower bound technique. We show that the feasible interpolation technique applies to all resolution-type QBF proof systems, whether expansion or CDCL based. This provides the first truly general lower bound technique for QBF proof systems, and—at the same time—hugely extends the scope of the feasible interpolation method. (We note that in recent work [BCMS16b], this technique has also been shown to apply to a QBF version of the cutting planes proof system.)

In a nutshell, feasible interpolation works for true implications $A(\vec{p}, \vec{q}) \rightarrow B(\vec{p}, \vec{r})$ (or, equivalently, false conjunctions $A(\vec{p}, \vec{q}) \wedge \neg B(\vec{p}, \vec{r})$), which by Craig’s interpolation theorem [Cra57] possess interpolants $C(\vec{p})$ in the common variables \vec{p} . Such interpolants, even though they exist, may not be of polynomial size [Mun84]. However, it may be the case that we can always efficiently extract such interpolants from a proof of the implication in a particular proof system P , and in this case, the system P is said to admit feasible interpolation. If we know that a particular class of formulas does not admit small interpolants (either unconditional or under suitable assumptions), then there cannot exist small proofs of the formulas in the system P . Here we show that this feasible interpolation theorem holds for arbitrarily quantified formulas $A(\vec{p}, \vec{q})$ and $B(\vec{p}, \vec{r})$ above, when the common variables \vec{p} are existentially quantified before all other variables.

2. New lower bounds for QBF systems. As our second contribution we exhibit new hard formulas for QBF resolution systems. Of course, exponential lower bounds for these systems follow immediately from the known lower bounds for resolution (in these systems, refuting a totally quantified false sentence that uses only existential quantifiers degenerates to classical resolution). However, we can better understand the power of such systems to handle arbitrary QBFs if we have more examples of false QBFs that use existential and universal quantifiers in non-trivial ways and that are hard to refute in these systems. It is fair to say that we are currently quite short of hard examples: research so far has mainly concentrated on formulas of Kleine Büning et al. [KKF95] and their modifications [BCJ15, BWJ14], a principle by Janota and Marques-Silva [JM15], and a class of parity formulas recently introduced by Beyersdorff et al. [BCJ15]. This again is in sharp contrast with classical proof complexity where a wealth of different combinatorial principles as well as random formulas are known to be hard for resolution.

Our new hard formulas are QBF contradictions formalising the easy and appealing fact that a graph cannot both have and not have a k -clique. The trick is that in our formulation, each interpolant for these formulas has to solve the k -clique problem. Using our interpolation theorem together with the exponential lower bound for the monotone circuit complexity of clique [AB87], we obtain exponential lower bounds for the clique-no-clique formulas in all CDCL and expansion-based QBF resolution systems.

We remark that conceptually our clique-no-clique formulas are different from and indeed simpler than the clique-colour formulas used for the interpolation technique in classical proof systems. This is due to the more succinct expressivity of QBF. Indeed it is not clear how the clique-no-clique principle could even be formulated succinctly in propositional logic.

3. Comparison to strategy extraction. On a conceptual level, we uncover a tight relationship between feasible interpolation and strategy extraction. Strategy extraction is a very desirable property of QBF proof systems and is known to hold for the main resolution-based systems. From a refutation of a false QBF, a winning strategy for the universal player can be efficiently extracted.

Like feasible interpolation, the lower bound technique based on strategy extraction by Beyersdorff et al. [BCJ15, BBC16] also transfers circuit lower bounds to proof size bounds. However, instead of monotone circuit bounds as in the case of feasible interpolation, the strategy extraction technique imports AC^0 circuit lower bounds (or further circuit bounds for circuit classes directly corresponding to the lines in the proof system [BBC16]). Here we

show that each feasible interpolation problem can be transformed into a strategy extraction problem, where the interpolant corresponds to the winning strategy of the universal player on the first universal variable. This clarifies that indeed feasible interpolation can be viewed as a special case of strategy extraction.

Organisation of the paper. The remaining part of this article is organised as follows. In Section 2 we review the definitions and relations of relevant QBF proof systems. In Section 3 we start by recalling the overall idea for feasible interpolation and show interpolation theorems for the strongest CDCL-based system **LQU⁺-Res** as well as the strongest expansion-based proof system **IRM-calc**. This implies feasible interpolation for all QBF resolution-based systems. Further we show that all these systems even admit monotone feasible interpolation. In Section 4 we obtain the new lower bounds for the clique-no-clique formulas. Section 5 reformulates interpolation as a strategy extraction problem.

2. PRELIMINARIES

A literal is a boolean variable or its negation. We say a literal x is complementary to the literal $\neg x$ and vice versa. A *clause* is a disjunction of literals. For notational convenience, we sometimes also refer to a clause as a set of literals. The empty clause is denoted by \square , and is semantically equivalent to false. We denote true by 1 and false by 0. A formula in *conjunctive normal form* (CNF) is a conjunction of clauses. For a literal $l = x$ or $l = \neg x$, we write $\text{var}(l)$ for x and extend this notation to $\text{var}(C)$ for a clause C . Let α be any partial assignment. For a clause C , we write $C|_\alpha$ for the clause obtained after applying the partial assignment α to C . For example, applying $\alpha : x_1 \leftarrow 0$ to the clause $C \equiv (x_1 \vee x_2 \vee x_3)$ yields $C|_\alpha \equiv (x_2 \vee x_3)$, and applying $\alpha' : x_1 \leftarrow 1$ to the same clause gives $C|_{\alpha'} \equiv 1$. In the former case, we say that C evaluates to the clause $(x_2 \vee x_3)$ under the assignment α , and in the latter case, it evaluates to 1 under the assignment α' . Similarly, for a formula F , we write $F|_\alpha$ for the restriction of the formula to the partial assignment.

Quantified Boolean Formulas (QBFs) extend propositional logic with the boolean quantifiers \forall and \exists . They have the standard semantics that $\forall x.F$ is satisfied by the same truth assignments to its free variables as $F|_{x=0} \wedge F|_{x=1}$, and $\exists x.F$ as $F|_{x=0} \vee F|_{x=1}$. We assume that QBFs are fully quantified (no free variables), in *closed prenex form*, and with a CNF matrix, i.e, we consider the form $Q_1x_1 \dots Q_nx_n.\phi$, where $Q_i \in \{\exists, \forall\}$, and the formula ϕ is in CNF and is defined on the set of variables $X = \{x_1, \dots, x_n\}$. Further, we assume that complementary literals do not appear in the same clause; that is, no clause in the matrix is tautological. The propositional part ϕ is called the *matrix* and the rest the *prefix*. We abbreviate the prefix by the notation Qx . The *index* $\text{ind}(x)$ of a variable is its position in the prefix; thus $\text{ind}(x_i) = i$. When $Q_i = \exists$ ($Q_i = \forall$, respectively), we say that x_i is an existential variable (a universal variable, resp.). A literal l is said to be existential (universal) if $\text{var}(l)$ is existential (universal, resp.). For a literal l , we write $\text{ind}(l)$ for $\text{ind}(\text{var}(l))$.


Often it is useful to think of a QBF $Qx.\phi$ as a *game* between the *universal* and the *existential player*. In the i -th step of the game, player Q_i assigns a value to the variable 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. Let u be a universal variable u with index i . At the i th step of the game, when the universal player has to decide what value to assign to u , all variables with index less than i already have values assigned to them. A *strategy for u* is thus a function from the set of assignments to

the variables with index $< i$ to $\{0, 1\}$. A strategy for the universal player is a collection of strategies, one for each universal variable. A strategy is a *winning strategy* for the universal player if, using it, the universal player can win any possible game, irrespective of the strategy used by the existential player. A QBF is false iff there exists a *winning strategy* for the universal player ([GVB11], [AB09, Sec. 4.2.2], [Pap94, Chap. 19]).

Resolution-based calculi for QBF. We now give a brief overview of the main existing resolution-based calculi for QBF. For the technical proofs in this paper, full details are needed only for two systems, **LQU⁺-Res** and **IRM-calc**, both of which are included in the overview.

Recall that resolution for propositional proofs (where all variables are existential) operates by inferring clauses, starting from the clauses of the given formula (axioms), until the empty clause is derived. From clauses $C \vee x$ and $D \vee \neg x$ that have been already inferred, it can infer the clause $C \vee D$, by resolving on the variable x . Here, x is referred to as the pivot, and $C \vee D$ is the resolvent. The clauses $C \vee x$ and $D \vee \neg x$ are referred to as parents of the clause $C \vee D$ in the proof. In a representation of a proof as a graph, each clause in the proof is a node, and edges are directed from parent to child. This system can be augmented in various ways to handle QBFs with universal variables.

