Hardness Characterisations and Size-Width Lower Bounds for QBF Resolution*

OLAF BEYERSDORFF, Friedrich-Schiller-Universität Jena, Germany
JOSHUA BLINKHORN, Friedrich-Schiller-Universität Jena, Germany
MEENA MAHAJAN, The Institute of Mathematical Sciences, HBNI, Chennai, India
TOMÁŠ PEITL, TU Wien, Austria

We provide a *tight characterisation of proof size in resolution for quantified Boolean formulas (QBF) via circuit complexity.* Such a characterisation was previously obtained for a hierarchy of QBF Frege systems [16], but leaving open the most important case of QBF resolution. Different from the Frege case, our characterisation uses a new version of decision lists as its circuit model, which is stronger than the CNFs the system works with. Our decision list model is well suited to compute countermodels for QBFs. Our characterisation works for both *Q-Resolution and QU-Resolution.*

Using our characterisation we obtain a *size-width relation for QBF resolution* in the spirit of the celebrated result for propositional resolution [4]. However, our result is not just a replication of the propositional relation – intriguingly ruled out for QBF in previous research [12] – but shows a different dependence between size, width, and quantifier complexity. An essential ingredient is an improved relation between the size and width of term decision lists; this may be of independent interest.

We demonstrate that *our new technique elegantly reproves known QBF hardness results* and unifies previous lower-bound techniques in the QBF domain.

CCS Concepts: • **Theory of computation → Logic**; **Proof complexity**; **Circuit complexity**; **Automated reasoning**.

Additional Key Words and Phrases: quantified Boolean formulas, proof complexity, size-width tradeoff

ACM Reference Format:
Olaf Beyersdorff, Joshua Blinkhorn, Meena Mahajan, and Tomáš Peitl. 2023. Hardness Characterisations and Size-Width Lower Bounds for QBF Resolution. *ACM Trans. Comput. Logic* 24, 2, Article 10 (January 2023), 30 pages. https://doi.org/10.1145/3565286

1 INTRODUCTION

Proof complexity is a field at the intersection of logic and complexity that studies the difficulty of proving formal theorems, where difficulty of proving is associated with the size of proofs in different proof calculi. Obtaining lower bounds to the size of proofs is the central and most challenging goal in proof complexity, and the endeavour bears tight relations to central questions in computational complexity [24, 35] and first-order logic [5, 23]. In addition to this foundational quest, proof complexity has become the main theoretical tool for the analysis of powerful SAT solvers that routinely solve huge industrial instances of the NP-complete SAT problem [19, 41, 49].

*A preliminary version of this article appeared in the proceedings of the conference LICS 2020 [7].

Authors' addresses: Olaf Beyersdorff, Institut für Informatik, Friedrich-Schiller-Universität Jena, Ernst-Abbe-Platz 2, 07743, Jena, Thüringen, Germany, olaf.beyersdorff@uni-jena.de; Joshua Blinkhorn, Institut für Informatik, Friedrich-Schiller-Universität Jena, Ernst-Abbe-Platz 2, 07743, Jena, Thüringen, Germany, joshua.blinkhorn@uni-jena.de; Meena Mahajan, The Institute of Mathematical Sciences, HBNI, Chennai, CIT Campus, Chennai, India, 600 113, meena@imsc.res.in; Tomáš Peitl, Institute of Logic and Computation, TU Wien, Favoritenstraße 9-11, 1040, Vienna, Austria, peitl@ac.tuwien.ac.at.

Many conceptually different proof systems have been studied, but the *resolution system* [17, 45] – operating on clauses and using just one rule – has received by far the greatest attention. This is because resolution is a foundational system from the theoretical point of view [46], but also because resolution (and its subsystems) underpin modern SAT solving [19, 41], whereby lower bounds on resolution proof size provide lower bounds on solving time.

In the past two decades, researchers have tried to lift the successes of SAT solving and propositional proof complexity to even more computationally challenging settings, with *quantified Boolean formulas* (QBF) receiving key attention. As a PSPACE-complete problem, QBF widely generalises SAT and encompasses the polynomial hierarchy, a source of many practical problems [26, 34, 40] that are efficiently tackled by modern QBF solvers. As in the propositional case, QBF resolution systems play a key role in understanding the efficiency and limits of current solving. Arguably, the simplest QBF resolution system is QU-Res, augmenting propositional resolution by just one universal reduction rule [27, 33].

There is a long-standing belief in the proof complexity community (cf. [3]) that there exist strong connections between *the logical problem* of determining the size of the shortest proof for a given formula (proof size bounds) and *the complexity problem* of finding small circuits for explicit functions corresponding to the formula (circuit bounds).

While such a formal connection has so far appeared elusive for central propositional proof systems such as resolution or Frege systems, some connections are known, for example between algebraic proof systems and algebraic circuit complexity [28]. Arguably, the clearest such connection has been shown in the QBF domain, between the hierarchy of QBF Frege systems and the corresponding circuit classes. For QBF Frege (where lines are propositional formulas, i.e. NC^1 circuits) the connection manifests as follows: there are QBFs that require superpolynomial-size proofs in QBF Frege if, and only if, there are functions requiring superpolynomial-size NC^1 circuits or there are propositional formulas requiring superpolynomial-size propositional Frege proofs [16]. This characterisation unites central problems from circuit complexity (NC^1 lower bounds) with central problems from proof complexity (Frege lower bounds). However, such a connection has remained open for resolution systems (either QBF or propositional), which are of prime importance, theoretically and practically.

1.1 Our contributions

A. Characterising QU-Res **hardness on bounded alternation.** We obtain a *tight characterisation of* QU-Res *hardness in terms of circuit lower bounds.* More precisely, we show that a sequence of QBFs Q_n of bounded quantifier complexity requires superpolynomial QU-Res proofs if and only if each countermodel for Q_n requires superpolynomial circuit size (in a natural circuit model defined on decision lists as explained below) or if Q_n exhibits propositional resolution hardness (defined in a precise sense, Theorem 4.18). We thus identify *a dichotomy for* QU-Res *hardness:* it either rests on circuit lower bounds or on propositional resolution lower bounds. We note that the second case is inevitable: each propositional resolution lower bound (e.g. for the pigeonhole principle [29]) can be easily turned into a QU-Res lower bound. The surprising insight is that 'genuine QBF hardness' (cf. [14, 20]) can be completely characterised by circuit hardness.

Our result is best obtained in a model of QBF systems that 'filters out' propositional hardness (the second case above). For this we use the model of oracle QBF proof systems defined in [14], which employs an NP oracle to perform arbitrary propositional entailments in one inference step. For example, in the oracle system $\mathsf{QU}\text{-}^{\mathsf{NP}}\mathsf{Res}$, propositional resolution derivations of arbitrary size can be performed in just one step. The use of an NP oracle in $\mathsf{QU}\text{-}^{\mathsf{NP}}\mathsf{Res}$ is akin to the use of SAT solvers as oracles in QBF solving [39].

The hardness characterisation we obtain for QU^{NP}-Res is in terms of *unified decision lists* (UDL). This is a natural adaptation of the classical model of decision lists [44], which computes functions $\{0,1\}^n \to \{0,1\}$, to multi-output functions $\{0,1\}^n \to \{0,1\}^m$. Our first main result (Theorem 4.2) shows that for bounded-alternation QBFs, proof size in QU^{NP}-Res is polynomially related to the size of UDLs computing countermodels of the QBF.

Technically, this result is shown via *two simulations*. The first efficiently extracts UDLs from QU^{NP}-Res proofs (Theorem 4.5). Single-output decision lists have been used before to extract winning strategies for QBFs [2, 8, 10]. Here we show that winning strategies can also be extracted via multi-output decision lists, and these can be combined via a direct product construction (Definition 4.3) into one single UDL that computes the countermodel. We argue that representing the countermodel by just one function (computed by the UDL) is quite natural. However, it differs from the conventional approach, which represents the countermodel as a collection of Herbrand functions, one for each universal variable.

The *second simulation* turns a UDL into a QU^{NP}-Res refutation (Theorem 4.11). This is *conceptually novel,* as – to the best of our knowledge – the efficient construction of proofs from countermodels has not been considered before. In the course of the simulation, we obtain a normal form for proofs via the *entailment sequence* associated with a UDL (Definition 4.8). Inference steps in this entailment sequence also allow us to pinpoint sources for propositional hardness that arise when replacing NP oracle calls with actual resolution derivations. This way we obtain the dichotomy for QU-Res explained above (Theorem 4.18).

B. QU-Resolution and Q-Resolution. While QU-Res is arguably the simplest QBF resolution system from a logical perspective (it just adds the universal reduction rule to propositional resolution), there are other QBF resolution systems that better correspond to ideas in QBF solving. A core system among these is Q-Resolution (Q-Res), which is also historically the first QBF resolution system [33]. Q-Res is a restriction of QU-Res in which resolution pivots must be existential. This corresponds to techniques in QCDCL solving [38] (even though Q-Res does not capture QCDCL precisely [31]).

The system QU-Res is exponentially stronger than Q-Res [27], the separation provided by the prominent KBKF_n formulas [33]. These formulas use unbounded quantifier alternations, and indeed, we show that every separation must be of this form. We obtain the surprising result that Q-Res and QU-Res are *polynomially equivalent* on QBFs of bounded quantifier alternation (Theorem 5.4). This simulation is shown by a direct construction.

As a consequence, our hardness characterisation in terms of UDLs transfers directly to Q-Res (Corollary 5.6).

C. Size and width for QBF Resolution. Our new connection between QBF resolution and UDLs not only provides a tight characterisation of QBF resolution hardness, it also paves the way towards *a powerful lower-bound method.* We show that lower bounds on resolution width – defined as the (existential) size of the largest clause in the proof – directly imply lower bounds for proof size. The celebrated result of Ben-Sasson & Wigderson [4] provides such a size-width result for propositional resolution. Indeed, the vast majority of resolution hardness results are nowadays shown via this method.

Here we provide *the first size-width result for QBF* (Theorem 6.2). In a nutshell it says that each short QU-Res proof can be transformed into a narrow proof, where a proof is narrow if it does not contain a clause with many existential literals. What is perhaps most surprising is that the authors of [12, 22] have previously ruled out a similar size-width result for Q-Res and QU-Res. Not only did they show that the proof method of [4] does not lift to QBF, they also provided concrete QBF counterexamples to their size-width relation.

Two ingredients are required for Theorem 6.2: our UDL characterisation of proofs; and a size-width transfer for decision lists. The second ingredient, a decision-list size-width relation indeed already exists in the literature, due to Bshouty [18, repeated here as Lemma 6.3]. While this would suffice to obtain superpolynomial lower bounds on proof size, by careful analysis, essentially replicating the proof of the size-width transfer for resolution of Ben-Sasson and Wigderson [4, Theorem 3.5] for decision lists, we are able to improve Bshouty's result and shave off a factor of $\log n$ (Lemmas 6.4 and 6.5). Thanks to this improvement, we obtain a size-width result for QU-Res (indeed even for the model of $\mathsf{QU}^{\mathsf{NP}}$-Res, yielding stronger size lower bounds) which can deliver exponential QU-Res lower bounds of the form $\exp(\Omega(n))$ compared to $\exp(\Omega(n/\log n))$ obtained by using Bshouty's original result (here n is the number of existential variables).[1]

Our result is not a mere QBF replication of Ben-Sasson & Wigderson's result [4]. There are two crucial differences. First, in contrast to [4] our size-width result does not depend on the initial width of the formula. This makes the technique easier to apply and avoids the need for Tseitin transformations, which are often required in the propositional domain [4]. Second, our size bound depends on the number of quantifier alternations of the QBF. Crucially, the counterexamples of [12, 22] use unbounded alternations, thus ruling out the relation of [4], but not contradicting our Theorem 6.2.

D. Unification of previous lower-bound techniques. Our hardness characterisation in terms of UDLs together with the size-width method *encompasses and extends previous lower bound methods for QBF resolution.* In addition to lifted propositional techniques [11, 13], there exist *two genuine QBF techniques:* strategy extraction [8, 9] and the size-cost-capacity technique [6]. These techniques are orthogonal in the sense that each yields hardness results that cannot be shown by the other. Here we demonstrate that UDL hardness captures both.

In the *strategy extraction method* [8, 9], lower bounds are shown by extracting strategies in terms of a collection of single-output decision lists, which can be turned into bounded-depth circuits. The authors of [8, 9] then construct QBFs with a single universal variable whose unique Herbrand function is hard to compute by bounded-depth circuits (such as the parity function [30]). Such functions are also hard for UDLs (Section 4.5). Moreover, we show that width bounds for QBFs based on the parity and majority functions are easy to obtain (Section 6.2). We thus *elegantly reprove previous hardness results* for parity and majority formulas [8, 9] with our technique, without the need to import substantial circuit complexity results [30, 43, 47].

The *size-cost-capacity technique* [6] establishes hardness for QBFs where countermodels might be easy to compute by single-output decision lists, but must have large range. The large range immediately implies large UDLs (Section 4.5), hence again we can show the hardness results with our new technique. We illustrate this with the equality formulas (Theorem 6.7).

Organisation. The remainder of this article is organised as follows. In Section 2 we review notions from logic. Section 3 introduces our UDL model and explains how UDLs compute countermodels. In Section 4 we show our characterisation of QU-Res proof size by UDL size, which is extended to Q-Res in Section 5. Section 6 contains the size-width relation together with a number of applications. We conclude in Section 7 with a discussion and open problems.

[1]The conference version of this paper [7] only states such subexponential lower bounds of the form $\exp(\Omega(n/\log n))$. Here we improve this to truly exponential lower bounds by showing a strengthening of Bshouty's result and an improved size-width relation for QU-Res.

2 PRELIMINARIES

Propositional logic. $\mathcal{V}$ is a countable set of Boolean *variables.* A *literal* is a variable z in $\mathcal{V}$ or its negation $\overline{z}$, with $\text{var}(z) = \text{var}(\overline{z}) = z$. The literals z and $\overline{z}$ are *complementary.* For any literal a, the complementary literal is denoted $\overline{a}$.

A *clause* is a disjunction $c := a_1 \vee \cdots \vee a_k$ of pairwise non-complementary literals, with $\text{vars}(c) := \{\text{var}(a_i) : i \in [k]\}$. We often remove the disjunction symbols from a written clause, for example we write $z_1\overline{z_2}z_3$ for $z_1 \vee \overline{z_2} \vee z_3$. Given a set Z of Boolean variables, $c\restriction_Z$ is the disjunction of literals a appearing in c with $\text{var}(a) \in Z$.

A *conjunctive normal form* formula (CNF) is a conjunction $F := c_1 \wedge \cdots \wedge c_k$ of clauses, with $\text{vars}(F) := \bigcup_{i=1}^{k} \text{vars}(c_i)$.

