Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 188 (2020)

Hard QBFs for Merge Resolution*

OLAF BEYERSDORFF, Institut für Informatik, Friedrich-Schiller-Universität Jena, Germany
JOSHUA BLINKHORN, Institut für Informatik, Friedrich-Schiller-Universität Jena, Germany
MEENA MAHAJAN, The Institute of Mathematical Sciences (CI of Homi Bhabha National Institute), India
TOMÁŠ PEITL, Institute of Logic and Computation, TU Wien, Austria
GAURAV SOOD†, The Institute of Mathematical Sciences (CI of Homi Bhabha National Institute), India

We prove the first genuine QBF proof size lower bounds for the proof system Merge Resolution (MRes [7]), a refutational proof system for prenex quantified Boolean formulas (QBF) with a CNF matrix. Unlike most QBF resolution systems in the literature, proofs in MRes consist of resolution steps *together* with information on countermodels, which are syntactically stored in the proofs as merge maps. As demonstrated in [7], this makes MRes quite powerful: it has strategy extraction by design and allows short proofs for formulas which are hard for classical QBF resolution systems.

Here we show the first genuine QBF *exponential lower bounds for MRes*, thereby uncovering limitations of MRes. Technically, the results are either transferred from bounds from circuit complexity (for restricted versions of MRes) or directly obtained by combinatorial arguments (for full MRes). Our results imply that the MRes approach is *largely orthogonal to other QBF resolution models* such as the QCDCL resolution systems QRes and QURes and the expansion systems ∀Exp + Res and IR.

CCS Concepts: • **Theory of computation → Proof complexity**.

Additional Key Words and Phrases: QBF, resolution, proof complexity, lower bounds

1 INTRODUCTION

Proof complexity aims to provide a theoretical understanding of the ease or difficulty of proving statements formally. It also aims to explain the success stories of, as well as the obstacles faced by, algorithmic approaches to hard problems such as satisfiability (SAT) and Quantified Boolean Formulas (QBF) [22, 36]. While propositional proof complexity, the study of proofs of unsatisfiability of propositional formulas, has been around for decades [25, 32], the area of *QBF proof complexity* is relatively new, with theoretical studies gaining traction only in the last decade or so [2, 3, 8, 12, 13]. While inheriting and using a wealth of techniques from propositional proof complexity [14, 16, 30], QBF proof complexity has also given several new perspectives specific to QBF [6, 29, 42], and these perspectives and their connections to QBF solving [10, 18, 39, 45] as well as their practical applications [41] have driven the search for newer proof systems [1, 13, 27, 33, 37].

Many of the currently known QBF proof systems are built on the best-studied propositional proof system *resolution* [20, 40]. Broadly speaking, resolution has been adapted to handle and eliminate the universal variables in QBFs in two intrinsically different ways. The first is an *expansion-based approach:* universal variables are eliminated by implicitly expanding the universal quantifiers into conjunctions, creating annotated copies of existential variables. Universal

*A preliminary version of this article appeared in the proceedings of the 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science – FSTTCS 2020 [9].
†Current affiliation: Department of Computer Science, University of Haifa, Israel

Authors' addresses: Olaf Beyersdorff, olaf.beyersdorff@uni-jena.de, Institut für Informatik, Friedrich-Schiller-Universität Jena, Ernst-Abbe-Platz 2, 07743, Jena, Germany; Joshua Blinkhorn, joshua.blinkhorn@uni-jena.de, Institut für Informatik, Friedrich-Schiller-Universität Jena, Ernst-Abbe-Platz 2, 07743, Jena, Germany; Meena Mahajan, meena@imsc.res.in, The Institute of Mathematical Sciences (CI of Homi Bhabha National Institute), IV Cross Road, CIT Campus, Taramani, Chennai, India; Tomáš Peitl, peitl@ac.tuwien.ac.at, Institute of Logic and Computation, TU Wien, Favoritenstraße 9-11, 1040, Vienna, Austria; Gaurav Sood, gaurav.sood.work@gmail.com, The Institute of Mathematical Sciences (CI of Homi Bhabha National Institute), IV Cross Road, CIT Campus, Taramani, Chennai, India.

ISSN 1433-8092

variables thus appear in the proofs only in the annotations. The systems ∀Exp + Res, IR, and IRM [13, 29] are of this type. The second is a *reduction-rule approach:* under certain conditions, resolution may be blocked, and also under certain conditions, universal variables can be deleted from clauses. The conditions are formulated to preserve soundness, ensuring that if a QBF is true, then so is the QBF resulting from adding a derived clause. The systems QRes, QURes, CP + ∀Red [15, 31, 44] are of this type.

A central role in QBF proof complexity is played by the *two-player evaluation game* on QBFs, and the existence of winning strategies for the universal player in false QBFs. For many QBF resolution systems, such strategies were used to construct proofs and demonstrate completeness, and soundness was demonstrated by extracting such strategies from proofs [1, 13, 26]. The *strategy extraction* procedures build partial strategies at each line of the proof, with the strategies at the final line forming a complete countermodel. These extraction procedures are based on the fact that in each application of a rule in the proof system, any winning strategies of the existential player are not destroyed.

In the systems QRes [31] and QURes [44], the soundness of the resolution rule is ensured by enforcing a very simple side-condition: variables other than the resolved variable (referred to henceforth as the pivot) cannot appear in both polarities in the antecedents. It was observed early on that this is often too restrictive. The *long-distance resolution proof system* LD-QRes [1, 45] arose from efforts to have less restrictive but still sound rules. In this system, a universal variable could appear in both polarities and get merged in the consequent, provided it was to the right of the pivot in the quantifier prefix. This preserves soundness, but the strategy extraction procedures become notably more complex.

The system LD-QRes, while provably better than QRes [26], is still needlessly restrictive in some situations. In particular, by checking a very simple syntactic prefix-ordering condition, it fails to exploit the fact that soundness is not lost even if universal variables to the left of the pivot are merged in both antecedents, provided the partial strategies built for them in both antecedents are identical. A *new system Merge Resolution* (*MRes*) was introduced recently [7] by a subset of the current authors, precisely to address this point. In MRes, partial strategies are explicitly represented within the proof, in a particular representation format called merge maps – these are essentially deterministic branching programs (DBPs). In this format, isomorphism checking can be done efficiently, and this opens the way for enabling sound applications of resolution that would have been blocked in LD-QRes (and QRes). In [7], it was shown that this permitted a simulation of reductionless LD-QRes, denoted rLD-QRes, the fragment of LD-QRes where all reductions are postponed to the very end (no reduction step is followed by a resolution step). (This fragment was identified as interesting in [19]; see also [38].) More importantly, it was also shown in [7] that enabling resolution steps blocked in LD-QRes brought a rich pay-off: there are families of formulas, the Equality and the SquaredEquality formulas, with short (linear-size) proofs in MRes, even in its tree-like and regular versions, but requiring exponential size in QRes, QURes, CP + ∀Red, ∀Exp + Res, and IR. It is notable that the hardness of Equality (and also Squared Equality) in these systems stems from a certain semantic cost associated with these formulas and a corresponding lower bound [5, 6]. Thus the results of [7] show that such semantic costs are not a barrier for MRes.

In this paper, we explore the price paid for overcoming the semantic cost barrier. We show that (expectedly) MRes is not an unconditional success story. Building strategies into proofs via merge maps, and screening out unsoundness only through isomorphism tests, comes at a fairly heavy price: exponentially long proofs for various formulas.

It may be noted that for existentially quantified QBFs, all the QBF proof systems mentioned in this paper coincide with Resolution (or in case of CP + ∀Red, with Cutting Planes). Therefore lower bounds for these propositional proof systems trivially lift to the corresponding QBF proof system. In particular, the separations of tree-like and regular MRes from MRes and other systems follow directly from the propositional case. However, such lower bounds do not tell us much about the limitations of the QBF proof system other than what is known from the underlying propositional proof system. Therefore, in QBF proof complexity, we are interested in 'genuine' QBF lower bounds, i.e. lower bounds that do not follow from propositional lower bounds (cf. [17] on how to formally define the notion of 'genuine' lower bounds in many QBF proof systems). The lower bounds we establish here are of this nature. Specifically, we may consider an MRes derivation of a line from a given set of lines to be purely propositional if at each step, each merge map appearing in the consequent line already appears (in an isomorphic form) in at least one of the antecedent lines. The derivation thus does not contribute to building up the strategies. Collapsing such derivations to single steps (say, by accessing an NP oracle) leaves behind a proof in which purely propositional hardness has been removed. Our arguments show that even such proofs must be large; in this sense, our bounds are genuine QBF lower bounds.

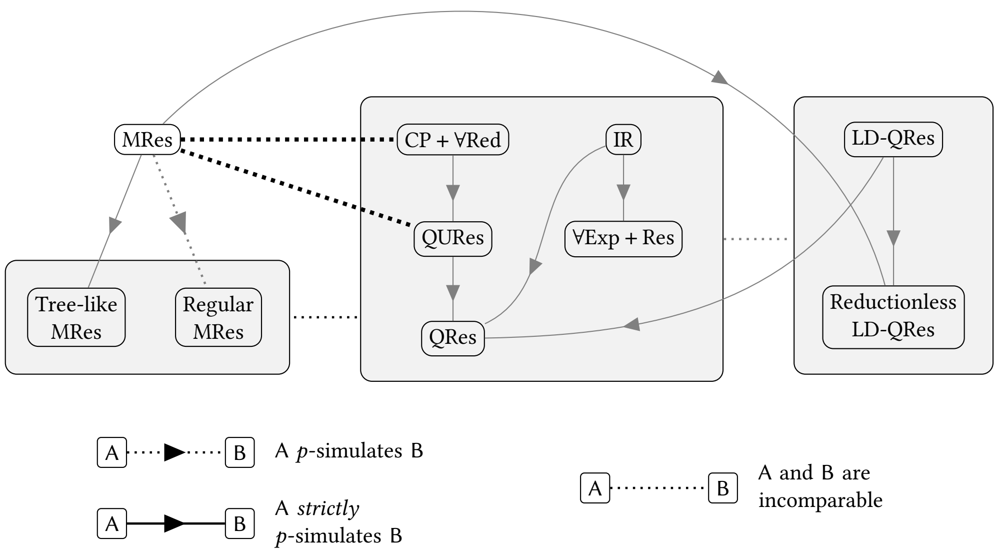

Fig. 1. Visual summary of the proof complexity landscape, with new results shown in bold. Lines from/to a big grey box mean that the line is from/to every proof system within the box. New separations are summarised in Theorems 4.1 and 4.2.

(A) Lower bounds from circuit complexity for restricted versions of MRes. Since the strategies are explicitly represented inside the proofs, computational hardness of strategies immediately translates to proof size lower bounds. While computational hardness of strategies is a known source of hardness in all reduction-based proof systems admitting efficient strategy extraction [11, 13], the computational model relevant for MRes is one for which no unconditional lower bounds are known. For tree-like and regular MRes, the relevant models are decision trees and read-once DBPs, where lower bounds are known. Using this approach, we show:

(1) Tree-like MRes does not simulate regular and general MRes, in terms of genuine size. The QParity formulas witness the separation (Theorem 3.3) as their unique countermodel is the parity function which requires large decision trees.
Note: unlike in the propositional setting, we do not know whether regular MRes simulates tree-like MRes.

