Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 44 (2017)

Reasons for Hardness in Quantified Boolean Formulas*

Olaf Beyersdorff[1], Luke Hinde[1], and Ján Pich[2]

[1] School of Computing, University of Leeds
{o.beyersdorff,sclpeh}@leeds.ac.uk
[2] Kurt Gödel Research Center, University of Vienna
jan.pich@univie.ac.at

Abstract. We aim to understand inherent reasons for lower bounds for QBF proof systems, and revisit and compare two previous approaches in this direction.

The first of these relates size lower bounds for strong QBF Frege systems to circuit lower bounds via *strategy extraction* (Beyersdorff & Pich, LICS'16). Here we show a refined version of strategy extraction and thereby for any QBF proof system obtain a trichotomy for hardness: (1) via circuit lower bounds, (2) via propositional Resolution lower bounds, or (3) 'genuine' QBF lower bounds.

The second approach tries to explain QBF lower bounds through *quantifier alternations* in a system called relaxing QU-Res (Chen, ICALP'16). We prove a strong lower bound for relaxing QU-Res, which at the same time exhibits significant shortcomings of that model. Prompted by this we propose an alternative, improved version, allowing more flexible oracle queries in proofs. We show that lower bounds in our new model correspond to the trichotomy obtained via strategy extraction.

1 Introduction

Proof complexity studies the question of how difficult it is to prove theorems in different formal proof systems. The main question is thus: for a given theorem ϕ and proof system P, what is the size of the shortest proof of ϕ in P? This research has strong and productive connections to several other areas, most notably to computational complexity, with the aim of separating complexity classes through Cook's programme [10,13], and to first-order logic (theories of bounded arithmetic [12,25]). In recent years, progress in practical SAT- and QBF-solving has been a major motivation for proof complexity, as runs of SAT-solvers correspond to proofs of (un)satisfiability of CNFs. Analysis of the corresponding proof system provides the framework for understanding the power and the limitations of the solver [10].

The majority of work in proof complexity has been focused on *propositional proof complexity*, on proof systems for classical propositional logic. In particular,

* Supported by grant no. 48138 from the John Templeton Foundation and by the Austrian Science Fund (FWF) under project number 28699.

ISSN 1433-8092

Resolution [30] has received much attention as it models the approach taken by many modern SAT-solvers.

QBF proof complexity is a comparatively young field, studying proof systems for quantified Boolean formulas. Determining the truth of a QBF is PSPACE-complete, and so has wider ranging applications than SAT-solving, extending to fields such as formal verification and planning [3, 14, 29]. Similarly to the propositional case, several Resolution-based QBF proof systems have been suggested and analysed [1,5–7,16,21,23,33] to model the approaches taken by QBF solvers. Of particular importance are Q-Resolution [23] and universal Q-Resolution (QU-Res) [16], which as analogues of propositional Resolution form the base systems for conflict-driven clause learning (CDCL) QBF solving [17].

Stronger systems in the form of QBF Frege systems were developed recently [4]. As in the propositional framework, by restricting the lines in Frege to a circuit class $\mathcal{C}$ we obtain a hierarchy of (QBF) $\mathcal{C}$-Frege systems, corresponding to the hierarchy of circuit classes.

A conceptually simple but powerful technique for constructing QBF proof size lower bounds from Boolean circuit lower bounds was developed in [4, 6]. This *strategy extraction technique* employs the complexity of Herbrand functions witnessing the universal quantifiers. In [4] the technique was used to show strong lower bounds for QBF Frege systems, including exponential lower bounds for QBF $AC^0[p]$-Frege (which is in stark contrast to the situation in propositional Frege, where lower bounds for $AC^0[p]$-Frege are wide open).

Recent work has tightened the connection to circuit complexity further. In [8] it has been shown that for natural circuit classes $\mathcal{C}$, a lower bound for proof size in QBF $\mathcal{C}$-Frege corresponds to either a lower bound for propositional $\mathcal{C}$-Frege, or a lower bound for the circuit class $\mathcal{C}$. This characterisation points to a distinction between lower bounds derived from lower bounds on propositional proof systems, and 'genuine' QBF lower bounds.

More widely, *understanding the reasons of hardness* for QBF proof systems and solving constitutes a major challenge, which at current is only insufficiently mastered. Most QBF proof systems use a propositional system such as Resolution or Frege as their core, implying that on existentially quantified formulas the QBF system coincides with its classical core system. This leads to the somewhat disturbing fact that lower bounds for e.g. propositional Resolution trivially lift to any of the studied QBF Resolution systems.

Motivated by this observation, Chen [11] introduced a new notion of *proof system ensemble*, in particular for QU-Res called *relaxing QU-Res*, with the aim to distinguish between lower bounds lifted from propositional Resolution and 'genuine' QBF lower bounds arising from quantifier alternation of the QBFs. Quantifier alternation as also been empirically observed as a source of hardness [26, 27], making this a very interesting direction for theoretical study.

Our Contributions The main aim of this paper is to gain a refined understanding of the reasons for QBF hardness, both following the strategy extraction paradigm [8] and the paradigm via quantifier alternation [11]. We revisit both models and relate them in their explanatory power.

A. Refinement of formalised strategy extraction. We describe a decomposition of QBF solving into SAT solving and a search for small circuits witnessing a given QBF. This relies on an improvement of the strategy extraction theorem from [8] which says that, given polynomial-size QBF $\mathcal{C}$-Frege proofs of QBFs ψ_n, one can construct small $\mathcal{C}$ circuits witnessing the existential quantifiers in ψ_n in such a way that the resulting 'witnessed' propositional formulas have polynomial-size proofs in $\mathcal{C}$-Frege. Here, we show that in fact the witnessed formulas have polynomial-size proofs even in tree-like Resolution (Theorem 1).

Applying a similar decomposition, we observe that polynomial-size lower bounds on a sequence of QBFs in any QBF proof system can be categorized as either (1) a circuit lower bound, (2) a Resolution lower bound, or (3) a genuine QBF lower bound (Theorem 2).

B. Lower bounds for relaxing QU-Res. We revisit relaxing QU-Res, introduced in [11] with the aim of distinguishing propositional bounds from QBF bounds arising from quantifier alternation. In particular, Chen [11] gives an exponential lower bound for relaxing QU-Res that applies to quantified Boolean circuits, however with no small CNF representations (Appendix A). As this is a somewhat atypical feature in proof complexity, we improve this by presenting QBFs with CNF matrices that require exponential-size relaxing QU-Res proofs (Theorem 9). Our formulas use a new construction that combines two false QBFs Φ and Ψ into their product formula $\Phi \otimes \Psi$ such that each short QU-Res proof must refute Ψ before it refutes Φ.

These product formulas have another compelling feature: their hardness for relaxing QU-Res (and QU-Res) rests on the hardness of the pigeonhole principle for propositional Resolution. Our lower bound therefore suggests that relaxing QU-Res does not capture 'genuine' hardness of QBFs due to quantifier alternation.

C. New systems for 'genuine' QBF hardness. Noting this situation, we propose new QBF proof systems, Σ_k^p-QU-Res (Def. 15). The systems bear similarities to relaxing QU-Res, particularly in the use of relaxations of quantifiers and a proof checking algorithm with access to a Σ_k^p-oracle. The major difference is that oracle queries in our algorithm may appear at any point in the proof.

It is interesting to relate lower bounds in Σ_1^p-QU-Res to our trichotomy shown in A. In this direction, we prove that Σ_1^p-QU-Res admits strategy extraction by depth-3 Boolean circuits (Lemma 18). Hence QU-Res lower bounds stemming from circuit lower bounds (case (1) in the trichotomy in A) translate to lower bounds in Σ_1^p-QU-Res. Further, if a QBF is hard for QU-Res due to a Resolution lower bound (case (2) in A), it has short proofs in Σ_1^p-QU-Res. We also demonstrate that a variant of the prominent formulas of Kleine Büning et al. [23] simultaneously has genuine QBF lower bounds as per case (3) in A (Theorem 4) and is hard for Σ_k^p-QU-Res proofs for any constant k (Theorem 22).

Organisation. In Sec. 2 we detail necessary background. Section 3 refines formalised strategy extraction and the characterisation of QBF lower bounds from [8]. In Sec. 4 we show the lower bound for relaxing QU-Res. Section 5 contains the definition of Σ_k^p-QU-Res and a comparison of lower bounds in these systems with the characterisation in Sec. 3. In Sec. 6, we analyse the hardness of several QBF families in these proof systems.

2 Preliminaries

Quantified Boolean Formulas. A (prenex normal form) *quantified Boolean formula* (QBF) $\Phi = \mathcal{Q}_1 x_1 \dots \mathcal{Q}_n x_n . \phi(x_1, \dots, x_n)$ consists of a propositional formula ϕ, usually expressed as a CNF, and a quantifier prefix $\mathcal{Q}_1 x_1 \dots \mathcal{Q}_n x_n$, where each $\mathcal{Q}_i \in \{\exists, \forall\}$ ranges over $\{0, 1\}$.