A *term* is a finite conjunction $t := a_1 \wedge \cdots \wedge a_k$ of non-complementary literals, with $\text{vars}(t) := \{\text{var}(a_i) : i \in [k]\}$. $t\restriction_Z$ is defined similarly as for clauses. The negation of t is the clause $\overline{t} := \overline{a}_1 \vee \cdots \vee \overline{a}_k$. The negation of a clause c is the unique term $\overline{c}$ whose negation is c. The width of a clause or term is the number of its literals.

An *assignment* τ to a set Z of Boolean variables is a function from Z into the set of *Boolean constants* $\{0, 1\}$. The set of all assignments to Z is denoted $\{0, 1\}^Z$. A partial assignment to Z is an assignment to a subset of Z. We often represent assignments as terms, as there is a natural one-one correspondence between the two. The term t with $\text{vars}(t) = Z$ represents the assignment $\tau : Z \rightarrow \{0, 1\}$ which maps $z \in Z$ to 0 if, and only if, $\overline{z}$ is a conjunct in t.

The *restriction* of a literal, clause, CNF or term ϕ by τ, denoted $\phi\,[\tau]$, is the result of substituting each variable z in Z by $\tau(z)$, followed by applying the standard simplifications for Boolean constants, i.e. $\overline{0} \mapsto 1$, $\overline{1} \mapsto 0$, $c \vee 0 \mapsto c$, $c \vee 1 \mapsto 1$, $t \wedge 1 \mapsto t$, and $t \wedge 0 \mapsto 0$. We say that τ *satisfies* ϕ when $\phi\,[\tau] = 1$, and falsifies ϕ when $\phi\,[\tau] = 0$.

Otherwise, a *formula*, and *substitution* of formulas for variables, is defined in the standard way for propositional logic (cf. [48]). A formula F *entails* another formula G (written $F \vDash G$) when every assignment to $\text{vars}(F) \cup \text{vars}(G)$ satisfying F also satisfies G. Formulas F and G are *logically equivalent* (written $F \equiv G$) when they entail one another.

Quantified Boolean formulas. A *quantified Boolean formula* (QBF) Q of alternation depth d is a formula of the form $P \cdot F$, where $P := \exists X_1 \forall U_1 \cdots \exists X_d \forall U_d \exists X_{d+1}$ is called the *quantifier prefix* and F is a CNF called the *matrix.*[2] The X_i, U_i are pairwise-disjoint sets of Boolean variables called the *blocks* of Q.

The sets $\text{vars}_\exists(Q) := \bigcup_{i=1}^{d+1} X_i$ and $\text{vars}_\forall(Q) := \bigcup_{i=1}^{d} U_i$ are referred to as the *existential variables* and *universal variables* of Q, respectively, and their union $\text{vars}(Q)$ as the *variables* of Q. We say that an assignment τ to a set $Z \subseteq \text{vars}(Q)$ is *existential* if $Z \subseteq \text{vars}_\exists(Q)$, and *universal* if $Z \subseteq \text{vars}_\forall(Q)$. Given two variables z, z' in $\text{vars}(Q)$, we say that z is *left of* z' (written $z <_P z'$) when z belongs to a block quantified before that of z'. We deal only with *closed* QBFs, i.e. those for which $\text{vars}(F) \subseteq \text{vars}(Q)$. The *restriction* of Q by an assignment τ is $Q\,[\tau] := P\,[\tau] \cdot F\,[\tau]$, where $P\,[\tau]$ is obtained from P by deleting each variable in $\text{vars}(\tau)$ (along with its quantifier).

A set of QBFs has *bounded alternation* if each has alternation depth at most d, for some constant d.

QBF resolution proof systems. We work with *refutational* QBF proof systems, i.e. systems proving the falsity of a given QBF. We call a refutational QBF proof system P *sound* when there is no P-refutation of a true QBF, and *complete* when every false QBF has a P-refutation. Given two refutational QBF proof systems P and Q, we say that P *p-simulates* Q (written Q $\leq_p$ P) when there exists a polynomial-time computable translation mapping Q-refutations into P-refutations,

[2]Note that our definition of alternation depth differs slightly from the usual way of counting alternations, by counting only alternations from universal to existential blocks, effectively counting the number of universal blocks.

while preserving the refuted QBF [24]; we say that P *p-simulates* Q *on bounded alternation* if the translation of a Q-refutation π runs in time $O(|\pi|^{f(d)})$ for some computable function f, where d is the alternation depth of the formula refuted by π (in other words, in polynomial time whenever we restrict ourselves to formulas of bounded alternation, though possibly not in polynomial time on all QBFs). We say that P and Q are *p-equivalent* (written P $\equiv_p$ Q) when they p-simulate one another; and analogously for bounded alternation.

QU-Resolution (QU-Res) is the QBF analogue of propositional resolution [17, 45], defined as follows.

*Definition 2.1 (*QU-Res *[27, 33]).* A QU-Res derivation from a QBF $P \cdot F$ is a sequence of clauses $\pi := c_1, \ldots, c_s$ in which each c_i is derived by one of the following rules:

- *Axiom:* c_i is a clause in the matrix F;
- *Resolution:* $c_i = a \vee b$, where $c_r = a \vee z$ and $c_s = b \vee \bar{z}$ for some $r, s < i$ and variable z.
- *Weakening:* $c_i = c_r \vee b$ for some $r < i$ and clause b.
- *Universal reduction:* $c_i = c_r[\mu]$ for some $r < i$ and some universal assignment μ with $\mathrm{vars}_\exists(c_r) <_P \mathrm{vars}(\mu)$.[3]

The size of π is $|\pi| = s$, and π is a *refutation* when $c_s = \bot$. We say that a clause is *fully universal*, if it can be reduced to the empty clause $\bot$, i.e. if it consists of universal variables only.

The axiom, resolution and weakening rules together are *propositionally implicationally complete*; that is, if $F \vDash c$, then there exists a derivation of c from F. The refutational QBF proof system $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ allows any such correct *propositional* implication to be derived in a single step, eliminating all hardness due to propositional resolution.[4]

Definition 2.2 ($\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ *[14]).* $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ is defined as for QU-Res, except that the resolution and weakening rules are replaced by the following single rule, which requests that c_i be implied by the clauses derived earlier:

- Σ_1*-rule:* $\bigwedge_{j=1}^{i-1} c_j \vDash c_i$.

We can assume that there are no *postponed reductions*; universal variables are reduced as soon as the reduction is permitted. Corollary 1 in [1] establishes that postponing reductions does not shorten proofs in QU-Res. For the oracle system $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ too, postponed reductions can be eliminated with at most polynomial size blow-up as follows: let $\Pi = c_1, c_2, \ldots, c_t$ be a $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ refutation. Construct the sequence $\Pi' = d_1, e_1, d_2, e_2, \ldots, d_t, e_t$ where each e_i is obtained from d_i by applying all enabled reductions, and if c_i is obtained by reduction on c_j, then $d_i = e_j$ (and hence $e_i = e_j$ as well), otherwise $d_i = c_i$. Inductively, we see that each d_i and e_i is a subclause of c_i. Hence if $\bigwedge_{j=1}^{i-1} c_j \vDash c_i$, then $\bigwedge_{j=1}^{i-1}(d_j \wedge e_j) \vDash c_i$. Thus Π' is also a $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ refutation, with no postponed reductions. If Π has A, R, E axiom, reduction, and entailment steps respectively, then Π' has A axiom steps, at most E entailment steps, and at most $A + E$ reduction steps.

In the following, we will assume that any universal reduction step is due to a total assignment to exactly one block, i.e. $\mathrm{vars}(\mu) = U_j$ for some $j \in \{1, \ldots, d\}$. This restriction simplifies some of the arguments, while maintaining a p-equivalent proof system. Since more than one block may be reducible at some stage, the proof size can indeed blow up, though only by a factor d; we simply reduce all reducible blocks one by one rather than at once.

[3]Some definitions of QU-Res disallow deriving tautological clauses [33]. The definition of universal reduction chosen here eliminates this restriction.

[4]Note that proofs in $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ cannot necessarily be checked in polynomial time, hence $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ is not a proof system in the sense of [24], but conforms to our definition of proof system above (cf. also [15] for a formal definition of oracle proof systems).

3 COUNTERMODELS AS DECISION LISTS

A *countermodel* witnesses the falsity of a QBF. In the literature, countermodels are usually defined in one of two equivalent ways (under various names): either as a collection of functions, one for each universal variable (called here *distributed countermodel*), or as a single function (*unified countermodel*). In this section, we recall the definitions of distributed and unified countermodels. We show that distributed countermodels represented by term decision lists are unsuitable for characterising hardness in $\mathsf{QU}^{\mathsf{NP}}$-Res (Subsection 3.1) and propose a model for multi-output term decision lists which serves as a natural representation for unified countermodels (Subsection 3.2).

3.1 Distributed countermodels

A distributed countermodel defines a set of formulas which, when substituted for the universal variables, leaves the matrix unsatisfiable. In order to respect the variable dependencies imposed by the order of quantification, each function must depend only on the preceding existential variables.[5]

Definition 3.1 (distributed countermodel). Let Q be a QBF with $\text{vars}_\forall(Q) = u_1, \dots, u_m$, and let D_i denote the union of the existential blocks preceding u_i in the prefix. A *distributed countermodel* for Q is a collection of functions $\{f_i\}_{i\in[m]}$ of the form $f_i : \{0,1\}^{D_i} \to \{0,1\}$, such that the substitution of formula representations of $f_1, \dots, f_m$ for the universal variables $u_1, \dots, u_m$ in F yields an unsatisfiable formula.

We illustrate this concept with the equality formulas, which we will use as a running example.

Definition 3.2 (equality [6]). The n^{th} *equality formula* is

$$Q_n^{\text{EQ}} := \exists x_1 \cdots x_n \forall u_1 \cdots u_n \exists z_1 \cdots z_n \cdot (\overline{z_1} \vee \cdots \vee \overline{z_n}) \wedge \bigwedge_{i=1}^{n} \Big((\overline{x_i} \vee \overline{u_i} \vee z_i) \wedge (x_i \vee u_i \vee z_i) \Big).$$

Example 3.3. The n^{th} equality formula has the unique distributed countermodel $\{f_i\}_{i\in[n]}$, where

$$\begin{aligned} f_i \;:\; \{0,1\}^X &\to \{0,1\} \\ \tau &\mapsto \begin{cases} 0 & \text{if } \tau(x_i) = 0\,, \\ 1 & \text{if } \tau(x_i) = 1\,, \end{cases} \end{aligned}$$

where $X = \{x_1, \dots, x_n\}$. Here, each function f_i is represented by the atomic formula x_i. It is easy to see that substituting each u_i for x_i in the matrix of Q_n^{EQ} yields an unsatisfiable formula. ■

Particularly in the context of strategy extraction, whereby one translates QBF refutations into countermodels, it is quite natural to represent a distributed countermodel as a set of term decision lists, one for each individual function [8]. Let us recall the traditional definition of a term decision list.

Definition 3.4 (decision list [44]). Given a set X of variables, a *decision list* is a sequence of pairs $L := (\varepsilon_1, b_1), \dots, (\varepsilon_s, b_s)$ where

- the ε_i are terms with $\text{vars}(\varepsilon_i) \subseteq X$ and $\bigvee_{i=1}^{s} \varepsilon_i \equiv \top$,
- the b_i are Boolean constants, i.e. 0 or 1.

L computes the function from $\{0,1\}^X$ into $\{0,1\}$ mapping τ to b_i, where i is the least natural number for which τ satisfies ε_i. The size of L, denoted by $|L|$, is s.

[5]Preceding universals can also be included as dependencies (cf. [8]), producing a potentially stronger model.

As far as characterising QU-Res hardness is concerned, the problem with this computation model – distributed countermodels represented as decision lists – is that it is too strong, even for bounded alternation depth. For example, the distributed countermodel $\{f_i\}_{i\in[n]}$ from Example 3.3 can be computed by n constant-size decision lists, namely

$$L_i := (x_i, u_i), (\overline{x_i}, \overline{u_i}), \quad i \in [n],$$

but the equality formulas require exponential-size $\mathsf{QU}^{\mathsf{NP}}$-Res refutations [6].

3.2 Unified countermodels

A unified countermodel is a single function which *simultaneously* represents the individual functions of a distributed countermodel. Formally, there are two differences. First, the output of the function is not a $\{0, 1\}$ value, but a total assignment to the universal variables, giving a $\{0, 1\}$ value for *each* universal variable. Second, the prefix dependencies, which are implicit in the function signatures of a distributed countermodel, must be explicitly enforced.

Definition 3.5 (unified countermodel). Let $Q := P \cdot F$ be a QBF of alternation depth d. A *unified countermodel* for Q is a function $f : \{0,1\}^{\mathrm{vars}_\exists(Q)} \to \{0,1\}^{\mathrm{vars}_\forall(Q)}$ satisfying two conditions:

(a) for each $\tau \in \mathrm{dom}(f)$, $\tau \wedge f(\tau)$ falsifies F;
(b) for each $\tau, \rho \in \mathrm{dom}(f)$ and each $i \in [d]$, if τ, ρ agree on the first i existential blocks, then $f(\tau), f(\rho)$ agree on the first i universal blocks.

Example 3.6. The n^{th} equality formula has the unique unified countermodel

$$f_{\mathsf{EQ}} \quad : \quad \{0,1\}^X \quad \to \quad \{0,1\}^U$$

where $X = \{x_1, \dots, x_n\}$, $U = \{u_1, \dots, u_n\}$, and $f_{\mathsf{EQ}}(\tau) : U \to \{0, 1\}$ is the assignment mapping each u_i to $\tau(x_i)$. It is easy to see that f_{EQ} is a single-function representation of the distributed countermodel from Example 3.3, and readily verified that conditions (a) and (b) of Definition 3.5 are satisfied. ■

In order to represent a unified countermodel as a decision list, we specify a new format to allow simultaneous output for multiple Boolean variables. This is achieved in the most natural way, specifying a term over the universal variables which represents the desired output assignment.

Definition 3.7 (multi-output decision list). Given sets X and U of Boolean variables, a *multi-output term decision list (MDL)* is a sequence of pairs $L := (\varepsilon_1, \mu_1), \dots, (\varepsilon_s, \mu_s)$ where

- the ε_i are terms with $\mathrm{vars}(\varepsilon_i) \subseteq X$ and $\bigvee_{i=1}^{s} \varepsilon_i \equiv \top$,
- the μ_i are terms with $\mathrm{vars}(\mu_i) = U$.

L computes the function from $\{0,1\}^X$ into $\{0,1\}^U$ mapping τ to μ_i, where i is the least natural number for which τ satisfies ε_i. The size of L, denoted by $|L|$, is s; we call $\varepsilon_1, \dots, \varepsilon_s$ the *input terms* of L, and $\mu_1, \dots, \mu_s$ the *output terms* of L. The *input width* of L is the maximum width of any of its input terms.

Definition 3.8 (unified decision list). An MDL computing a unified countermodel for a QBF Q is called a *unified decision list (UDL)* for Q.

Without ambiguity, we will use the same symbol (e.g. L) to represent both the UDL and its computed function.