(2) Tree-like MRes is incomparable with the dag-like and tree-like versions of QRes, QURes, CP + ∀Red, ∀Exp + Res and IR (Theorem 4.1).
One direction was shown in [7] via the Equality formulas: these formulas are easy for tree-like MRes but hard for dag-like QRes, QURes, CP + ∀Red, ∀Exp + Res, IR. The other direction is witnessed by the Completion Principle formulas, easy in tree-like versions of QRes and ∀Exp + Res [28, 29], but exponentially hard for tree-like MRes (Theorem 3.12). Unlike the QParity formulas, these formulas do not have unique countermodels. However, we show that every countermodel requires large decision-tree size, and hence obtain the lower bound for tree-like MRes.

(B) Combinatorial lower bounds for MRes. Even when winning strategies are easy to compute by DBPs, the formulas can be hard for MRes. We establish such hardness in three cases, obtaining more incomparabilities.

(1) The LQParity formulas, easy in ∀Exp + Res [13], are exponentially hard for regular MRes (Theorem 3.6). Hence regular MRes is incomparable with ∀Exp + Res and IR.
(2) The Completion Principle formulas, easy in tree-like versions of QRes and ∀Exp + Res [28, 29], are exponentially hard for regular MRes (Theorem 3.13). Hence regular MRes is incomparable with the dag-like and tree-like versions of QRes, QURes, CP + ∀Red, ∀Exp + Res and IR (Theorem 4.1).
(3) The KBKF-lq formulas, easy in QURes [2], are exponentially hard for MRes (Theorem 3.17). Hence MRes is incomparable with QURes and CP + ∀Red (Theorem 4.2).

The third hardness result above for the KBKF-lq formulas provides the first genuine lower bound for the full system of MRes, for which previously no such lower bounds were known.

Figure 1 depicts the *simulation order and incomparabilities* we establish involving MRes and its refinements. Amongst the five systems in the big grey box, all relationships not directly implied by depicted connections are known to be incomparabilities [13, 15, 29].

More recently, upper bounds for the proof system MRes have been established, in [24], and variants of MRes have been explored, in [23].

Organization of this paper. We define QBFs and MRes in Section 2. In Section 3, we prove lower bounds for many formula families. Finally, in Section 4, we give the resulting separations among QBF proof systems.

2 PRELIMINARIES

Let $[n] = \{1, 2, \ldots, n\}$ and $[m, n] = \{m, \ldots, n\}$.

Variables take Boolean values, and a literal ℓ is a variable x or its negation $\neg x$ (also denoted $\bar{x}$). We say that $x = \mathrm{var}(\ell)$. A clause is a disjunction of literals, and a conjunctive-normal-form (CNF) formula is a conjunction of clauses. We represent clauses interchangeably as disjunctions of literals and sets of literals. Similarly, we represent CNF formulas interchangeably as conjunctions of clauses and sets of clauses.

The *resolution rule* derives, from clauses $C \vee \ell$ and $D \vee \neg\ell$ for some literal ℓ, the clause $C \vee D$. We say that $C \vee D$ is the resolvent, $x = \mathrm{var}(\ell)$ is the pivot, and denote this by $C \vee D = \mathrm{res}(C \vee \ell, D \vee \neg\ell, x)$. Representing clauses as sets of literals, we say that $C \cup D$ is the resolvent of $C \cup \{\ell\}$ and $D \cup \{\bar{\ell}\}$ on pivot x, and denote this by $C \cup D = \mathrm{res}(C \cup \{\ell\}, D \cup \{\bar{\ell}\}, x)$.

The *propositional proof system Resolution* proves that a CNF formula F is unsatisfiable by deriving the empty clause through repeated applications of the resolution rule.

2.1 Quantified Boolean formulas

A *Quantified Boolean Formula* (QBF) in *prenex conjunctive normal form* is denoted $\Phi \coloneqq Q \cdot \phi$, where

- $Q = Q_1 Z_1 Q_2 Z_2 \ldots Q_k Z_k$ is the quantifier prefix, in which Z_i are pairwise disjoint finite sets of Boolean variables, $Q_i \in \{\exists, \forall\}$ for each $i \in [k]$ and $Q_i \neq Q_{i+1}$ for each $i \in [k-1]$, and
- the matrix ϕ is a CNF over $\text{vars}(\Phi) \coloneqq \cup_{i \in [k]} Z_i$.

The existential (resp. universal) variables of Φ, typically denoted X or $X_\exists$ (resp. U or $X_\forall$) is the set obtained as the union of Z_i for which $Q_i = \exists$ (resp. $Q_i = \forall$). The prefix Q defines a binary relation $<_Q$ on $\text{vars}(\Phi)$, such that $z <_Q z'$ holds iff $z \in Z_i$, $z' \in Z_j$, and $i < j$, in which case we say that z' is right of z and z is left of z'. For each $u \in U$, we define $L_Q(u) \coloneqq \{x \in X \mid x <_Q u\}$, i.e. the existential variables left of u.

For a set of variables Z, let $\langle Z \rangle$ denote the set of assignments to Z. A *strategy* h for a QBF Φ is a set $\{h^u \mid u \in U\}$ of functions $h^u \colon \langle L_Q(u) \rangle \to \{0, 1\}$ (for each $\alpha \in \langle X \rangle$, $h^u(\alpha{\restriction}_{L_Q(u)})$ and $h(\alpha)$ should be interpreted as a Boolean assignment to the variable u and the variable set U respectively). Additionally h is *winning* if, for each $\alpha \in \langle X \rangle$, the restriction of ϕ by the assignment $(\alpha, h(\alpha))$ is false. We use the terms "winning strategy" and "countermodel" interchangeably. A QBF is called false if it has a countermodel, and true if it does not.

The semantics of QBFs is also explained by a *two-player evaluation game* played on a QBF. In a run of the game, two players, the existential and the universal player, assign values to the variables in the order of quantification in the prefix. The existential player wins if the assignment so constructed satisfies all the clauses of ϕ; otherwise the universal player wins. Assigning values according to a countermodel guarantees that the universal player wins no matter how the existential player plays; hence the term "winning strategy".

2.2 The Merge Resolution proof system

We first describe the idea behind the Merge Resolution (MRes) proof system. MRes is a line-based proof system. A refutation in Merge Resolution is a sequence of lines. Each line L consists of a clause C with only existential literals, and a partial strategy h^u for each universal variable u. The idea is to maintain the invariant that for each existential assignment α, if α falsifies C, then α extended by the partial universal assignment setting each u to $h^u(\alpha)$ falsifies at least one of the clauses used to derive L. Thus the set of functions $\{h^u\}$ gives a partial strategy (for the universal player) that wins whenever the existential player plays from the set of assignments falsifying C. The goal is to derive a line with the empty clause; the corresponding strategy at that line will be a complete winning strategy for the universal player, i,e. a countermodel. Along the way, resolution is used on the clauses. If the pivot is x, then for universal variables u right of x, the partial strategies can be combined with a branching decision on x. However, for u left of x, in the evaluation game, the value of u is already set when x is to be assigned. Thus already existing non-trivial partial strategies for u cannot be combined with a branching decision, and so this resolution step is blocked. However, if both the strategies are identical, or if one of them is trivial (unspecified), then the non-trivial strategy can be carried forward while maintaining the desired invariant. Checking whether strategies are identical can itself be hard, making verification of the proof difficult. In MRes, this is handled by choosing a particular representation called merge maps, where isomorphism checks are easy.

Now we can describe the proof system itself. First we describe *merge maps*.

Definition 2.1. Merge maps are deterministic branching programs, specified by a sequence of instructions of one of the following two forms:

- $\langle \texttt{line } \ell \rangle : b$, where $b \in \{*, 0, 1\}$.[1]
  Merge maps containing a single such instruction are called simple. In particular, if $b = *$, then they are called trivial.
- $\langle \texttt{line } \ell \rangle$: `If` $x = 0$ `then go to` $\langle \texttt{line } \ell_1 \rangle$ `else go to` $\langle \texttt{line } \ell_2 \rangle$, for some $\ell_1, \ell_2 < \ell$. In a merge map M for u, all queried variables x must precede u in the quantifier prefix.
  Merge maps with such instructions are called complex.

(All line numbers are natural numbers.) The merge map M^u computes a partial strategy for the universal variable u starting at the largest line number (the leading instruction) and following the instructions in the natural way. The value $*$ denotes an undefined value.

Definition 2.2. Two merge maps M_1 and M_2 are said to be *consistent*, denoted $M_1 \bowtie M_2$, if for every line number i appearing in both M_1, M_2, the instructions with line number i are identical.

When two merge maps, M_1 and M_2, are consistent, it is possible to build the merge map: `If` $x = 0$ `then go to` M_1 `else go to` M_2 without repeating the common parts of M_1 and M_2. To be more precise, the new merge map will contain all instructions of M_1 and M_2 and the following additional instruction: `If` $x = 0$ `then go to` $\langle$`leading instruction of` $M_1\rangle$ `else go to` $\langle$`leading instruction of` $M_2\rangle$.

Definition 2.3. Two merge maps M_1, M_2 are said to be isomorphic, denoted $M_1 \simeq M_2$, if there is a bijection between the line numbers in M_1 and M_2 that transforms M_1 to M_2 in the natural way.

For the remainder of this section let $\Phi = Q \cdot \phi$ be a QBF with existential variables X and universal variables U.

Definition 2.4. The *proof system MRes* has the following rules:

(1) *Axiom:* For a clause A in the matrix ϕ, let C be the existential part of A. For each universal variable u, let b_u be the value u must take to falsify A; if $u \notin \text{var}(A)$, then $b_u = *$. For any natural number i, the line $(C, \{M^u : u \in U\})$ where each M^u is the simple merge map $\langle i \rangle : b_u$ can be derived in MRes.
(2) *Resolution:* From lines $L_a = (C_a, \{M^u_a : u \in U\})$ for $a \in \{0, 1\}$, in MRes, the line $L = (C, \{M^u : u \in U\})$ can be derived, where for some $x \in X$,
   - $C = \text{res}(C_0, C_1, x)$, and
   - for each $u \in U$; either M^u_a is trivial and $M^u = M^u_{1-a}$ for some a; or $M^u = M^u_0 \simeq M^u_1$; or x precedes u, $M_1 \bowtie M_2$ and M^u has all the instructions of M^u_1 and M^u_2 in addition to the following instruction:
     `If` $x = 0$ `then go to` $\langle$`leading instruction of` $M^u_1\rangle$ `else go to` $\langle$`leading instruction of` $M^u_2\rangle$.
     The line number of this leading instruction is the number (position) of the line L in the derivation.