We start by describing the proof systems modelling *CDCL-based QBF solving*; their rules are summarized in Figure 1. The most basic and important system is *Q-resolution (Q-Res)* by Kleine Büning et al. [KKF95]. It is a resolution-like calculus that operates on QBFs in prenex form with CNF matrix. The lines in a **Q-Res** proof are clauses. In addition to the axioms, **Q-Res** comprises the resolution rule **S \exists R** and universal reduction \forall -Red (cf. Figure 1). Note that the conditions in **S \exists R** explicitly disallow inferring a tautology; this is syntactically essential for soundness.



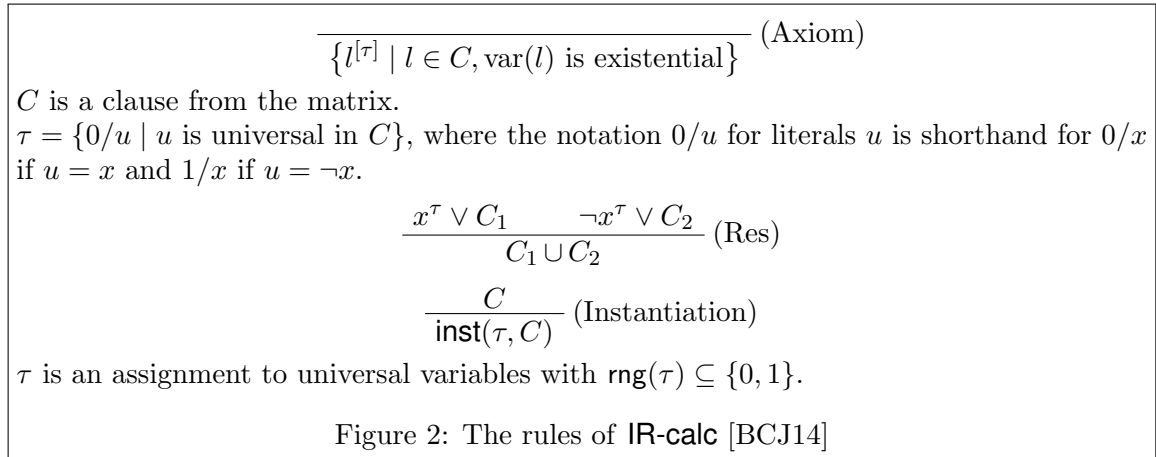
Long-distance resolution (LD-Q-Res) appears originally in the work of Zhang and Malik [ZM02] and was formalized into a calculus by Balabanov and Jiang [BJ12]. It allows resolving clauses $C \vee x$ and $D \vee \neg x$ on an existential variable x even if C and D contain complementary literals (where $C \vee D$ would be a tautology), provided the complementary literals correspond to universal variables with index greater than the index of the pivot variable x . It merges complementary literals of a universal variable u into the special literal u^* which then appears in the resolvent. We define $\text{ind}(u^*) = \text{ind}(u)$. **LD-Q-Res** uses the rules $\text{L}\exists\text{R}$, $\forall\text{-Red}$ and $\forall\text{-Red}^*$ (cf. Figure 1).

QU-resolution (QU-Res) by Van Gelder [VG12] removes the restriction from **Q-Res** that the resolved variable must be an existential variable and allows resolution of universal variables. The rules of **QU-Res** are $\text{S}\exists\text{R}$, $\text{S}\forall\text{R}$ and $\forall\text{-Red}$ (cf. Figure 1).


LQU⁺-Res by Balabanov et al. [BWJ14] extends **LD-Q-Res** by allowing short and long distance resolution pivots to be universal. However, the pivot is never a merged literal z^* . **LQU⁺-Res** uses the rules $\text{L}\exists\text{R}$, $\text{L}\forall\text{R}$, $\forall\text{-Red}$ and $\forall\text{-Red}^*$ (cf. Figure 1).

The second type of calculi models *expansion-based QBF solving*. These calculi are based on *instantiation* of universal variables: $\forall\text{Exp+Res}$ by Janota and Marques-Silva [JM15], **IR-calc**, and **IRM-calc** by Beyersdorff et al. [BCJ14]. All these calculi operate on clauses that comprise only existential variables from the original QBF, which are additionally *annotated* by a substitution to some universal variables, e.g. $\neg x^{0/u_1 1/u_2}$. For any annotated literal l^σ , the substitution σ must not make assignments to variables at a higher quantification level than l , i.e. if $u \in \text{dom}(\sigma)$, then u is universal and $\text{ind}(u) < \text{ind}(l)$. To preserve this invariant, we use the *auxiliary notation* $l^{[\sigma]}$, which for an existential literal l and an assignment σ to the universal variables filters out all assignments that are not permitted, i.e. $l^{[\sigma]} = l\{c/u \in \sigma \mid \text{ind}(u) < \text{ind}(l), c \in \{0,1\}\}$.

As annotations can be partial assignments, we use auxiliary operations of *completion* and *instantiation*. For assignments τ and μ , we write $\tau \circ \mu$ for the assignment σ defined as follows: $\sigma(x) = \tau(x)$ if $x \in \text{dom}(\tau)$, otherwise $\sigma(x) = \mu(x)$ if $x \in \text{dom}(\mu) \setminus \text{dom}(\tau)$. The operation $\tau \circ \mu$ is called *completion* because μ provides values for variables not defined in τ . The operation is associative and therefore we can omit parentheses. For an assignment τ and an annotated clause C , the function $\text{inst}(\tau, C)$ returns the annotated clause $\{l^{[\sigma \circ \tau]} \mid l^\sigma \in C\}$. The system **IR-calc** is defined in Figure 2.



The calculus **IRM-calc** further extends **IR-calc** by enabling annotations containing $*$. The rules of the calculus **IRM-calc** are presented in Figure 3. The symbol $*$ may be introduced by the merge rule, e.g. by collapsing $x^{0/u} \vee x^{1/u}$ into $x^{*/u}$.



The simulation order of QBF resolution systems is shown in Figure 4. All proof systems have been exponentially separated (cf. [BCJ15]).




Figure 4: The simulation order of QBF resolution systems. Systems on the left correspond to expansion-based solving, whereas the systems on the right are CDCL based.

We end this section with a definition that is used in later sections. It generalises the notion of weakening. For clauses containing only literals of the form $x_i, \neg x_i$ (no l^*), clause D weakens clause C if every literal in C is also present in D ; i.e. $C \subseteq D$. With merged literals, the analogous notion of weakening is as defined below.

Definition 2.1. For clauses C, D we write $C \preceq D$ if for any literal $l \in C$ we have $l \in D$ or $l^* \in D$ and for any $l^* \in C$ we have $l^* \in D$.

For annotations τ and σ we say that $\tau \preceq \sigma$ if $\text{dom}(\tau) = \text{dom}(\sigma)$ and for any $c/u \in \tau$ we have $c/u \in \sigma$ or $*/u \in \sigma$ and for any $*/u \in \tau$ we have $*/u \in \sigma$. If C, D are annotated clauses, we write $C \preceq D$ if there is an injective function $f : C \hookrightarrow D$ such that for all $l^\tau \in C$ we have $f(l^\tau) = l^\sigma$ with $\tau \preceq \sigma$.

Note: the requirement above that f is injective ensures that $x^{0/u} \vee x^{1/v} \not\preceq x^{0/u, 1/v}$.

3. FEASIBLE INTERPOLATION AND FEASIBLE MONOTONE INTERPOLATION

In this section we show that feasible interpolation and feasible monotone interpolation hold for **LQU⁺-Res** and **IRM-calc**. We adapt the technique first used by Pudlák [Pud97] to re-prove and generalise the result of Krajíček [Kra97].

3.1. The setting. Consider a false QBF \mathcal{F} of the form

$$\exists \vec{p} \mathcal{Q} \vec{q} \mathcal{Q} \vec{r} [A(\vec{p}, \vec{q}) \wedge B(\vec{p}, \vec{r})],$$

where, \vec{p} , \vec{q} , and \vec{r} are mutually disjoint sets of propositional variables, $A(\vec{p}, \vec{q})$ is a CNF formula on variables \vec{p} and \vec{q} , and $B(\vec{p}, \vec{r})$ is a CNF formula on variables \vec{p} and \vec{r} . Thus \vec{p} contains all the common variables between them. The \vec{q} and \vec{r} variables can be quantified arbitrarily, with any number of alternations between quantifiers. The QBF is equivalent to the following, not in prenex form

$$\exists \vec{p} [\mathcal{Q} \vec{q}. A(\vec{p}, \vec{q}) \wedge \mathcal{Q} \vec{r}. B(\vec{p}, \vec{r})].$$

Let \vec{a} denote an assignment to the \vec{p} variables. We denote $A(\vec{p}, \vec{q})|_{\vec{a}}$ by $A(\vec{a}, \vec{q})$ and $B(\vec{p}, \vec{r})|_{\vec{a}}$ by $B(\vec{a}, \vec{r})$.