The semantics of such a QBF can be considered as a game between players $\exists$ and $\forall$. On the ith turn, the player corresponding to $\mathcal{Q}_i$ assigns a 0/1 value to x_i. After all the variables have been assigned, the $\exists$ player (resp. $\forall$ player) wins the game if ϕ evaluates to 1 (resp. 0).

Given a variable x_i, a *strategy* for x_i is a function $\sigma_i : \{x_1, \dots, x_{i-1}\} \to \{0, 1\}$. A *winning strategy* for the $\exists$ (resp. $\forall$) player, consists of a strategy for each existential (resp. universal) variable which wins all possible games on Φ. A QBF is false (resp. true) if and only if there is a winning strategy for the $\forall$ player (resp. $\exists$ player).

The quantifier complexity of a QBF is described by inductively defined classes Σ_i^b and Π_i^b, counting the number of quantifier alternations. By Σ_i^p (resp. Π_i^p) we denote the i^{th} level of the polynomial hierarchy, for which deciding truth of Σ_i^b (resp. Π_i^b) formulas is complete.

Proof Complexity. A *proof system* for a language $\mathcal{L}$ is a polynomial-time computable surjective function $f : \{0, 1\}^* \to \mathcal{L}$ [13]. If $f(\pi) = \phi$, we say π is an f-proof of ϕ. Given proof systems P and Q for $\mathcal{L}$, P *p-simulates* Q if there is a polynomial-time function t with $P(t(\pi)) = Q(\pi)$ for any π. Two proof systems are *p-equivalent* if they p-simulate each other.

Here we consider proof systems for propositional tautologies and fully quantified true QBFs. We also consider proof systems for unsatisfiable formulas and false QBFs and use the words *proof* and *refutation* interchangeably.

Resolution [30] is one of the best studied propositional proof systems [32]. Given two clauses $C \vee x$ and $D \vee \neg x$, Resolution can derive the clause $C \vee D$. A Resolution proof that a CNF ϕ is unsatisfiable is a derivation of the empty clause $\bot$ using the resolution rule.

QU-Resolution (QU-Res) [16] is a natural extension of Resolution to QBFs. Given a QBF $\Phi = \mathcal{Q}_1 x_1 \dots \mathcal{Q}_n x_n . \phi$, where ϕ is a CNF, a QU-Res refutation of Φ is a derivation of $\bot$ from the clauses of ϕ. It uses the Resolution rule (with the extra condition that deriving tautological clauses is not allowed) and the $\forall$*-reduction rule*, which from a clause $C \vee l$ with literal l on universal variable x_i (i.e., $l = x_i$ or $l = \neg x_i$) can derive the clause C provided C contains no literals on $x_{i+1}, \dots, x_n$.

A proof in Resolution (and QU-Res, and other proof systems) can be represented as a directed acyclic graph (dag) with a root labelled by $\bot$, and input vertices labelled with clauses from the CNF. If we restrict the dag to be a tree, we define *tree-like Resolution*, which we denote by R^*. Tree-like Resolution is known to be weaker than Resolution [9].

Frege Systems. Frege systems are common 'textbook' proof systems comprised of a set of axiom schemes and inference rules [13]. Lines of a Frege proof are formulas in propositional variables and Boolean connectives $\wedge, \vee, \neg$. A Frege proof of ϕ is a sequence of formulas, ending with ϕ, in which each formula is either a substitution instance of an axiom, or is inferred from previous formulas by a valid inference rule. We also consider refutational Frege systems, in which we start with the formula $\neg\phi$ and derive a contradiction.

For a given circuit class $\mathcal{C}$, we define $\mathcal{C}$-Frege, as in [22], to be a Frege system which works with lines consisting of circuits in $\mathcal{C}$ and a finite set of derivation rules. If $\mathcal{C}$ consists of all Boolean circuits, then $\mathcal{C}$-Frege is p-equivalent to extended Frege (EF). If $\mathcal{C}$ is restricted to Boolean formulas, i.e. $\mathcal{C} = NC^1$, then NC^1-Frege is Frege as defined above.

An elegant method for extending $\mathcal{C}$-Frege systems to QBF was shown in [4]. The QBF proof system $\mathcal{C}$-Frege+$\forall$-red is a refutational proof system working with circuits from $\mathcal{C}$. The inference rules of $\mathcal{C}$-Frege+$\forall$-red are those of $\mathcal{C}$-Frege, along with the $\forall$-red rule $\frac{L_j(u)}{L_j(u/B)}$ where u is quantified innermost among the variables of the proof line L_j with respect to the quantifier prefix, and the circuit B does not contain any variables to the right of u. Restricting the circuit B in the $\forall$-red rule to the constants $0, 1$ results in a p-equivalent system [8].

3 Strategy extraction and reasons for hardness

A QBF proof system P has the *strategy extraction property* if for any P-proof π of a QBF ψ of the general form $\forall x_1 \exists y_1 \ldots \forall x_n \exists y_n.\, \phi(x_1, \ldots, x_n, y_1, \ldots, y_n)$, where ϕ is a propositional formula, there are $|\pi|^{O(1)}$-size circuits C_i witnessing the existential quantifiers in ψ, i.e.

$$\bigwedge_{i=1}^{n} (y_i \leftrightarrow C_i(x_1, \ldots, x_i, y_1, \ldots, y_{i-1})) \rightarrow \phi(x_1, \ldots, x_n, y_1, \ldots, y_n). \qquad (1)$$

The strategy extraction is *Q-formalised* if, in addition, the propositional formulas (1) have $|\pi|^{O(1)}$-size proofs in a propositional proof system Q.

For any QBF ψ, either there is a propositional formula as in (1) equivalent to ψ, or there are no (small) circuits C_i witnessing the existential variables, and so no QBF proof system with the strategy extraction property can prove ψ feasibly.

The task of *QBF solving based on proof systems admitting strategy extraction is thus reducible to the task of finding the witnessing circuits C_i, and then SAT solving of the witnessed formula.* Alternatively, we can speak about a reduction of QBF solving to Σ_2^b-formulas with existentially quantified witnessing circuits: $\exists C_1 \ldots C_n \forall x_1 \ldots x_n\ y_1 \ldots y_n.\ \bigwedge_{i=1}^{n} (y_i \leftrightarrow C_i(x_1, \ldots, x_i, y_1, \ldots, y_{i-1})) \rightarrow \phi(x_1, \ldots, x_n, y_1, \ldots, y_n)$

We will show that all QBF proof systems P p-simulated by EF+$\forall$-red[3] have R^*-formalized strategy extraction. More precisely, we improve the formalised strategy extraction for EF+$\forall$-red from [8] by observing that the witnessing circuits can encode extension variables, which allows us to replace the EF proof of the witnessed formula with an R^* proof.

Consequently, instead of determining whether there is a short P-proof of ψ, one can solve the equivalent problem of whether there are small circuits C_i and a short R^*-proof of (1). As R^* is quasi-automatisable (i.e., R^* refutations for a given CNF can be constructed in quasi-polynomial time in the size of the smallest R^* proof [2]), the problem is essentially reduced to the search for the right witnessing circuits C_i.

Theorem 1. *Let $\mathcal{C}$ be the circuit class NC^1 or $P/poly$.*[4] *Given a $\mathcal{C}$-Frege+$\forall$-red refutation π of a QBF $\exists x_1 \forall y_1 \ldots \exists x_n \forall y_n .\, \phi(x_1, \ldots, x_n, y_1, \ldots, y_n)$ where $\phi \in \Sigma_0^b$, we can construct in time $|\pi|^{O(1)}$ an R^* refutation of*

$$\bigwedge_{i=1}^{n} (y_i \leftrightarrow C_i(x_1, \ldots, x_i, y_1, \ldots, y_{i-1})) \wedge \phi(x_1, \ldots, x_n, y_1, \ldots, y_n) \tag{2}$$

for some circuits $C_i \in \mathcal{C}$.

Proof. By the formalised strategy extraction theorem for $\mathcal{C}$-Frege systems [8], there is a $\mathcal{C}$-Frege proof of the witnessed formula (2). This means there is an R^* refutation of

$$\mathit{Ext} \wedge \bigwedge_{i=1}^{n} (y_i \leftrightarrow C_i(x_1, \ldots, x_i, y_1, \ldots, y_{i-1})) \wedge \phi(x_1, \ldots, x_n, y_1, \ldots, y_n)$$

where *Ext* is a set of extension axioms defining $\mathcal{C}$ formulas built on variables $x_1, \ldots, x_n, y_1, \ldots, y_n$. With the exception of those depending on y_n, these axioms can be encoded into circuits C_i with each extension variable represented by a possibly redundant gate of a circuit C_i. In order to remove the extension variables depending on y_n, we construct two independent R^* refutations, one with all occurrences of y_n in clauses of *Ext* substituted by 0 and the other with occurrences of y_n in *Ext* substituted by 1. This results in two R^* derivations, both at most as large as the original, one concluding with $\{y_n\}$ and the other with $\{\neg y_n\}$. Resolving on these two clauses we obtain the needed R^* derivation without the extension variables depending on y_n. □

The reduction of QBF solving to SAT solving presented above is also of use for proving QBF proof complexity lower bounds. In [8] it was shown that any super-polynomial lower bound on EF+$\forall$-red is either a super-polynomial circuit

[3] This includes all commonly studied Resolution-based QBF systems.