   With slight abuse of notation, we will call L the resolvent of L_0 and L_1 with pivot x, and denote this by $L = \text{res}(L_0, L_1, x)$.
   Note that [7] also requires that the positive literal of the pivot appears in the first argument, so $x \in C_0$ (i.e. the clause at line L_0) and $\overline{x} \in C_1$ (the clause at line L_1). However, this was only for syntactic convenience, and the way we formulate our arguments, this is not necessary.)

Note that the entire merge maps are not stored at each line, only the leading instruction specific to the line. Due to consistency, this is enough information to build the entire map from the derivation. As noted in [7] (Proposition 19),

[1] In [7], the notation used is $b \in \{*, u, \overline{u}\}$; $u, \overline{u}, *$ denote $u = 1, u = 0$, undefined respectively.

for lines within the same derivation, the corresponding merge maps are always consistent. Therefore, in the above definition, we don't have to explicitly do a consistency check.

Definition 2.5. A *refutation* is a derivation using these rules and ending in a line with the empty existential clause. The size of the refutation is the number of lines.

In the rest of this paper, we will denote refutations by the Greek letter Π. A refutation can be represented as a graph (with edges directed from the antecedents to the consequent, hence from the axioms to the final line). We denote the graph corresponding to refutation Π by G_Π. The lines of Π will be denoted by L, L_1, L_2, L', L'' etc. For lines L, L_i and L', and universal variable z, the respective clause, merge map and the function computed by the merge map will be denoted by C, M^z, h^z, C_i, M_i^z, h_i^z and C', $(M')^z$, $(h')^z$ respectively.

Definition 2.6. Let Y be a subset of the existential variables X of Φ. We say that an MRes refutation Π of Φ is Y*-regular* if for each $y \in Y$, there is no leaf-to-root path in G_Π that uses y as pivot more than once. An X-regular proof is simply called a *regular proof.* If G_Π is a tree, then we say that Π is a *tree-like proof.*

Example 2.7. We reproduce from [7] a small example to illustrate how MRes operates. The formulas to be refuted are the Equality formulas from [6], defined as follows: The *Equality family* is the QBF family whose nth instance has the prefix $\exists x_1, \ldots, x_n, \forall u_1, \ldots, u_n, \exists t_1, \ldots, t_n$ and the matrix consisting of the clauses $\{x_i, u_i, t_i\}$, $\{\overline{x}_i, \overline{u}_i, t_i\}$ for $i \in [n]$, and $\{\overline{t}_1, \ldots, \overline{t}_n\}$.

In [7] (Example 3), these formulas are shown to have linear-size refutations in the system rLD-QRes denoting reductionless LD-QRes, the fragment of LD-QRes where all reductions are postponed to the very end (no reduction step is followed by a resolution step). Later in [7] (Theorem 22), MRes is shown to simulate reductionless LD-QRes. Hence these formulas are easy to refute in MRes. On the other hand, these formulas are known to require exponential-size refutations in QRes, QURes, CP + $\forall$Red [6], $\forall$Exp + Res and IR [5] (cf. [4] on how to apply the lower bound technique from [5] to the Equality formulas).

Here, we directly present the implied linear-size MRes refutations (in fact, these refutations are also tree-like and regular) for the Equality formulas.

First, we download the axioms. Line 0 downloads the long clause $\{\overline{t}_1, \ldots, \overline{t}_n\}$, with all trivial merge maps. The next $2n$ lines download the short axiom clauses. Letting $i \in [n]$, we define these lines as follows: Line $2i-1$ is the clause $\{x_i, t_i\}$ with merge map 0 for u_i and all other merge maps are trivial. Line $2i$ is the clause $\{\overline{x}_i, t_i\}$ with merge map 1 for u_i and all other merge maps are trivial.

For $i \in [n]$, line $2n+i$ is obtained by applying the merge resolution rule on lines $2i-1$ and $2i$. This gives the clause $\{t_i\}$; the merge maps for $j \neq i$ are trivial, and the merge map for u_i has the instruction: `If` $x_i = 0$ `then go to` $\langle$`line` $2i-1\rangle$ `else go to` $\langle$`line` $2i\rangle$.

At line $3n+1$, applying merge resolution on lines 0 and $2n+1$, we obtain the clause $\{\overline{t}_2, \ldots, \overline{t}_n\}$. The merge map for u_1 is taken from line $2n+1$, since at line 0 it is trivial.

Now for $i \in [2, n]$, line $3n+i$ is obtained by applying merge resolution on lines $2n+i$ and $3n+i-1$. This gives the clause $\{\overline{t}_{i+1}, \ldots, \overline{t}_n\}$. The merge map for u_i is taken from line $2n+i$ since at line $3n+i-1$ it is trivial. For $j < i$, the merge map for u_j is taken from line $3n+i-1$ since at line $2n+i$ it is trivial. Effectively, at this line, for all $j \leq i$, the merge map for u_j is from line $2n+j$, and for all $j > i$, the merge map for u_j is trivial.

Line $4n$ derives the empty clause and the strategy computing, for each $i \in [n]$, $u_i = x_i$. This completes the refutation.

A crucial fact about the proof system MRes, shown in [7], is that the merge maps at the final line of an MRes refutation compute a countermodel for the QBF. To establish this fact, some stronger properties of MRes refutations are established and will be useful to us. We restate the relevant properties here.

Lemma 2.8 (Extracted/adapted from [7] Section 4.3, (Proof of Lemma 21)). *Let $\Phi = Q \cdot \phi$ be a QBF with existential variables X and universal variables U. Let $\Pi \stackrel{\text{def}}{=} L_1, \ldots, L_m$ be an MRes refutation of Φ, where each $L_i = (C_i, \{M_i^u \mid u \in U\})$. Further, for each $i \in [m]$,*

- *let α_i be the minimal partial assignment falsifying C_i,*
- *let A_i be the set of assignments to X consistent with α_i,*
- *for each $u \in U$, let h_i^u be the function computed by M_i^u,*
- *for each $\alpha \in A_i$, let $h_i(\alpha)$ be the partial assignment which sets variable u to $h_i^u(\alpha\restriction_{L_Q(u)})$ if $h_i^u(\alpha\restriction_{L_Q(u)}) \neq *$, and leaves it unset otherwise.*

Then for each $\alpha \in A_i$, the (partial) assignment $(\alpha, h_i(\alpha))$ falsifies at least one clause of ϕ used in the sub-derivation of L_i.

Proposition 2.9 ([7]). *Let Φ and Π be as defined in Lemma 2.8. Then, for all $u \in U$, M_m^u is isomorphic to a subgraph of G_Π (up to path contraction).*

3 LOWER BOUNDS

We now start to explore lower bounds for MRes. In Section 3.1, we show how to construct generic hard formulas for tree-like and regular MRes. In Sections 3.2 to 3.5, we prove lower bounds for specific QBF formulas.

3.1 Lower bounds for generic formulas

The following theorem, implicit in [7], is an immediate consequence of Lemma 2.8 and Proposition 2.9.

Theorem 3.1. *Let $\Phi = Q \cdot \phi$ be a false QBF with existential variables X and universal variables U. If, for every countermodel of Φ, the function for some universal variable u requires size at least s to compute by branching programs (resp. decision trees, read-once branching programs), then every MRes (resp. tree-like MRes, regular MRes) refutation of Φ has size at least s.*

Currently, no exponential lower bounds are known for general branching programs. Therefore, we cannot use the above theorem to prove lower bounds for general MRes. However, we can use it to prove exponential lower bounds for tree-like and regular MRes. To do so, we need a QBF whose countermodel requires exponential decision-trees (resp. read-once branching programs). We now show how to construct such QBFs. This follows the method used, for instance, in [13, Sec. 4.1] and [38, Sec. 6].

Let $f : X \to \{0, 1\}$ be a Boolean function, let C_f be a Boolean circuit encoding f, and let u be a variable not in X. Using the Tseitin transformation [43], we can construct a CNF formula $\phi(X, u, Y)$ such that $\exists Y.\phi(X, u, Y)$ is logically equivalent to $C_f(X) \neq u$. Then, the QBF formula $\Phi \coloneqq \exists X \forall u \exists Y.\phi(X, u, Y)$, called the QBF encoding of f, is a false QBF formula with f as the unique winning strategy. Moreover, the size of Φ is polynomial in the size of C_f. This is the desired QBF formula.

3.2 The QParity formulas

We now turn our attention to lower bounds for specific formulas. We start with the QParity formulas in this section. These are the formulas obtained by the Tseitin transformation described above, using a linear-size read-once branching program computing the parity function. These formulas were defined in [13] where they were shown to be hard for QRes and QURes. It was also shown that these formulas are easy for the expansion-based systems $\forall$Exp + Res, IR and IRM. It was hence concluded that QRes and QURes do not simulate $\forall$Exp + Res, IR and IRM.

Before we define the formulas, we set up some notation. For variables o, o_1, o_2, let $\mathrm{xor}(o_1, o)$ and $\mathrm{xor}(o_1, o_2, o)$ be the following sets of clauses:

$$\mathrm{xor}(o_1, o) = \{\overline{o_1} \vee o, o_1 \vee \overline{o}\},$$
$$\mathrm{xor}(o_1, o_2, o) = \{\overline{o_1} \vee \overline{o_2} \vee \overline{o}, \overline{o_1} \vee o_2 \vee o, o_1 \vee \overline{o_2} \vee o, o_1 \vee o_2 \vee \overline{o}\}$$

Note that xor on a set of variables is just the CNF representation of the constraint that the number of variables set to true is even. That is, $\mathrm{xor}(o_1, o)$ is satisfied iff $o \equiv o_1 \pmod 2$, and $\mathrm{xor}(o_1, o_2, o)$ is satisfied iff $o \equiv o_1 + o_2 \pmod 2$.

Definition 3.2. The $\mathrm{QParity}_n$ formula [13] is the QBF $\exists x_1, \ldots, x_n, \forall z, \exists t_1, \ldots, t_n.\ \left(\bigwedge_{i\in[n+1]} \phi_n^i\right)$ where

$$\begin{aligned}
\phi_n^1 &= \mathrm{xor}(x_1, t_1);\\
\phi_n^i &= \mathrm{xor}(t_{i-1}, x_i, t_i), \qquad \forall i \in [2, n];\\
\phi_n^{n+1} &= \{t_n \vee z, \overline{t_n} \vee \overline{z}\}.
\end{aligned}$$

The QBFs are false: they claim that there exist $x_1, \ldots, x_n$ such that $x_1 + \cdots + x_n$ is neither congruent to 0 nor 1 modulo 2. Note that the only winning strategy for the universal player is to play z satisfying $z \equiv x_1 + \cdots + x_n \pmod 2$.

Theorem 3.3. *Every tree-like MRes refutation of $\mathrm{QParity}_n$ has size at least 2^n.*

Proof. It is a folklore fact that the n-input parity function requires decision-trees of size at least 2^n. From Theorem 3.1, we obtain the desired lower bound. □

3.3 The LQParity formulas

We now turn our attention to the LQParity formulas. These formulas are variants of the QParity formulas, and were originally defined in [13]. The variant is designed to allow the QRes lower bound arguments for QParity to be adapted also to LD-QRes. Like the QParity formulas, these formulas are easy for several proof systems, including $\forall$Exp + Res, IR and IRM, but were shown to be hard forQURes and LD-QRes. This then established that LD-QRes does not simulate $\forall$Exp + Res, IR and IRM [13].