Definition 3.1. Let \mathcal{F} be a false QBF of the form $\exists \vec{p} \mathcal{Q} \vec{q} \mathcal{Q} \vec{r}. [A(\vec{p}, \vec{q}) \wedge B(\vec{p}, \vec{r})]$. An *interpolation circuit* for \mathcal{F} is a boolean circuit G such that on every 0, 1 assignment \vec{a} for \vec{p} we have

$$\begin{aligned} G(\vec{a}) = 0 &\implies \mathcal{Q} \vec{q}. A(\vec{a}, \vec{q}) \text{ is false, and} \\ G(\vec{a}) = 1 &\implies \mathcal{Q} \vec{r}. B(\vec{a}, \vec{r}) \text{ is false.} \end{aligned}$$

We say that a QBF proof system S has *feasible interpolation* if there is an effective procedure that, given any S -proof π of a QBF \mathcal{F} of the form above, outputs an interpolation circuit for \mathcal{F} of size polynomial in the size of π .

We say that the procedure *extracts* the circuit from the proof.

We say that S has *monotone feasible interpolation* if the following holds: in the same setting as above, if \vec{p} appears only positively in $A(\vec{p}, \vec{q})$, then the interpolation circuit for \mathcal{F} extracted from π is monotone.

As our main results, we show that both **LQU⁺-Res** and **IRM-calc** have monotone feasible interpolation.

Before proving the interpolation theorems, we first outline the general idea:

Proof idea. Fix a proof system $S \in \{\mathbf{LQU}^+\text{-Res}, \mathbf{IRM-calc}\}$ and an S -proof π of \mathcal{F} . Consider the following definition of a \vec{q} -clause and an \vec{r} -clause.

Definition 3.2. We call a clause C in π a \vec{q} -clause (resp. \vec{r} -clause), if C contains only variables \vec{p}, \vec{q} (resp. \vec{p}, \vec{r}). We also call C a \vec{q} -clause (resp. \vec{r} -clause), if C contains only \vec{p} variables, but all its descendant clauses in the proof π (all clauses with a directed path to C in π) are \vec{q} (resp. \vec{r})-clauses. In the case of **IRM-calc** the variables appearing in the annotations are irrelevant and can be from either set.

From π we construct a circuit C_π with the \vec{p} -variables as inputs: for each node u with clause C_u in the proof π , associate a gate g_u (or a constant-size circuit) in the circuit C_π . Next, we inductively construct, for any assignment \vec{a} to the \vec{p} variables, another proof-like structure $\pi'(\vec{a})$. For each node u with clause C_u in the proof π , associate a clause $C'_{u, \vec{a}}$

in the structure $\pi'(\vec{a})$. Finally, we obtain $\pi''(\vec{a})$ from the structure $\pi'(\vec{a})$ by instantiating \vec{p} variables to the assignment \vec{a} (that is, $C''_{u,\vec{a}} = C'_{u,\vec{a}}|_{\vec{a}}$ for each node u) and doing some pruning, and show that $\pi''(\vec{a})$ is a valid proof in S . We then find that if $C_\pi(\vec{a}) = 0$, then $\pi''(\vec{a})$ uses only \vec{q} -clauses and thus is a refutation of $\mathcal{Q}\vec{q}.A(\vec{a}, \vec{q})$, and if $C_\pi(\vec{a}) = 1$, then $\pi''(\vec{a})$ uses only \vec{r} -clauses and thus is a refutation of $\mathcal{Q}\vec{r}.B(\vec{a}, \vec{r})$. Thus C_π is the desired interpolant circuit.

More precisely, we show by induction on the height of u in π (that is, the length of the longest path to u from a source node in π) that:

- (1) $C'_{u,\vec{a}} \preceq C_u$.
- (2) $g_u(\vec{a}) = 0 \implies C''_{u,\vec{a}}$ is a \vec{q} -clause and can be obtained from the clauses of $A(\vec{a}, \vec{q})$ alone using the rules of S .
- (3) $g_u(\vec{a}) = 1 \implies C''_{u,\vec{a}}$ is an \vec{r} -clause and can be obtained from the clauses of $B(\vec{a}, \vec{r})$ alone using the rules of S .

From the above, we have the following conclusion. Let r be the root of π . Then on any assignment \vec{a} to the \vec{p} variables we have:

- (1) $C'_{r,\vec{a}} \preceq C_r = \square$, so $C'_{r,\vec{a}} = \square$. Therefore, $C''_{r,\vec{a}} = C'_{r,\vec{a}}|_{\vec{a}} = \square$.
- (2) $g_r(\vec{a}) = 0 \implies \square$ is a \vec{q} -clause and can be obtained from the clauses of $A(\vec{a}, \vec{q})$ alone using the rules of system S . Hence by soundness of S , $\mathcal{Q}\vec{q}.A(\vec{a}, \vec{q})$ is false.
- (3) $g_r(\vec{a}) = 1 \implies \square$ is an \vec{r} -clause and can be obtained from the clauses of $B(\vec{a}, \vec{r})$ alone using the rules of system S . Hence by soundness of S , $\mathcal{Q}\vec{r}.B(\vec{a}, \vec{r})$ is false.

Thus g_r , the output gate of the circuit, computes an interpolant.

When \mathcal{F} has only existential quantification, π is a classical resolution proof, and this is exactly the interpolant computed by Pudlák's method [Pud97]. The challenge here is to construct π' and π'' appropriately when the stronger proof systems are used for general QBF, while maintaining the inductive invariants.

3.2. Interpolants from LQU⁺-Res proofs. We now implement the idea described above for LQU⁺-Res.

Theorem 3.3. LQU⁺-Res has feasible interpolation.

Proof. As mentioned in the proof idea, for an LQU⁺-Res proof π of \mathcal{F} we first describe the circuit C_π with input \vec{p} .

Construction of the circuit C_π : The DAG underlying the circuit is exactly the same as the DAG underlying the proof π . For each node u with clause C_u in π we associate a gate g_u as follows:

u is a leaf node: If $C_u \in A(\vec{p}, \vec{q})$ then g_u is a constant 0 gate. If $C_u \in B(\vec{p}, \vec{r})$ then g_u is a constant 1 gate.

u is an internal node: We distinguish four cases.

- (1) u was derived by a universal reduction step. In this case put a no-operation gate (identity gate) for g_u .

- (2) u corresponds to a resolution step with an existential variable $x \in \vec{p}$ as pivot. Nodes v and w are its two parents, i.e.

$$\frac{\overbrace{C_1 \vee x}^{\text{node } v} \quad \overbrace{C_2 \vee \neg x}^{\text{node } w}}{\underbrace{C}_{\text{node } u}}$$

In this case, put a selector gate $\text{sel}(x, g_v, g_w)$ for g_u . Here, $\text{sel}(x, a, b) = a$, when $x = 0$ and $\text{sel}(x, a, b) = b$, when $x = 1$. That is, $\text{sel}(x, a, b) = (\neg x \wedge a) \vee (x \wedge b)$. Note that all the variables in \vec{p} are existential variables without annotations (equivalently, with empty annotations).

- (3) u corresponds to a resolution step with an existential or universal variable $x \in \vec{q}$ as pivot. Put an OR gate for g_u .
- (4) u corresponds to a resolution step with an existential or universal variable $x \in \vec{r}$ as pivot. Put an AND gate for g_u .

This completes the description of the circuit C_π .

Construction of π' and π'' : Following our proof idea, we now describe, for each node u in π with clause C_u , the associated clause $C'_{u,\vec{a}}$ in $\pi'(\vec{a})$. Once $\pi'(\vec{a})$ is defined, the structure $\pi''(\vec{a})$ is obtained by instantiating \vec{p} variables by the assignment \vec{a} in each clause of $\pi'(\vec{a})$, cutting away any edge out of a node where the clause evaluates to 1, and deleting nodes which now have no path to the root node. That is, for each node u , if $C'_{u,\vec{a}}|_{\vec{a}} = 1$, then the node u is removed, and otherwise the node u survives and the associated clause $C''_{u,\vec{a}}$ is equal to $C'_{u,\vec{a}}|_{\vec{a}}$.

We show (by induction on the height of u in π) that:

- (1) $C'_{u,\vec{a}} \preceq C_u$.
- (2) $g_u(\vec{a}) = 0 \implies C''_{u,\vec{a}}$ is a \vec{q} -clause and can be obtained from the clauses of $A(\vec{a}, \vec{q})$ alone using the rules of system **LQU⁺-Res**.
- (3) $g_u(\vec{a}) = 1 \implies C''_{u,\vec{a}}$ is a \vec{r} -clause and can be obtained from the clauses of $B(\vec{a}, \vec{r})$ alone using the rules of system **LQU⁺-Res**.

As described in the proof outline, this suffices to conclude that C_π computes an interpolant. We now present the construction details.

At leaf level: Let node u be a leaf in π . Then $C'_{u,\vec{a}} = C_u$; that is, we copy the clause as it is. Trivially, we have $C'_{u,\vec{a}} \preceq C_u$. By construction of C_π , the conditions concerning $g_u(\vec{a})$ and $C''_{u,\vec{a}}$ are satisfied.