[4] The result easily generalises to further 'natural' circuit classes $\mathcal{C}$ such as AC^0 or TC^0, but we will focus here on the two most interesting cases NC^1 and $P/poly$ leading to Frege and EF systems, respectively.

lower bound or a super-polynomial lower bound on EF. Here we generalise this phenomenon to other QBF proof systems.

Let P be a refutational QBF proof system operating on clauses of matrices of QBFs (given in a prenex form with CNF matrices) which contains a resolution rule that allows resolution on both existential and universal variables. We say that a set of clauses C defines a formula $C_i(\boldsymbol{x}) = z$ for a circuit C_i with input variables $\boldsymbol{x}$ and output variable z if z appears in a literal of some clause in C and for any assignment of the input variables there is exactly one assignment of the remaining variables satisfying all clauses in C.

Whenever a QBF ψ as above is hard for a QBF proof system P it is for one of the following reasons:

1. the existential quantifiers in ψ cannot be witnessed by circuits C_i such that formulas $\bigwedge_i C_i(x_1, \dots, x_i, y_1, \dots, y_{i-1}) = y_i$ have $|\phi|^{O(1)}$-size P-derivations from $\neg\phi$.
2. the existential quantifiers in ψ are witnessable as in 1. but the witnessed formula $\bigwedge_{i=1}^{n} (y_i \leftrightarrow C_i(x_1, \dots, x_i, y_1, \dots, y_{i-1})) \wedge \neg\phi(x_1, \dots, x_n, y_1, \dots, y_n)$ is hard for Resolution.

This characterisation can be specified further.

Theorem 2. *Let P be a refutational QBF proof system as above admitting strategy extraction by $\mathcal{C}$ circuits. If $\psi_n = \forall x_1 \exists y_1 \dots \forall x_n \exists y_n.\, \phi_n(x_1, \dots, x_n, y_1, \dots, y_n)$ are QBFs with propositional CNF ϕ_n, which do not have polynomial-size proofs in P, then one of the following holds:*

1. ***Circuit lower bound.*** *The existential variables in ψ_n are not witnessable by $\mathcal{C}$ circuits.*
2. ***Resolution lower bound.*** *Condition 1. does not hold, but for all $\mathcal{C}$ circuits witnessing ψ_n, the witnessed formulas require super-polynomial size Resolution refutations.*
3. ***Genuine QBF hardness.*** *There are circuits $C_i \in \mathcal{C}$ witnessing ψ_n so that the witnessed formulas have polynomial-size Resolution refutations, but for all such circuits C_i it is hard to derive $\bigwedge_i C_i(x_1, \dots, x_i, y_1, \dots, y_{i-1}) = y_i$ from $\neg\phi_n$ in P.* □

Proof. If the existential variables in ψ_n are not witnessable by $\mathcal{C}$ circuits, we are done. We therefore assume that there are $\mathcal{C}$ circuits witnessing the existential variables.

Suppose that there are some circuits $C_i \in \mathcal{C}$ such that the witnessed formula (2) has a polynomial-size Resolution refutation. If this is not the case, we are done as we are in case 2.

We can construct a refutation of $\neg\psi_n$ in P by first deriving $\bigwedge_i C_i(x_1, \dots, x_i, y_1, \dots, y_{i-1}) = y_i$ from $\neg\phi_n$, and then refuting $\bigwedge_i (C_i \leftrightarrow y_i) \wedge \neg\phi_n$. Since the refutation of $\bigwedge_i (C_i \leftrightarrow y_i) \wedge \neg\phi$ is assumed to have a polynomial-size refutation, but any refutation of $\neg\psi_n$ requires super-polynomial-size, it must be the case that for the circuits $C_i \in \mathcal{C}$, the derivation of $\bigwedge_i (C_i \leftrightarrow y_i)$ from $\neg\phi_n$ requires super-polynomial size (case 3). □

This means that any QBF lower bound on P is either a circuit lower bound, a propositional proof complexity lower bound, or it is a ‘genuine’ QBF proof complexity lower bound in the sense that P cannot derive efficiently some circuits witnessing the existential quantifiers in the original formula and whenever it can do that for some other witnessing circuits, the witnessed formula is hard for Resolution.

The last possibility does not happen in the case of strong systems like EF+∀-red [8]. The situation is, however, more delicate with weaker systems, where we can indeed encounter ‘genuine’ QBF lower bounds. We give an example.

Definition 3 (Kleine Büning et al. [23]). *The QBFs* KBKF_n *are defined as* $\exists y_0 y_1 y_1' \forall x_1 \dots \exists y_k y_k' \forall x_k \dots \forall x_n \exists y_{n+1} \dots y_{n+n}. \bigwedge_{i=1}^{2n} C_i \wedge C_i'$*, where*

$$\begin{aligned} C_0 &= \{\neg y_0\} & C_0' &= \{y_0, \neg y_1, \neg y_1'\} \\ C_k &= \{y_k, \neg x_k, \neg y_{k+1} \neg y_{k+1}'\} & C_k' &= \{y_k', x_k, \neg y_{k+1}, \neg y_{k+1}'\} \\ C_n &= \{y_n, \neg x_n, \neg y_{n+1}, \dots, \neg y_{n+n}\} & C_n' &= \{y_n', x_n, \neg y_{n+1}, \dots, y_{n+n}\} \\ C_{n+t} &= \{x_t, y_{n+t}\} & C_{n+t}' &= \{\neg x_t, y_{n+t}\} \end{aligned}$$

This family of QBFs is known to require proofs of size $2^{\Omega(n)}$ in Q-Resolution [6,23]. Although KBKF_n have polynomial-size refutations in QU-Res, the exponential-size lower bound can be lifted to QU-Res using the formulas KBKF_n', obtained by adding new universal variables z_k, quantified at the same level as x_k, and adding the literal z_k or $\neg z_k$ to each clause containing x_k or $\neg x_k$, respectively [1]. This lower bound is a ‘genuine’ QBF proof complexity lower bound.

Theorem 4. *The formulas* KBKF_n' *are hard for QU-Res due to genuine QBF hardness (case 3 in Theorem 2).*

Proof. It is clear that playing the variables x_k and z_k identical to y_k' is a winning strategy for the universal player, and so there are circuits C_i as described in Theorem 2 which are of constant size.

Looking now at the witnessed formula $\bigwedge_{i=1}^{n}((x_i \leftrightarrow y_i') \wedge z_i \leftrightarrow y_i')) \wedge \phi$, we show this can be refuted by a linear-size proof. By resolving on each x_i and z_i to replace these with the relevant literal on y_i', we obtain the clauses $y_i' \vee y_{n+i}$ and $\neg y_i' \vee y_{n+i}$. Resolving on each y_{n+i} gives y_n', $y_n \vee \neg y_n'$ and consequently y_n. For each i, we use y_i and y_i' to deduce y_{i-1} and y_{i-1}' and finally y_0, completing the refutation.

Since KBKF_n' is known to require exponential size proofs in QU-Res [1], by Theorem 2, it must satisfy one of the three conditions given. We have established that there are small witnessing circuits, and that the witnessed formula is easy to refute, and so it must be the case that it is hard to derive the witnessing circuits. □

4 Hardness due to quantifier alternation

The characterisation of QBF proof system lower bounds given above is a very natural one. We now show that other suggested reasons for hardness correspond with it.

An alternative characterisation of QBF lower bounds that has previously been suggested is based on the alternation of quantifiers in the quantifier prefix. Most studied QBF proof systems build on a propositional proof system (e.g. Resolution) and on Σ_1^b formulas just coincide with the propositional base system. Therefore we can obtain QBF lower bounds directly from the propositional lower bounds. Characterising lower bounds by quantifier alternation aims to distinguish between such propositional lower bounds and 'genuine' QBF lower bounds arising from the alternation of quantifiers. Relaxing QU-Res has been put forward as a proof system to determine hardness due to quantifier alternation [11].

Definition 5 (Chen [11]). *A* Relaxing QU-Res *proof of a QBF Φ uses the same deduction rules as QU-Res, but can introduce any axiom from the set $H(\Phi, \Pi_k^b)$, defined below, for some constant k.*

For a quantifier prefix $\Pi = \mathcal{Q}_1 x_1 \dots \mathcal{Q}_n x_n$, if π is a permutation such that $\pi(i) < \pi(j)$ whenever $i < j$ and $\mathcal{Q}_i = \forall$ and $\mathcal{Q}_j = \exists$, then the prefix $\Pi' = \mathcal{Q}_{\pi(1)} x_{\pi(1)} \dots \mathcal{Q}_{\pi(n)} x_{\pi(n)}$ is a relaxation*. Intuitively, a relaxation involves 'moving $\forall$-variables to the left'. If Π' is a Σ_k^b-prefix, we call Π' a Σ_k^b-*relaxation.

Let $\Phi = \Pi.\phi$ be a QBF. For a clause A, let α be the minimal assignment falsifying A. Construct $\Pi[\alpha]$ by removing all variables in α, and replacing any $\forall$-quantifers left of a variable in α by $\exists$. If there is some Π_k^b-relaxation $\Pi'[\alpha]$ of $\Pi[\alpha]$ such that $\Pi'[\alpha].\phi[\alpha]$ is false, then $A \in H(\Phi, \Pi_k^b)$.