We now describe the formulas. They are obtained from the QParity formulas by duplicating each clause except those in ϕ_n^{n+1}, and inserting the universal variable z in one copy and its negation $\overline{z}$ in the other. Formally, they can be defined as follows: For variables o, o_1, o_2, z, let $\mathrm{xor}_l(o_1, o, z)$ and $\mathrm{xor}_l(o_1, o_2, o, z)$ be the following sets of clauses:

$$\mathrm{xor}_l(o_1, o, z) = \{\overline{o_1} \vee o \vee z, o_1 \vee \overline{o} \vee z\},$$
$$\mathrm{xor}_l(o_1, o_2, o, z) = \{\overline{o_1} \vee \overline{o_2} \vee \overline{o} \vee z, \overline{o_1} \vee o_2 \vee o \vee z, o_1 \vee \overline{o_2} \vee o \vee z, o_1 \vee o_2 \vee \overline{o} \vee z\}$$

Definition 3.4. The $\mathrm{LQParity}_n$ formula [13] is the QBF $\exists x_1, \ldots, x_n, \forall z, \exists t_1, \ldots, t_n.\ \left(\bigwedge_{i\in[n+1]} \phi_n^i\right)$ where

$$\begin{aligned}
\phi_n^1 &= \mathrm{xor}_l(x_1, t_1, z) \cup \mathrm{xor}_l(x_1, t_1, \overline{z}),\\
\phi_n^i &= \mathrm{xor}_l(t_{i-1}, x_i, t_i, z) \cup \mathrm{xor}_l(t_{i-1}, x_i, t_i, \overline{z}) \qquad \forall i \in [2, n],\\
\phi_n^{n+1} &= \{t_n \vee z, \overline{t_n} \vee \overline{z}\}.
\end{aligned}$$

For $i, j \in [n+1], i \leq j$, let $\phi_n^{[i,j]}$ denote $\bigwedge_{k\in[i,j]} \phi_n^k$. Also, let $X = \{x_1, \ldots, x_n\}$ and $T = \{t_1, \ldots, t_n\}$.

Observation 3.5. *(a) For each $i \in [n]$, and each $C \in \phi_n^i$, $\{x_i, t_i\} \subseteq var(C)$; and (b) for each $i \in [n+1] \setminus \{1\}$, and each $C \in \phi_n^i$, $\{t_{i-1}\} \subseteq var(C)$.*

We will now show that LQParity formulas require exponential-size refutations in regular MRes.

Theorem 3.6. *Every T-regular refutation of LQParity$_n$ in MRes, and hence any regular MRes refutation, has size at least 2^n.*

The proof proceeds as follows: Let Π be a T-regular MRes refutation of LQParity$_n$. Since every axiom has a variable from T while the final clause in Π is empty, there is a maximal "component" (say $\mathcal{S}$) of the proof leading to and including the final line, where all clauses are T-free. The clauses in this component involve only the X variables. We show that the "boundary" ($\partial\mathcal{S}$) of this component is large, by showing in Lemma 3.9 that each clause at the boundary must be wide. (This idea was used in [38] to show that CR is hard for reductionless LD-QRes.) To establish the width bound, we note that no lines have trivial strategies. Since the pivots at the boundary are variables from T, the merge maps incoming into each boundary resolution must be isomorphic. By carefully analysing which axiom clauses can and must be used to derive lines just above the boundary (Lemma 3.8), we conclude that the merge maps must be simple, yielding the lower bound. To fill in all the details, we first describe some properties (Lemma 3.7) of Π that will be used in obtaining this result.

Recall that the lines of Π have mergemaps for the universal variable. Since LQParity formulas have a single universal variable, we avoid the superscript z. Thus a line L (resp. L_1, L_2, L', L'' etc) has a merge map M (resp. M_1, M_2, M', M'', etc) for z and computes the function h (resp, h_1, h_2, h', h'', etc.).

Let G_Π be the derivation graph corresponding to Π (with edges directed from the antecedents to the consequent, hence from the axioms to the final line). We will refer to the nodes of this graph by the corresponding line. For $L, L' \in \Pi$, we will say $L \rightsquigarrow L'$ if there is a path from L to L' in G_Π.

For a line $L \in \Pi$, let Π_L be the minimal sub-derivation of L, and let G_{Π_L} be the corresponding subgraph of G_Π with sink L. Define $\mathrm{UC}(\Pi_L) = \{\phi_n^i \mid i \in [n+1], \mathrm{leaves}(G_{\Pi_L}) \cap \phi_n^i \neq \emptyset\}$, and $\mathrm{UCI}(\Pi_L) = \{i \in [n+1] \mid \phi_n^i \in \mathrm{UC}(\Pi_L)\}$. (The abbreviations UC and UCI stand for UsedConstraints and UsedConstraintsIndex respectively.) Note that for any leaf L, $\mathrm{UCI}(\Pi_L)$ is a singleton.

Define $\mathcal{S}$ to be the set of those lines in Π where the clause part has no T variable and furthermore there is a path in G_Π from the line to the final empty clause via lines where all the clauses also have no T variables. Let $\partial\mathcal{S}$, called the boundary of $\mathcal{S}$, denote the set of leaves in the subgraph of G_Π restricted to $\mathcal{S}$; these are lines that are in $\mathcal{S}$ but their parents are not in $\mathcal{S}$. Note that no leaf of Π is in $\mathcal{S}$ because all leaves of G_Π contain a variable in T.

Lemma 3.7. *Let $L = (C, M)$ be a line of Π. Then $\mathrm{UCI}(\Pi_L)$ is an interval $[i, j]$ for some $1 \leq i \leq j \leq n+1$. Furthermore, (below i, j refer to the endpoints of this interval)*

(1) For all $k \in [i, j-1]$, $t_k \notin var(C)$.
(2) If $i > 1$, then $t_{i-1} \in var(C)$.
(3) If $j \leq n$, then $t_j \in var(C)$.
(4) $|var(C) \cap T| = 1$ iff $[i, j]$ contains exactly one of $1, n+1$.
$var(C) \cap T = \emptyset$ iff $[i, j] = [1, n+1]$.
(5) For all $k \in [i, j] \cap [1, n]$, $x_k \in var(C) \cup var(M)$.

Proof. Let $I = \mathrm{UCI}(\Pi_L)$. Assume, to the contrary, that I is not an interval; for some $k \in [2, n]$, I contains an index $i < k$ and an index $j > k$, but does not contain k. Let L' be the first line in Π such that $\mathrm{UCI}(\Pi_{L'})$ intersects both $[1, k-1]$ and $[k+1, n+1]$. Since leaves have singleton UCI sets, L' is not a leaf. Say $L' = \text{res}(L'', L''', v)$. Assume that $\text{UCI}(\Pi_{L''}) \subseteq [1, k-1]$ and $\text{UCI}(\Pi_{L'''}) \subseteq [k+1, n+1]$; the argument for the other case is identical. So $v \in \text{var}_\exists(\text{UC}(\Pi_{L''})) \subseteq \text{var}_\exists(\phi_n^{[1,k-1]})$, and $v \in \text{var}_\exists(\text{UC}(\Pi_{L'''})) \subseteq \text{var}_\exists(\phi_n^{[k+1,n+1]})$. But $\text{var}_\exists(\phi_n^{[1,k-1]})$ and $\text{var}_\exists(\phi_n^{[k+1,n+1]})$ are disjoint, a contradiction.

Fixing i, j so that $I = \text{UCI}(\Pi_L) = [i, j]$, we now prove the remaining statements in the Lemma.

(1) Fix any $k \in [i, j-1]$. Note that $\{k, k+1\} \subseteq \text{UCI}(\Pi_L)$. Let L' be the first line in Π_L such that $\{k, k+1\} \subseteq \text{UCI}(\Pi_{L'})$. Say L' is obtained as $\text{res}(L'', L''', v)$. Assume that $\text{UCI}(\Pi_{L''})$ contributes k and $\text{UCI}(\Pi_{L''})$ contributes $k+1$; the other case is symmetric. Since $\text{UCI}(\Pi_{L''})$ must also be an interval, and since it contains k but not $k+1$, $\text{UCI}(\Pi_{L''}) \subseteq [1, k] \cap \text{UCI}(\Pi_L) = [i, k]$. Similarly, $\text{UCI}(\Pi_{L'''}) \subseteq [k+1, j]$. The pivot variable v must thus belong to both $\phi_n^{[i,k]}$ and $\phi_n^{[k+1,j]}$; the only such existential variable is t_k. Hence each t_k is used as a pivot in Π_L. Since Π is T-regular, and since t_k is used as a pivot to derive L' inside Π_L, it cannot reappear in any line on any path from (including) L' to the final clause. Hence it does not appear in L.

(2) Let $i > 1$. By Observation 3.5, t_{i-1} appears in at least one axiom used in Π_L. Assume to the contrary that $t_{i-1} \notin \text{var}(C)$. Let ρ_C be the minimal partial assignment falsifying C. By assumption, ρ_C does not set t_{i-1}, and by item (1) above, ρ_C does not set any variable t_k with $i \le k < j$. Extend ρ_C arbitrarily to all unassigned variables in $(X \cup T) \setminus \{t_{i-1}, \ldots, t_{j-1}\}$ to get ρ_1. Since the merge map M does not depend on variables in T, the partial assignment ρ_1 is sufficient to evaluate M and h. Define the value y as follows:

$$y = \begin{cases} \rho_1(t_j) & \text{if } j \le n, \\ h(\rho_1) & \text{if } j = n+1. \end{cases}$$

For $b \in \{0, 1\}$, let ρ_1^b denote the extension of ρ_1 by $t_{i-1} = b$. Exactly one of ρ_1^0, ρ_1^1 satisfies the equation $t_{i-1} + x_i + x_{i+1} + \ldots + x_j + y \equiv 0 \bmod 2$; let this extension be ρ_2. Then there is a unique extension α of ρ_2 to $X \cup T$ such that

- if $j \le n$, then α satisfies the existential part of all clauses in $\phi_n^{[i,j]}$;
- if $j = n+1$, then $(\alpha, h(\rho_1))$ satisfies all clauses in $\phi_n^{[i,j]}$. (That is, assigning $X \cup T$ according to α and assigning z the value $h(\rho_1)$ satisfies $\phi_n^{[i,j]}$.)

(To find α, work backwards from y to determine the appropriate values of $t_{j-1}, t_{j-2}, \ldots, t_i$ to satisfy ϕ_n^j, ϕ_n^{j-1}, $\ldots$, ϕ_n^i.)

Note that $h(\rho_1) = h(\rho_2) = h(\alpha)$. So $(\alpha, h(\alpha))$ falsifies C (since it extends ρ_C) and satisfies all axiom clauses used to derive L. This contradicts Lemma 2.8.