At an internal node we distinguish four cases based on the rule that was applied.

At an internal node with universal reduction: Let node u be an internal node in π corresponding to a universal reduction step on some universal literal x or x^* . Let node v be its only parent. Here we consider only the case where the universal literal is x . The case of x^* is identical. We have

$$\frac{C_v = \overbrace{D_v \vee x}^{\text{node } v}}{C_u = \underbrace{D_v}_{\text{node } u}}, \quad x \text{ is a universal literal, } \forall \text{ existential literal } l \in D_v, \text{ind}(l) < \text{ind}(x).$$

In this case, define $C'_{u,\vec{a}} = C'_{v,\vec{a}} \setminus \{x, \neg x, x^*\}$. By induction, $C'_{v,\vec{a}} \preceq C_v = D_v \vee x$. Therefore, $C'_{u,\vec{a}} = C'_{v,\vec{a}} \setminus \{x, \neg x, x^*\} \preceq D_v = C_u$.

If $g_u(\vec{a}) = 0$, then we know that $g_v(\vec{a}) = 0$ as $g_u(\vec{a}) = g_v(\vec{a})$. By the induction hypothesis, we know that $C''_{v,\vec{a}} = C'_{v,\vec{a}}|_{\vec{a}}$ is a \vec{q} -clause and can be derived using $A(\vec{a}, \vec{q})$ alone via **LQU⁺-Res**. Recall that $C'_{u,\vec{a}} = C'_{v,\vec{a}} \setminus \{x, \neg x, x^*\}$ in this case. Since \vec{a} is an assignment to the \vec{p} variables and $x \notin \vec{p}$, $C'_{u,\vec{a}}|_{\vec{a}} = C''_{u,\vec{a}}$ is also a \vec{q} -clause and can be derived using $A(\vec{a}, \vec{q})$ alone via **LQU⁺-Res**. (Either $C''_{u,\vec{a}}$ already equals $C''_{v,\vec{a}}$, or x needs to be dropped. In the latter case, the condition on $\text{ind}(x)$ is satisfied at $C''_{u,\vec{a}}$ because it is satisfied at C_v in π and $C'_{v,\vec{a}} \preceq C_v$. So we can drop x from $C''_{v,\vec{a}}$ to get $C''_{u,\vec{a}}$.)

The situation is dual for the case when $g_u(\vec{a}) = 1$; we get \vec{r} -clauses.

At an internal node with \vec{p} -resolution: Let node u in the proof π correspond to a resolution step with pivot $x \in \vec{p}$. Note that x is existential, as \vec{p} variables occur only existentially in \mathcal{F} . We have

$$\frac{C_v = \overbrace{C_1 \vee U_1 \vee x}^{\text{node } v} \quad \overbrace{C_2 \vee U_2 \vee \neg x}^{\text{node } w}}{C_u = \underbrace{C_1 \vee C_2 \vee U}_{\text{node } u}}.$$

In the assignment \vec{a} , if $x = 0$, then define $C'_{u,\vec{a}} = C'_{v,\vec{a}} \setminus \{x\}$ and if $x = 1$ then define $C'_{u,\vec{a}} = C'_{w,\vec{a}} \setminus \{\neg x\}$. By induction, we have $C'_{v,\vec{a}} \preceq C_v$ and $C'_{w,\vec{a}} \preceq C_w$. So, if $x = 0$, we have $C'_{u,\vec{a}} = C'_{v,\vec{a}} \setminus \{x\} \preceq C_1 \vee U_1 \preceq C_u$. If $x = 1$, we have $C'_{u,\vec{a}} \preceq C'_{w,\vec{a}} \setminus \{\neg x\} \preceq C_2 \vee U_2 \preceq C_u$.

In this case g_u is a selector gate. If $x = 0$ in the assignment \vec{a} , then $g_u(\vec{a}) = g_v(\vec{a})$ and $C''_{u,\vec{a}} = C''_{v,\vec{a}}$. Since the conditions concerning $g_v(\vec{a})$ and $C''_{v,\vec{a}}$ are satisfied by induction, the conditions concerning $g_u(\vec{a})$ and $C''_{u,\vec{a}}$ are satisfied as well. Similarly, if $x = 1$, then $g_u(\vec{a}) = g_w(\vec{a})$ and $C''_{u,\vec{a}} = C''_{w,\vec{a}}$, and the statements that are inductively true at w hold at u as well.

At an internal node with \vec{q} -resolution: Let node u in the proof π correspond to a resolution step with pivot $x \in \vec{q}$. Note that x may be existential or universal. We have

$$\frac{C_v = \overbrace{C_1 \vee U_1 \vee x}^{\text{node } v} \quad \overbrace{C_2 \vee U_2 \vee \neg x}^{\text{node } w}}{C_u = \underbrace{C_1 \vee C_2 \vee U}_{\text{node } u}}, \quad x \in \vec{q}.$$

If $g_v(\vec{a}) = 1$ then define $C'_{u,\vec{a}} = C'_{v,\vec{a}}$. By induction, we know that $C''_{u,\vec{a}} = C''_{v,\vec{a}}$ is an \vec{r} -clause. Since x is a \vec{q} -variable and is not instantiated by \vec{a} , it must be the case that $x \notin C'_{v,\vec{a}}$. Thus $C'_{u,\vec{a}} = C'_{v,\vec{a}} \preceq C_v \setminus \{x\} \preceq C_u$.

Else if $g_w(\vec{a}) = 1$, define $C'_{u,\vec{a}} = C'_{w,\vec{a}}$. By a similar analysis as above, $C'_{u,\vec{a}} = C'_{w,\vec{a}} \preceq C_w \setminus \{\neg x\} \preceq C_u$.

If $g_v(\vec{a}) = g_w(\vec{a}) = 0$, and if $x \notin C'_{v,\vec{a}}$, define $C'_{u,\vec{a}} = C'_{v,\vec{a}}$. Otherwise, if $\neg x \notin C'_{w,\vec{a}}$, define $C'_{u,\vec{a}} = C'_{w,\vec{a}}$. It follows from induction that $C'_{u,\vec{a}} \preceq C_u$.

Else, define $C'_{u,\vec{a}}$ to be the resolvent of $C'_{v,\vec{a}}$ and $C'_{w,\vec{a}}$ on x . By induction, we know that $C'_{v,\vec{a}} \setminus \{x\} \preceq C_1 \vee U_1$ and $C'_{w,\vec{a}} \setminus \{\neg x\} \preceq C_2 \vee U_2$. Hence $C'_{u,\vec{a}} \preceq C_1 \vee C_2 \vee U = C_u$.

We need to verify the conditions on $g_u(\vec{a})$ and $C''_{u,\vec{a}}$. The case when $g_u(\vec{a}) = 1$ is immediate, since $C''_{u,\vec{a}}$ copies a clause known by induction to be an \vec{r} -clause. So now consider

the case when $g_u(\vec{a}) = 0$. By induction, we know that both $C''_{v,\vec{a}} = C'_{v,\vec{a}}|\vec{a}$ and $C''_{w,\vec{a}} = C'_{w,\vec{a}}|\vec{a}$ are \vec{q} -clauses and can be derived using $A(\vec{a}, \vec{q})$ alone via **LQU⁺-Res**.

We have three cases. If $C'_{u,\vec{a}} = C'_{v,\vec{a}}$ or $C'_{u,\vec{a}} = C'_{w,\vec{a}}$, then by induction we are done. Otherwise, $C'_{u,\vec{a}}$ is obtained from $C'_{v,\vec{a}}$ and $C'_{w,\vec{a}}$ via a resolution step on pivot x . Since \vec{a} is an assignment to the \vec{p} variables and $x \notin \vec{p}$, $C''_{u,\vec{a}}$ can be derived from $C''_{v,\vec{a}}$ and $C''_{w,\vec{a}}$ via the same resolution step.

Note: A simple observation is that $C'_{u,\vec{a}}$ is always a subset of C_u with only one exception, which is that some special symbol u^* in C_u may be converted into u in $C'_{u,\vec{a}}$. This leads us to define the relation \preceq . Also, the resolution step in $\pi''(\vec{a})$ is applicable in **LQU⁺-Res** because (1) every mergable universal variable in $C''_{v,\vec{a}}$ and $C''_{w,\vec{a}}$ was also mergable earlier in C_v and C_w in π . (2) Every common non-mergable existential variable in $C''_{v,\vec{a}}$ and $C''_{w,\vec{a}}$ was also a non-mergable existential variable in C_v and C_w . (3) Every non-mergable universal variable in $C''_{v,\vec{a}}$ and $C''_{w,\vec{a}}$ was also a non-mergable universal pair in C_v and C_w . (4) The operations do not disturb the indices of variables, therefore if variable x satisfies the index condition in π it satisfies it in $\pi''(\vec{a})$ as well.