For some families of QBFs, such as the pigeonhole principle, other propositional formulas or indeed any QBF with a prefix with constant alternation, relaxing QU-Res has polynomial-size proofs, whereas QU-Res may require exponential-size proofs.

However, lower bounds for both tree-like and dag-like relaxing QU-Res were also shown in [11]. The lower bound for dag-like relaxing QU-Res in [11] is rather unconventional as the proof system works with clauses, whereas the lower bound applies to circuits without polynomial-size CNF representations (cf. Appendix A). Here we present formulas with polynomially many clauses that require exponential-size proofs in relaxing QU-Res.

Furthermore, the lower bounds we show on the size of QU-Res proofs of these formulas are clearly due to lower bounds on Resolution proofs of the pigeonhole principle, rather than alternation of quantifiers, or any other 'genuine' QBF reasons. It follows that this is the case for relaxing QU-Res as well. This demonstrates that relaxing QU-Res is not an adequate formalism to distinguish propositional lower bounds from genuine QBF lower bounds.

To begin, we present a method of combining two false QBFs to produce another false QBF. This method might also be of independent interest for the creation of hard QBFs.

Definition 6. *Let $\Phi = \Lambda(\boldsymbol{x}) \cdot \bigwedge_{i=1}^{n} C_i(\boldsymbol{x})$ and $\Psi = \Pi(\boldsymbol{z}) \cdot \bigwedge_{j=1}^{m} D_j(\boldsymbol{z})$ be QBFs consisting of quantifier prefixes Λ and Π over the disjoint sets of variables $\boldsymbol{x}$ and $\boldsymbol{z}$ respectively, and of clauses C_i and D_j over the corresponding variables.*

Then define

$$\Phi \otimes \Psi := \Lambda(\boldsymbol{x})\Pi(\boldsymbol{z}_1)\dots\Pi(\boldsymbol{z}_n) \cdot \bigwedge_{i=1}^{n}\bigwedge_{j=1}^{m}(C_i(\boldsymbol{x}) \vee D_j(\boldsymbol{z}_i))$$

where each $\boldsymbol{z}_i$ is a fresh copy of the variables $\boldsymbol{z}$, distinct from $\boldsymbol{x}$, for each $i = 1, \dots, n$.

The new formula $\Phi \otimes \Psi$ is false if and only if Φ and Ψ are both false. We can combine a winning strategy for the universal variables of Φ with a winning strategy for the universal variables of Ψ to construct a strategy which must falsify some $C_i(\boldsymbol{x})$ and, for each i, will falsify some $D_j(\boldsymbol{z}_i)$. It is therefore the case that the strategy will falsify some $C_i(\boldsymbol{x}) \vee D_j(\boldsymbol{z}_i)$. Similarly, a winning strategy for the existential player for either Φ or Ψ will give a winning strategy for $\Phi \otimes \Psi$.

The proof size for $\Phi \otimes \Psi$ is bounded by the size of proofs required by Φ and Ψ.

Lemma 7. *Let $\Phi = \boldsymbol{\mathcal{Q}}. \bigwedge_{i=1}^{n} C_i$ and $\Psi = \boldsymbol{\mathcal{S}}. \bigwedge_{j=1}^{m} D_j$ be minimally unsatisfiable QBFs. Let $S_P(\Phi)$ be the size of the smallest P-proof for Φ (and similarly for other formulas). Then*

$$\max(S_P(\Phi), S_P(\Psi)) \leq S_P(\Phi \otimes \Psi) \leq S_P(\Phi) + n \cdot S_P(\Psi).$$

Moreover, if P is QU-Res, then $S_P(\Phi \otimes \Psi) = S_P(\Phi) + n \cdot S_P(\Psi)$.

Proof. All clauses of $\Phi \otimes \Psi$ are necessary for a refutation. By assigning variables from Φ or the copies of Ψ appropriately, the lines in the proof can be restricted to a refutation of Φ or Ψ, and so $max(S_P(\Phi), S_P(\Psi)) \leq S_P(\Phi \otimes \Psi)$. Since $\Phi \otimes \Psi$ can be refuted by first deriving each clause C_i from $\bigwedge_{j=1}^{m}(C_i(\boldsymbol{x}) \vee D_j(\boldsymbol{z}_i))$, which can be done in $S_P(\Psi)$, and then refuting $\bigwedge_{i=1}^{n} C_i(\boldsymbol{x})$ with size $S_P(\Phi)$, we can find a refutation of $\Phi \otimes \Psi$ of size $S_P(\Phi) + n \cdot S_P(\Psi)$.

As noted, by restricting the variables we can construct a refutation of $\Phi(\boldsymbol{x})$ and each $\Psi(\boldsymbol{z}_i)$ assigning variables. In QU-Res, each resolution step or ∀-reduction step can only be performed on one variable, and so will only remain in one of these proofs, being replaced by a weakening or trivial step in all others. Any QU-Res proof of $\Phi \otimes \Psi$ must therefore have size at least $S_P(\Phi) + n \cdot S_P(\Psi)$. Equality comes from the upper bound above. □

We use this method to construct a family of false QBFs that require exponential-size proofs in QU-Res. These QBFs are the product of propositional formulas hard for Resolution and of QBFs easy for QU-Res, so the hardness of the product is clearly derived from the propositional lower bound. Yet, these product formulas are also hard for relaxing QU-Res. The QBF is obtained by taking the product of the pigeonhole principle, defined below, and the formulas by Kleine Büning et al. [23] as defined in Definition 3 above.

Definition 8. *The* pigeonhole principle *for m pigeons and n holes, denoted PHP_n^m, is the CNF*

$$\bigwedge_{i=1}^{m}(x_{i,1} \vee \dots \vee x_{i,n}) \wedge \bigwedge_{j=1}^{m}\bigwedge_{1 \leq i_1 < i_2 \leq n}(\neg x_{i_1,j} \vee \neg x_{i_2,j})$$

.

For $m > n$, this is unsatisfiable, and for $m = n + 1$ it has been shown that $2^{\Omega(n)}$ clauses are required to refute it in Resolution, and indeed in any constant-depth Frege system [18, 24, 28].

Theorem 9. *The QBFs* $\Phi_n := \mathrm{PHP}_n^{n+1} \otimes \mathrm{KBKF}_n$ *require relaxing QU-Res proofs of size* $2^{\Omega(n)}$.

Since QU-Res when restricted to a propositional formula is equivalent to Resolution, and PHP_n^{n+1} requires proofs of size $2^{\Omega(n)}$ in Resolution [18], we know that PHP_n^{n+1} requires QU-Res proofs of size at least $2^{\Omega(n)}$. In QU-Res, it is known that the formulas KBKF_n have linear-size proofs [16]. Given the proof size bounds on Φ_n given by Lemma 7, this QU-Res lower bound for Φ_n is unambiguously due to the lower bound for PHP_n^{n+1} in Resolution.

We first show that any relaxation of the quantifier prefix of KBKF_n is true.

Lemma 10. *Any relaxation of the quantifier prefix of* KBKF_n *to a* Π_t^b*-prefix results in a true QBF, for any* $t < n$.

Proof. To produce a Π_t^b-relaxation of the quantifier prefix, for $t < n$, there must be some k such that either x_k is quantified existentially, or x_k is quantified to the left of y_k and y'_k. In either case, we can construct a winning strategy for the existential player.

If some x_k is now quantified existentially, then a winning strategy for the existential player is to play $y_i = 0$, $y'_i = 1$ for each $i \leq k$, and to play $y_j = y'_j = 1$ for each $j > k$. Finally, playing $y_{n+i} = 1$ for each i then satisfies every clause apart from $y_{k-1} \vee \neg x_k \vee \neg y_{k+1} \vee \neg y'_{k+1}$, which can be satisfied by playing $x_k = 0$.

If some x_k is universally quantified to the left of y_k, y'_k, then the strategy for the existential variables is as above, except on the variables y_k and y'_k. When assigning these variables, the existential strategy looks at the value of x_k. If $x_k = 0$, then play $y_k = 0$, $y'_k = 1$. If $x_k = 1$, then play $y_k = 1$, $y'_k = 0$. This strategy will then satisfy all clauses. □

Any clause in the variables of Φ_n can be written as $X \vee Z_1 \vee \cdots \vee Z_m$ where X is a clause in the variables of $\boldsymbol{x}$, and Z_i is a clause in the variables of $\boldsymbol{z}_i$. We use the terms Z-variables and X-variables to refer to any variables in $\boldsymbol{z}_1, \ldots, \boldsymbol{z}_m$ and $\boldsymbol{x}$ respectively. Similarly, given a clause C, we use *X-clause* and *Z-clause* to refer to the restriction of C to the X-variables and Z-variables, and denote these restrictions by C^X and C^Z.

To prove Theorem 9, we first show that, for any clause A derived as an axiom by relaxing QU-Res, if A^X requires at least c clauses from PHP_n^{n+1} to prove, then it must also contain at least c existentially quantified Z-variables (Lemma 11).