Note that the insistence on a single function suitably reduces the strength of the computational model, in terms of representation size. For example, UDLs for the equality formulas must have exponential size, matching the exponential-size $\mathsf{QU}^{\mathsf{NP}}$-Res refutations. This is due to the fact that the range of the unique unified countermodel, which is the complete set of universal assignments, has cardinality 2^n.

Obviously, this holds generally. Since every entry of a UDL produces exactly one output, there must be at least as many entries in the UDL as there are different outputs in the countermodel. Hence the minimal range cardinality of a unified countermodel for a QBF Q is a lower bound to the size of a UDL for Q.

4 CHARACTERISING HARDNESS IN QU-Res ON BOUNDED ALTERNATION

In this section, we demonstrate that UDLs have *exactly* the right strength to characterise $\text{QU}^{\text{NP}}_{-}\text{Res}$ refutation size on bounded alternation QBFs. For this, we cast UDLs as a refutational QBF proof system.

*Definition 4.1 (*UDL*).* A UDL-refutation of a QBF Q is a UDL $L := (\varepsilon_1, \mu_1), \dots, (\varepsilon_s, \mu_s)$ for Q. The size of the UDL-refutation L is the size $|L| = s$ of L as a UDL.

Our central result is the following.

Theorem 4.2. $\text{QU}^{\text{NP}}_{-}\text{Res} \equiv_p$ UDL *on bounded-alternation QBFs.*

The two individual p-simulations are shown in Subsection 4.1 (Corollary 4.6) and Subsection 4.2 (Corollary 4.12). In Subsection 4.3 we demonstrate that the equivalence cannot be extended to unbounded alternation depth.

In Subsection 4.4 we characterise bounded-alternation hardness in QU-Res, insofar as superpolynomial QU-Res lower bounds come either from large UDLs or from an embedded propositional resolution lower bound. Finally, in Subsection 4.5, we discuss how UDL lower bounds encompass both the strategy extraction [8, 9] and size-cost techniques for QU-Res [6].

4.1 From $\text{QU}^{\text{NP}}_{-}\text{Res}$ to unified decision lists

In this subsection, we show an efficient transformation from $\text{QU}^{\text{NP}}_{-}\text{Res}$ refutations into unified decision lists. The transformation is a two-step process.

In the *first step,* we transform the refutation π into a collection of multi-output term decision lists, each of which computes the countermodel for just a single universal block, based on assignments to *all* previous blocks (including previous universal blocks). This constitutes a modification of the strategy extraction procedure from [2, 9], which works per universal variable, rather than per universal block. The size of each of the lists we obtain this way is bounded by the size $|\pi|$ of the refutation, and their number is equal to the alternation depth d of the refuted formula.

In the *second step,* we transform the collection into a single unified decision list, substituting all dependence on universal variables with their own decision lists and merging the entire collection. This involves taking a kind of 'direct product' of multi-output term decision lists (defined shortly), the size of which is the product of the sizes of the operands. Thus, with d lists of size at most $|\pi|$, we obtain a UDL of size at most $|\pi|^d$, which is a polynomial quantity as long as alternation depth d remains bounded.

In the rest of this subsection, we prove this formally. We turn first to the definition of our direct-product-like operation; the full transformation is described in Theorem 4.5 and its proof.

Definition 4.3 (direct product). Let X_1, U_1, X_2 and U_2 be pairwise-disjoint sets of Boolean variables, and let $L := (\varepsilon_1, \mu_1), \dots, (\varepsilon_s, \mu_s)$ and $M := (\delta_1, \nu_1), \dots, (\delta_t, \nu_t)$ be multi-output term decision lists with

$$\begin{aligned} &\text{vars}(\varepsilon_i) \subseteq X_1 \text{ and } \text{vars}(\mu_i) = U_1\,, && \text{for } i \in [s]\,, \\ &\text{vars}(\delta_j) \subseteq X_1 \cup U_1 \cup X_2 \text{ and } \text{vars}(\nu_j) = U_2\,, && \text{for } j \in [t]\,. \end{aligned}$$

The *direct product* $L \times M$ is the decision list

$$(\varepsilon_1 \wedge \delta_1 [\mu_1], \mu_1 \wedge \nu_1), \ldots, (\varepsilon_s \wedge \delta_1 [\mu_s], \mu_s \wedge \nu_1),$$
$$\vdots$$
$$(\varepsilon_1 \wedge \delta_t [\mu_1], \mu_1 \wedge \nu_t), \ldots, (\varepsilon_s \wedge \delta_t [\mu_s], \mu_s \wedge \nu_t).$$

The direct product $L \times M$ computes a function based on M, which first queries L for the assignment to U_1. Informally, the U_1 variables in M are substituted for the function computed by L, while U_1 is moved from the domain to the codomain. This is stated formally as follows.

Proposition 4.4. *Let X_1, U_1, X_2 and U_2 be pairwise-disjoint Boolean variable sets, and let L and M be multi-output decision lists computing $L : \{0,1\}^{X_1} \to \{0,1\}^{U_1}$ and $M : \{0,1\}^{X_1 \cup U_1 \cup X_2} \to \{0,1\}^{U_2}$. Then $L \times M$ computes the function*

$$\begin{array}{rcccl} L \times M & : & \{0,1\}^{X_1 \cup X_2} & \to & \{0,1\}^{U_1 \cup U_2} \\ & & \tau & \mapsto & L(\tau\restriction_{X_1}) \wedge M(\tau \wedge L(\tau\restriction_{X_1})). \end{array}$$

Proof. Let $L := (\varepsilon_1, \mu_1), \ldots, (\varepsilon_s, \mu_s)$, $M := (\delta_1, \nu_1), \ldots, (\delta_t, \nu_t)$. Let $\tau \in \{0,1\}^{X_1 \cup X_2}$, and let a and b be the least natural numbers such that $\tau\restriction_{X_1}$ satisfies ε_a and τ satisfies $\delta_b [\mu_a]$. By definition of decision list (Definition 3.7),

$$(L \times M)(\tau) = \mu_a \wedge \nu_b.$$

Clearly, $L(\tau\restriction_{X_1}) = \mu_a$ by definition of decision list, therefore $\tau \wedge L(\tau\restriction_{X_1}) = \tau \wedge \mu_a$. Aiming for contradiction, suppose that $M(\tau \wedge \mu_a) \neq \nu_b$. Since τ satisfies $\delta_b [\mu_a]$, $\tau \wedge \mu_a$ satisfies δ_b. Therefore $\tau \wedge \mu_a$ satisfies some $\delta_{b'}$ with $b' < b$. It follows that τ satisfies $\delta_{b'} [\mu_a]$, contradicting the minimality of b. □

Note that the size of a direct product is indeed the product of the sizes of the original decision lists.

Theorem 4.5. *A $\mathrm{QU}^{\mathrm{NP}}$-Res refutation π of a QBF Q of alternation depth d can be transformed into a UDL $t(\pi)$ for Q, where $|t(\pi)| \leq |\pi|^d$. The transformation t is computable in time $O(|\pi|^d)$.*

Proof. Let $Q := P := \exists X_1 \forall U_1 \cdots \exists X_d \forall U_d \exists X_{d+1} \cdot F$ be a QBF, and let $\pi := c_1, \ldots, c_s$ be a $\mathrm{QU}^{\mathrm{NP}}$-Res refutation of Q. We assume without loss of generality that each universal reduction step in π is due to a total assignment to a single universal block (as discussed in the preliminaries).

For each $i \in [d]$ and $j \in [s+1]$, we define a collection of multi-output term decision lists recursively as follows: $L_i^{s+1} := (\top, \alpha_i)$, where α_i is some fixed assignment to U_i; for each $j \in [s]$,

$$L_i^j := \begin{cases} (\overline{c_j}, \mu), L_i^{j+1} & \text{if } c_j \text{ was derived by universal reduction due to } \mu \in \{0,1\}^{U_i}, \\ L_i^{j+1} & \text{otherwise.} \end{cases}$$

The intuition behind these lists is the same as in the original strategy-extraction algorithm [2]. L_i^j approximate L_i^1, which finds the first clause of π which was derived by universal reduction and is falsified by the input assignment, and sets the assignment to the i-th universal block to match the reduction assignment μ – the meaning of the reduction assignment is exactly that it should be picked if the reduced clause is falsified. Because of how earlier clauses imply later clauses in the proof, if we identify the first reduction step whose conclusion is falsified and match its reduction assignment, we ultimately obtain that axioms restricted by the combined assignment under consideration jointly imply the empty clause; and that can only happen if one of the axioms is already falsified. It then remains to merge all these lists together to obtain a single UDL. We now show how this is done formally.

By backwards induction on $j \in [s+1]$, we show that the combined direct product of these lists

$$L^j := L_1^j \times \left(L_2^j \times \cdots \times \left(L_{d-1}^j \times L_d^j\right) \cdots\right)$$

is a UDL for $P \cdot F \wedge \bigwedge_{k=1}^{j-1} c_k$. We therefore prove the theorem, i.e. that L^1 is a UDL for Q of size at most $|\pi|^d$, that can clearly be constructed in time $O(|\pi|^d)$.

It is clear by construction that each L_i^j computes a function

$$L_i^j : \{0,1\}^{X_1 \cup \cdots \cup X_i \cup U_1 \cup \cdots \cup U_{i-1}} \to \{0,1\}^{U_i} .$$

Hence, by definition of direct product (Definition 4.3), L^j computes a function

$$L^j : \{0,1\}^{\text{vars}_\exists(Q)} \to \{0,1\}^{\text{vars}_\forall(Q)}$$

satisfying condition (b) for a unified countermodel (Definition 3.5). It remains to show that condition (a) is satisfied; that is, for each $\tau \in \{0,1\}^{\text{vars}_\exists(Q)}$, we must show that $\tau \wedge L^j(\tau)$ falsifies $F \wedge \bigwedge_{k=1}^{j-1} c_k$.

Base case $j = s+1$. Since c_s is the empty clause, $\tau \wedge L^{s+1}(\tau)$ always falsifies $F \wedge \bigwedge_{k=1}^{s} c_k$.

Inductive step $j \in [s]$. We consider two cases, based on how c_j was derived.

Suppose that c_j was introduced as an axiom, or derived by the Σ_1-rule. In either case, $L^j = L^{j+1}$ and $F \wedge \bigwedge_{k=1}^{j-1} c_k \vDash c_j$. By the inductive hypothesis we know that $\tau \wedge L^{j+1}(\tau)$ falsifies $F \wedge \bigwedge_{k=1}^{j} c_k$. It follows that $\tau \wedge L^j(\tau)$ falsifies $F \wedge \bigwedge_{k=1}^{j-1} c_k$.

On the other hand, suppose that c_j was derived by universal reduction from c_r due to the assignment $\mu \in U_i$. In this case, $L_k^j = L_k^{j+1}$ for each $k \neq i$. We consider two cases.

(a) Suppose that $\tau \wedge L^{j+1}(\tau)$ falsifies c_j. Consider the direct product of lists up to, but not including L_i^j, namely

$$M^j := L_1^j \times \left(L_2^j \times \cdots \times \left(L_{i-2}^j \times L_{i-1}^j\right) \cdots\right) ,$$

and let D_i and D_{i-1} denote the union of existential blocks preceding U_i and U_{i-1} respectively. It is easy to see that

$$\tau\restriction_{D_i} \wedge M^j(\tau\restriction_{D_{i-1}}) \text{ satisfies } \overline{c_j} ,$$

from which it follows that

$$L_j^i\left(\tau\restriction_{D_i} \wedge M^j(\tau\restriction_{D_{i-1}})\right) = \mu .$$

As a result, $L^j(\tau)$ extends μ. Therefore $\tau \wedge L^j(\tau)$ falsifies c_r, which belongs to $F \wedge \bigwedge_{k=1}^{j-1} c_k$.

(b) On the other hand, suppose that $\tau \wedge L^{j+1}(\tau)$ satisfies c_j. Then the addition of $(\overline{c_j}, \mu)$ to L_i^{j+1} has no effect on L^{j+1}, so that $L^j(\tau) = L^{j+1}(\tau)$. Hence $\tau \wedge L^j(\tau)$ falsifies $F \wedge \bigwedge_{k=1}^{j-1} c_k$ by the inductive hypothesis.

□

Corollary 4.6. $\mathsf{QU}^{\mathsf{NP}}$-Res $\leq_p$ UDL *on bounded alternation.*

4.2 From unified decision lists to $\mathsf{QU}^{\mathsf{NP}}$-Res

In this subsection, we show an efficient translation from UDLs back into $\mathsf{QU}^{\mathsf{NP}}$-Res refutations. The transformation uses a notion of restriction for UDLs.

Definition 4.7 (restriction of a UDL). Given an assignment α and a multi-output decision list $L := (\varepsilon_1, \mu_1), \ldots, (\varepsilon_s, \mu_s)$, the *restriction of L by α* is

$$L[\alpha] := (\varepsilon_1[\alpha], \mu_1[\alpha]), \ldots, (\varepsilon_s[\alpha], \mu_s[\alpha]) .$$

The entailment sequence. We summarise our method as follows: we transform a UDL L for a QBF Q into a sequence of clauses $\mathcal{E}(L)$. Each clause in the sequence is entailed by Q *and the universal reduction of the previous clauses in the sequence.* The final clause is fully universal, yielding a refutation. We refer to the sequence $\mathcal{E}(L)$ as the *entailment sequence* for L.

First, some extra notation and nomenclature. Given a clause b and a sequence of clauses $\pi := c_1, \dots, c_s$, we define

$$b \otimes \pi := b \vee c_1, \dots, b \vee c_s \,.$$

Given a UDL $L := (\varepsilon_1, \mu_1), \dots, (\varepsilon_s, \mu_s)$ for a QBF Q and block Z of Q, the Z-component of (ε_i, μ_i) is $(\varepsilon_i \wedge \mu_i)\restriction_Z$.

Also, we note the following: without loss of generality we can assume that rightmost existential variables (on which no universal variable can depend) do not appear in a UDL. That is, given a QBF with prefix

$$P := \exists X_1 \forall U_1 \cdots \exists X_d \forall U_d \exists X_{d+1} \,,$$

the X_{d+1}-components in any UDL for Q can be deleted while preserving the computed countermodel. This is an easy consequence of condition (b) in the definition of unified countermodel (Definition 3.5).

Definition 4.8 (entailment sequence). Given a UDL $L := (\varepsilon_1, \mu_1), \dots, (\varepsilon_s, \mu_s)$ for a QBF Q, the *entailment sequence* $\mathcal{E}(L)$ is defined recursively on the alternation depth d of Q.