At an internal node with \vec{r} -resolution: Let node u in π correspond to a resolution step with pivot $x \in \vec{r}$. This is dual to the case above. \square

3.3. Interpolants from IRM-calc proofs. We now establish the interpolation theorem for the expansion-based calculi, following the same overall idea described in Section 3.1.

Theorem 3.4. IRM-calc has feasible interpolation.

Proof. This proof closely follows that of Theorem 3.3, but with several changes in the proof details. We describe the changes here.

Construction of the circuit C_π : The circuit construction is very similar to that for **LQU⁺-Res**. Leaves and resolution nodes are treated as before. Instantiation and merging nodes are treated as the universal reduction nodes were; that is, the corresponding gates are no-operation (identity) gates.

Construction of π' and π'' : As before we construct a proof-like structure $\pi'(\vec{a})$, which depends on the assignment \vec{a} to the \vec{p} variables, the proof π of \mathcal{F} , and the circuit C_π . For each node u in π , with clause C_u , we associate a clause $C'_{u,\vec{a}}$ in $\pi'(\vec{a})$, and let $C''_{u,\vec{a}}$ be the instantiation of $C'_{u,\vec{a}}$ by the assignment \vec{a} . We show (by induction on the height of u in π) that:

- (1) $C'_{u,\vec{a}} \preceq C_u$.
- (2) $g_u(\vec{a}) = 0 \implies C''_{u,\vec{a}}$ is a \vec{q} -clause and can be obtained from the clauses of $A(\vec{a}, \vec{q})$ alone using the rules of system **IRM-calc**.
- (3) $g_u(\vec{a}) = 1 \implies C''_{u,\vec{a}}$ is a \vec{r} -clause and can be obtained from the clauses of $B(\vec{a}, \vec{r})$ alone using the rules of system **IRM-calc**.

Once again, as described in the proof outline, this suffices to conclude that the circuit C_π computes an interpolant.

Recall that for annotated clauses, the meaning of \preceq is slightly different and is given in Definition 2.1.

At a leaf level: Let node u be a leaf in π . Then $C'_{u,\vec{a}} = C_u$; that is, copy the clause as it is. Trivially, $C'_{u,\vec{a}} \preceq C_u$. By construction of C_π , the conditions concerning $g_u(\vec{a})$ and $C''_{u,\vec{a}}$ are satisfied.

At an internal node with instantiation: Let node u be an internal node in π corresponding to an instantiation step by τ . And let node v be its only parent. We know $C_u = \text{inst}(\tau, C_v)$.

Suppose $l^{\sigma'} \in \text{inst}(\tau, C'_{v,\vec{a}})$. Then for some $\xi', l^{\xi'} \in C'_{v,\vec{a}}$, and $l^{\sigma'} = l^{[\xi' \circ \tau]}$; hence σ' is a subset of ξ' completed with τ . By induction we know that $C'_{v,\vec{a}} \preceq C_v$. We have an injective function $f : C'_{v,\vec{a}} \hookrightarrow C_v$ that demonstrates this. Let $f(l^{\xi'}) = l^\xi$. Hence $l^\xi \in C_v$ for some $\xi' \preceq \xi$. So $l^\sigma = l^{[\xi \circ \tau]} \in C_u$. Since the annotations introduced by instantiation match, $\sigma' \preceq \sigma$. We use this to define a function $g : \text{inst}(\tau, C'_{v,\vec{a}}) \rightarrow C_u$ where $g(l^{\sigma'}) = l^\sigma$. Now we find any l^{τ_1}, l^{τ_2} where $g(l^{\tau_1}) = g(l^{\tau_2}) = l^\tau$ and perform a merging step on l^{τ_1} and l^{τ_2} ; note that the resulting literal $l^{\tau'}$ will still satisfy $\tau' \preceq \tau$. Eventually we get a clause which we define as $\text{instmerge}(\tau, C'_{v,\vec{a}}, C_u) = C'_{u,\vec{a}}$ where this function is injective. We will use this notation to refer to this process of instantiation and then deliberate merging to get $\preceq C_u$.

Therefore $C'_{u,\vec{a}} \preceq C_u$.

If the node u is not pruned out in $\pi''(\vec{a})$, then $C''_{u,\vec{a}}$ contains no satisfied \vec{p} literals; hence neither does $C'_{v,\vec{a}}$. Therefore $C''_{u,\vec{a}}$ is derived from $C''_{v,\vec{a}}$; this is a valid step in the proof system.

Because we only use instantiation and merging or a dummy step, $C''_{u,\vec{a}}$ is a \vec{q} -clause if and only if $C''_{v,\vec{a}}$ is a \vec{q} -clause. Therefore the no-operation (identity) gate g_u gives a valid result by induction.

At an internal node with merging: Let node u be an internal node in π corresponding to a merging step. Let node v be its only parent. We have

$$\frac{C_v = D_v \vee b^\mu \vee b^\sigma}{C_u = D_v \vee b^\xi}$$

where $\text{dom}(\mu) = \text{dom}(\sigma)$ and ξ is obtained by merging the annotations μ, σ . That is, $\xi = \text{AMerge}(\mu, \sigma) = \{c/u \mid c/u \in \mu, c/u \in \sigma\} \cup \{*/u \mid c/u \in \mu, d/u \in \sigma, c \neq d\}$. Note that $\mu, \sigma \preceq \text{AMerge}(\mu, \sigma)$.

Note that from the induction hypothesis, $C'_{v,\vec{a}} \preceq C_v$, so there is an injective function $f : C'_{v,\vec{a}} \hookrightarrow C_v$. Suppose $C'_{v,\vec{a}}$ contains two distinct literals $b^{\mu'}$ and $b^{\sigma'}$ where $f(b^{\mu'}) = b^\mu$ and $f(b^{\sigma'}) = b^\sigma$. So $C'_{v,\vec{a}} = D'_v \vee b^{\mu'} \vee b^{\sigma'}$. Then let $C'_{u,\vec{a}} = D'_v \vee b^{\xi'}$, where $\xi' = \text{AMerge}(\mu', \sigma')$. Otherwise let $C'_{u,\vec{a}} = C'_{v,\vec{a}}$.

We first observe whenever we do actual merging, if $c/u \in \xi'$ then one of the following holds:

- (1) $c/u \in \sigma'$. Then $c/u \in \sigma$ or $*/u \in \sigma$, and so $c/u \in \xi$ or $*/u \in \xi$.
- (2) $c/u \in \mu'$. Then $c/u \in \mu$ or $*/u \in \mu$, and so $c/u \in \xi$ or $*/u \in \xi$.
- (3) $e/u \in \mu', d/u \in \sigma', e \neq d$, in which case $*/u \in \xi$.

Since all other annotated literals are unaffected, $C'_{u,\vec{a}} \preceq C_u$. We never merge \vec{p} literals as they have no annotations, so if $C''_{u,\vec{a}}$ is not pruned away, then $C''_{u,\vec{a}}$ is derived from $C''_{v,\vec{a}}$ via merging.

In case we do not merge, there might be some $b^{\sigma'} \in C'_{v,\vec{a}}$ with $\sigma' \preceq \sigma$, which is not removed by merging. However $\sigma' \preceq \sigma \preceq \xi$, so $C'_{u,\vec{a}} = C'_{v,\vec{a}} \preceq C_u$. As $C''_{u,\vec{a}} = C''_{v,\vec{a}}$, this is a valid inference step (in fact, a dummy step).

Because we only use merging or a dummy step, $C''_{u,\vec{a}}$ is a \vec{q} -clause if and only if $C''_{v,\vec{a}}$ is a \vec{q} -clause, therefore the no-operation (identity) gate g_u gives a valid result by induction.

At an internal node with \vec{p} -resolution: We do not have any annotations on \vec{p} -literals. So in this case we construct C'_u and C''_u exactly as we would for an LQU⁺-Res proof.

At an internal node with \vec{q} -resolution: When we have a resolution step between nodes v and w on a \vec{q} pivot to get node u , we have

$$\frac{C_v = x^{\tau \cup \xi} \vee D_v \quad C_w = \neg x^{\tau \cup \sigma} \vee D_w}{C_u = \text{inst}(\sigma, D_v) \cup \text{inst}(\xi, D_w)}$$

where $\text{dom}(\tau)$, $\text{dom}(\xi)$ and $\text{dom}(\sigma)$ are mutually disjoint, and $\text{rng}(\tau) \subseteq \{0, 1\}$.

In order to do dummy instantiations we will need to define a $\{0, 1\}$ version of ξ and σ . So we define $\xi' = \{c/u \mid c/u \in \xi, c \in \{0, 1\}\} \cup \{0/u \mid */u \in \xi\}$, $\sigma' = \{c/u \mid c/u \in \sigma, c \in \{0, 1\}\} \cup \{0/u \mid */u \in \sigma\}$. This gives us the desirable property that $\xi' \preceq \xi$, $\sigma' \preceq \sigma$.