We then establish an upper bound on the size of a proof of an X-clause derived from c axioms of PHP_n^{n+1} which depends only on c (Lemma 12). Using this, we conclude that any relaxing QU-Res axiom where the corresponding X-clause requires proofs of size 2^k must contain $\Omega(k)$ Z-variables (Corollary 13).

Lastly, we show that given any relaxing QU-Res proof, for each assignment to the Z-variables, we can find an axiom containing $\Omega(n)$ Z-variables which agrees with the given Z-assignment (Lemma 14). From this, we conclude that the proof must contain $2^{\Omega(n)}$ axioms.

Lemma 11. *Suppose that the clause $A = A^X \vee A^Z$ is derived as an axiom of Φ_n by relaxing QU-Res. Let $Z_{i_1}, \ldots, Z_{i_l}$ be such that all the existential variables in A^Z are in some Z_{i_j}. Then the clause A^X is a semantic consequence of the pigeonhole principle axioms $C_{i_1}, \ldots, C_{i_l}$, i.e. $C_{i_1} \wedge \cdots \wedge C_{i_l} \models A^X$.*

Proof. Suppose that $C_{i_1} \wedge \cdots \wedge C_{i_l} \not\models A^X$. Let α be an assignment to the X-variables which falsifies A^X but satisfies each C_{i_j}. We can extend α to the minimal assignment α' which falsifies A. We show that for any Π^b_t-relaxation of Φ_n, for $t < n$, we can extend α' to a winning strategy for the existential player.

Given a Π^b_t-relaxation of Φ_n, with quantifier prefix $\boldsymbol{\mathcal{Q}}'$, we show by induction that for each k, we can construct a strategy σ_k on the existential variables of X and $Z_1, \ldots, Z_k$ which extends α' and is a winning strategy for

$$\boldsymbol{\mathcal{Q}}' \cdot \bigwedge_{i=1}^{k} \bigwedge_{j=1}^{m} (C_i(\boldsymbol{x}) \vee D_j(\boldsymbol{z}_i))$$

Let $\sigma_0 := \alpha'$. This clearly satisfies the empty conjunction. For each k, we extend the strategy σ_{k-1} which satisfies $\bigwedge_{i=1}^{k-1} \bigwedge_{j=1}^{m} (C_i(\boldsymbol{x}) \vee D_j(\boldsymbol{z}_i))$. It therefore suffices to find a strategy for the unassigned Z_k variables which satisfies $\bigwedge_{j=1}^{m} (C_k(\boldsymbol{x}) \vee D_j(\boldsymbol{z}_k))$. We divide into two possible cases:

- Suppose $k = i_j$ for some $1 \leq j \leq l$. Then α', and hence σ_{k-1}, already satisfies $C_k(\boldsymbol{x})$. Therefore σ_{k-1} also satisfies each clause $C_k(\boldsymbol{x}) \vee D(\boldsymbol{z}_i)$ for any D, and we define $\sigma_k = \sigma_{k-1}$ on the variables where σ_{k-1} is defined. We can define σ_k arbitrarily on the remaining variables of Z_k.
- Suppose $k \neq i_j$ for any $1 \leq j \leq l$. Then A^Z does not contain any existential variables in Z_k so α', and hence σ_{k-1}, are not defined on any existential variables in Z_k. Any Π^b_t-relaxation of KBKF$_n$ is true, by Lemma 10. Let τ_k be a strategy for the existential variables of Z_k which is a winning strategy for $\boldsymbol{\mathcal{Q}}' \cdot \bigwedge_{j=1}^{m} D_j(\boldsymbol{z}_k)$, and so also for $\boldsymbol{\mathcal{Q}}' \cdot \bigwedge_{j=1}^{m} C_k \vee D_j(\boldsymbol{z}_k)$.
  As σ_{k-1} is not defined on any existential variables from Z_k, τ_k and σ_{k-1} are strategies for disjoint sets of variables. We extend our strategy σ_{k-1} with τ_k to give σ_k, a winning strategy for $\boldsymbol{\mathcal{Q}}' \cdot \bigwedge_{i=1}^{k} \bigwedge_{j=1}^{m} (C_i(\boldsymbol{x}) \vee D_j(\boldsymbol{z}_i))$.

The final strategy σ_n is therefore a winning strategy for the existential variables of the Π^b_t-relaxation of Φ_n, and σ_n extends the assignment α'. This suffices to show that the relaxation of $\Phi_n[\alpha']$ is true. Since α' extends β, the minimal assignment falsifying A, with assignments to existential variables only, the procedure detailed here will construct a winning existential strategy for any Π^b_t-relaxation of $\Phi[\beta]$, and so any Π^b_t-relaxation of $\Phi[\beta]$ is true. This does not satisfy the axiom derivation rules of relaxing QU-Res, and so A cannot be derived as an axiom in this system. □

This is enough to show that if we use relaxing QU-Res to derive an axiom A, and A^X requires at least l axioms from PHP_n^{n+1} in any proof, then A must contain existential variables from at least l different Z_i. In particular, A contains at least l distinct Z-variables.

The next lemma gives an upper bound for the size of Resolution proofs from a fixed number of axioms from PHP_n^{n+1}. This upper bound also applies to the length of a Resolution proof of the X-clause of an axiom containing a small number of Z-variables.

Lemma 12. *Suppose C is a clause derived by Resolution from PHP_n^{n+1}, and there exist axioms $C_1, \dots, C_t$ from PHP_n^{n+1} such that $C_1 \wedge \dots \wedge C_t \models C$. Then there is a Resolution proof of C of size at most 18^t.*

Combining this with Lemma 11 shows that any relaxing QU-Res axiom A for which A^X requires a large QU-Res derivation from the axioms of the pigeonhole principle must also contain a large number of Z-variables.

Corollary 13. *Let A be an axiom derived from Φ_n by relaxing QU-Res. Let $S(A^X)$ be the size of the smallest Resolution derivation of A^X from PHP_n^{n+1}. Then A must contain at least $\frac{1}{\log 18} \log S(A^X)$ existential Z-variables.*

Proof (Lemma 12). We show that without weakening, which can be done in one step at the end if needed, there are at most 18^t clauses that can be derived by Resolution from t axioms of PHP_n^{n+1}. This upper bound is far from tight, but is sufficient for the proof of Theorem 9.

Given t clauses from PHP_n^{n+1}, all negative literals are in clauses of size 2. Thus there are at most $2t$ variables x_i which appear in both positive and negative literals in the clauses $C_1, \dots, C_t$. There remain at most t blocks Y_j of pure positive literals or pure negative literals, at most one corresponding to each C_i. Any clause derived by Resolution from $C_1, \dots, C_t$ must contain each variable x_i as a positive literal, a negative literal or not at all, and must contain some subset of the blocks of pure literals. Thus the total number of clauses derivable in Resolution from $C_1, \dots, C_t$ is at most $3^{2t} \cdot 2^t = 18^t$. Any Resolution derivation of C from $C_1, \dots, C_t$ therefore has size at most 18^t. □

The last result we need to prove Theorem 9 is to show that for any existential Z-assignment α, the restriction of a proof of Φ_n by α results in a refutation of (X-axioms derived from) PHP_n^{n+1}.

Lemma 14. *Given a relaxing QU-Res proof π and an assignment α to the existential Z-variables of Φ_n, $\pi|_\alpha^X$ contains a sound Resolution refutation of the X-axioms corresponding to axioms agreeing with α.*

Proof. Consider $\pi|_\alpha$, the result of restricting π to those clauses which agree with α. We show by induction that ${\pi|_\alpha}^X$ is a Resolution refutation from the X-axioms, of size at most $f(n)$.

- The empty clause is the root of the Resolution proof on the X-variables, and clearly agrees with α.

- Suppose a clause C is derived by a $\forall$-red step on a Z-variable u. Then clearly $C \vee u$ agrees with α if C agrees with α, since α does not assign u. Also $C^X = (C \vee u)^X$, so this is a sound step in a Resolution refutation.
- Suppose C agrees with α and C is derived from C_1 and C_2 by resolving on an X-variable x. Then $C_1^Z, C_2^Z \subseteq C^Z$, and so both C_1 and C_2 agree with α since C does so. Observe also that C^X is derived from C_1^X and C_2^X by a single Resolution step on x.
- Suppose C agrees with α and C is derived from C_1 and C_2 by resolving on a Z-variable z. Then at least one of C_1 and C_2 must agree with α, depending on the value of $\alpha(z)$. As $C_1^X, C_2^X \subseteq C^X$, we can derive C^X by a weakening step from whichever agrees with the Z-assignment, or both if z is universally quantified.

This completes our induction, and proves that the X-clauses of the clauses in π which agree with α are a valid Resolution proof. □

We are now ready to use these lemmas to prove our lower bound for relaxing QU-Res.