- if $d = 1$, $\mathcal{E}(L) := \overline{\varepsilon_1} \vee \overline{\mu_1}, \dots, \overline{\varepsilon_s} \vee \overline{\mu_s}$,
- if $d \geq 2$, for each $i \in [s]$ define L_i as the list obtained from L by replacing the first $i-1$ existential terms by their X_1 components, and setting all U_1 components to $\mu_i\restriction_{U_1}$. We define $\mathcal{E}(L)$ as the sequence $\pi_1, \dots, \pi_s$, where

$$\pi_i := (\overline{\varepsilon_i}\restriction_{X_1} \vee \overline{\mu_i}\restriction_{U_1}) \otimes \mathcal{E}(L_i\left[\varepsilon_i\restriction_{X_1} \wedge \mu_i\restriction_{U_1}\right]) \,.$$

The size of $\mathcal{E}(L)$, denoted $|\mathcal{E}(L)|$, is the number of clauses in the sequence; the *existential width* of $\mathcal{E}(L)$ is the maximum number of existential literals in any of its clauses.

As an exercise to absorb the definition better, let us prove the following lemma about the size and the width of entailment sequences, which will come in handy later.

Lemma 4.9. *Let L be a UDL for a QBF Q of alternation depth d. Then $|\mathcal{E}(L)| \leq |L|^d$ and the width of $\mathcal{E}(L)$ is at most d times the width of L.*

Proof. Both parts follow by induction on d; the difference comes from how the two attributes aggregate over $\otimes$ and concatenation of sequences: size multiplies, while width adds up. □

The intuition behind the construction of the entailment sequence, in particular when the alternation depth exceeds 1, is not obvious. We will elaborate upon this later. For now, the important property is the fulfilment of Lemma 4.10. For a clause c of a QBF $Q = P \cdot F$ we take the assignment $\nu : \{u \in \text{vars}_\forall(c) : \text{vars}_\exists(c) <_P u\} \to \{0, 1\}$ which maps u to 1 if, and only if, $\overline{u}$ is in c_i, i.e. the assignment that reduces every universal literal that can be reduced, and we write $\text{red}(c) := c\,[\nu]$ for the clause obtained from c by 'maximum reduction.'

Lemma 4.10. *Let L be a unified decision list for a QBF $Q := P \cdot F$, and let $\mathcal{E}(L) = c_1, \dots, c_r$. Then c_r is fully universal, and, for each $i \in [r]$,*

$$F \wedge \bigwedge_{j=1}^{i-1} \text{red}(c_j) \vDash c_i \,.$$

We defer the proof of this lemma to the end of the subsection. The entailment of each clause by the universal reduction of its predecessors (in conjunction with the matrix F) gives rise to a straightforward QU^{NP}-Res refutation.

Theorem 4.11. *A UDL L for a QBF Q of alternation depth d can be transformed into a* $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ *refutation $t(L)$ for Q, where $|t(L)| \leq O(|L|^d)$. The transformation t is computable in time $O(|L|^d)$.*

Proof. Let $\mathcal{E}(L) = c_1, \dots, c_r$. By Lemma 4.10, the sequence π, consisting of the clauses of the matrix of Q followed by

$$c_1, \mathrm{red}(c_1), \dots, c_r, \mathrm{red}(c_r)\,,$$

is a $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ refutation of Q, and by Lemma 4.9 $r \leq |L|^d$. By a simple induction on alternation depth d, one verifies that π can be constructed in time $O(r)$. □

Corollary 4.12. UDL $\leq_p$ $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ *on bounded alternation.*

Intuition and example. The $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ refutation obtained from a UDL L consists of the entailment sequence interleaved with reduction steps. The clauses in the entailment sequence are intended to witness the fact that L is indeed a UDL for the formula, i.e., the lines in the entailment sequence describe the correctness of the outputs produced in the lines of the UDL (this is the intuition behind Lemma 4.10). The $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ proof constructed from the entailment sequence can then be understood as a formal proof of correctness of the UDL. Note that we start with a correct UDL, but the entailment sequence models this correctness formally within $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$.

For this, the idea is to 'unpack' the countermodel represented by L into its round-by-round responses in the 2-player game corresponding to the standard QBF semantics. Each L_i as defined, upon suitable restriction, is a UDL for the formula Q_i defined in the formal proof below; Q_i is the restriction of the formula obtained after one round of the game (an existential player move followed by a universal player move), where the restriction leads the UDL computation to line i or later. Recursively obtaining $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ proofs for each Q_i from the corresponding L_i, the combination into a single proof expresses the fact that the ith refutation should be used on partial assignments inconsistent with the first $i-1$ lines of L; this is achieved by the direct product operation.

In the simplest case, with alternation depth $d = 1$, the entailment sequence is composed merely of the negations of the combined existential and universal terms in the UDL (i.e. $\overline{\varepsilon_i} \vee \overline{\mu_i}$). The universal reduction of each clause is merely $\overline{\varepsilon_i}$, the negation of the corresponding existential term. In this case, the fact that each clause is entailed by the universal reductions of its predecessors in conjunction with the matrix (Lemma 4.10) follows straightforwardly from the definition of UDL. In fact, in this case the entailments stated in Lemma 4.10 can be easily seen to exactly model the correctness of the outputs produced in each line of the UDL.

This forms the base case for a general argument by induction, when the alternation depth exceeds 1. In the entailment sequence definition, the lists L_i are defined so that $L_i\left[\varepsilon_i\!\restriction_{X_1} \wedge \mu_i\!\restriction_{U_1}\right]$ is a UDL for the QBF

$$\left(P \cdot F \wedge \bigwedge_{k=1}^{i-1} \overline{c_k}\!\restriction_{X_1}\right)\left[\varepsilon_i\!\restriction_{X_1} \wedge \mu_i\!\restriction_{U_1}\right]\,. \tag{1}$$

Note that each of the negated X_1-components $\overline{c_k}\!\restriction_{X_1}$ is the universal reduction of a clause already appearing in $\mathcal{E}(L)$ before π_i. This is not obvious; it relies on the fact that the final clause of each $\mathcal{E}(L_k\left[\varepsilon_k\!\restriction_{X_1} \wedge \mu_k\!\restriction_{U_1}\right])$ is fully universal.

The addition of these negated X_1-components to the matrix is the reason why the first $i-1$ existential terms in L_i are replaced by their X_1 components. Assignments satisfying the i^{th} term are guaranteed to falsify one of these clauses. One might suspect that the first $i-1$ lines could be removed altogether, somewhat simplifying the definition of $\mathcal{E}(L)$. Unfortunately, it is not clear that such a construction would produce a UDL for the QBF in (1). The assignments satisfying the removed lines are distributed arbitrarily across the remaining ones, so that the computed function may not satisfy the proper dependencies (condition (b) of Definition 3.5).

Note that the U_1-components in L_i are set uniformly to $\mu_i\restriction_{U_1}$ merely so that restriction by that assignment deletes them all.

Construction of the entailment sequence, along with the corresponding $\mathsf{QU}^{\mathsf{NP}}$-Res refutation, is illustrated by the following example.

Example 4.13. We will construct an entailment sequence for the QBF

$$\exists x_1 \forall u_1 \exists z_1 \exists x_2 \forall u_2 \exists z_2 \cdot \overline{x_1}\overline{u_1}z_1 \wedge x_1 u_1 z_1 \wedge \overline{x_2}\overline{u_2}z_2 \wedge x_2 u_2 z_2 \wedge \overline{z_1}\overline{z_2}\,.$$

This QBF is Q_2^{INT}, the second instance of the *interleaved equality family*, which we will meet in the following subsection. We write the blocks of Q_2^{INT} as follows: $X_1 := \{x_1\}$, $U_1 := \{u_1\}$, $X_2 := \{z_1, x_2\}$, $U_2 := \{u_2\}$, and $X_3 := \{z_2\}$. Note that the alternation depth of Q_2^{INT} is 2.

Similar to the original equality formulas, a unified countermodel for this QBF sets each u_i equal to the corresponding x_i, with the values of the z_i essentially ignored. This countermodel is computed by the following UDL L:

$$(x_1 \wedge x_2, u_1 \wedge u_2), (x_1 \wedge \overline{x_2}, u_1 \wedge \overline{u_2}), (x_2, \overline{u_1} \wedge u_2), (\top, \overline{u_1} \wedge \overline{u_2}).$$

We now construct the entailment sequence $\mathcal{E}(L)$. First we obtain the lists L_1, L_2, L_3, L_4 and their appropriate restrictions. These restrictions are easily transformed (they have alternation depth 1), and pieced together to obtain the complete entailment sequence.

L_1 is obtained from L by replacing each U_1-component by the U_1-component of the first line, namely the term u_1. So the restriction of L_1 by the X_1- and U_1-components of the first line, namely the assignment $x_1 \wedge u_1$, is

$$(x_2, u_2), (\overline{x_2}, \overline{u_2}), (x_2, u_2), (\top, \overline{u_2})\,.$$

Since the final two lines are redundant, this simplifies to $L_1\,[x_1 \wedge u_1] = (x_2, u_2), (\top, \overline{u_2})$. Hence we have

$$\begin{aligned} \mathcal{E}(L_1\,[x_1 \wedge u_1]) &= \overline{x_2}\overline{u_2}, u_2\,, \\ \pi_1 &= \overline{x_1}\overline{u_1} \otimes \mathcal{E}(L_1\,[x_1 \wedge u_1]) \\ &= \overline{x_1}\overline{u_1}\overline{x_2}\overline{u_2}, \overline{x_1}\overline{u_1}u_2\,. \end{aligned}$$

L_2 is obtained from L by replacing the first existential term by its X_1-component x_1, then replacing each U_1-component by the U_1-component of the second line, namely the term u_1:

$$(x_1, u_1 \wedge u_2), (x_1 \wedge \overline{x_2}, u_1 \wedge \overline{u_2}), (x_2, u_1 \wedge u_2), (\top, u_1 \wedge \overline{u_2})\,.$$

Restriction of L_2 by the X_1- and U_1-components of the second line, namely $x_1 \wedge u_1$, yields

$$(\top, u_2), (\overline{x_2}, \overline{u_2}), (x_2, u_2), (\top, \overline{u_2})\,.$$

Every line except the first is redundant, so this simplifies to $L_2\,[x_1 \wedge u_1] = (\top, u_2)$. In this case we get

$$\begin{aligned} \mathcal{E}(L_2\,[x_1 \wedge u_1]) &= \overline{u_2}\,, \\ \pi_2 &= \overline{x_1}\overline{u_1} \otimes \mathcal{E}(L_2\,[x_1 \wedge u_1]) \\ &= \overline{x_1}\overline{u_1}\overline{u_2}\,. \end{aligned}$$

Continuing in this way for L_3 and L_4, one verifies that

$$\begin{aligned} L_3\,[\overline{u_1}] \;=\; L_4\,[\overline{u_1}] &= (x_1, u_2), (x_2, u_2), (\top, \overline{u_2})\,, \\ \pi_3 \;=\; \pi_4 &= \overline{x_1}u_1\overline{u_2}, u_1\overline{x_2}\overline{u_2}, u_1 u_2\,. \end{aligned}$$

The fact that $\pi_3 = \pi_4$ is coincidental (note that the X_1-components of the third and fourth lines are both empty, and both U_1-components are $\overline{u_1}$).

Piecing together the π_i, the entailment sequence for L is

$$\begin{aligned}\mathcal{E}(L) &= \pi_1, \pi_2, \pi_3, \pi_4\\ &= \overline{x_1u_1x_2u_2}, \overline{x_1u_1}u_2, \overline{x_1u_1u_2}, \overline{x_1}u_1\overline{u_2}, u_1\overline{x_2u_2}, u_1u_2,\\ &\quad\ \overline{x_1}u_1\overline{u_2}, u_1\overline{x_2u_2}, u_1u_2\ .\end{aligned}$$

We can now illustrate how the entailment sequence gives rise to a QU$^{\text{NP}}$-Res refutation. In fact, several clauses in this particular entailment sequence are superfluous and can be ignored, so we work with the subsequence

$$\overline{x_1u_1x_2u_2}, \overline{x_1u_1}u_2, u_1\overline{x_2u_2}, u_1u_2\ .$$

The essential point is that each clause in the sequence is entailed by the matrix of Q_2^{INT} in conjunction with the universal reduction of the preceding clauses. For example, the first clause is entailed by the matrix of Q_2^{INT} alone; in fact

$$\overline{x_1u_1}z_1 \wedge \overline{x_2u_2}z_2 \wedge \overline{z_1z_2} \vDash \overline{x_1u_1x_2u_2}\ .$$

An easy way to verify this is to construct a resolution derivation:

$$\dfrac{\dfrac{\overline{x_1u_1}z_1 \qquad \overline{z_1z_2}}{\overline{x_1u_1z_2}} \qquad \overline{x_2u_2}z_2}{\overline{x_1u_1x_2u_2}}$$

The second clause in the sequence is entailed by the matrix of Q_2^{INT} and the universal reduction of the first clause ($\overline{x_1u_1x_2}$):

$$\overline{x_1u_1}z_1 \wedge x_2u_2z_2 \wedge \overline{z_1z_2} \wedge \overline{x_1u_1x_2} \vDash \overline{x_1u_1}u_2\ .$$

Again, we can verify this with a resolution derivation:

$$\dfrac{\dfrac{\dfrac{x_2u_2z_2 \qquad \overline{x_1u_1x_2}}{\overline{x_1u_1}u_2z_2} \qquad \overline{z_1z_2}}{\overline{x_1u_1}u_2\overline{z_1}} \qquad \overline{x_1u_1}z_1}{\overline{x_1u_1}u_2}$$

Similarly the third clause is entailed by the matrix and the universal reductions of the first two clauses (strictly, only the reduction of the second ($\overline{x_1}$) is required)

$$x_1u_1z_1 \wedge \overline{x_2u_2}z_2 \wedge \overline{z_1z_2} \wedge \overline{x_1} \vDash u_1\overline{x_2u_2}\,,$$

and the pattern continues for the final clause:

$$x_1u_1z_1 \wedge x_2u_2z_2 \wedge \overline{z_1z_2} \wedge \overline{x_1} \wedge u_1\overline{x_2} \vDash u_1u_2\ .$$

Resolution derivations verifying these steps can be found easily.

Each individual entailment can be derived immediately using the Σ_1-rule. As the final clause u_1u_2 is fully universal, its universal reduction is the empty clause, yielding a refutation of Q_2^{INT}. ■

Proof of Lemma 4.10. Let $L := (\varepsilon_1, \mu_1), \dots, (\varepsilon_s, \mu_s)$, and let

$$P := \exists X_1 \forall U_1 \cdots \exists X_d \forall U_d \exists X_{d+1}\ .$$

Without loss of generality, we can assume that the X_{d+1}-components of L are all empty, and that the final existential term is $\top$. We proceed by induction on the alternation depth d of Q. Let $i \in [r]$.