(3) Let $j \le n$. Assume to the contrary that $t_j \notin \text{var}(C)$. The argument is identical to that in item 2 (only the indices differ): ρ_C falsifies C; ρ_1 extends it arbitrarily to all unassigned variables in $(X \cup T) \setminus \{t_i, \ldots, t_j\}$; ρ_2 is the extension of ρ_1 obtained by setting t_j so as to satisfy the equation $t_{i-1} + x_i + x_{i+1} + \ldots + x_j + t_j \equiv 0 \bmod 2$; (Here, if $i = 1$, discard t_0 from the equation; i.e. assume $t_0 = 0$); α is the unique extension of ρ_2 to $X \cup T$ satisfying $\phi_n^{[i,j]}$ (To obtain α, work forwards obtaining $t_i, t_{i+1}, \ldots, t_{j-1}$). Now $(\alpha, h(\alpha))$ contradicts Lemma 2.8.

(4) Since $\text{UCI}(\Pi_L) = [i, j]$, variables t_k for $k \notin [i-1, j]$ do not appear in any of the used axioms (Observation 3.5) and hence do not appear in C. By the preceding three items, $\text{var}(C) \cap T$ does not include any t_k with $k \in [i, j-1]$, includes t_{i-1} whenever $i > 1$, and includes t_j whenever $j < n+1$. The claim follows.

(5) Assume to the contrary that for some $k \in [i, j]$, $x_k \notin \text{var}(C) \cup \text{var}(M)$. The argument is similar to that in item (2): ρ_C falsifies C; ρ_1 extends it arbitrarily to all unassigned variables in $(X \setminus \{x_k\}) \cup (T \setminus \{t_i, \ldots, t_{j-1}\})$; y is the value of t_j if $j \le n$ and the value of h otherwise (since $x_k \notin \text{var}(M)$, ρ_1 is sufficient to evaluate h); ρ_2 is the

extension of ρ_1 obtained by setting x_k so as to satisfy the equation $t_{i-1} + x_i + x_{i+1} + \ldots + x_j + y \equiv 0 \bmod 2$; (Here, if $i = 1$, discard t_0 from the equation; i.e. assume $t_0 = 0$); α is the unique extension of ρ_2 to $X \cup T$ satisfying $\phi_n^{[i,j]}$ (To obtain α, work forwards from t_i towards t_{j-1}). Now $(\alpha, h(\alpha))$ contradicts Lemma 2.8. □

Lemma 3.8. *Let $L \in \partial\mathcal{S}$ be derived in Π as $L = res(L', L'', t_k)$. Then* $\mathrm{UCI}(\Pi_L) = [1, n+1]$*, and* $\mathrm{UCI}(\Pi_{L'}), \mathrm{UCI}(\Pi_{L''})$ *partition* $[1, n+1]$ *into* $[1, k], [k+1, n+1]$.

Proof. Since $L \in \partial\mathcal{S}$, L has no variable from T. By Lemma 3.7(4), $\mathrm{UCI}(\Pi_L) = [1, n+1]$.

Since $L = \mathrm{res}(L', L'', t_k)$, we see that $\mathrm{var}(C') \cap T = \mathrm{var}(C'') \cap T = \{t_k\}$. By Lemma 3.7(2,3,4), we see that $\mathrm{UCI}(\Pi_{L'}), \mathrm{UCI}(\Pi_{L''}) \in \{[1, k], [k+1, n+1]\}$. If both $\mathrm{UCI}(\Pi_{L'}), \mathrm{UCI}(\Pi_{L''})$ equal $[k+1, n+1]$, then $\mathrm{UCI}(\Pi_L) = [k+1, n+1]$, contradicting $\mathrm{UCI}(\Pi_l) = [1, n+1]$. On the other hand, if both $\mathrm{UCI}(\Pi_{L'}), \mathrm{UCI}(\Pi_{L''})$ equal $[1, k]$, then $\mathrm{UCI}(\Pi_L) = [1, k]$. Since t_k is a pivot variable, $k \le n$, contradicting $\mathrm{UCI}(\Pi_L) = [1, n+1]$. Hence one each of $\mathrm{UCI}(\Pi_{L'}), \mathrm{UCI}(\Pi_{L''})$ equals $[1, k]$ and $[k+1, n+1]$ as claimed. □

Lemma 3.9. *For all $L \in \partial\mathcal{S}$, $width(C) = n$.*

Proof. Let $L \in \partial\mathcal{S}$ be derived in Π as $L = \mathrm{res}(L', L'', t_k)$. Since all axioms create non-trivial strategies, neither M' nor M'' equals $*$. By the rules of MRes, $M' = M'' = M \neq *$. We will show that in fact M must be a constant strategy, $M \in \{0, 1\}$.

By definition of $\partial\mathcal{S}$, $\mathrm{var}(C) \cap T = \emptyset$, and hence $\mathrm{var}(C') \cap T = \mathrm{var}(C'') \cap T = \{t_k\}$. By Lemma 3.8, $\mathrm{UCI}(\Pi_L) = [1, n+1]$ is partitioned by $\mathrm{UCI}(\Pi_{L'})$ and $\mathrm{UCI}(\Pi_{L''})$ into $[1, k], [k+1, n+1]$.

Assume $\mathrm{UCI}(\Pi_{L'}) = [1, k]$, $\mathrm{UCI}(\Pi_{L''}) = [k+1, n+1]$; the argument in the other case is identical. Then $\mathrm{var}(M) = \mathrm{var}(M') \subseteq \mathrm{var}(\phi^{[1,k]}) \cap X = \{x_1, \ldots, x_k\}$, and $\mathrm{var}(M) = \mathrm{var}(M'') \subseteq \mathrm{var}(\phi^{[k+1,n+1]}) \cap X = \{x_{k+1}, \ldots, x_n\}$. The only way both these conditions can be satisfied is if $\mathrm{var}(M) = \emptyset$; that is, M is a constant strategy.

Since $\mathrm{UCI}(\Pi_L) = [1, n+1]$ and $\mathrm{var}(M) = \emptyset$, Lemma 3.7(5) implies that $X \subseteq \mathrm{var}(C)$. Therefore $\mathrm{width}(C) = n$. □

Now we can put together the proof of Theorem 3.6.

Proof of Theorem 3.6. Let Π be a T-regular refutation of $\mathrm{LQParity}_n$ in MRes. Let $\mathcal{S}$ and $\partial\mathcal{S}$ be as defined in the beginning of this sub-section. By definition, for each $L = (C, M) \in \mathcal{S}$, $\mathrm{var}(C) \subseteq X$. Let $\widehat{\Pi} = \{C \mid L = (C, M) \in \mathcal{S}\}$. Then $\widehat{\Pi}$ contains a propositional resolution refutation of $\mathcal{C} = \{C \mid L = (C, M) \in \partial\mathcal{S}\}$. Therefore $\mathcal{C}$ is an unsatisfiable CNF formula over the n variables in X. By Lemma 3.9, each clause in $\mathcal{C}$ has width n and so is falsified by exactly one assignment. Therefore, to ensure that each of the 2^n assignments falsifies some clause, (at least) 2^n clauses are required. Therefore $|\mathcal{C}| \geqslant 2^n$. Hence $|\Pi| \geqslant 2^n$. □

3.4 The Completion Principle formulas

We now move to the Completion Principle (CR_n) formulas introduced in [29]. From a proof-complexity viewpoint, these formulas are very simple: they have polynomial-size, in fact linear-size, refutations in QRes, and hence in QURes, CP + ∀Red, ∀Exp + Res and IR [28, 29]. The QRes refutations are even tree-like; [34]. They are known to be hard for QRes if the resolution pivots must respect the quantifier ordering (level-ordered QRes); [28, 29].

In this section, we prove that CR_n requires exponential-size proofs in tree-like and regular MRes. Recall that no simulation is known between tree-like and regular MRes, so these statements require separate proofs. We believe that CR_n requires exponential size refutations in general MRes as well, but we have not been able to prove this.

We first define the formulas:

Definition 3.10. The Completion Principle formulas CR_n [29] are defined as follows:

$$\mathrm{CR}_n = \underset{i,j\in[n]}{\exists} x_{ij}, \forall z, \underset{i\in[n]}{\exists} a_i, \underset{j\in[n]}{\exists} b_j. \left(\bigwedge_{i,j\in[n]} (A_{ij} \wedge B_{ij}) \right) \wedge L_A \wedge L_B$$

where $A_{ij} = x_{ij} \vee z \vee a_i$, $B_{ij} = \overline{x_{ij}} \vee \overline{z} \vee b_j$, $L_A = \overline{a_1} \vee \cdots \vee \overline{a_n}$, and $L_B = \overline{b_1} \vee \cdots \vee \overline{b_n}$.

Let X, A, B denote the variable sets $\{x_{ij} : i, j \in [n]\}$, $\{a_i : i \in [n]\}$, and $\{b_j : j \in [n]\}$. It is convenient to think of the X variables as arranged in an $n \times n$ matrix.

Intuitively, the formulas describe a completion game, played on the matrix

$$\begin{pmatrix} a_1 & \dots & a_1 & \dots & a_n & \dots & a_n \\ b_1 & \dots & b_n & \dots & b_1 & \dots & b_n \end{pmatrix}$$

where the $\exists$-player first deletes exactly one cell per column and the $\forall$-player then chooses one row. The $\forall$-player wins if his row contains all of A or all of B (cf. [29]). In the formalisation, an assignment to the variable x_{ij} indicates the choices of the first player in the column containing a_i and b_j.

3.4.1 Lower bound for tree-like MRes. For the QBF CR_n, the winning strategy for the universal player (countermodel) is not unique. However, we show that all countermodels require large decision trees.

LEMMA 3.11. *Every countermodel for CR_n has decision tree size complexity at least 2^n.*

PROOF. We prove the size bound by showing that in every decision tree for every countermodel, all root-to-leaf paths query at least n variables, and hence the decision tree has at least 2^n nodes.

Assume to the contrary that some countermodel h is computed by a decision tree M that has a root-to-leaf path p querying less than n variables. Then there exist $k, \ell \in [n]$ such that no variable from Row k and no variable from Column ℓ is on this path. Let ρ_p be the minimal partial assignment that takes this path in M, and let ρ' be an arbitrary extension of ρ_p to variables in $\{x_{ij} \mid i \neq k, j \neq \ell\}$. Consider the following extension of ρ' to variables in $(X \setminus \{x_{k\ell}\}) \cup T$, giving assignment σ:
Set all variables in row k (other than $x_{k,\ell}$) to 1.
Set all variables in column ℓ (other than $x_{k,\ell}$) to 0.
Set a_k and b_ℓ to 0 and all other a_i, b_j variables to 1.

For $n \geq 2$, σ satisfies all the clauses of CR_n except $A_{k\ell}$ and $B_{k\ell}$, which get restricted to $x_{k\ell} \vee z$ and $\overline{x_{k\ell}} \vee \overline{z}$ respectively.