Now resuming the construction of C'_u , we use information from the circuit to construct this. If $g_v(\vec{a}) = 1$, then we define $C'_{u,\vec{a}} = \text{instmerge}(\sigma', C'_{v,\vec{a}}, C_u)$. Otherwise, if $g_w(\vec{a}) = 1$, then we define $C'_{u,\vec{a}} = \text{instmerge}(\xi', C'_{w,\vec{a}}, C_u)$. In these cases, we know by the inductive claim that $C'_{u,\vec{a}}$ does not contain any \vec{q} literals. Therefore $C'_{u,\vec{a}}$ is the correct instantiation (as $\xi' \preceq \xi$, $\sigma' \preceq \sigma$) of some subset of D_v or D_w . Hence $C'_{u,\vec{a}} \preceq C_u$. Furthermore since g_u is an OR gate evaluating to 1 and since $C''_{u,\vec{a}}$, an \vec{r} -clause, can be obtained by an instantiation step, our inductive claim is true.

Now suppose $g_v(\vec{a}) = 0$ and $g_w(\vec{a}) = 0$. If there is no $x^\mu \in C'_{v,\vec{a}}$ such that $\mu \preceq \tau \cup \xi$, then define $C'_{u,\vec{a}} = \text{instmerge}(\sigma', C'_{v,\vec{a}}, C_u)$. Else, if there is no $\neg x^\mu \in C'_{w,\vec{a}}$ such that $\mu \preceq \tau \cup \sigma$, then define $C'_{u,\vec{a}} = \text{instmerge}(\xi', C'_{w,\vec{a}}, C_u)$. In these cases we know that $C'_{u,\vec{a}}$ is the correct instantiation (as $\xi' \preceq \xi$, $\sigma' \preceq \sigma$) of some subset of D_v or D_w ; hence $C'_{u,\vec{a}} \preceq C_u$. Furthermore, since g_u is an OR gate evaluating to 0, and since $C''_{u,\vec{a}}$, a \vec{q} -clause, can be obtained by an instantiation step, our inductive claim is true.

The final case is when $g_v(\vec{a}) = g_w(\vec{a}) = 0$ and $x^{\tau \cup \xi_1} \in C'_{v,\vec{a}}$ for some $\xi_1 \preceq \xi$ and $\neg x^{\tau \cup \sigma_1} \in C'_{w,\vec{a}}$ for some $\sigma_1 \preceq \sigma$. Here, because $\text{dom}(\tau)$, $\text{dom}(\xi)$ and $\text{dom}(\sigma)$ are mutually disjoint, $\text{dom}(\tau)$, $\text{dom}(\xi_1)$ and $\text{dom}(\sigma_1)$ are also mutually disjoint. Thus we can do the resolution step

$$\frac{C'_{v,\vec{a}} = x^{\tau \cup \xi_1} \vee D'_v \quad C'_{w,\vec{a}} = \neg x^{\tau \cup \sigma_1} \vee D'_w}{\text{inst}(\sigma_1, D'_v) \cup \text{inst}(\xi_1, D'_w)}.$$

Since $\text{instmerge}(\sigma_1, D'_v, C_u) \preceq \text{inst}(\sigma, D_v)$ and $\text{instmerge}(\xi_1, D'_w, C_u) \preceq \text{inst}(\xi, D_w)$, we can follow up $\text{inst}(\sigma_1, D'_v) \cup \text{inst}(\xi_1, D'_w)$ with sufficient merging steps to get a clause $C' \preceq C_u$; we define this clause to be the clause $C'_{u,\vec{a}}$. By the inductive claim, both $C''_{v,\vec{a}}$ and $C''_{w,\vec{a}}$ are \vec{q} -clauses; hence $C''_{u,\vec{a}}$ is also a \vec{q} -clause and is obtained via a valid resolution step.

At an internal node with \vec{r} -resolution: When we have a resolution step between nodes u and v on an \vec{r} -literal, this is the dual of the previous case. \square

3.4. Monotone Interpolation. To transfer known circuit lower bounds into size of proof bounds, we need a monotone version of the previous interpolation theorems, which we prove next.

Theorem 3.5. LQU⁺-Res and IRM-calc have monotone feasible interpolation.

Proof. In previous subsections, we have shown that the circuit $C_\pi(\vec{p})$ is a correct interpolant for the QBF \mathcal{F} . That is, if $C_\pi(\vec{p}) = 0$ then $\mathcal{Q}\vec{q}.A(\vec{a}, \vec{q})$ is false, and if $C_\pi(\vec{p}) = 1$ then $\mathcal{Q}\vec{r}.B(\vec{a}, \vec{r})$ is false.

However, if \vec{p} occurs only positively in $A(\vec{p}, \vec{q})$ then we construct a monotone circuit $C_\pi^{mon}(\vec{p})$ such that, on every 0, 1 assignment \vec{a} to \vec{p} we have

$$\begin{aligned} C_\pi^{mon}(\vec{a}) = 0 &\implies \mathcal{Q}\vec{q}.A(\vec{a}, \vec{q}) \text{ is false, and} \\ C_\pi^{mon}(\vec{a}) = 1 &\implies \mathcal{Q}\vec{r}.B(\vec{a}, \vec{r}) \text{ is false.} \end{aligned}$$

We obtain $C_\pi^{mon}(\vec{p})$ from $C_\pi(\vec{p})$ by replacing all selector gates $g_u = \text{sel}(x, g_v, g_w)$ by the following monotone ternary connective: $g_u = (x \vee g_v) \wedge g_w$ where nodes v and w are the parents of u in π . We also change the proof-like structure $\pi'(\vec{a})$; the construction is the same as before except that at \vec{p} -resolution nodes, the rule for fixing $C'_{u,\vec{a}}$ is also changed to reflect the monotone function used instead.

More precisely, the functions $\text{sel}(x, g_v, g_w)$ and $g_u = (x \vee g_v) \wedge g_w$ differ only when $x = 0$, $g_v(\vec{a}) = 1$, and $g_w(\vec{a}) = 0$. We set $C'_{u,\vec{a}}$ to $C'_{w,\vec{a}} \setminus \{\neg x\}$ if $x = 1$ or if $x = 0$, $g_v(\vec{a}) = 1$ and $g_w(\vec{a}) = 0$, and to $C'_{v,\vec{a}} \setminus \{x\}$ otherwise.

It suffices to verify the inductive statements in the case when $x = 0$, $g_v(\vec{a}) = 1$, and $g_w(\vec{a}) = 0$. We have to show that $C'_{u,\vec{a}} \preceq C_u$; this holds by induction. We also have to show that $C''_{u,\vec{a}}$ is a \vec{q} -clause, and can be derived using $A(\vec{a}, \vec{q})$ clauses alone via the appropriate proof system. By induction, since $g_w(\vec{a}) = 0$, we conclude that $C''_{w,\vec{a}}$ cannot contain $\neg x$: it can be derived from the clauses of $A(\vec{p}, \vec{q})$ alone, by the positivity constraint, these clauses do not contain $\neg x$, and the derivation cannot introduce literals. Hence $C''_{w,\vec{a}} = C'_{w,\vec{a}} \setminus \{\neg x\}|_{\vec{a}}$, which is $C''_{u,\vec{a}}$. \square

4. NEW EXPONENTIAL LOWER BOUNDS FOR IRM-CALC AND LQU⁺-RES

We now apply our interpolation theorems to obtain new exponential lower bounds for a new class of QBFs. The lower bound will be directly transferred from the following monotone circuit lower bound for the problem CLIQUE(n, k), asking whether a given graph with n nodes has a clique of size k .

Theorem 4.1 (Alon & Boppana [AB87]). All monotone circuits that compute CLIQUE($n, n/2$) are of exponential size.

We now build the QBF. Fix an integer n (indicating the number of vertices of the graph) and let \vec{p} be the set of variables $\{p_{uv} \mid 1 \leq u < v \leq n\}$. An assignment to \vec{p} picks a set of edges, and thus an n -vertex graph. Let \vec{q} be the set of variables $\{q_{iu} \mid i \in [\frac{n}{2}], u \in [n]\}$. We use the following clauses.

$$\begin{aligned} C_i &= q_{i1} \vee \cdots \vee q_{in} && \text{for } i \in [\frac{n}{2}] \\ D_{i,j,u} &= \neg q_{iu} \vee \neg q_{ju} && \text{for } i, j \in [\frac{n}{2}], i < j \text{ and } u \in [n] \\ E_{i,u,v} &= \neg q_{iu} \vee \neg q_{iv} && \text{for } i \in [\frac{n}{2}] \text{ and } u, v \in [n], u < v \\ F_{i,j,v} &= \neg q_{iv} \vee \neg q_{jv} \vee p_{uv} && \text{for } i, j \in [\frac{n}{2}], i \neq j \text{ and } u, v \in [n], u < v. \end{aligned}$$

(For notational convenience, we interpret the assignment $p_{uv} = 0$ to mean that the edge uv is present in the graph.)