Base case $d = 1$. In this case $r = s$, $c_i = \overline{\varepsilon_i} \vee \overline{\mu_i}$, and $\mathsf{red}(c_i) = \overline{\varepsilon_i}$. Let τ be a total assignment falsifying $\overline{\varepsilon_i} \vee \overline{\mu_i}$. If the existential part $\tau_\exists$ satisfies $\bigvee_{k=1}^{i-1} \varepsilon_k$, then it falsifies

$$\bigwedge_{k=1}^{i-1} \overline{\varepsilon_k} = \bigwedge_{k=1}^{i-1} \mathsf{red}(c_k)\ .$$

Otherwise, since $\tau_\exists$ satisfies ε_i, and the universal part $\tau_\forall$ is equal to μ_i, τ falsifies F by definition of countermodel. Since $\varepsilon_s = \top$, $c_s = \bot \vee \overline{\mu_s}$ is fully universal.

Inductive step $d \geq 2$. For each $j \in [s]$, we put

$$\alpha_j := \varepsilon_j \restriction_{X_1} \wedge \mu_j \restriction_{U_1},$$

and claim that $L_j\left[\alpha_j\right]$ is a unified decision list for

$$Q_j := P\left[\alpha_j\right] \cdot \left(F \wedge \bigwedge_{k=1}^{j-1} \overline{\varepsilon_k} \restriction_{X_1}\right)\left[\alpha_j\right],$$

which is a QBF of alternation depth $d-1$. We prove the claim later.

Let p and q be natural numbers such that

$$c_i = \overline{\varepsilon_p} \restriction_{X_1} \vee \overline{\mu_p} \restriction_{U_1} \vee b_q$$

where $\mathcal{E}(L_p\left[\alpha_p\right]) = b_1, \dots, b_{s_p}$. By the inductive hypothesis,

$$\left(F \wedge \bigwedge_{k=1}^{p-1} \overline{\varepsilon_k} \restriction_{X_1}\right)\left[\alpha_p\right] \wedge \bigwedge_{k=1}^{q-1} \mathrm{red}(b_k) \vDash b_q,$$

from which it follows that

$$F \wedge \bigwedge_{k=1}^{p-1} \overline{\varepsilon_k} \restriction_{X_1} \wedge \bigwedge_{k=1}^{q-1} \mathrm{red}(\overline{\varepsilon_p} \restriction_{X_1} \vee \overline{\mu_p} \restriction_{U_1} \vee b_k) \tag{2}$$

entails $\overline{\varepsilon_p} \restriction_{X_1} \vee \overline{\mu_p} \restriction_{U_1} \vee b_q = c_i$.

We show that each conjunct in (2) besides F is $\mathrm{red}(c)$ for some c appearing in $\mathcal{E}(L)$ before c_i. For each $k \in [q-1]$, the clause $\overline{\varepsilon_p} \restriction_{X_1} \vee \overline{\mu_p} \restriction_{U_1} \vee b_k$ appears in $\mathcal{E}(L)$ before c_i by definition. For each $k \in [p-1]$,

$$\overline{\varepsilon_k} \restriction_{X_1} = \mathrm{red}(\overline{\varepsilon_k} \restriction_{X_1} \vee \overline{\mu_k} \restriction_{U_1} \vee f_k)$$

where f_k is the final clause of $\mathcal{E}(L_k\left[\alpha_k\right])$, which is fully universal by the inductive hypothesis, and the clause $\overline{\varepsilon_k} \restriction_{X_1} \vee \overline{\mu_k} \restriction_{U_1} \vee f_k$ appears in L before c_i.

Since $\varepsilon_s = \top$, $c_r = \bot \vee \overline{\mu_s} \restriction_{U_1} \vee f_s$ is fully universal. This completes the inductive step.

Proof of claim. Fixing $j \in [s]$, we show that $L_j\left[\alpha_j\right]$ computes a unified countermodel for Q_j by checking both conditions in Definition 3.5.

(a) Let $\tau \in \{0,1\}^{\mathrm{vars}_\exists(Q_j)}$, and let

$$\sigma := \varepsilon_j \wedge \tau \restriction_{\mathrm{vars}(\tau) \setminus \mathrm{vars}(\varepsilon_j)}.$$

If τ falsifies $\bigwedge_{k=1}^{j-1} \overline{\varepsilon_j} \restriction_{X_1}\left[\alpha_j\right]$, then $\tau \wedge L_j\left[\alpha_j\right](\tau)$ already falsifies the matrix of Q_j, so we assume otherwise. Then $L(\sigma) = \mu_j$, and since $\varepsilon_j \restriction_{X_1} \wedge \tau$ agrees with σ on X_1, $L(\varepsilon_j \restriction_{X_1} \wedge \tau)$ agrees with μ_j on U_1. It follows that

$$L(\varepsilon_j \restriction_{X_1} \wedge \tau) = \mu_j \restriction_{U_1} \wedge L_j\left[\alpha_j\right](\tau),$$

whereby $\alpha_j \wedge \tau \wedge L_j\left[\alpha_j\right](\tau)$ falsifies F, by definition of countermodel. Hence $\tau \wedge L_j\left[\alpha_j\right](\tau)$ falsifies $F\left[\alpha_j\right]$, and therefore falsifies the matrix of Q_j.

(b) Let $\tau, \rho \in \{0,1\}^{\mathrm{vars}_\exists(Q_j)}$, and suppose that τ and ρ agree on the first r existential blocks of Q_j for some $r \in [d-1]$. Since τ and ρ agree on X_1 in particular, if either of them satisfies $\bigwedge_{k=1}^{j-1} \overline{\varepsilon_k} \restriction_{X_1}\left[\alpha_j\right]$, then we have $L_j\left[\alpha_j\right](\tau) = L_j\left[\alpha_j\right](\rho)$ satisfying the condition trivially, so we assume otherwise. Notice that $L_j\left[\alpha_j\right](\tau)$ is $L(\varepsilon_j \restriction_{X_1} \wedge \tau)$ with the U_1-component removed, and likewise for ρ. Since $\varepsilon_j \restriction_{X_1} \wedge \tau$ and $\varepsilon_j \restriction_{X_1} \wedge \rho$ agree on the first $r+1$ existential blocks of Q,

$L(\varepsilon_j\restriction_{X_1} \wedge \tau)$ and $L(\varepsilon_j\restriction_{X_1} \wedge \rho)$ agree on the first $r+1$ universal blocks of Q, thus $L_j\left[\alpha_j\right](\tau)$ and $L_j\left[\alpha_j\right](\rho)$ agree on the first r universal blocks of Q_j.

□

4.3 Unbounded alternation

Theorem 4.2 does not extend to QBFs in general; UDLs prove to be too weak for QBFs of unbounded alternation depth. To show this, we consider a version of the equality formulas with an unbounded, 'interleaved' prefix.

Definition 4.14 (interleaved equality). The n^{th} *interleaved equality formula* Q_n^{INT} is obtained from Q_n^{EQ} by replacing the prefix with $\exists x_1 \forall u_1 \exists z_1 \cdots \exists x_n \forall u_n \exists z_n$.

Recall that the countermodel range for the original equality formulas is the complete set of universal assignments. In fact, this remains true under the interleaved prefix.

Proposition 4.15. *If f is a unified countermodel for Q_n^{INT}, then* $\operatorname{rng}(f) = \{0,1\}^U$, *where* $U = \{u_1, \ldots, u_n\}$.

Proof. The idea of the proof is to show that any countermodel must copy the value of x_i into u_i. Because with the interleaved prefix u_i additionally has access to the values of z_j for $j < i$, we must rule out a larger number of candidate countermodels, which requires some attention to technical detail. The formal proof follows.

For each $i \in [n]$, let D_i denote the existential variables appearing before u_i in the prefix of Q_n^{INT}. We show that the range of any countermodel for Q_n^{INT} is $\{0,1\}^U$, and the proposition follows.

Let f be a countermodel for Q_n^{INT}, and let μ be an arbitrary total assignment to the universal variables. We prove that $\mu = f(\varepsilon)$, where

$$\begin{aligned} \varepsilon(x_i) &:= \begin{cases} 0 & \text{if } \mu(u_i) = 0 \\ 1 & \text{if } \mu(u_i) = 1 \end{cases}, && \text{for } i \in [n], \\ \varepsilon(z_i) &:= 1, && \text{for } i \in [n]. \end{aligned}$$

Aiming for contradiction, let j be the least natural number for which $f(\varepsilon)\restriction_{\{u_j\}} \neq \mu\restriction_{\{u_j\}}$. The matrix of $Q_n^{\mathsf{INT}}[\varepsilon\restriction_{D_j}]$ is

$$a z_j \wedge \overline{z_j} \cdots \overline{z_n} \wedge \bigwedge_{i=j+1}^{n} (\overline{x_i u_i} z_i \wedge x_i u_i z_i)$$

where a is the literal represented by the assignment $f(\varepsilon)\restriction_{\{u_j\}}$. This matrix is satisfied by the assignment

$$f(\varepsilon)\restriction_{\{u_j\}} \wedge \overline{z_j} \wedge z_{j+1} \wedge \cdots \wedge z_n .$$

Now, let δ be any total existential assignment that extends

$$\varepsilon\restriction_{D_j} \wedge \overline{z_j} \wedge z_{j+1} \wedge \cdots \wedge z_n .$$

Since ε and δ agree on D_j, the assignments $f(\varepsilon)\restriction_{\{u_j\}}$ and $f(\delta)\restriction_{\{u_j\}}$ are identical. It follows that the assignment $\delta \cup f(\delta)$ satisfies the matrix of Q_n^{INT}, contradicting the fact that f is a countermodel for Q_n^{INT}. □

As a consequence, the interleaved equality family requires UDLs of exponential size. However, they also admit short QU-Res refutations. As shown in Figure 1, Q_n^{INT} can be reduced to Q_{n-1}^{INT} in a constant-size derivation.

Proposition 4.16. *The interleaved equality formulas admit linear-size* QU-Res *refutations.*

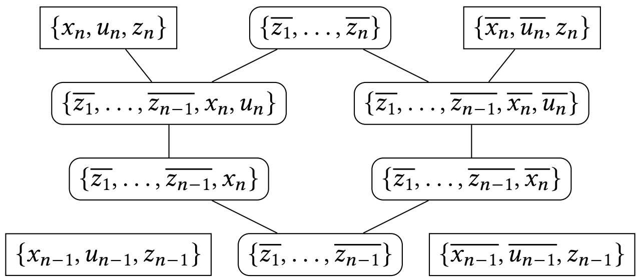

Fig. 1. First portion of a QU-Res refutation of Q_n^{INT}.

Thus distributed decision lists are unsuitable for characterising $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ refutation size when the alternation depth is unbounded.

Corollary 4.17. $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res} \not\leq_p$ UDL *on unbounded alternation.*

4.4 Characterisation of hardness for QU-Res on bounded alternation

If we consider only families of bounded alternation QBFs, given the equivalence between UDLs and the oracle system $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ (Theorem 4.2), there can be only two reasons for hardness in the classical system QU-Res: either

(a) the family requires large UDLs, or
(b) the family harbours propositional resolution hardness.

The main question here is regarding case (b), and what it really means for a QBF family to 'harbour' propositional hardness. In fact, we can give a precise answer: for every family of small UDLs, some steps in the entailment sequences are hard for resolution. This gives rise to a hard sequence of unsatisfiable CNFs for each small family of UDLs.

The result, stated in the following theorem, is a complete characterisation of QU-Res hardness (on bounded alternation), analogous to the hardness characterisations for Frege+∀red and EF+∀red from [16].

Theorem 4.18. *Given a bounded-alternation QBF family* $\{P_n \cdot F_n\}_{n\in\mathbb{N}}$ *requiring superpolynomial-size* QU-Res *refutations, either*

(a) $\{P \cdot F\}_{n\in\mathbb{N}}$ *requires superpolynomial-size UDLs, or*
(b) *for each family of polynomial-size UDLs* $\{L_n\}_{n\in\mathbb{N}}$ *for* $P_n \cdot F_n$ *with entailment sequences* $\mathcal{E}(L_n) = c_1^n, \ldots, c_{r_n}^n$*, there exist natural numbers* $i_n \in [r_n]$ *such that the CNF family*

$$\left\{\left(F_n \wedge \bigwedge_{k=1}^{i_n-1} \mathrm{red}(c_k^n)\right)\left[\overline{c_{i_n}^n}\right]\right\}_{n\in\mathbb{N}}$$

requires superpolynomial-size resolution refutations.

Proof. For $n \in \mathbb{N}$ and $i_n \in [r_n]$, we put

$$\phi_{i_n} := \left(F_n \wedge \bigwedge_{k=1}^{i_n-1} \mathrm{red}(c_k^n)\right)\left[\overline{c_{i_n}^n}\right] .$$

Note that ϕ_{i_n} is unsatisfiable by Lemma 4.10.

Suppose now that neither condition (a) nor condition (b) holds. Then there exist polynomials $p(n), q(n)$ and a family of UDLs $\{L_n\}_{n\in\mathbb{N}}$ with $|L_n| \leq p(n)$ and with entailment sequences $\mathcal{E}(L_n) =$

$c_1^n, \dots, c_{r_n}^n$, such that for all $i_n \in [r_n]$ the CNFs ϕ_{i_n} have resolution refutations of size at most $q(n)$. Let $i_n \in [r_n]$.

By assumption, the alternation depth of each $P_n \cdot F_n$ is bounded above by a constant d, and so Lemma 4.9 provides the bound $|\mathcal{E}(L_n)| = r_n \leq p(n)^d$. Given an arbitrary CNF G and clause b, it is easy to see that a resolution refutation π of $G\,[b]$ can be transformed into a resolution derivation of b from G of size $|\pi|+1$ (it may be necessary to add a weakening step). Hence, there exist derivations of $c_{i_n}^n$ from $F_n \wedge \bigwedge_{k=1}^{i_n-1} \mathrm{red}(c_k^n)$ of size $q'(n) = q(n) + 1$.

Now, beginning with the axiom clauses F_n, and successively deriving and reducing the clauses in $\mathcal{E}(L_n)$, we obtain QU-Res refutations of $P_n \cdot F_n$ of size $O(|P_n \cdot F_n| + p(n)^d q'(n))$. Hence $P_n \cdot F_n$ has polynomial-size QU-Res refutations. □

4.5 Unification of lower-bound techniques

The *two main existing lower-bound techniques* for resolution-based QBF proof systems are *strategy extraction* [8, 9] and *size-cost-capacity* [6]. As far as proof-size lower bounds for bounded-alternation QBFs are concerned, our hardness characterisation (Theorem 4.18) encompasses both.

Indeed, the exact lower bounds for all known bounded-alternation hardness results (all of which have alternation depth 1) can be shown as the result of a UDL lower bound. For QBFs with a single universal block, we have the following immediate corollary to Theorems 4.5 and 4.11.