Let $\alpha_0 = \sigma \cup \{x_{k\ell} = 0\}$ and $\alpha_1 = \sigma \cup \{x_{k\ell} = 1\}$. Since both α_0 and α_1 extend ρ_p, they follow path p, therefore $h(\alpha_0) = h(\alpha_1)$. If $h(\alpha_0) = h(\alpha_1) = 0$, then $(\alpha_1, h(\alpha_1))$ satisfies all clauses of CR_n. On the other hand, if $h(\alpha_0) = h(\alpha_1) = 1$, then $(\alpha_0, h(\alpha_0))$ satisfies all clauses of CR_n. Thus in either case, h is not a countermodel for CR_n. □

From Theorem 3.1 and Lemma 3.11, we obtain the desired lower bound.

THEOREM 3.12. *Every tree-like MRes refutation of CR_n formulas has size at least 2^n.*

3.4.2 Lower bound for regular MRes. We now show that these formulas require exponential-size refutations in regular MRes. The idea of using Theorem 3.1 does not work here, since there is a winning strategy for the universal player with a small read-once branching program. (The strategy is to choose $z = 0$ exactly if X has an all-0s column.) We therefore directly analyse the combinatorial structure of a refutation to show that it must be large.

THEOREM 3.13. *Every $(A \cup B)$-regular refutation of CR_n in MRes, and hence every regular MRes refutation, has size at least 2^{n-1}.*

The high level idea is the same as the LQParity lower bound in Section 3.3. Let Π be a $(A\cup B)$-regular MRes refutation of CR_n. Since every axiom has a variable from $A\cup B$ while the final clause in Π is empty, there is a maximal "component" (say $\mathcal{S}$) of the proof leading to and including the final line, where all clauses are $(A \cup B)$-free. The clauses in this component involve only the X variables. We show that the "boundary" $\partial\mathcal{S}$ of this component is large, by showing in Lemma 3.14 that each clause here must be wide. (This idea was used in [38] to show that CR is hard for reductionless LD-QRes.)

To establish the width bound, we first note that except for the axioms L_A, L_B, no lines have trivial strategies. Since the pivots at the boundary are variables from $A \cup B$, which are all to the right of z, the merge maps incoming into each boundary resolution must be isomorphic. By analysing what axiom clauses cannot be used to derive lines just above the boundary, we show that many variables are absent in the corresponding merge maps, and invoking soundness of MRes, we show that they must then be present in the boundary clause, making it wide.

Proof of Theorem 3.13. Let Π be an $(A \cup B)$-regular refutation of CR_n (for $n \geq 2$) in MRes.

The lines of Π will be denoted by L, L_1, L_2, L', L'' etc. For lines L, L_i and L', and universal variable z; the respective clause, merge map and the function computed by the merge map will be denoted by $C, M^z, h^z, C_i, M_i^z, h_i^z$ and $C', (M')^z, (h')^z$ respectively. However, since CR_n has a single universal variable, we avoid the superscript z for the merge-maps and the corresponding functions. So, for these formulas, they will be denoted simply by M M_i, M', h, h_i and h'.

Define $\mathcal{S}$ to be the set of those lines in Π where the clause part has no variable from $A \cup B$, and furthermore there is a path in G_Π from the line to the final empty clause via lines where all the clauses also have no variables from $A \cup B$. Let $\partial\mathcal{S}$, called the boundary of $\mathcal{S}$, denote the set of leaves in the subgraph of G_Π restricted to $\mathcal{S}$; these are lines that are in $\mathcal{S}$ but their parents are not in $\mathcal{S}$. Note that no leaf of Π is in $\mathcal{S}$ because all leaves of G_Π contain a variable in $A \cup B$.

By definition, for each $L = (C, M) \in \mathcal{S}$, we have $\mathrm{var}(C) \subseteq X$. The sub-derivation $\widehat{\Pi} = \{C \mid L = (C, M) \in \mathcal{S}\}$ contains a propositional resolution refutation of the conjunction of clauses $\mathcal{C} = \{C \mid L = (C, M) \in \partial\mathcal{S}\}$. Hence $\mathcal{C}$ is an unsatisfiable CNF formula over the n^2 variables in X. We show below, in Lemma 3.14, that each clause in $\mathcal{C}$ has width at least $n - 1$. Hence it is falsified by at most $2^{n^2-(n-1)}$ assignments. Therefore, to ensure that each of the 2^{n^2} assignments falsifies some clause, at least 2^{n-1} clauses are required. Therefore $|\mathcal{C}| \geqslant 2^{n-1}$. Hence $|\Pi| = 2^{\Omega(n)}$. ◻

To conclude the proof of Theorem 3.13, it remains to prove the following lemma:

Lemma 3.14. *For all $L = (C, M) \in \partial\mathcal{S}$, $\mathit{width}(C) \geq n - 1$.*

Proof. Since $\mathrm{var}(C) \cap (A \cup B) = \emptyset$, we know that L is not a leaf of Π. Say $L = \mathrm{res}(L_1, L_2, v)$ where $L_1 = (C_1, M_1)$ and $L_2 = (C_2, M_2)$. Since $\mathrm{var}(C_1) \cap (A \cup B) \neq \emptyset$ and $\mathrm{var}(C_2) \cap (A \cup B) \neq \emptyset$, we have $v \in A \cup B$. Consider the case when $v \in A$; the argument for the case when $v \in B$ is symmetric. Without loss of generality, assume that $v = a_n$, and $a_n \in C_1$ and $\overline{a_n} \in C_2$.

Since Π is $(A \cup B)$-regular, a_n does not occur as a pivot in the sub-derivation Π_{L_1}. Therefore $L_A \notin \mathrm{leaves}(G_{\Pi_{L_1}})$ (otherwise $\overline{a_n} \in C_1$, and therefore C_1 would be tautological clause, a contradiction). This implies that the sub-derivation Π_{L_1} cannot use any axiom that contains a positive literal in A other than a_n, since such a literal would have to be eliminated by resolution before reaching C_1, requiring the corresponding negated literal, and L_A is the only axiom with negated literals from A. That is, Π_{L_1} does not use any of the axioms A_{ij} for $i \in [n-1]$. The positive literal x_{ij} appears only in A_{ij}. Hence for $i \in [n-1]$, $j \in [n]$, x_{ij} is not a pivot in Π_{L_1} and hence does not appear in M_1. On the other hand, M_1 is not trivial since some A_{nj} clause is used.

The clause C_2 contains $\overline{a_n}$, but no other $\overline{a_i}$. So C_2 is not the axiom L_A. Hence M_2 is not trivial.

Since the pivot a_n at the step obtaining line L is to the right of z, by the rules of MRes, M_1 and M_2 are isomorphic. Hence for each $i \in [n-1]$, and each $j \in [n]$, $x_{ij} \notin \text{var}(M_2)$. We claim the following:

Claim 3.15. *Either for all $i \in [n-1]$, C_2 has a variable of the form x_{i*}, or for all $j \in [n]$, C_2 has a variable of the form x_{*j}. In either case, C_2 has at least $n-1$ variables.*

We know that $\overline{a_n} \in C_2$, and for all $i \in [n-1]$, for all $j \in [n]$, $x_{ij} \notin \text{var}(M_2)$. Aiming for contradiction, suppose that there exist $i \in [n-1]$ and $j \in [n]$ such that for all $\ell \in [n]$, $x_{i\ell} \notin \text{var}(C_2)$, and for all $k \in [n]$, $x_{kj} \notin \text{var}(C_2)$. Fix such an i, j.

Let ρ be the minimum partial assignment falsifying C_2. Then

- ρ sets $a_n = 1$, leaves all other variables in $A \cup B$ unset.
- ρ does not set any $x_{i\ell}$ or x_{kj}.

For $c \in \{0, 1\}$, extend ρ to α_c as follows: Set $a_i = 0, b_j = 0$, set all other unset variables from $A \cup B$ to 1. Set $x_{ij} = c$. All $x_{i\ell}$ other than x_{ij} set to 1. All x_{kj} other than x_{ij} set to 0. Set remaining variables arbitrarily (but in the same way in α_0 and α_1).

The common part of α_0 and α_1 satisfies all axiom clauses except A_{ij} and B_{ij}, and does not falsify any axiom. The extensions α_c satisfy one more axiom, and still do not falsify the remaining axiom (it has a universal literal z or $\overline{z}$). They both falsify C_2, since they extend ρ.

Since α_0 and α_1 agree everywhere except on x_{ij}, and since $x_{ij} \notin \text{var}(M_2)$, it follows that $M_2(\alpha_0) = M_2(\alpha_1) = d$, say.

By Lemma 2.8, both (α_0, d) and (α_1, d) should falsify some axiom. However, $(\alpha_{\bar{d}}, d)$ actually satisfies all axioms, a contradiction. This completes the proof of the claim, and hence of the lemma as well. ◻

3.5 The KBKF-lq formulas

In this section, we turn towards the KBKF-lq formulas, defined in [2]. These formulas are variants of the KBKF formulas defined in [31]. The KBKF formulas and their variants are significant in QBF proof complexity. The KBKF formulas were used to prove the first lower bound in QBF proof complexity, for the proof system QRes [31]. They have short refutations in LD-QRes [26] (that paper uses the name φ_t) and in QURes [44]. Variants of this formula were then used to show non-simulations beteen QURes, LD-QRes, IR, and others. The specific variant of interest to us here is KBKF-lq. This variant was constructed in [2] to obtain formulas hard for LD-QRes. It is known that the KBKF-lq formulas are hard for LD-QRes [2] and for IRM [13] but have polynomial-size refutations in QURes [2].

Here we show that the KBKF-lq formulas are hard for the full system of Merge Resolution, thus making it our strongest lower bound in the paper. This constitutes the first genuine-to-QBF lower bound for unrestricted MRes in the literature.

Definition 3.16. The KBKF-lq$_n$ formulas [2] consist of the quantifier prefix

$$\exists d_1, e_1, \forall x_1, \exists d_2, e_2, \forall x_2, \ldots, \exists d_n, e_n, \forall x_n, \exists f_1, f_2, \ldots, f_n$$

and the clauses

$$
\begin{array}{llll}
A_0 = \{\overline{d_1}, \overline{e_1}, \overline{f_1}, \ldots, \overline{f_n}\} & & & \\
A_i^d = \{d_i, x_i, \overline{d_{i+1}}, \overline{e_{i+1}}, \overline{f_1}, \ldots, \overline{f_n}\} & A_i^e = \{e_i, \overline{x_i}, \overline{d_{i+1}}, \overline{e_{i+1}}, \overline{f_1}, \ldots, \overline{f_n}\} & & \forall i \in [n-1] \\
A_n^d = \{d_n, x_n, \overline{f_1}, \ldots, \overline{f_n}\} & A_n^e = \{e_n, \overline{x_n}, \overline{f_1}, \ldots, \overline{f_n}\} & & \\
B_i^0 = \{x_i, f_i, \overline{f_{i+1}}, \ldots \overline{f_n}\} & B_i^1 = \{\overline{x_i}, f_i, \overline{f_{i+1}}, \ldots \overline{f_n}\} & & \forall i \in [n-1] \\
B_n^0 = \{x_n, f_n\} & B_n^1 = \{\overline{x_n}, f_n\} & &
\end{array}
$$

Note that the existential part of each clause in KBKF-lq$_n$ is a Horn clause (at most one positive literal), and except A_0, is even strict Horn (exactly one positive literal).