Theorem 9. *The QBFs $\Phi_n := \mathrm{PHP}_n^{n+1} \otimes \mathrm{KBKF}_n$ require relaxing QU-Res proofs of size $2^{\Omega(n)}$.*

Proof. Suppose that π is a relaxing QU-Res proof of Φ_n with $|\pi| = f(n)$. Given an assignment α to the existential Z-variables, $\pi|_\alpha^X$ is a sound Resolution refutation of the X-axioms (Lemma 14), and has at most $f(n)$ axioms. Since any Resolution refutation of PHP_n^{n+1} requires proofs of size at least 2^{kn} for some constant k, some X-axiom B in $\pi|_\alpha^X$ requires a Resolution derivation of size at least $\frac{2^{kn}-f(n)}{f(n)} = \frac{2^{kn}}{f(n)} - 1$. By Corollary 13, there is an axiom A in π such that $A^X = B$, and so A contains at least $c(kn - \log f(n)) =: g(n)$ existential Z-variables, which agree with α.

For every assignment α to the existential Z-variables, we can find such an axiom containing at least $g(n)$ existential Z-variables and agreeing with a. As each of these axioms can agree with at most a $2^{-g(n)}$ proportion of the possible assignments α, π must contain at least $2^{g(n)}$ axioms. As a proof cannot contain more axioms than its length, we conclude that $2^{g(n)} \leq f(n)$, i.e.

$$2^{ckn} \leq f(n) 2^{c \log f(n)} = f(n)^{c+1}$$

and so $f(n) = 2^{\Omega(n)}$. □

We have shown that $\mathrm{PHP}_n^{n+1} \otimes \mathrm{KBKF}_n$ requires proofs of size $2^{\Omega(n)}$ in relaxing QU-Res, despite consisting of a propositional formula which is hard for Resolution combined with a QBF which is easy for QU-Res.

5 An alternative definition of hardness from alternation

In this section, we define a new set of proof systems which better characterise whether a QBF lower bound is due to the alternation of quantifiers, or due to a propositional lower bound. In these proof systems, $\mathrm{PHP}_n^{n+1} \otimes \mathrm{KBKF}_n$ has linear-size proofs.

Definition 15. *A Σ_k^p-QU-Res proof of a QBF Φ is a derivation of the empty clause by any of the rules of QU-Res, or the Σ_k^p-*derivation *rule*

$$\frac{C_1 \ \dots \ C_l}{D}$$

for any l, where there is some Σ_k^b-relaxation Π' of the quantifier prefix Π such that $\Pi'. \bigwedge_{i=1}^{l} C_i \models \Pi'.D \wedge \bigwedge_{i=1}^{l} C_i$.

In the context of these proof systems, we define a Σ_k^b-relaxation of a quantifier prefix as in Definition 5, i.e. any movement of universally quantified variables to the left. We also allow replacing any $\forall$ quantifier by $\exists$. Allowing this replacement is not necessary, but as shown in Lemma 19, it allows us to restrict our attention to Σ_{2k+1}^p-QU-Res, eliminating the need for a similarly defined Π_m^p-QU-Res.

It is straightforward to define Σ_k^p-P similarly for any QBF proof system P which works with proof lines, and several of the following results will hold for any suitable P. For simplicity, we state and prove these results only for QU-Res.

The completeness of Σ_k^p-QU-Res is clear since any QU-Res proof is also a Σ_k^p-QU-Res proof. To demonstrate the soundness, note that QU-Res (with weakening) is both sound and inferentially complete [16]. Thus we can replace any Σ_k^p-derivation with a QU-Res derivation consistent with the Σ_k^b-relaxation. This QU-Res derivation will therefore also be consistent with the original quantifier prefix, and so from any Σ_k^p-QU-Res refutation, we can construct a QU-Res refutation. Since QU-Res is known to be sound, Σ_k^p-QU-Res is also sound.

We can now suggest our definition of hardness due to quantifier alternation.

Definition 16. *A family of QBFs is* hard due to quantifier alternation *if it requires superpolynomial-size Σ_1^p-QU-Res refutations.*

A QBF family has alternation hardness Σ_k^p *if it has polynomial-size proofs in Σ_k^p-QU-Res, but requires superpolynomial-size proofs in Σ_{k-1}^p-QU-Res.*

The proof complexity of formulas in Σ_1^p-QU-Res is of particular interest, as recent success in SAT solving has resulted in some QBF solvers embedding a SAT solver as a black box [20, 31]. The oracle access to Σ_1^p models this technique, and may provide some insight as to the power and limitations of such QBF solvers.

As noted in Section 4, the formulas $\mathrm{PHP}_n^{n+1} \otimes \mathrm{KBKF}_n$ require QU-Res proofs of size $2^{\Omega(n)}$ due to the lower bound on Resolution. Here we show that these formulas have polynomial-size proofs in Σ_1^p-QU-Res, even using only a single Σ_1^p-derivation, and so are not hard for QU-Res due to quantifier alternation. This is in sharp contrast with the lower bound shown in Theorem 9 for relaxing QU-Res, despite this proof system also making use of oracles for Σ_k^p.

Theorem 17. $\mathrm{PHP}_n^{n+1} \otimes \mathrm{KBKF}_n$ *have Σ_1^p-QU-Res proofs of length $O(n^3)$.*

Proof. Define the clauses C_i and D_j such that $\mathrm{PHP}_n^{n+1} = \bigwedge_i C_i$ and $\mathrm{KBKF}_n = \Pi \cdot \bigwedge_j D_j$, and so the clauses of $\mathrm{PHP}_n^{n+1} \otimes \mathrm{KBKF}_n$ are $C_i(\boldsymbol{x}) \vee D_j(\boldsymbol{z}_i)$ for all i, j.

Since there is an $O(n)$-length refutation of KBKF_n in QU-Res, we know that QU-Res can derive $C_i(\boldsymbol{x})$ from $\bigwedge_j C_i(\boldsymbol{x}) \vee D_j(\boldsymbol{z}_i)$ in $O(n)$ lines. There are $O(n^2)$ clauses C_i in PHP_n^{n+1}, so there is a QU-Res derivation of $\bigwedge_i C_i(\boldsymbol{x})$ in $O(n^3)$ lines. All the variables in $\boldsymbol{x}$ are existentially quantified, and PHP_n^{n+1} is false, thus from $\bigwedge_i C_i(\boldsymbol{x})$, the empty clause can be derived in a single Σ_1^p-derivation step. □

Any clause derived as an axiom of Φ using a Σ_k^p-oracle by relaxing QU-Res can also be derived from the clauses of Φ by a single Σ_k^p-derivation in Σ_k^p-QU-Res. It is easy to see from this that Σ_k^p-QU-Res p-simulates relaxing QU-Res. Tbeorem 17 shows an exponential separation.

In order to compare this characterisation of lower bounds by quantifier alternation with the characterisation given in Section 3, we first show that Σ_1^p-QU-Res still has the same strategy extraction property as QU-Res. Analogous results apply for $\mathcal{C}$-Frege+∀-Red systems with strategy extraction in the appropriate circuit classes.

Lemma 18. *Σ_1^p-QU-Res has strategy extraction by depth-3 Boolean circuits.*

Proof. QU-Res is known to have strategy extraction by depth-3 Boolean circuits [4]. We extend this result to Σ_1^p-QU-Res by showing that Σ_1^p-derivations do not contain any information on the strategy for the universal player.

From any Σ_k^p-QU-Res proof we can construct a QU-Res proof by replacing the Σ_k^p-derivation steps with a QU-Res derivation of the clauses. This is possible by the inferential completeness of QU-Res, and furthermore each Σ_k^p-derivation can be replaced by a QU-Res derivation consistent with the Σ_k^b-relaxation.

In the case of Σ_1^b, the relaxation of the prefix treats all variables as existential. A QU-Res proof constructed in this way, while potentially much larger than the Σ_1^p-QU-Res proof, does not contain any additional ∀-reduction steps that were not in the Σ_1^p-QU-Res proof. Strategy extraction for QU-Res constructs a strategy which is polynomial in the number of ∀-reduction steps of the proof, as noted in [4]. Given any Σ_1^p-QU-Res proof, it is therefore possible to extract a strategy for the universal variables as a depth-3 Boolean circuit with size polynomial in the length of the proof. □

As a consequence of Lemma 18, QBFs hard for QU-Res by item 1 of Theorem 2 (hardness due to strategy extraction) are therefore still hard for Σ_1^p-QU-Res. Intuitively, we expect lower bounds due to strategy extraction to also be lower bounds due to alternation, as strategy extraction is a technique which inherently relies on universally quantified variables and the order of the quantification.

Consider now QBFs hard for QU-Res by item 2 in Theorem 2. There are polynomial-size strategies for the universal variables, but for all of these, the witnessed formulas require superpolynomial-size proofs in Resolution. Using the normal form for proofs described in [8], we can construct short proofs of these QBFs in Σ_1^p-QU-Res, deriving the witnessed formula, and then using a Σ_1^p-derivation to derive $\bot$. This demonstrates that QBFs in the second category are not hard due to alternation of quantifiers.

For sufficiently strong proof systems, such as Frege+$\forall$-red, these are the only two possible reasons for hardness [8]. As Lemma 18 extends naturally to Σ_1^p-Frege+$\forall$-red, the characterisation of hardness for QBF Frege systems in [8] (circuit lower bounds vs propositional Frege lower bounds) therefore coincides with our characterisation via quantifier alternation.