Corollary 4.19. *Let $\{Q_n\}_{n\in\mathbb{N}}$ be a QBF family of alternation depth 1. Then the following are equivalent statements:*

- *$\{Q_n\}_{n\in\mathbb{N}}$ admits UDLs of size $O(s(n))$;*
- *$\{Q_n\}_{n\in\mathbb{N}}$ admits $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ refutations of size $O(s(n))$.*

Lower bounds by strategy extraction. In [8, 9], a general method was exhibited for forming a QBF Q_f whose unique countermodel is a given Boolean function f. Proof-size lower bounds were shown via strategy extraction, instantiating the function f by PARITY [9, Thm. 14], MAJORITY [8, Cor. 5.7] and SIPSER_d [8, Cor. 5.12], and importing known hardness results for these functions from circuit complexity [30, 43, 47]. In all three cases, the resulting QBF family has a single universal variable, and the imported circuit lower bound holds also for UDLs. As such, all three lower bounds for $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ follow from Corollary 4.19.

Lower bounds by size-cost-capacity. A largely orthogonal technique was proposed in [6]. Here it was shown that the so-called *cost* of a QBF is an absolute lower bound on its $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ refutation size.[6]

In fact, for alternation depth 1, the cost of a QBF is equal to the minimal cardinality of countermodel range, which in turn is a trivial lower bound on UDL size. As such, the lower bounds for equality [6, Thm. 3.5] and random QBFs [6, Thm. 7.9], both of which have alternation depth 1, follow from Corollary 4.19 once the exponential countermodel-range lower bound is established.

5 EQUIVALENCE OF QU-Res AND Q-Res ON BOUNDED ALTERNATION

The natural follow-up question, prompted by our work in Section 4, is whether our results also hold for Q-Resolution (QU-Res without universal pivots). In particular, does the UDL characterisation (Theorem 4.2) continue to hold? In this section, we show that the answer is yes. An immediate corollary is that $\mathrm{Q}\overset{\mathrm{NP}}{-}\mathrm{Res}$ and $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ are p-equivalent on bounded-alternation QBFs.

Perhaps the most obvious approach would be to show that our transformations between $\mathrm{QU}\overset{\mathrm{NP}}{-}\mathrm{Res}$ and UDL go through without resolution on universal pivots. However, we choose another approach.

[6]This is actually shown in the proof of Theorem 4.5. The cost of Q is equal to the maximum, over the individual lists L_i, of the minimal list size (cf. [6]).

We show directly that $\mathrm{Q}^{\mathrm{NP}}_{-}\mathrm{Res}$ is equivalent to $\mathrm{QU}^{\mathrm{NP}}_{-}\mathrm{Res}$, and therefore to UDL. This approach throws up a further interesting result, namely that the classical systems Q-Res and QU-Res are also p-equivalent on bounded alternation.

Definitions of Q-Res **and** $\mathrm{Q}^{\mathrm{NP}}_{-}\mathrm{Res}$. Q-Res is identical to QU-Res, except that resolution pivots must be existential variables.

*Definition 5.1 (*Q-Res *[33]).* A Q-Res derivation from a QBF $P \cdot F$ is a sequence of clauses $\pi := c_1, \dots, c_s$ in which each c_i is derived by one of the following rules:

- *Axiom:* c_i is a clause in the matrix F;
- *∃-Resolution:* $c_i = a \vee b$, where $c_r = a \vee x$ and $c_s = b \vee \overline{x}$ for some $r, s < i$ and some existential variable x.
- *Weakening:* $c_i = c_r \vee b$ for some $r < i$ and clause b.
- *Universal reduction:* $c_i = c_r[\mu]$ for some $r < i$ and some universal assignment μ with $\mathrm{vars}_\exists(c_r) <_P \mathrm{vars}(\mu)$.

The size of π is $|\pi| = s$, and π is a *refutation* when $c_s = \bot$.

For the oracle version of Q-Res, we want to specify a rule which allows a propositional derivation to be collapsed into a single inference. This is complicated by the fact that Q-Res is not propositionally implicationally complete; that is, from $F \vDash c$ it does not follow that c can be derived from F using the axiom, ∃-resolution and weakening rules. As such we do not reuse the Σ_1-rule from $\mathrm{QU}^{\mathrm{NP}}_{-}\mathrm{Res}$, but rather define a new version capturing the insistence on existential pivots.

Definition 5.2 ($\mathrm{Q}^{\mathrm{NP}}_{-}\mathrm{Res}$*).* $\mathrm{Q}^{\mathrm{NP}}_{-}\mathrm{Res}$ is defined as Q-Res, except that the resolution and weakening rules are replaced by the following rule:

- $\Sigma_1^\exists$*-rule:* For some $G \subseteq \{c_1, \dots, c_{i-1}\}$,
  (a) $\bigwedge_{b \in G} b^\exists \vDash c_i^\exists$, and
  (b) for each $b \in G$, $b^\forall$ is a subclause of $c_i^\forall$,

where $c^\exists$ and $c^\forall$ denote the existential and universal subclauses of any clause c.

Equivalences on bounded alternation depth. Both of the p-equivalences that we want to show can be proved constructively, and the essential observation is the following: all of the universal resolutions from a single block can be removed from a QU-Res refutation in quadratic time.

It is also important that the number of universal reduction steps grows only quadratically during the transformation. We denote the number of universal reduction steps in a refutation π by $|\pi|_\forall$.

Lemma 5.3. *Let π be a* QU-Res *refutation of a QBF Q of alternation depth d. For each $i \in [d]$, π can be transformed into a refutation $t(\pi)$ of Q with $|t(\pi)| = O(|\pi|^2)$ and $|t(\pi)|_\forall = O(|\pi|_\forall^2)$ in which there are no resolutions on the i^{th} universal block. The transformation is computable in time $O(|\pi|^2)$.*

Proof. Let $c_1, \dots, c_s$ be a QU-Res refutation of a QBF $\exists X_1 \forall U_1 \cdots \exists X_d \forall U_d \exists X_{d+1} \cdot F$, and let $i \in [d]$. We describe the transformation t recursively on the number r of U_i reductions in π.

If $r = 0$, we obtain $t(\pi)$ from π by removing all U_i resolutions in the following way: we delete all clauses containing a positive U_i literal, and add the empty clause at the end of the refutation. The negative U_i literals, which are no longer resolved away, accumulate through the refutation, and are removed at the conclusion by the addition of a single universal reduction step (hence the addition of the empty clause).

If $r \geq 1$, we find the first U_i reduction step c_j appearing in π, and consider the subderivation π_j ending in c_j. Suppose that the antecedent of c_j is $c_j \vee R$. Now we remove all U_i resolutions from π_j, obtaining a new sequence π'_j, as follows: for each U_i literal in R, we remove all clauses containing the complementary literal; for each variable in U_i not appearing in R, we remove all clauses containing the positive literal. Once again, all U_i literals that are no longer resolved away accumulate through the derivation, and are universally reduced at the conclusion. Then we define $t(\pi) := \pi'_j, t(\pi')$, where π' is identical to π, except that c_j is introduced as an axiom, rather than derived by universal reduction.

It is clear that $|t(\pi)| = O(|\pi|^2)$ and $|t(\pi)|_\forall = O(|\pi|^2_\forall)$, and that t can be computed in time $O(|\pi|^2)$. It remains to prove that $t(\pi)$ is a valid QU-Res refutation of Q with no U_i resolutions. We do this by induction on r.

The base case $r = 0$ is clear. For the inductive step $r \geq 1$, it is clear that π'_j is a valid QU-Res derivation of c_j with no U_i resolutions. Since π' is a QU-Res refutation of $P \cdot F \wedge c_j$ with $r - 1$ U_i reductions, $t(\pi')$ is a valid QU-Res refutation of $P \cdot F \wedge c_j$ with no U_i resolutions, by the inductive hypothesis. The inductive step follows, as c_j is the conclusion of π'_j. □

Now we show the p-equivalence of the classical systems, which is an easy consequence of Lemma 5.3.

Theorem 5.4. Q-Res $\equiv_p$ QU-Res *on bounded alternation.*

Proof. Since QU-Res trivially p-simulates Q-Res, we need only show the reverse simulation. By repeated application of Lemma 5.3, QU-Res refutations π of QBFs of alternation depth d can be transformed into Q-Res refutations of size $O(|\pi|^{2^d})$ in time $O(|\pi|^{2^d})$. Hence Q-Res p-simulates QU-Res when d is bounded above by a constant. □

Next, we show the p-equivalence of the oracle systems.

Theorem 5.5. $\text{Q}^{\text{NP}}_{-}\text{Res} \equiv_p \text{QU}^{\text{NP}}_{-}\text{Res}$ *on bounded alternation.*

Proof. $\text{QU}^{\text{NP}}_{-}\text{Res}$ trivially p-simulates $\text{Q}^{\text{NP}}_{-}\text{Res}$, so we need only show the reverse simulation. Let π be a $\text{QU}^{\text{NP}}_{-}\text{Res}$ refutation of a QBF Q of alternation depth d. We transform π into a $\text{Q}^{\text{NP}}_{-}\text{Res}$ refutation $t(\pi)$ of size $O(|\pi|^{2^d})$.

Since resolution is implicationally complete, whenever the Σ_1-rule is applied, the consequent can be derived by resolution from the antecedents. Hence we can obtain a QU-Res refutation π_0 from π by replacing each entailment step with a resolution derivation. Moreover, $|\pi_0|_\forall = |\pi|_\forall$.

Next we remove the universal resolution steps from π_0 by applying Lemma 5.3 for each $i \in [d]$. We obtain a Q-Res refutation π_1 with $|\pi_1|_\forall = O(|\pi|^{2^d}_\forall)$.

Finally, we transform π_1 into a $\text{Q}^{\text{NP}}_{-}\text{Res}$ refutation $t(\pi)$ as follows. Call a clause in π_1 *surplus* if it is neither an axiom, nor the conclusion, nor the antecedent of a reduction step. We obtain $t(\pi)$ from π_1 by deleting all surplus clauses.

To see that $t(\pi)$ is indeed a $\text{QU}^{\text{NP}}_{-}\text{Res}$ refutation, observe that the removal of surplus clauses from the antecedents preserves $\exists$-entailment steps (realised by the $\Sigma_1^\exists$-rule), since surplus clauses are already $\exists$-entailed by the preceding clauses. As $t(\pi)$ contains only axioms, reduction steps, and antecedents of reduction steps, its size is at most

$$|Q| + 2(|\pi_1|_\forall) = |Q| + O(|\pi|^{2^d}) .$$

Assuming without loss of generality that $|Q| \leq |\pi|$, we have $|t(\pi)| = O(|\pi|^{2^d})$. □

As a corollary of Theorems 4.2 and 5.5, UDLs characterise $\text{Q}^{\text{NP}}_{-}\text{Res}$ refutation size on bounded QBFs.

Corollary 5.6. $\text{Q}^{\text{NP}}_{-}\text{Res} \equiv_p$ UDL *on bounded alternation.*

Unbounded alternation depth. The equivalences in Theorems 5.4 and 5.5 cannot be extended to QBFs in general. The former case is ruled out by the fact that Q-Res does not simulate QU-Res [27], the separation being shown by the QBFs $\{\mathrm{KBKF}_n\}_{n\in\mathbb{N}}$ introduced by Kleine Büning, Karpinski and Flögel [33], which have unbounded alternation depth. Indeed, Theorem 5.4 shows that any such constructive separation *must* be due to a QBF family with unbounded alternation.

The latter case is ruled out by the same QBFs. It is clear that the exponential Q-Res lower bound for KBKF_n [10, 33] is due to exponentially many universal reduction steps (see the proof by size-cost in [6]), giving rise to an exponential lower bound for $\mathrm{Q}^{\underline{\mathrm{NP}}}\mathrm{Res}$. The existence of short (i.e. polynomial-size) $\mathrm{QU}^{\underline{\mathrm{NP}}}\mathrm{Res}$ refutations follows from the existence of short QU-Res refutations. So $\mathrm{Q}^{\underline{\mathrm{NP}}}\mathrm{Res}$ does not simulate $\mathrm{QU}^{\underline{\mathrm{NP}}}\mathrm{Res}$ on unbounded alternation.

6 SIZE-WIDTH FOR QBF RESOLUTION

The seminal paper of Ben-Sasson and Wigderson [4] introduced the celebrated size-width relations, equations which show that short resolution refutations must also be *narrow*. This powerful technique allows resolution size lower bounds to be obtained via *width* lower bounds, the point being that width lower bounds are often much easier to show.

Let us first recall the size-width relation for (general) resolution.[7] The width of a clause is the number of literals it contains, and the width of a resolution refutation is the maximal width of a clause in the sequence. The initial width of a CNF is the maximal width amongst its clauses.

Theorem 6.1 ([4]). *Let F be a CNF with n variables, let $w(F)$ denote the initial width of F, and let $s(F \vdash \bot)$ and $w(F \vdash \bot)$ denote the minimal size and minimal width of a resolution refutation of F. Then*

$$s(F \vdash \bot) = \exp\left(\Omega\left(\frac{(w(F \vdash \bot) - w(F))^2}{n}\right)\right).$$

Size-width is arguably *the* main lower-bound technique for propositional resolution, and its applicability to QBFs has already been investigated [12, 22]. Unfortunately, only negative results were obtained, ruling out the exact relations of Ben-Sasson and Wigderson for various width measures.

In this section, we use the connection to UDLs to show the first positive results, and we apply our new size-width relation to reprove some superpolynomial lower bounds.

6.1 A size-width relation for $\mathrm{QU}^{\underline{\mathrm{NP}}}\mathrm{Res}$

Previous work [12] considered two natural width measures for QBF refutations:

(a) the *standard notion of width*, i.e. the maximal number of literals appearing in a single clause;
(b) *existential width*, i.e. the maximal number of existential literals appearing in a single clause.

We argue that the *correct measure of width for a* $\mathrm{QU}^{\underline{\mathrm{NP}}}\mathrm{Res}$ *refutation is existential width with the axiom clauses not considered.* Thus, we define the existential width of a $\mathrm{QU}^{\underline{\mathrm{NP}}}\mathrm{Res}$ refutation as the maximal number of existential literals appearing in a non-axiom clause.[8] With this definition of existential width, the following size-width relation holds.

Theorem 6.2. *Let $Q = \exists X_1 \forall U_1 \cdots \exists X_d \forall U_d \exists X_{d+1} \cdot F$ be a QBF of alternation depth d, let $v := \sum_{i=1}^{d} |X_i|$ (i.e. the number of existential variables excluding those in the last block), and let $s(F \vdash \bot)$*

[7] There is a separate relation for tree-like resolution [4].
[8] With this definition, the width of an axiom clause c implicitly enters the calculation of the width of a proof in case there is a universal reduction step performed on c.

and $w_\exists(F \vdash \bot)$ *denote the minimal size and minimal existential width of a* $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ *refutation of* Q. *Then*

$$s(F \vdash \bot) = \exp\left(\Omega\left(\frac{(w_\exists(Q \vdash \bot))^2}{d^3 v}\right)\right).$$

Before we proceed to prove Theorem 6.2, a couple of remarks are in order, by way of comparison with the original relation of Ben-Sasson and Wigderson [4].