We can now express $\text{CLIQUE}(n, n/2)$ as a polynomial-size QBF $\exists \vec{q}. A_n(\vec{p}, \vec{q})$, where

$$A_n(\vec{p}, \vec{q}) = \bigwedge_{i \in [\frac{n}{2}]} C_i \wedge \bigwedge_{i < j, u \in [n]} D_{i,j,u} \wedge \bigwedge_{i \in [\frac{n}{2}], u < v} E_{i,u,v} \wedge \bigwedge_{i < j, u \neq v} F_{i,j,u,v}.$$

Here the edge variables \vec{p} appear positively in $A_n(\vec{p}, \vec{q})$.

Likewise $\text{no-CLIQUE}(n, n/2)$ can be written as a polynomial-size QBF $\forall r_1^1 \exists r_2^2. B_n(\vec{p}, r_1^1, r_2^2)$. To construct this we use a polynomial-size circuit that checks whether the nodes specified by r_1^1 fail to form a clique in the graph given by \vec{p} . We then use existential variables r_2^2 for the gates of the circuit and can then form a CNF $B_n(\vec{p}, r_1^1, r_2^2)$ that represents the circuit computation.

Now we can form a sequence of false QBFs, stating that the graph encoded in \vec{p} both has a clique of size $n/2$ (as witnessed by \vec{q}) and likewise does not have such a clique as expressed in the B part:

$$\Phi_n = \exists \vec{p} \exists \vec{q} \forall r_1^1 \exists r_2^2. A_n(\vec{p}, \vec{q}) \wedge B_n(\vec{p}, r_1^1, r_2^2).$$

This formula has the unique interpolant $\text{CLIQUE}(n, n/2)(\vec{p})$. But since all monotone circuits for this are of exponential size by Theorem 4.1, and since monotone circuits of size polynomial in IRM-calc and $\text{LQU}^+\text{-Res}$ proofs can be extracted by Theorem 3.5, all such proofs must be of exponential size, yielding:

Theorem 4.2. The QBFs $\Phi_n(\vec{p}, \vec{q}, \vec{r})$ require exponential-size proofs in IRM-calc and $\text{LQU}^+\text{-Res}$.

Note: A slightly different, and arguably more transparent, way of encoding $\text{no-CLIQUE}(n, n/2)$ is described in [BCMS16b].

5. FEASIBLE INTERPOLATION VS. STRATEGY EXTRACTION

Recall the two player game semantics of a QBF explained in Section 2. Every false QBF has a winning strategy for the universal player, where the strategy value for each variable depends only on the values of the variables played before. We now explain the relation between strategy extraction — one of the main paradigms for QBF systems — and feasible interpolation. In Section 3 we studied QBFs of the form $\mathcal{F} = \exists \vec{p} \mathcal{Q} \vec{q} \mathcal{Q} \vec{r}. [A(\vec{p}, \vec{q}) \wedge B(\vec{p}, \vec{r})]$. If we add a common universal variable b we can change it to an equivalent QBF

$$\mathcal{F}^b = \exists \vec{p} \forall b \mathcal{Q} \vec{q} \mathcal{Q} \vec{r}. [(A(\vec{p}, \vec{q}) \vee b) \wedge (B(\vec{p}, \vec{r}) \vee \neg b)].$$

This can be expressed with a CNF matrix by inserting the literal b into each clause of $A(\vec{p}, \vec{q})$ and the literal $\neg b$ into each clause of $B(\vec{p}, \vec{r})$. Let \mathcal{F}^b also denote this equivalent QBF.

If \mathcal{F} is false, then also \mathcal{F}^b is false and thus the universal player has a winning strategy, including a strategy for $b = \sigma(\vec{p})$ for the common universal variable b .

Remark 5.1. Every winning strategy $\sigma(\vec{p})$ for b is an interpolant for \mathcal{F} , i.e., for every 0, 1 assignment \vec{a} of \vec{p} we have

$$\begin{aligned} \sigma(\vec{a}) = 0 &\implies \mathcal{Q} \vec{q}. A(\vec{a}, \vec{q}) \text{ is false, and} \\ \sigma(\vec{a}) = 1 &\implies \mathcal{Q} \vec{r}. B(\vec{a}, \vec{r}) \text{ is false.} \end{aligned}$$

Proof. Suppose not. Then there are two possibilities.

- There is some \vec{a} where $\sigma(\vec{a}) = 0$ and $\mathcal{Q}\vec{q}.A(\vec{a}, \vec{q})$ is true. Then setting $b = 0$ would satisfy $\mathcal{Q}\vec{r}.B(\vec{p}, \vec{r}) \vee \neg b$. But $\mathcal{Q}\vec{q}.A(\vec{a}, \vec{q}) \vee b$ is also satisfied. Hence this cannot be part of the winning strategy for the universal player.
- There is some \vec{a} where $\sigma(\vec{a}) = 1$ and $\mathcal{Q}\vec{r}.B(\vec{a}, \vec{r})$ is true. This is the dual of the above. \square

This observation means that every interpolation problem can be reformulated as a strategy extraction problem. We will now show that from proofs of these reformulated interpolation problems we can extract a (monotone) Boolean circuit for the winning strategy on the new universal variable b .

Strategy extraction was recently shown to be a powerful lower bound technique for QBF resolution systems. In strategy extraction, from a refutation of a false QBF, winning strategies for the universal player for all universal variables can be efficiently extracted. Devising QBFs that require computationally hard strategies then leads to lower bounds for QBF proof systems. This technique applies both to **Q-Res** [BCJ15], where AC^0 lower bounds for e.g. parity are used, as well as to much stronger QBF Frege systems where the full spectrum of current (and conjectured) lower circuit bounds is employed [BBC16]. In fact, Beyersdorff and Pich [BP16] show that lower bounds for QBF Frege systems can only come either (a) from circuit lower bounds via the strategy extraction technique or (b) from lower bounds for classical proposition Frege. This picture is reconfirmed here as well: QBF resolution lower bounds via feasible interpolation fall under paradigm (a) as they are in fact lower bounds via strategy extraction.

To show this we now prove how to extract strategies for interpolation problems, first for **LQU⁺-Res** and then for **IRM-calc**.

Theorem 5.2.

- (1) From each **LQU⁺-Res** refutation π of \mathcal{F}^b we can extract in polynomial time a boolean circuit for $\sigma(\vec{p})$, i.e., the part of the winning strategy for variable b .
- (2) If in the same setting as above for \mathcal{F}^b , the variables \vec{p} appear only positively in $A(\vec{p}, \vec{q})$, then we can extract a monotone boolean circuit for $\sigma(\vec{p})$ from a **LQU⁺-Res** refutation π of \mathcal{F}^b in polynomial time (in the size of π).

Proof. As we can compute the (monotone) interpolant when b is absent, we use the same proof with a few modifications for the new formula.

We first change the definition of \vec{q} and \vec{r} -clauses to allow for b and $\neg b$ literals.

Definition 5.3. We call any clause in the proof a \vec{q} -clause (resp. \vec{r} -clause) if it contains only variables \vec{p}, \vec{q} or literal b (resp. \vec{p}, \vec{r} or literal $\neg b$). We retain the inheritance property for clauses only containing \vec{p} variables.

Construction of the circuit C_π : When constructing the circuit, we now also need to consider a resolution step on the common universal variable b :

$$\frac{C_v = \overbrace{C_1 \vee U_1 \vee b}^{\text{node } v} \quad \overbrace{C_2 \vee U_2 \vee \neg b}^{\text{node } w}}{C_u = \underbrace{C_1 \vee C_2 \vee U}_{\text{node } u}} = C_w.$$

Here we can arbitrarily pick one of v or w . For example here we pick v and let g_u be wired to g_v with the no-operation (identity) gate, disregarding the input from g_w .

Construction of π' and π'' : We slightly modify the invariants to include the new definitions. Additionally we make a small change to the first invariant.

- (1) $C'_{u,\vec{a}} \setminus \{b, \neg b\} \preceq C_u$.
- (2) $g_u(\vec{a}) = 0 \implies C''_{u,\vec{a}}$ is a \vec{q} -clause and can be obtained from the clauses of $A(\vec{a}, \vec{q})$ alone using the rules of **LQU⁺-Res**.
- (3) $g_u(\vec{a}) = 1 \implies C''_{u,\vec{a}}$ is a \vec{r} -clause and can be obtained from the clauses of $B(\vec{a}, \vec{r})$ alone using the rules of **LQU⁺-Res**.

Notice also that $b^* \notin C''_{u,\vec{a}}$ as b^* can only arise from a long distance resolution step on a \vec{p} variable but these are instantiated and so never occur as pivots in the proof π'' assuming the induction hypothesis.