We use the following shorthand notation. Sets of variables: $D = \{d_1, \ldots, d_n\}$, $E = \{e_1, \ldots, e_n\}$, $F = \{f_1, \ldots, f_n\}$, and $X = \{x_1, \ldots, x_n\}$. Sets of literals: For $Y \in \{D, E, X, F\}$, set $Y^1 = \{u \mid u \in Y\}$ and $Y^0 = \{\overline{u} \mid u \in Y\}$. Sets of clauses:

$$
\begin{array}{lllllll}
\mathcal{A}_0 & = & \{A_0\} & & & & \\
\mathcal{A}_i & = & \{A_i^d, A_i^e\} \quad \forall i \in [n] & & \mathcal{B}_i & = & \{B_i^0, B_i^1\} \quad \forall i \in [n] \\
\mathcal{A}_{[i,j]} & = & \cup_{k \in [i,j]} \mathcal{A}_k \quad \forall i, j \in [0, n], i \le j & & \mathcal{B}_{[i,j]} & = & \cup_{k \in [i,j]} \mathcal{B}_k \quad \forall i, j \in [n], i \le j \\
\mathcal{A} & = & \mathcal{A}_{[0,n]} & & \mathcal{B} & = & \mathcal{B}_{[1,n]}
\end{array}
$$

Theorem 3.17. *Every MRes refutation of KBKF-lq$_n$ has size at least 2^n.*

This proof follows the same high-level idea as the proofs of Theorems 3.6 and 3.13. Namely, in any refutation, a maximal component is identified (in this case, the F-free component, where no clause has a variable from F) and its boundary is shown to be large. However, the idea for showing that the boundary is large is completely different. The proofs of Theorems 3.6 and 3.13 established that the boundary clauses must be wide. Here, the complexity measure for the boundary is not width but the nature of the merge maps, or equivalently, of the partial strategies. We identify a property called *self-dependence* which captures the right complexity; a merge map for x_i has this self-dependence property if it depends on at least one of d_i, e_i. We show that all merge-maps at the final line must have self-dependence, whereas at the boundary lines none of the merge maps have self-dependence. We use this to then conclude that there must be exponentially many lines.

To show that self-dependence is not possible outside the F-free component, we show that from a line with F-variables and at least one self-dependent strategy, the F-variables can never be removed.

Elaborating on the roadmap of the argument: Let Π be an MRes refutation of KBKF-lq$_n$. Each line in Π has the form $L = (C, M^{x_1}, \ldots, M^{x_n})$ where C is a clause over D, E, F, and each M^{x_i} is a merge map computing a strategy for x_i.

Define $\mathcal{S}$ to be the set of those lines in Π where the clause part has no F variable and furthermore the line has a path in G_Π to the final empty clause via lines where all the clauses also have no F variables. Let $\partial\mathcal{S}$, called the boundary of $\mathcal{S}$, denote the set of leaves in the subgraph of G_Π restricted to $\mathcal{S}$; these are lines that are in $\mathcal{S}$ but their parents are not in $\mathcal{S}$. Note that by definition, for each $L = (C, \{M^{x_i} \mid i \in [n]\}) \in \mathcal{S}$, $\mathrm{var}(C) \subseteq D \cup E$. No line in $\mathcal{S}$ (and in particular, no line in $\partial\mathcal{S}$) is an axiom since all axiom clauses have variables from F.

Recall that the variables of KBKF-lq$_n$ can be naturally grouped based on the quantifier prefix: for $i \in [n]$, the ith group has d_i, e_i, x_i, and the $(n+1)$th group has the F variables. By construction, the merge map for x_i does not depend on variables in later groups, as is indeed required for a countermodel. We say that a merge map for x_i has *self-dependence* if it does depend on d_i and/or e_i.

We show that every merge map at every line in $\mathcal{S}$ is non-trivial (Lemma 3.22). Further, we show that at every line on the boundary of $\mathcal{S}$, i.e. in $\partial\mathcal{S}$, no merge map has self-dependence (Lemma 3.23). Using this, we conclude that $\partial\mathcal{S}$ must be exponentially large, since in every countermodel the strategy of each variable must have self-dependence (Proposition 3.24).

In order to show that lines in $\partial\mathcal{S}$ do not have self-dependence, we first establish several properties of the sets of axiom clauses used in a sub-derivation (Lemmas 3.18 to 3.21).

For a line $L \in \Pi$, let Π_L be the minimal sub-derivation of L, and let G_{Π_L} be the corresponding subgraph of G_Π with sink L. Let $\mathrm{UCI}(\Pi_L) = \{i \in [0, n] \mid \mathrm{leaves}(G_{\Pi_L}) \cap \mathcal{A}_i \neq \emptyset\}$. (UCI stands for UsedConstraintsIndex). Note that we are only looking at the clauses in $\mathcal{A}$ to define UCI.

Lemma 3.18. *For every line $L = (C, \{M^{x_i} \mid i \in [n]\})$ of Π, $|C \cap F^1| \leq 1$. Furthermore, $\mathrm{UCI}(\Pi_L) = \emptyset \Leftrightarrow C \cap F^1 \neq \emptyset$. (Recall that $F^1 = \{f_1, f_2, \dots, f_n\}$, i.e. the set of positive literals over the variable set $\{f_1, f_2, \dots, f_n\}$.)*

Proof. Since the existential part of each clause in KBKF-lq$_n$ is a Horn clause, and since the resolvent of Horn clauses is also Horn, $|C \cap F^1| \leq 1$ for each line of Π. It thus suffices to prove that $\forall L \in \Pi,\ \mathrm{UCI}(\Pi_L) = \emptyset \iff C \cap F^1 \neq \emptyset$.

($\Rightarrow$): For an arbitrary line $L \in \Pi$, suppose $\mathrm{UCI}(\Pi_L) = \emptyset$, so L is derived from $\mathcal{B}$. Since $\mathrm{var}_\exists(\mathcal{B}) = F$, $\mathrm{var}(C) \subseteq F$. The existential part of these clauses is strict Horn, and the resolvent of strict Horn clauses is also strict Horn, so C is strict Horn. So $C \cap F^1 \neq \emptyset$.

($\Leftarrow$): The statement $C \cap F^1 \neq \emptyset \Rightarrow \mathrm{UCI}(\Pi_L) = \emptyset$ holds at all axioms. Assume to the contrary that it does not hold everywhere in Π. Pick a highest L (closest to the axioms) for which this statement fails. That is, $C \cap F^1 \neq \emptyset$, and $\mathrm{UCI}(\Pi_L) \neq \emptyset$. Let L', L'' be the parents of L in Π; by choice of L, both L' and L'' satisfy the statement. Let f_j be the positive literal in C (unique, because C is Horn). Without loss of generality, $f_j \in C'$. Since L' satisfies the statement, $\mathrm{UCI}(\Pi_{L'}) = \emptyset$. So $\mathrm{var}(C') \subseteq F$, and since C' is Horn, $C' \setminus \{f_j\} \subseteq F^0$. Since $f_j \in C$, the pivot at this step is not f_j, so it must be an f_k for some $\overline{f_k} \in C'$. So $f_k \in C''$. Since L'' satisfies the statement, $\mathrm{UCI}(\Pi_{L''}) = \emptyset$. But then $\mathrm{UCI}(\Pi_L) = \mathrm{UCI}(\Pi_{L'}) \cup \mathrm{UCI}(\Pi_{L''}) = \emptyset$, contradicting our choice of L. Hence our assumption was wrong, and the statement holds for all L in Π. □

Lemma 3.19. *A line $L = (C, \{M^{x_i} \mid i \in [n]\})$ of Π with $\mathrm{UCI}(\Pi_L) = \emptyset$ has these properties:*

(1) $\mathrm{var}(C) \subseteq F$; for all $i \in [n]$, $M^{x_i} \in \{*, 0, 1\}$;
(2) For some $j \in [n]$, $f_j \in C$ and $M^{x_j} \in \{0, 1\}$; such a j is unique;
(3) For the unique j from (2), for $1 \leq i < j$, $f_i \notin \mathrm{var}(C)$ and $M^{x_i} = *$;
(4) For $j < i \leq n$, if $f_i \notin \mathrm{var}(C)$, then $M^{x_j} \in \{0, 1\}$.

Proof. (1) Since $\mathrm{UCI}(\Pi_L) = \emptyset$, $\mathrm{var}(C) \subseteq \mathrm{var}_\exists(\mathcal{B}) = F$.
All pivots in Π_L are from F, and all universal variables are left of F in the quantifier prefix. So no step in Π_L can use the merge operation to update merge maps; all steps in Π_L use only the select operation, which does not create any branching.

(2) By Lemma 3.18, $|C \cap F^1| = 1$, so there is a unique j with the literal $f_j \in C$. This literal appears only in the clauses of $\mathcal{B}_j$, both of which create a non-trivial strategy for x_j. So $M^{x_j} \neq *$. By item (1) proven above, $M^{x_j} \in \{0, 1\}$.

(3) Let k be the least index such that Π_L uses an axiom from $\mathcal{B}_k$. Since the positive literal f_j is in C and appears only in $\mathcal{B}_j$, $k \leq j$. Assume $k < j$. The axiom from $\mathcal{B}_k$ introduces the positive literal f_k into Π_L, and by choice of k, no axiom in Π_L has the literal $\overline{f_k}$. Hence f_k cannot be removed by resolution, and so $f_k \in C$, contradicting the fact that C is Horn. So in fact $k = j$. This means that no axiom introduces the variables f_i, $i < j$, into Π_L, so

$f_i \notin \mathrm{var}(C)$. Furthermore, amongst all the axioms in $\mathcal{B}$, only the axioms in $\mathcal{B}_i$ have a non-trivial merge map for x_i. Hence for $i < j$, no non-trivial merge map for x_i is created.

(4) Since $f_j \in C$, Π_L uses an axiom from $\mathcal{B}_j$. This axiom introduces the literals $\overline{f_i}$, for $j < i \le n$, into Π_L. If $\overline{f_i}$ is removed (by resolution) in Π_L, then an axiom from $\mathcal{B}_i$ must be used to introduce the positive literal f_i. This axiom created a non-trivial merge map for x_i, so the merge map for x_i at L is also non-trivial. □

Lemma 3.20. *Let $L = (C, \{M^{x_i} \mid i \in [n]\})$ be a line of Π with $\mathrm{UCI}(\Pi_L) \neq \emptyset$. Then $\mathrm{UCI}(\Pi_L)$ is an interval $[a, b]$ for some $0 \le a \le b \le n$. Furthermore, (in the items below, a, b refer to the endpoints of this interval), it has the following properties:*

(1) For $k \in [n] \cap [a, b]$, $M^{x_k} \neq *$.
(2) If $a \ge 1$, then $|\{d_a, e_a\} \cap C| = 1$. If $a = 0$, then C does not have any positive literal.
(3) If $b < n$, then $\overline{d_{b+1}}, \overline{e_{b+1}} \in C$.
(4) For all $k \in [n] \setminus [a, b]$, (i) $d_k, e_k \notin var(M^{x_k})$, and (ii) if $M^{x_k} = *$ then $\overline{f_k} \in C$.