The first notable difference is the absence of an initial width term, and the related switch from counting the total number of variables n, to the number v of existential variables outside the last block. These both arise from the fact that we apply our narrowing transformation (Lemma 6.5) to UDLs, and those have no concept of initial width; and never even see the variables from the last block. In a way, this highlights that the variables from the last block, and in particular Tseitin variables, are irrelevant for hardness in $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$. Ignoring the last block will turn out crucial later, when proving the lower bound for the majority formulas (Corollary 6.11).

The second obvious difference is in the denominator of the exponent. Here we incur a factor of d^3, related to alternation depth. Hence our relation works best when the alternation depth is bounded, or at least grows very modestly. In Subsection 6.3, we show that in this kind of size-width relation some dependence on the alternation depth is unavoidable.

Note that Theorem 6.2 is not a direct generalization of Theorem 6.1; the propositional case in Theorem 6.2 would have to be obtained by setting $d = 0$, in which case $v = 0$ as well, and in fact $w_\exists(Q \vdash \bot) = 0$ too, because every propositional formula can be refuted in one 0-width $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ step. Thus, we would obtain a meaningless expression containing zeros in both the numerator and the denominator. We also note that the straightforward generalisation of Theorem 6.1 to QBFs does not hold as shown in earlier work [12, 22] (cf. the discussion in Section 6.3).

Proof of the QBF size-width relation. We prove Theorem 6.2 via a transformation from $\mathsf{QU}^{\mathsf{NP}}_{-}\mathsf{Res}$ to UDL and back. A central step in the transformation is based on an adaptation of the following Lemma of Bshouty [18]. It states a size-width relation for (single-output) term decision list. Here, the *width of a decision list* is the maximal width of a term in the list.

Lemma 6.3 ([18]). *Let* $f : \{0,1\}^Z \to \{0,1\}$ *be a function, where* Z *is a set of* n *Boolean variables. If* f *is computed by a decision list of size* s, *then it is also computed by a decision list of width* $O(\sqrt{n \log n \log s})$.

However, UDLs are multi-output term decision lists, so we need to generalise this result for multiple outputs. This is actually quite straightforward, and we could simply copy Bshouty's proof to obtain a generalized version of the lemma for MDLs. [9] However, we take a different approach, and prove an MDL version of Lemma 6.3 following the proof of Ben-Sasson and Wigderson of the size-width transfer for resolution [4, Theorem 3.5]. In this way we obtain a better bound than in Lemma 6.3, by a factor of $\log n$. As a corollary, we obtain a strengthened version of Bshouty's lemma for ordinary decision lists as well.

We split the proof into two parts: Lemma 6.4 states the narrowing transformation of MDLs for arbitrary target width (under appropriate conditions), and Lemma 6.5 plugs in the right parameters to get the optimal bound, resulting in a strengthening of Bshouty's Lemma 6.3.

Recall that the *input width of an MDL* is defined as the maximal width of an input term in the list.

Lemma 6.4. $\forall d \geq 0\ \forall n \geq 0\ \forall b \geq 0$ *if an MDL* L *on* n *input variables has fewer than* $a(n,d)^b$ *terms of input width greater than* d, *then it can be transformed into an equivalent MDL* M *of input width at most* $d + b$, *where* $a(n,d) = (1 - \frac{d}{2n})^{-1}$.

[9] We have done this in the conference version of this paper [7].

Proof. We prove the statement for every d by double induction on n and b. Let $d \geq 0$ be fixed. For an MDL L, let $L^\star_d$ denote the set of *fat* terms of L, i.e. those of (input) width greater than d.

The base case $b = 0$ is trivial, as the condition $|L^\star_d| < a(n,d)^b = 1$ ensures that L already has width at most $d + b = d$. Similarly, the base case $n \leq d$ is also trivial, as the width of a term cannot be larger than the number of variables.

For the inductive step, consider an MDL L for which $|L^\star_d| < a(n,d)^b$. Since the number of occurrences of literals in the terms of $L^\star_d$ is greater than $d|L^\star_d|$ and there are $2n$ literals, by the pigeonhole principle there is a literal c which occurs in more than $d|L^\star_d|/2n$ terms of $L^\star_d$. Therefore, the list $L\restriction_{\overline{c}}$ has fewer than

$$|L^\star_d| - \frac{d|L^\star_d|}{2n} = \frac{|L^\star_d|}{a(n,d)}$$

fat terms. In other words,

$$\left|(L\restriction_{\overline{c}})^\star_d\right| < \frac{|L^\star_d|}{a(n,d)} < a(n,d)^{b-1} < a(n-1,d)^{b-1}\,.$$

Thus, by the induction hypothesis, $L\restriction_{\overline{c}}$ can be transformed into an equivalent MDL L_1 of width at most $d + b - 1$. On the other hand, the list $L\restriction_c$ has $n-1$ variables, and it holds that

$$|(L\restriction_c)^\star_d| \leq |L^\star_d| < a(n,d)^b < a(n-1,d)^b\,,$$

so by the induction hypothesis it can be transformed into an equivalent list L_2 of width at most $d + b$.

Now consider the list M which consists of $\overline{c} \otimes L_1$ followed by L_2. Because L_1 is equivalent to $L\restriction_{\overline{c}}$ and L_2 is equivalent to $L\restriction_c$, M is equivalent to L. The width of M is at most $d + b$. □

The next lemma both improves Bshouty's lemma and generalises it from decision lists to MDLs.

Lemma 6.5. *Let f be a multi-output Boolean function. with n input variables. If f is computed by an MDL of size s, then it is also computed by an MDL of input width $O(\sqrt{n \log s})$.*

Proof. We will apply Lemma 6.4 with $d = b = \left\lceil \sqrt{2n \ln s} \right\rceil$.[10] Let $a := a(n,d)$. We will show that $s < a^b$, which allows us to use Lemma 6.4 to obtain the statement. By $\ln(1+x) \leq x$ we have

$$\ln a = -\ln\left(1 - \frac{d}{2n}\right) \geq \frac{d}{2n}$$

and hence

$$\log_a s = \frac{\ln s}{\ln a} \leq \frac{2n \ln s}{d} < \sqrt{2n \ln s} < b.$$

□

We may now proceed to prove Theorem 6.2, applying Lemma 6.5 to UDLs. In the context of UDLs, since their input variables are the existential variables of a QBF, we speak of *existential width* when referring to their input width.

[10]To be irritatingly pedantic, note that $d > 0$, and $a > 1$, as long as $n > 0$ and $s > 1$. The remaining case is when the MDL computes a constant and has width 0 anyway. Also, note that $\sqrt{2n \ln s} < d$ (the inequality is strict) because e is a transcendental number. Of course, neither of this has any bearing on the asymptotics.

Proof. Let $Q = \exists X_1 \forall U_1 \cdots \exists X_d \forall U_d \exists X_{d+1} \cdot F$ be a QBF of alternation depth d, let $X := \bigcup_{i=1}^{d} X_i$, $v := |X|$, and let π be a *shortest* $\text{QU}^{\text{NP}}_{-}\text{Res}$ refutation of Q, i.e. $s(Q \vdash \bot) = |\pi|$. By Theorem 4.5, π can be transformed into a UDL L of size at most $|\pi|^d$. Since L only uses the variables from X as input variables, by Lemma 6.5, L can be transformed into a UDL M of existential width

$$w_{\exists}(M) = O\left(\sqrt{v \log(|\pi|^d)}\right) = O\left(\sqrt{dv \log |\pi|}\right) .$$

From Lemma 4.9 it follows that the $\text{QU}^{\text{NP}}_{-}\text{Res}$ refutation ρ of Q based on $\mathcal{E}(M)$ (i.e. $t(M)$ as described in the proof of Theorem 4.11) has existential width at most $d \cdot w_{\exists}(M)$. Therefore

$$w_{\exists}(Q \vdash \bot) = O\left(d \cdot \sqrt{dv \log |\pi|}\right) ,$$

and solving for $|\pi|$ yields the theorem statement. □

6.2 $\text{QU}^{\text{NP}}_{-}\text{Res}$ lower bounds by size-width

We illustrate the application of the QBF size-width relation by reproving three exponential QU-Res lower bounds from the literature.

A useful feature of our translation via UDLs is that UDL width lower bounds imply $\text{QU}^{\text{NP}}_{-}\text{Res}$ width lower bounds. Indeed, it is readily verified that the translation in Theorem 4.11 (from UDL to $\text{QU}^{\text{NP}}_{-}\text{Res}$) preserves existential width when the alternation depth is 1.

Proposition 6.6. *A UDL for a QBF Q of alternation depth 1 can be transformed into a* $\text{QU}^{\text{NP}}_{-}\text{Res}$ *refutation of Q with no increase in existential width.*

In the forthcoming examples, linear lower bounds on the existential width of UDLs can be shown with relative ease, whereby application of Proposition 6.6 and Theorem 6.2 yields a size lower bound of $\exp(\Omega(n))$. This is in contrast to the application of size-width relations for propositional resolution, where showing width lower bounds still entails quite some work (cf. [4]).

The equality family. We first show that UDLs for the equality formulas require linear existential width.

Theorem 6.7. *Any UDL for Q_n^{EQ} has existential width n.*

Proof. Let $L := (\varepsilon_1, \mu_1), \ldots, (\varepsilon_s, \mu_s)$ be a UDL for Q_n^{EQ}, and note that L computes the unique countermodel

$$\begin{array}{rcccl} f_{\text{EQ}} & : & \{0,1\}^X & \to & \{0,1\}^U \\ & & \tau & \mapsto & f_{\text{EQ}}(\tau) , \end{array}$$

where $X = \{x_1, \ldots, x_n\}$, $U = \{u_1, \ldots, u_n\}$, and $f_{\text{EQ}}(\tau)(u_i) = \tau(x_i)$ for each $i \in [n]$. Note that the countermodel f_{EQ} amounts to setting each $u_i = x_i$.

Aiming for contradiction, suppose that L has existential width $w < n$. In particular, ε_1 is a term of width less than n, so there exists some variable x_i that does not appear in ε_1. It follows that there exist two assignments $\tau, \rho \in \{0,1\}^X$, both of which satisfy ε_1, with $\tau(x_i) \neq \rho(x_i)$. We deduce that $f_{\text{EQ}}(\tau) = f_{\text{EQ}}(\rho)$, but also that $\tau(x_i) \neq \rho(x_i)$, in contradiction with the definition of f_{EQ}. □

The parity family. Arguing along the same lines, we obtain a linear lower bound on the existential width of UDLs for the parity formulas.

Definition 6.8 (parity [9]). The n^{th} *parity formula* is

$$Q_n^{\mathsf{PAR}} := \exists x_1 \cdots x_n \forall u \exists z_1 \cdots z_n \cdot (x_1 \vee \overline{z}_1) \wedge (\overline{x}_1 \vee z_1) \wedge$$
$$(\overline{u} \vee \overline{z}_n) \wedge (u \vee z_n) \wedge \bigwedge_{i=1}^{n-1} \oplus(x_{i+1}, z_i, z_{i+1}),$$

where $\oplus(x_{i+1}, z_i, z_{i+1})$ consists of the four clauses

$$(x_{i+1} \vee z_i \vee \overline{z_{i+1}}) \wedge (\overline{x_{i+1}} \vee \overline{z_i} \vee \overline{z_{i+1}}) \wedge$$
$$(x_{i+1} \vee \overline{z_i} \vee z_{i+1}) \wedge (\overline{x_{i+1}} \vee z_i \vee z_{i+1}).$$

Theorem 6.9. *Any UDL for Q_n^{PAR} has existential width n.*

Proof. Let $L := (\varepsilon_1, \mu_1), \dots, (\varepsilon_s, \mu_s)$ be a UDL for Q_n^{PAR}, and note that L computes the unique countermodel

$$f_{\mathsf{PAR}} : \{0,1\}^X \to \{0,1\}^{\{u\}}$$
$$\tau \mapsto \left(u \mapsto \left(\Sigma_{i=1}^n \tau(x_i)\right) (\mathrm{mod}\ 2)\right),$$

where $X = \{x_1, \dots, x_n\}$, which amounts to $u = x_1 \oplus \cdots \oplus x_n$.

Similarly as for equality, if the width of ε_1 is strictly less than n, then there exist two assignments $\tau, \rho \in \{0,1\}^X$, both of which satisfy ε_1, and which disagree only at some variable x_i. It follows that $f_{\mathsf{PAR}}(\tau) = f_{\mathsf{PAR}}(\rho)$, and also that

$$\left(\Sigma_{i=1}^n \tau(x_i)\right) (\mathrm{mod}\ 2) \neq \left(\Sigma_{i=1}^n \rho(x_i)\right) (\mathrm{mod}\ 2),$$

contradicting the definition of the function f_{PAR}. □

The majority family. The majority function MAJ is defined as

$$\mathsf{MAJ}(x_1, \dots, x_n) = \left\lfloor \frac{1}{2} + \frac{\left(\Sigma_{i=1}^n x_i\right) - 1/2}{n} \right\rfloor.$$

For each $n \in \mathbb{N}$, let $Q_n^{\mathsf{MAJ}} := \exists x_1 \cdots x_n \forall u \exists z_1 \cdots z_m \cdot F_n$ denote a polynomial-size QBF whose unique countermodel f_{MAJ} amounts to $u = \mathsf{MAJ}(x_1, \dots, x_n)$; that is,

$$f_{\mathsf{MAJ}} : \{0,1\}^X \to \{0,1\}^{\{u\}}$$
$$\tau \mapsto \left(u \mapsto \mathsf{MAJ}(\tau(x_1), \dots, \tau(x_n))\right),$$

where $X = \{x_1, \dots, x_n\}$ (for an explicit construction of such formulas, see [8]). We can show straightforwardly that UDLs for $\{Q_n^{\mathsf{MAJ}}\}_{n \in \mathbb{N}}$ also require linear existential width.

Theorem 6.10. *Any UDL for Q_n^{MAJ} has existential width at least $\frac{n}{2}$.*

Proof. Let $L := (\varepsilon_1, \mu_1), \dots, (\varepsilon_s, \mu_s)$ be a UDL for Q_n^{MAJ}. If the width of ε_1 is strictly less than $\frac{n}{2}$, then there exist two assignments $\tau, \rho \in \{0,1\}^X$, both of which satisfy ε_1, such that

$$\mathsf{MAJ}\left(\tau(x_1), \dots, \tau(x_n)\right) \neq \mathsf{MAJ}\left(\rho(x_1), \dots, \rho(x_n)\right).$$

We reach a contradiction, since $L(\tau) = L(\rho)$, implying that L does not compute the unique countermodel f_{MAJ}. □

Application. Application of Proposition 6.6 and Theorem 6.2 gives the following refutation-size lower bounds.