We observe that the base cases work for the construction of π' and π'' . The only new part of the inductive step is when we have

$$\frac{C_v = \overbrace{C_1 \vee U_1 \vee b}^{\text{node } v} \quad \overbrace{C_2 \vee U_2 \vee \neg b}^{\text{node } w}}{C_u = \underbrace{C_1 \vee C_2 \vee U}_{\text{node } u}}.$$

To find $C'_{u,\vec{a}}$ we look at our choice of wiring in the circuit construction. If g_u is wired to g_v ($g_u = g_v$) then we take $C'_{u,\vec{a}}$ to equal $C'_{v,\vec{a}}$. Since $C'_{v,\vec{a}} \setminus \{b, \neg b\} \preceq C_v \setminus \{b, \neg b\} \preceq C_u$ we get $C'_{u,\vec{a}} \setminus \{b, \neg b\} \preceq C_u$. Since our choice of the clause is determined by our choice of wiring, then we retain our invariants in that way.

Notice that we never resolve a \vec{q} -clause with a \vec{r} clause in π'' so $b, \neg b$ will always be retained in their respective type of clauses.

From the above, we have the following conclusion. Let r be the root of π . Then on any assignment \vec{a} to the \vec{p} variables we have:

- (1): $C'_{r,\vec{a}} \setminus \{b, \neg b\} \preceq C_r = \square$. Therefore, $C''_{r,\vec{a}} \setminus \{b, \neg b\} = \square$. But $C''_{r,\vec{a}}$ can contain at most one of these literals, which can be universally reduced to complete a refutation.
- (2): $g_r(\vec{a}) = 0 \implies C''_{r,\vec{a}}$ is a \vec{q} -clause and can be obtained from the clauses of $A(\vec{a}, \vec{q})$ alone using the rules of system **LQU⁺-Res**. Hence by soundness of **LQU⁺-Res**, $\mathcal{Q}\vec{q}.A(\vec{a}, \vec{q})$ is false.
- (3): $g_r(\vec{a}) = 1 \implies C''_{r,\vec{a}}$ is an \vec{r} -clause and can be obtained from the clauses of $B(\vec{a}, \vec{r})$ alone using the rules of system **LQU⁺-Res**. Hence by soundness of **LQU⁺-Res**, $\mathcal{Q}\vec{r}.B(\vec{a}, \vec{r})$ is false.

Thus g_r , the output gate of the circuit, computes $\sigma(\vec{p})$. □

An analogous result to Theorem 5.2 also holds for **IRM-calc**.

Theorem 5.4.

- (1) From each **IRM-calc** refutation π of \mathcal{F}^b we can extract in polynomial time a boolean circuit for $\sigma(\vec{p})$, i.e., the part of the winning strategy for variable b .
- (2) If in the same setting as above for \mathcal{F}^b , the variables \vec{p} appear only positively in $A(\vec{p}, \vec{q})$, then we can extract a monotone boolean circuit for $\sigma(\vec{p})$ from a **IRM-calc** refutation π of \mathcal{F}^b in polynomial time (in the size of π).

Proof. We can use exactly the same constructions as in Theorem 3.4. The b literals do not affect the argument. □

As a corollary, the versions $\Phi_n^b(\vec{p}, \vec{q}, \vec{r})$ of the formulas from Section 4 also require exponential-size proofs in **IRM-calc** and **LQU⁺-Res**.

ACKNOWLEDGEMENTS

We thank Pavel Pudlák and Mikoláš Janota for helpful discussions on the relation between feasible interpolation and strategy extraction during the Dagstuhl Seminar ‘Optimal Algorithms and Proofs’ (14421).

REFERENCES

- [AB87] Noga Alon and Ravi B. Boppana. The monotone circuit complexity of boolean functions. *Combinatorica*, 7(1):1–22, 1987.
- [AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity – A Modern Approach*. Cambridge University Press, 2009.
- [BBC16] Olaf Beyersdorff, Ilario Bonacina, and Leroy Chew. Lower bounds: From circuits to QBF proof systems. In *Proc. ACM Conference on Innovations in Theoretical Computer Science (ITCS’16)*, pages 249–260. ACM, 2016.
- [BCJ14] Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. On unification of QBF resolution-based calculi. In *MFCS, II*, pages 81–93, 2014.
- [BCJ15] Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. Proof complexity of resolution-based QBF calculi. In *STACS*, pages 76–89, 2015.
- [BCMS15] Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Feasible interpolation for QBF resolution calculi. In *Proc. International Colloquium on Automata, Languages, and Programming (ICALP’15)*, pages 180–192. Springer, 2015.
- [BCMS16a] Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Are short proofs narrow? QBF resolution is not simple. In *Proc. Symposium on Theoretical Aspects of Computer Science (STACS’16)*, pages 15:1–15:14, 2016.
- [BCMS16b] Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Understanding Cutting Planes for QBFs. In *Proc. Foundations of Software Technology and Theoretical Computer Science (FSTTCS’16)*, pages 40:1–40:15, 2016. (Full version in ECCC TR 17-037.)
- [BCS15] Olaf Beyersdorff, Leroy Chew, and Karteek Sreenivasaiah. A game characterisation of tree-like Q-resolution size. In *LATA*, pages 486–498, 2015.
- [BDG⁺04] Maria Luisa Bonet, Carlos Domingo, Ricard Gavaldà, Alexis Maciel, and Toniann Pitassi. Non-automatizability of bounded-depth Frege proofs. *Computational Complexity*, 13(1–2):47–68, 2004.
- [BJ12] Valeriy Balabanov and Jie-Hong R. Jiang. Unified QBF certification and its applications. *Formal Methods in System Design*, 41(1):45–65, 2012.
- [BK14] Olaf Beyersdorff and Oliver Kullmann. Unified characterisations of resolution hardness measures. In *SAT*, pages 170–187, 2014.
- [BM08] Marco Benedetti and Hratch Mangassarian. QBF-based formal verification: Experience and perspectives. *JSAT*, 5(1-4):133–191, 2008.
- [BP16] Olaf Beyersdorff and Ján Pich. Understanding Gentzen and Frege systems for QBF. In *Proc. ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2016.
- [BPR00] Maria Luisa Bonet, Toniann Pitassi, and Ran Raz. On interpolation and automatization for Frege systems. *SIAM Journal on Computing*, 29(6):1939–1967, 2000.
- [BSW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow - resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001.
- [BWJ14] Valeriy Balabanov, Magdalena Widl, and Jie-Hong R. Jiang. QBF resolution systems and their proof complexities. In *SAT*, pages 154–169, 2014.
- [Cra57] William Craig. Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *The Journal of Symbolic Logic*, 22(3):269–285, 1957.
- [EKLP14] Uwe Egly, Martin Kronegger, Florian Lonsing, and Andreas Pfandler. Conformant planning as a case study of incremental QBF solving. In *Artificial Intelligence and Symbolic Computation (AISC’14)*, pages 120–131, 2014.

- [GVB11] Alexandra Goultiaeva, Allen Van Gelder, and Fahiem Bacchus. A uniform approach for generating proofs and strategies for both true and false QBF formulas. In *IJCAI*, pages 546–553, 2011.
- [Hru09] Pavel Hrubeš. On lengths of proofs in non-classical logics. *Annals of Pure and Applied Logic*, 157(2–3):194–205, 2009.
- [JM15] Mikolás Janota and Joao Marques-Silva. Expansion-based QBF solving versus Q-resolution. *Theor. Comput. Sci.*, 577:25–42, 2015.
- [KKF95] Hans Kleine Büning, Marek Karpinski, and Andreas Flögel. Resolution for quantified Boolean formulas. *Inf. Comput.*, 117(1):12–18, 1995.
- [KP98] Jan Krajíček and Pavel Pudlák. Some consequences of cryptographical conjectures for S_2^1 and *EF*. *Information and Computation*, 140(1):82–94, 1998.
- [Kra97] Jan Krajíček. Interpolation theorems, lower bounds for proof systems and independence results for bounded arithmetic. *J. Symb. Log.*, 62(2):457–486, 1997.
- [Kra11] Jan Krajíček. *Forcing with random variables and proof complexity*, volume 382 of *Lecture Note Series*. London Mathematical Society, 2011.
- [Mun84] Daniele Mundici. Tautologies with a unique Craig interpolant, uniform vs. nonuniform complexity. *Annals of Pure and Applied Logic*, 27:265–273, 1984.
- [Pap94] Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [Pud97] Pavel Pudlák. Lower bounds for resolution and cutting planes proofs and monotone computations. *The Journal of Symbolic Logic*, 62(3):981–998, 1997.
- [Pud00] Pavel Pudlák. Proofs as games. *American Math. Monthly*, pages 541–550, 2000.
- [Rin07] Jussi Rintanen. Asymptotically optimal encodings of conformant planning in QBF. In *AAAI*, pages 1045–1050. AAAI Press, 2007.
- [VG12] Allen Van Gelder. Contributions to the theory of practical quantified Boolean formula solving. In *CP*, pages 647–663, 2012.
- [ZM02] Lintao Zhang and Sharad Malik. Conflict driven learning in a quantified Boolean satisfiability solver. In *ICCAD*, pages 442–449, 2002.