6 Alternation Hardness of Specific Formulas

In this section we determine the precise alternation hardness of specific families of QBFs. We consider three different families, one from each of the categories defined in Theorem 2. While not all formulas from the same category will necessarily have the same alternation hardness, the bounds shown here reinforce the distinctions shown in Theorem 2.

The first step in establishing the alternation hardness of these formulas is to understand which levels are necessary to consider. Since the definition of relaxation allows replacing universal quantifiers with existential quantifiers, we can limit the proof systems under consideration to Σ_k^p-QU-Res for odd k.

Lemma 19. *If a family of QBFs has proofs of size $s(n)$ in Π_m^p-QU-Res or Σ_{2k}^p-QU-Res, then it has proofs of size $n \cdot s(n)$ in Σ_{m-1}^p-QU-Res or Σ_{2k-1}^p-QU-Res respectively.*

In particular, given a family of QBFs Φ_n, if the alternation hardness of Φ_n is precisely $\mathcal{C}$, then $\mathcal{C} = \Sigma_{2k+1}^b$ for some integer k.

Proof. We begin by demonstrating that from a Π_m^p-QU-Res refutation of Φ_n of size $s(n)$, we can construct a Σ_{m-1}^p-QU-Res refutation of size $O(s(n))$.

Consider the outermost block of universal variables in a Π_m^b-relaxation. A Σ_{m-1}^b-relaxation can be obtained by quantifying the variables in this block existentially. If a Π_m^p-derivation does not derive the empty clause, then all possible clauses derived by the Π_m^p-derivation contain at least one variable quantified existentially in the Π_m^b-relaxation. Thus we can still derive the same clauses using the Σ_{m-1}^b relaxation, as at no point would any QU-Res proof consistent with the Π_m^b-relaxation contain a $\forall$-reduction step on these universal variables. If the Π_m^p-derivation does derive the empty clause, then it is possible in the Σ_{m-1}^b-relaxation to derive a clause containing only variables which were universally quantified in the first block in the Π_m^b-relaxation. As these variables must be universally quantified in the original QBF, there is a proof using a Σ_{m-1}^b-relaxation of size $\leq p(n) + n$, which replaces the Π_m^p-deduction with a Σ_{m-1}^p-deduction and at most n $\forall$-reduction steps.

Given a Σ_{2k}^b-relaxation of the quantifier prefix, the innermost block of variables is universally quantified. By the definition of relaxation, these variables must also have been innermost in the original quantifier prefix. The first step in a Σ^p_{2k-1}-QU-Res proof is to $\forall$-reduce these variables in each axiom. The Σ^p_{2k}-QU-Res proof is then followed, with the innermost variables removed from the clauses. At each Σ^p_{2k}-derivation, the innermost variables are not present in any of the clauses, and so the Σ^b_{2k}-relaxation can be replaced by a Σ^b_{2k-1}-relaxation with these variables also existentially quantified. □

If the definition of relaxation were restricted to that of relaxing QU-Res, then the simulation of Π^p_m-QU-Res by Σ^p_{m-1}-QU-Res would not hold. With the exception of Σ^p_2-QU-Res, it would still be possible to reduce a Σ^p_{2k}-QU-Res proof to a Σ^p_{2k-1}-QU-Res proof by moving the innermost universal variables outwards to another block of universal quantifiers.

Lemmas 18 and 19 allow us to determine the precise alternation hardness of QParity_n, which were introduced in [6] as examples of formulas which are hard due to strategy extraction (item 1 in Theorem 2).

Definition 20 ([6]). *The formulas* QParity_n *consist of the quantifier prefix* $\exists x_1 \dots x_n \forall z \exists t_2 \dots t_n$ *and clauses expressing that* $t_2 \equiv x_1 \oplus x_2$*,* $t_k \equiv t_{k-1} \oplus x_k$ *for each* $3 \leq k \leq n$*, and* $z \equiv \neg t_n$*.*

The QBFs are false, and the only winning strategy for the $\forall$ player is to play $z \equiv \bigoplus_{i=1}^n x_i$. However, the parity function is hard to compute for depth-3 circuits [15,19], and so any QU-Res proof requires length $\Omega(2^n)$. The formulas QParity_n are therefore hard due to strategy extraction as defined in Theorem 2.

Corollary 21. *The formulas* QParity_n *have* Σ^b_3*-alternation hardness.*

The fact that the formulas QParity_n have Σ^b_3-alternation hardness shows that they are hard for QU-Res due to the alternation of quantifiers.

Proof. It is clear that QParity_n has short proofs in Σ^p_3-QU-Res, as their quantifier prefix is Σ^b_3. By Lemma 19, we need only show that QParity_n does not have polynomial size proofs in Σ^p_1-QU-Res. By Lemma 18, Σ^p_1-QU-Res has strategy extraction by depth-3 circuits. Since any depth-3 circuit for the parity function requires exponential size [15,19], any Σ^p_1-QU-Res refutation of QParity_n requires exponential size. □

It is clear that all formulas which fall into the second category of Theorem 2, of being hard only due to a lower bound on Resolution, have polynomial-size proofs in Σ_1-QU-Res, and so have Σ^p_1-alternation hardness.

The last family of QBFs we consider is KBKF'_n. By Theorem 4, the formulas KBKF'_n are hard for QU-Res due to a genuine QBF lower bound. As their hardness does not originate from a Resolution lower bound, we might expect them to be hard due to alternation. In fact, we can go further than this and show that the formulas KBKF'_n are hard for Σ^p_k-QU-Res for all k.

Theorem 22. *The formulas* KBKF'_n *require proofs of size* $2^{\Omega(n)}$ *in* Σ^p_k*-QU-Res for any constant* k*.*

Proof. Throughout this proof, we refer only to universal variables x_i. The variables x_i and z_i cannot be resolved until a $\forall$-reduction on the other has taken place, and both are quantified together. Thus whenever there is a $\forall$-reduction step on x_i, we can assume z_i is also $\forall$-reduced, and so x_i and z_i appear together with the same polarity. In the case of a relaxation of the quantifier prefix, x_i refers to the variable of x_i and z_i which is quantified furthest left.

QU-Res (with weakening) is implicationally complete, and so from any Σ_k^p-QU-Res proof we can construct a QU-Res proof by replacing each Σ_k-derivation with an appropriate series of QU-Res steps. The $\forall$-reduction steps replacing a given Σ_k^p-derivation are consistent with some Σ_k^b-relaxation of the quantifier prefix. We show a lower bound on the size of a Σ_k^p-QU-Res refutation of KBKF'_n by examining the QU-Res proof we obtain in this way. As all universal variables in KBKF'_n appear with another universal variable of the same polarity, at no point can there be a resolution step on universal variables. Thus once a clause contains a universal variable, the only way it can be removed from descendants of this clause is by $\forall$-reduction.

As observed in [23], before a $\forall$-reduction step on any clause is possible, the clause must contain a literal on all universal variables. Furthermore, all possible sets of literals on all universal variables are necessary for the QU-Res refutation. This observation can be extended to show that for the first $\forall$-reduction step on x_i, the clause contains literals on all previous universal literals.

Lemma 23. *If a clause derived from* KBKF'_n *contains a literal on* x_i*, and the derivation does not contain a* $\forall$*-reduction step on* x_i*, then it contains* y_j *or* y'_j *for some* $i \leq j \leq 2n$*.*

A further effect of Lemma 23 is to show that whenever the first $\forall$-reduction step is performed on the variable x_i, the clause must contain either y_i or y'_i.

Now suppose that π is a Σ_k^p-QU-Res proof of KBKF'_n. Let α be one of the 2^n possible assignments to the universal variables of KBKF'_n which the universal player may be required to play. We show that there is some clause in π which contains at least $n-k$ literals on universal variables and agrees with α.

Let π' be a QU-Res proof obtained by expanding the Σ_k^p-derivations of π. From the observations in [23], given an assignment α, there is some clause $C_{n-k} \in \pi'$ which is derived by a $\forall$-reduction step on x_{n-k}, such that C_{n-k} is not preceded by any $\forall$-reduction steps on $x_1, \ldots, x_{n-k}$, and the universal literals in C_{n-k} agree with α. In particular, this means that C_{n-k} contains literals on all universal variables left of x_{n-k}.

We look now at the derivation of C_{n-k} in π'. There must be some clause C_{n-k+1} derived by a $\forall$-reduction on x_{n-k+1} with no preceding such $\forall$-reduction. Construct clauses $C_{n-k+2}, \ldots, C_n$. Consider now the path through π' from C_n to C_{n-k} through each C_i. Since C_n contains literals on all universal variables, and the universal literals of C_{n-k} agree with α, all clauses on this path must contain literals on $x_1, \ldots, x_{n-k-1}$ agreeing with α.

We show that at least one clause in this path must also be in π. If this were not the case, then $C_n, \ldots, C_{n-k}$ are all in the expansion of a single Σ_k^p-derivation.

By the choice of C_i, each C_i contains a literal on y_i or y'_i by Lemma 23. The derivation of C_i by a $\forall$-reduction is therefore only possible if the corresponding Σ_k^b-relaxation quantifies x_i universally and to the right of y_i, y'_i. However if this were the case for each $n-k \leq i \leq n$, the relaxation would require at least $2k$ alternations of quantifiers. Thus there is some clause D on the path from C_n to C_{n-k} such that $D \in \pi$ and D contains literals on $x_1, \ldots, x_{n-k-1}$ agreeing with α.