Corollary 6.11. *$\{Q_n^{\mathsf{EQ}}\}_{n \in \mathbb{N}}$, $\{Q_n^{\mathsf{PAR}}\}_{n \in \mathbb{N}}$, and $\{Q_n^{\mathsf{MAJ}}\}_{n \in \mathbb{N}}$ require* $\mathsf{QU}^{\mathsf{NP}}$-Res *refutations of size $2^{\Omega(n)}$.*

Proof. For each of the three families we have that the number ν of existential variables outside the last block is n. With $d = 1$, Theorem 6.2 gives the result. □

We note that, in contrast to the original hardness proofs for the parity and majority families [8, 10], we obtained Corollary 6.11 without importing any lower bounds from circuit complexity. Also note that the majority formulas may have a quadratic number of variables in the last block [8], and if those were counted in Theorem 6.2, we would not obtain anything; thanks to ignoring them, the argument goes through smoothly.

6.3 Relation to previous work

As it was shown in [12, 22] that the propositional size-width relations (Theorem 6.1) do not lift to Q-Res or QU-Res, it is worthwhile taking a moment to see how those results are consistent with our size-width relation (Theorem 6.2).

The authors of [12, 22] showed that the 'existential-width analogue' of the propositional size-width relation, namely

$$s(Q \vdash \bot) = \exp\left(\Omega\left(\frac{(w_\exists(Q \vdash \bot) - w_\exists(Q))^2}{n}\right)\right), \tag{3}$$

does not hold in Q-Res or QU-Res. In particular, there exist QBFs $\{\phi_n\}_{n\in\mathbb{N}}$ (based on formulas from [32]) that

- have a linear number of variables: $|\mathrm{vars}(\phi_n)| = O(n)$;
- have constant initial existential width: $w_\exists(\phi_n) = O(1)$;
- require QU-Res refutations of linear existential width: $w_\exists(\phi_n \vdash \bot) = \Omega(n)$:
- admit QU-Res refutations of polynomial size: $s(\phi_n \vdash \bot) = n^{O(1)}$.

The QBFs $\{\phi_n\}_{n\in\mathbb{N}}$ clearly violate (3). However, no contradiction follows from Theorem 6.2. Since $\{\phi_n\}_{n\in\mathbb{N}}$ are unbounded alternation QBFs, the n^{th} instance having alternation depth n, Theorem 6.2 yields only a constant lower bound.

We can parameterize the expression in Theorem 6.2 by replacing the fixed exponent of d^3 with a variable c as follows:

$$s(F \vdash \bot) = \exp\left(\Omega\left(\frac{(w_\exists(Q \vdash \bot))^2}{d^c v}\right)\right).$$

It is clear that the smaller the c, the better the bound, and thus we can ask: what is $c^* = \min c$ such that the theorem still holds? Theorem 6.2 implies $c^* \leq 3$, while the formulas $\{\phi_n\}_{n\in\mathbb{N}}$ described above show $c^* \geq 1$, and by extension that the dependence on d cannot be removed, at least in this form. We leave closing the gap as an open problem for future work.

7 CONCLUSIONS

It is interesting to compare our characterisation of QBF resolution hardness with the characterisation of QBF Frege systems [16]. There the authors show a direct correspondence between $\mathcal{C}$-Frege (where lines in the system are $\mathcal{C}$-circuits) and the circuit class $\mathcal{C}$, e.g. hardness in QBF NC^1-Frege is characterised by NC^1 hardness. This is not the case in our results here. Resolution works with CNFs, i.e. formulas of depth 2. By a result of Krause [36], the complexity of decision lists (and hence of UDLs) is strictly intermediate between depth-2 and depth-3 circuits. Hence in QBF resolution, *our circuit model is strictly stronger than the model we use to represent the formulas.* This partly explains why ideas from [8, 16] do not suffice to characterise QBF resolution [14]. In addition to finding the right circuit model of UDLs, new technical ideas (such as the entailment sequence) are needed.

It is also clear from our results that UDLs do not characterise QU-Res hardness for QBFs of *unbounded* quantifier complexity. While QBFs of bounded quantification succinctly represent all problems from the polynomial hierarchy, which covers most applications of modern QBF solving and is prominently represented in QBF evaluation benchmarks [37, 42], we leave open the question of finding the right computational model to characterise QBF resolution for unbounded quantifier complexity.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their detailed and helpful feedback.

Research was supported by grants from the John Templeton Foundation (grant no. 60842), the DFG (grant BE 4209/3-1), the Austrian Science Fund FWF (grant no. J-4361), and the Carl Zeiss Foundation.

We acknowledge that Lemma 5.3 and Theorem 5.4 were independently proven by Judith Clymo [21]. We also acknowledge conversations with Judith Clymo at SAT'19 (Lisbon) on the topic.

REFERENCES

[1] Valeriy Balabanov, Hui-Ju Katherine Chiang, and Jie-Hong R. Jiang. 2014. Henkin quantifiers and Boolean formulae: A certification perspective of DQBF. *Theoretical Computer Science* 523 (2014), 86–100.

[2] Valeriy Balabanov and Jie-Hong R. Jiang. 2012. Unified QBF Certification and its Applications. *Formal Methods in System Design* 41, 1 (2012), 45–65.

[3] Paul Beame and Toniann Pitassi. 2001. Propositional Proof Complexity: Past, Present, and Future. In *Current Trends in Theoretical Computer Science: Entering the 21st Century*, G. Paun, G. Rozenberg, and A. Salomaa (Eds.). World Scientific Publishing, 71-75 Shelton Street, London, XGL WC2H 9JQ, UK, 42–70.

[4] Eli Ben-Sasson and Avi Wigderson. 2001. Short proofs are narrow - resolution made simple. *Journal of the ACM* 48, 2 (2001), 149–169.

[5] Olaf Beyersdorff. 2009. On the Correspondence Between Arithmetic Theories and Propositional Proof Systems – a Survey. *Mathematical Logic Quarterly* 55, 2 (2009), 116–137.

[6] Olaf Beyersdorff, Joshua Blinkhorn, and Luke Hinde. 2019. Size, Cost, and Capacity: A Semantic Technique for Hard Random QBFs. *Logical Methods in Computer Science* 15, 1 (2019).

[7] Olaf Beyersdorff, Joshua Blinkhorn, and Meena Mahajan. 2020. Hardness Characterisations and Size-Width Lower Bounds for QBF Resolution. In *LICS '20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science, Saarbrücken, Germany, July 8-11, 2020*, Holger Hermanns, Lijun Zhang, Naoki Kobayashi, and Dale Miller (Eds.). ACM, New York, 209–223. https://doi.org/10.1145/3373718.3394793

[8] Olaf Beyersdorff, Ilario Bonacina, and Leroy Chew. 2016. Lower Bounds: From Circuits to QBF Proof Systems. In *ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, Madhu Sudan (Ed.). ACM, New York, NY 10019-7434, USA, 249–260.

[9] Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. 2015. Proof Complexity of Resolution-based QBF Calculi. In *International Symposium on Theoretical Aspects of Computer Science (STACS) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 30)*, Ernst W. Mayr and Nicolas Ollinger (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Oktavie-Allee, 66687 Wadern, Deutschland, 76–89.

[10] Olaf Beyersdorff, Leroy Chew, and Mikolás Janota. 2019. New Resolution-Based QBF Calculi and Their Proof Complexity. *ACM Transactions on Computation Theory* 11, 4 (2019), 26:1–26:42.

[11] Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. 2017. Feasible Interpolation for QBF Resolution Calculi. *Logical Methods in Computer Science* 13 (2017). Issue 2.

[12] Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. 2018. Are Short Proofs Narrow? QBF Resolution is *not* so simple. *ACM Transactions on Computational Logic* 19, 1 (2018), 1:1–1:26.

[13] Olaf Beyersdorff, Leroy Chew, and Karteek Sreenivasaiah. 2019. A game characterisation of tree-like Q-Resolution size. *J. Comput. Syst. Sci.* 104 (2019), 82–101.

[14] Olaf Beyersdorff, Luke Hinde, and Ján Pich. 2020. Reasons for Hardness in QBF Proof Systems. *ACM Transactions on Computation Theory* 12, 2 (2020), 10:1–10:27.

[15] Olaf Beyersdorff, Johannes Köbler, and Sebastian Müller. 2011. Proof Systems that Take Advice. *Information and Computation* 209, 3 (2011), 320–332.

[16] Olaf Beyersdorff and Ján Pich. 2016. Understanding Gentzen and Frege Systems for QBF. In *Symposium on Logic in Computer Science (LICS)*, Martin Grohe, Eric Koskinen, and Natarajan Shankar (Eds.). ACM, New York, NY 10019-7434, USA, 146–155.

[17] A. Blake. 1937. *Canonical expressions in boolean algebra.* Ph. D. Dissertation. University of Chicago.

[18] Nader H. Bshouty. 1996. A Subexponential Exact Learning Algorithm for DNF Using Equivalence Queries. *Inform. Process. Lett.* 59, 1 (1996), 37–39.

[19] Samuel R. Buss. 2012. Towards NP-P via proof complexity and search. *Ann. Pure Appl. Logic* 163, 7 (2012), 906–917.

[20] Hubie Chen. 2017. Proof Complexity Modulo the Polynomial Hierarchy: Understanding Alternation as a Source of Hardness. *ACM Transactions on Computation Theory* 9, 3 (2017), 15:1–15:20.
[21] Judith Clymo. 2021. *Proof Complexity for Quantified Boolean Formulas*. Ph. D. Dissertation. School of Computing, University of Leeds.
[22] Judith Clymo and Olaf Beyersdorff. 2018. Relating size and width in variants of Q-resolution. *Inf. Process. Lett.* 138 (2018), 1–6.
[23] Stephen A. Cook and Phuong Nguyen. 2010. *Logical Foundations of Proof Complexity*. Cambridge University Press, Shaftesbury Road, Cambridge, CB2 8BS, UK.
[24] Stephen A. Cook and Robert A. Reckhow. 1979. The Relative Efficiency of Propositional Proof Systems. *Journal of Symbolic Logic* 44, 1 (1979), 36–50.
[25] Nadia Creignou and Daniel Le Berre (Eds.). 2016. *International Conference on Theory and Practice of Satisfiability Testing (SAT)*. Lecture Notes in Computer Science, Vol. 9710. Springer.
[26] Uwe Egly, Martin Kronegger, Florian Lonsing, and Andreas Pfandler. 2017. Conformant planning as a case study of incremental QBF solving. *Annals of Mathematics and Artificial Intelligence* 80, 1 (2017), 21–45.
[27] Allen Van Gelder. 2012. Contributions to the Theory of Practical Quantified Boolean Formula Solving. In *International Conference on Principles and Practice of Constraint Programming (CP) (Lecture Notes in Computer Science, Vol. 7514)*, Michela Milano (Ed.). Springer, 11 West 42nd Street 15th floor, New York, NY 10036, USA, 647–663.
[28] Joshua A. Grochow and Toniann Pitassi. 2018. Circuit Complexity, Proof Complexity, and Polynomial Identity Testing: The Ideal Proof System. *J. ACM* 65, 6 (2018), 37:1–37:59.
[29] Armin Haken. 1985. The intractability of Resolution. *Theoretical Computer Science* 39 (1985), 297–308.
[30] J. Håstad. 1987. *Computational Limitations of Small Depth Circuits*. MIT Press, Cambridge.
[31] Mikolás Janota. 2016. On Q-Resolution and CDCL QBF Solving, See [25], 402–418.
[32] Mikoláš Janota and João Marques-Silva. 2015. Expansion-based QBF solving versus Q-resolution. *Theoretical Computer Science* 577 (2015), 25–42.
[33] Hans Kleine Büning, Marek Karpinski, and Andreas Flögel. 1995. Resolution for Quantified Boolean Formulas. *Information and Computation* 117, 1 (1995), 12–18.
[34] Roman Kontchakov, Luca Pulina, Ulrike Sattler, Thomas Schneider, Petra Selmer, Frank Wolter, and Michael Zakharyaschev. 2009. Minimal Module Extraction from DL-Lite Ontologies Using QBF Solvers. In *International Joint Conference on Artificial Intelligence (IJCAI)*, Craig Boutilier (Ed.). AAAI Press, 2275 East Bayshore Road, Suite 160, Palo Alto, California 94303, USA, 836–841.
[35] Jan Krajíček. 1995. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*. Encyclopedia of Mathematics and Its Applications, Vol. 60. Cambridge University Press, Cambridge.
[36] Matthias Krause. 2006. On the computational power of Boolean decision lists. *Computational Complexity* 14, 4 (2006), 362–375.
[37] Florian Lonsing and Uwe Egly. 2018. Evaluating QBF Solvers: Quantifier Alternations Matter. In *International Conference on Principles and Practice of Constraint Programming (CP) (Lecture Notes in Computer Science, Vol. 11008)*, John N. Hooker (Ed.). Springer, 11 West 42nd Street 15th floor, New York, NY 10036, USA, 276–294.
[38] Florian Lonsing, Uwe Egly, and Allen Van Gelder. 2013. Efficient Clause Learning for Quantified Boolean Formulas via QBF Pseudo Unit Propagation. In *International Conference on Theory and Applications of Satisfiability Testing (SAT) (Lecture Notes in Computer Science, Vol. 7962)*, Matti Järvisalo and Allen Van Gelder (Eds.). Springer, 11 West 42nd Street 15th floor, New York, NY 10036, USA, 100–115.
[39] Florian Lonsing, Uwe Egly, and Martina Seidl. 2016. Q-Resolution with Generalized Axioms, See [25], 435–452.
[40] Hratch Mangassarian, Andreas G. Veneris, and Marco Benedetti. 2010. Robust QBF Encodings for Sequential Circuits with Applications to Verification, Debug, and Test. *IEEE Trans. Comput.* 59, 7 (2010), 981–994.
[41] Jakob Nordström. 2015. On the interplay between proof complexity and SAT solving. *SIGLOG News* 2, 3 (2015), 19–44.
[42] Luca Pulina and Martina Seidl. 2019. The 2016 and 2017 QBF solvers evaluations (QBFEVAL'16 and QBFEVAL'17). *Artif. Intell.* 274 (2019), 224–248.
[43] A. A. Razborov. 1987. Lower bounds for the size of circuits of bounded depth with basis $\{\wedge, \oplus\}$. *Math. Notes Acad. Sci. USSR* 41, 4 (1987), 333–338.
[44] Ronald L. Rivest. 1987. Learning Decision Lists. *Machine Learning* 2, 3 (1987), 229–246.
[45] John Alan Robinson. 1965. A Machine-Oriented Logic Based on the Resolution Principle. *Journal of the ACM* 12, 1 (1965), 23–41.
[46] Nathan Segerlind. 2007. The Complexity of Propositional Proofs. *Bulletin of Symbolic Logic* 13, 4 (2007), 417–481.
[47] R. Smolensky. 1987. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *Proc. 19th ACM Symposium on Theory of Computing*. ACM Press, New York, NY 10019-7434, USA, 77–82.
[48] Raymond M. Smullyan. 1995. *First-order Logic*. Dover Publications, 1196 S Little Creek Rd, Dover, DE 19901, USA.
[49] Moshe Y. Vardi. 2014. Boolean Satisfiability: Theory and Engineering. *Commun. ACM* 57, 3 (2014), 5.

Received October 2020; revised October 2021; accepted August 2022; published January 2023