There are 2^{n-k-1} possible assignments to $x_1, \ldots, x_{n-k-1}$ that α could define, and for each there is a clause in π which contains literals on all of these variables agreeing with α. The size of any Σ_k^p-QU-Res proof is therefore at least $2^{\Omega(n)}$. □

Proof (Lemma 23). No resolution steps on x_i are possible, so assume that the literal on x_i is a positive literal. The case for $\neg x_i$ is similar.

Suppose that the axiom introducing x_i is $x_i \vee y_{n+i}$. In this case, all existential variables are quantified to the right of x_i. No resolution steps on universal pivots are possible, and the only axiom which contains existential variables quantified both sides of x_i is $y'_i \vee x_i \vee \neg y_{i+1} \vee y'_{i+1}$. We can therefore assume that the literal x_i is introduced by this axiom.

Consider a clause derived from $y'_i \vee x_i \vee \neg y_{i+1} \vee \neg y'_{i+1}$. The only axioms which contain the literal $\neg y'_i$ also contain $\neg y_i$. The only axiom containing y_i also contains $\neg x_i$. Given a clause C containing x_i, derived from $y'_i \vee x_i \vee \neg y_{i+1} \vee \neg y'_{i+1}$ without any $\forall$-reduction steps on x_i, it cannot also be derived from $y_i \vee \neg x_i \vee \neg y_{i+1} \vee \neg y'_{i+1}$, as then C would contain both x_i and $\neg x_i$. Therefore, C must contain a literal on y_i or y'_i. □

7 Conclusion

We have undertaken an analysis of strategies and alternation as underlying reasons for the size of proofs in QBF proof systems. In the search for 'genuine' QBF lower bounds, these are the two characterisations which have received the most attention. We have shown that, for sufficiently strong proof systems (Frege and above), these two criteria are equivalent, and proposed a system for which all lower bounds are such proper QBF lower bounds.

A natural question is whether for weaker Resolution-based systems, QBFs from the third category of Theorem 2 are always hard due to alternation. Here we have only shown this for the special case of KBKF'_n. We also leave open the question of finding formulas which have alternation hardness precisely Σ_k^b for odd $k > 3$.

References

1. Valeriy Balabanov, Magdalena Widl, and Jie-Hong R. Jiang. QBF resolution systems and their proof complexities. In *Theory and Applications of Satisfiability Testing - SAT 2014*, pages 154–169, 2014.
2. Paul Beame and Toniann Pitassi. Simplified and improved resolution lower bounds. In *Symposium on Foundations of Computer Science (FOCS)*, pages 274–282, 1996.

3. Marco Benedetti and Hratch Mangassarian. QBF-based formal verification: Experience and perspectives. *JSAT*, 5(1-4):133–191, 2008.
4. Olaf Beyersdorff, Ilario Bonacina, and Leroy Chew. Lower bounds: From circuits to QBF proof systems. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 249–260, 2016.
5. Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. On unification of QBF resolution-based calculi. In *MFCS, II*, pages 81–93, 2014.
6. Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. Proof complexity of resolution-based QBF calculi. In *Proc. Symposium on Theoretical Aspects of Computer Science (STACS'15)*, pages 76–89. LIPIcs, 2015.
7. 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.
8. Olaf Beyersdorff and Ján Pich. Understanding Gentzen and Frege systems for QBF. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS*, pages 146–155, 2016.
9. Maria Luisa Bonet, Juan Luis Esteban, Nicola Galesi, and Jan Johannsen. On the relative complexity of resolution refinements and cutting planes proof systems. *SIAM Journal on Computing*, 30(5):1462–1484, 2000.
10. Samuel R. Buss. Towards NP-P via proof complexity and search. *Ann. Pure Appl. Logic*, 163(7):906–917, 2012.
11. Hubie Chen. Proof complexity modulo the polynomial hierarchy: Understanding alternation as a source of hardness. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP*, pages 94:1–94:14, 2016.
12. Stephen A. Cook and Phuong Nguyen. *Logical Foundations of Proof Complexity.* Cambridge University Press, 2010.
13. Stephen A. Cook and Robert A. Reckhow. The relative efficiency of propositional proof systems. *J. Symb. Log.*, 44(1):36–50, 1979.
14. 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.
15. Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, 1984.
16. Allen Van Gelder. Contributions to the theory of practical quantified Boolean formula solving. In *Principles and Practice of Constraint Programming - 18th International Conference, CP*, pages 647–663, 2012.
17. Enrico Giunchiglia, Paolo Marin, and Massimo Narizzano. Reasoning with quantified boolean formulas. In *Handbook of Satisfiability*, pages 761–780. IOS Press, 2009.
18. Armin Haken. The intractability of resolution. *Theor. Comput. Sci.*, 39:297–308, 1985.
19. Johan Håstad. Almost optimal lower bounds for small depth circuits. In S. Micali, editor, *Randomness and Computation, Advances in Computing Reasearch, Vol 5*, pages 143–170. JAI Press, 1989.
20. Mikolás Janota, William Klieber, Joao Marques-Silva, and Edmund M. Clarke. Solving QBF with counterexample guided refinement. *Journal of Artificial Intelligence*, 234:1–25, 2016.
21. Mikolás Janota and Joao Marques-Silva. Expansion-based QBF solving versus Q-resolution. *Theor. Comput. Sci.*, 577:25–42, 2015.
22. Emil Jerábek. Dual weak pigeonhole principle, Boolean complexity, and derandomization. *Ann. Pure Appl. Logic*, 129(1-3):1–37, 2004.

23. Hans Kleine Büning, Marek Karpinski, and Andreas Flögel. Resolution for quantified Boolean formulas. *Inf. Comput.*, 117(1):12–18, 1995.
24. Jan Krajícek, Pavel Pudlák, and Alan R. Woods. An exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. *Random Struct. Algorithms*, 7(1):15–40, 1995.
25. Jan Krajíček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*, volume 60 of *Encyclopedia of Mathematics and Its Applications*. Cambridge University Press, Cambridge, 1995.
26. Florian Lonsing and Uwe Egly. Evaluating QBF solvers: Quantifier alternations matter. *CoRR*, abs/1701.06612, 2017.
27. Florian Lonsing, Uwe Egly, and Martina Seidl. Q-resolution with generalized axioms. In *Theory and Applications of Satisfiability Testing (SAT'16)*, pages 435–452, 2016.
28. Toniann Pitassi, Paul Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3:97–140, 1993.
29. Jussi Rintanen. Asymptotically optimal encodings of conformant planning in QBF. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*, pages 1045–1050, 2007.
30. John Alan Robinson. A machine-oriented logic based on the resolution principle. *J. ACM*, 12(1):23–41, 1965.
31. Horst Samulowitz and Fahiem Bacchus. Using SAT in QBF. In *Principles and Practice of Constraint Programming, CP*, pages 578–592, 2005.
32. Nathan Segerlind. The complexity of propositional proofs. *Bulletin of Symbolic Logic*, 13(4):417–481, 2007.
33. Lintao Zhang and Sharad Malik. Conflict driven learning in a quantified Boolean satisfiability solver. In *Proceedings of the 2002 IEEE/ACM International Conference on Computer-aided Design, ICCAD*, pages 442–449, 2002.

A Chen's lower bound for relaxing QU-Res

Define $\Psi_n = \boldsymbol{\mathcal{Q}}_n \cdot \psi_n$ to be the quantified Boolean circuit consisting of the quantifier prefix $\boldsymbol{\mathcal{Q}}_n := \exists x_1 \forall y_1 \ldots \exists x_i \forall y_i \ldots \exists x_n \forall y_n$ and a (polynomial-sized) Boolean circuit ψ_n defined such that

$$\psi_n \leftrightarrow \sum_{i=1}^{n} (x_i + y_i) \not\equiv 0 \mod 3.$$

The quantified Boolean circuits Ψ_n then provide a lower bound for relaxing QU-Res.

Theorem 24 (Chen [11]). *Relaxing QU-Res requires proofs of size $\Omega(2^n)$ on Ψ_n.*

Lines in the relaxing QU-Res proof system are clauses, however there is no polynomial-size CNF equivalent to ψ_n.

Lemma 25. *Any CNF $\phi_n(\boldsymbol{x}, \boldsymbol{y})$ equivalent to $\psi_n(\boldsymbol{x}, \boldsymbol{y})$ must contain $\Omega(2^n)$ clauses.*

Proof. The circuit ψ_n has $2n$ input variables. For any assignment to $2n-1$ of these, the corresponding restriction of the circuit is not equivalent to 0. Any clause in an equivalent CNF must therefore contain literals on all $2n$ variables.

For each clause C in ϕ_n, there is therefore a unique assignment to $\boldsymbol{x}, \boldsymbol{y}$ which falsifies C. As each of the $\Omega(2^n)$ assignments on which ψ_n evaluates to 0 must falsify a clause, ϕ_n must contain $\Omega(2^n)$ clauses. □