Proof. Assume to the contrary that $\mathrm{UCI}(\Pi_L)$ is not an interval. Then there exist $0 \le a < c < b \le n$ such that $a, b \in \mathrm{UCI}(\Pi_L)$ but $c \notin \mathrm{UCI}(\Pi_L)$. Let L_1 be the first line in Π_L such that $\mathrm{UCI}(\Pi_{L_1})$ intersects both $[0, c-1]$ and $[c+1, n]$ (note that L_1 exists). Since leaves have singleton UCI sets, L_1 is not a leaf. Say $L_1 = \mathrm{res}(L_2, L_3, v)$. By our choice of L_1, exactly one each of $\mathrm{UCI}(\Pi_{L_2})$ and $\mathrm{UCI}(\Pi_{L_3})$ is a non-empty subset of $[0, c-1]$ and of $[c+1, n]$. So $v \in \mathrm{var}_\exists(\mathcal{A}_{[0,c-1]})$ and $v \in \mathrm{var}_\exists(\mathcal{A}_{[c+1,n]})$. But $\mathrm{var}_\exists(\mathcal{A}_{[0,c-1]}) \cap \mathrm{var}_\exists(\mathcal{A}_{[c+1,n]}) = F$, and by Lemma 3.18, both C_2 and C_3 contain variables of F only in negated form. So no variable from F can be a resolution pivot, a contradiction. It follows that $\mathrm{UCI}(\Pi_L)$ is an interval.

(1) For $k \in [n] \cap [a, b]$, some axiom from $\mathcal{A}_k$ has been used to derive L. Both these axioms create non-trivial strategies for x_k. Subsequent MRes steps cannot make a non-trivial strategy trivial.

(2) Consider first the case $a \ge 1$. Since C is a Horn clause, C can contain at most one of the literals d_a, e_a. Since $a \in \mathrm{UCI}(\Pi_L)$, at least one of A^d_a, A^e_a appears in leaves(Π_L), so at least one of the literals d_a, e_a is introduced into Π_L. Since A^d_{a-1} and A^e_{a-1} are the only axioms that contain $\overline{d_a}$ or $\overline{e_a}$, and since neither of these is used in Π_L, therefore the positive literals d_a, e_a, if introduced, cannot be removed through resolution. Hence at least one of them is in C. It follows that C has exactly one of d_a, e_a.
If $a = 0$, Π_L uses the clause A_0 which has only negative literals. The resolvent of such a clause and a Horn clause also has only negative literals. Following the sequence of resolutions on the path from a leaf using A_0 to C shows that C has only negative literals.

(3) Since $b < n$ and $b \in \mathrm{UCI}(\Pi_L)$, some clause from $\mathcal{A}_b$ is used in Π_L and introduces the literals $\overline{d_{b+1}}, \overline{e_{b+1}}$ into Π_L. Since $b + 1 \notin \mathrm{UCI}(\Pi_L)$, no leaf of Π_L contains the positive literals d_{b+1}, e_{b+1}. So $\overline{d_{b+1}}$ and $\overline{e_{b+1}}$ cannot be removed through resolution.

(4) For $k > b$, no leaf in Π_L contains the positive literals d_k, e_k. For $k < a$, no leaf in Π_L contains the negative literals $\overline{d_k}, \overline{e_k}$. Thus, for $k \notin [a, b]$, the variables d_k, e_k are not used as resolution pivots anywhere in Π_L, and hence are not queried in any of the merge maps.
Each negative literal $\overline{f_k}$ is present in every clause of $\mathcal{A}$, and hence is introduced into Π_L. If $M^{x_k} = *$, then $B^0_k, B^1_k \notin$ leaves(Π_L) (both of them have non-trivial merge maps for x_k). Since these are the only clauses with the positive literal f_k, the literal $\overline{f_k}$ cannot be removed in Π_L; hence $\overline{f_k} \in C$. □

Lemma 3.21. *For any line $L = (C, \{M^{x_i} \mid i \in [n]\})$ in Π, and any $k \in [n]$, if $\{d_k, e_k\} \cap var(M^{x_k}) \neq \emptyset$, then* $\text{UCI}(\Pi_L) = [a, n]$ *for some $a \leq k-1$.*

Proof. Since $\{d_k, e_k\} \cap \text{var}(M^{x_k}) \neq \emptyset$, either d_k or e_k must be used as a pivot in Π_L, and hence must appear in both polarities in Π_L. The variables d_k, e_k appear positively only in $\mathcal{A}_k$, and negatively only in $\mathcal{A}_{k-1}$. Hence $a \leq k-1$.

Suppose $b < n$. By Lemma 3.20 (3), both $\overline{d_{b+1}}$ and $\overline{e_{b+1}}$ are in C. Consider any path ρ in Π from L to the final line $L_\Box$. At every line on this path, the merge map for x_k queries at least one of d_k, e_k since it is at least as complex as the merge map M^{x_k}. Along this path, both d_{b+1} and e_{b+1} must appear as pivots, since the negated literals are eventually removed. Pick the first such step on ρ, and assume without loss of generality that the pivot is d_{b+1} (the other case is symmetric). So $\overline{d_{b+1}}$ is present in the line, say L_1, on ρ, and d_{b+1} is present in the clause L_2 with which it is resolved to obtain $L_3 = \text{res}(L_2, L_1, d_{b+1})$ on ρ. By Lemma 3.20 (2), $\text{UCI}(\Pi_{L_2}) = [b+1, b']$ for some $b' \geq b+1$. Hence by Lemma 3.20 (4), $d_k, e_k \notin \text{var}(M_2^{x_k})$. However, $\{d_k, e_k\} \cap \text{var}(M_1^{x_k}) \neq \emptyset$. Since this resolution on d_{b+1} is not blocked, it must be the case that $M_2^{x_k} = *$. Hence, by Lemma 3.20 (4), $\overline{f_k} \in C_2$ and so $\overline{f_k} \in C_3$. To remove this literal, at some later point along ρ, f_k must appear as pivot. However, at that point, the line from ρ has a complex merge map for x_k, while the line with the positive literal f_k has a non-trivial constant merge map (by Lemma 3.19 (2)). Hence the resolution on f_k is blocked, a contradiction. It follows that $b = n$. □

Lemma 3.22. *For all $L \in \mathcal{S}$, for all $k \in [n]$, $M^{x_k} \neq *$.*

Proof. Consider a line $L = (C, \{M^{x_i} \mid i \in [n]\}) \in \mathcal{S}$. Since $L \in \mathcal{S}$, it has no variables from F. So $C \cap F^1 = \emptyset$. (Recall that F^1 is the set of positive literals with variables from F; that is, $F^1 = \{f_1, f_2, \dots, f_n\}$. Similarly, $F^0 = \{\overline{f_1}, \overline{f_2}, \dots, \overline{f_n}\}$ is the set of negative literals over F.) By Lemma 3.18, $\text{UCI}(\Pi_L) \neq \emptyset$. Since every clause in $\mathcal{A}$ contains all literals in F^0, for each $k \in [n]$, Π_L has a leaf where the clause contains $\overline{f_k}$. This literal is removed in deriving L, so Π_L also has a leaf where the clause contains the positive literal f_k. That is, it uses an axiom from $\mathcal{B}_k$; this leaf has a non-trivial merge map for x_k. Since a step in MRes cannot make a non-trivial merge map trivial, the merge map for x_k at L is non-trivial. □

Lemma 3.23. *For all $L \in \partial\mathcal{S}$, for all $k \in [n]$, $d_k, e_k \notin var(M^{x_k})$.*

Proof. Consider a line $L \in \partial\mathcal{S}$; $L = (C, \{M^{x_i} \mid i \in [n]\})$. Assume to the contrary that for some $k \in [n]$, $\{d_k, e_k\} \cap \text{var}(M^{x_k}) \neq \emptyset$.

The line L is obtained by performing resolution on two non-$\mathcal{S}$ clauses with a pivot from F. Let $L = \text{res}(L', L'', f_\ell)$ for some $\ell \in [n]$; $f_\ell \in C'$ and $\overline{f_\ell} \in C''$. Since L has no variable in F, f_ℓ is the only variable from F in $\text{var}(C')$ and $\text{var}(C'')$.

Since C' has the literal $f_\ell \in F^1$, by Lemma 3.18, $\text{UCI}(\Pi_{L'}) = \emptyset$ and L' is derived exclusively from $\mathcal{B}$. Since $D \cup E$ and $\text{var}(\mathcal{B})$ are disjoint, all the merge maps in L' have no variable from $D \cup E$. So M^{x_k} gets its $D \cup E$ variables from $(M'')^{x_k}$. Since this does not block the resolution step, $(M')^{x_k}$ must be trivial and $M^{x_k} = (M'')^{x_k}$. Since $\text{var}(C') \cap F = f_\ell$, by Lemma 3.19 (2),(3),(4), $k < \ell$.

The line L'' has no literal from F^1, so by Lemma 3.18, $\text{UCI}(\Pi_{L''}) \neq \emptyset$. It has a merge map for x_k involving at least one of d_k, e_k, so by Lemma 3.21, $\text{UCI}(\Pi_{L''}) = [a, n]$ for some $a \leq k-1$. Thus we have $a \leq k-1 < k < \ell \leq n$.

Consider the resolution of L' with L''. By Lemma 3.19 (2), $(M')^{x_\ell} \in \{0, 1\}$, and by Lemma 3.20 (1), $(M'')^{x_\ell} \neq *$. To enable this resolution, $(M'')^{x_\ell} = (M')^{x_\ell}$. The clauses A_ℓ^d and A_ℓ^e give rise to different constant strategies for x_ℓ. So the derivation of L'' uses exactly one of these two clauses. Assume it uses A_ℓ^d; the other case is symmetric. Since $a < \ell$, the derivation of L'' uses a clause from $A_{\ell-1}$, introducing literals $\overline{d_\ell}$ and $\overline{e_\ell}$. Since the only clause containing positive literal e_ℓ is not used, $\overline{e_\ell}$ survives in C''. Going from L'' to L removes only $\overline{f_\ell}$, so $\overline{e_\ell} \in C$.

To summarize, at this stage we know that $L \in \partial\mathcal{S}$, $\overline{e_\ell} \in C$, $\{d_k, e_k\} \cap \text{var}(M^{x_k}) \neq \emptyset$, $M^{x_\ell} \in \{0, 1\}$ and $1 \leq k < \ell \leq n$.

Fix any path ρ in G_Π from L to $L_\Box$. Along this path, e_ℓ appears as the pivot somewhere, since the literal $\overline{e_\ell}$ is eventually removed. Consider the resolution step at that point, say $C_1 = \mathrm{res}(C_2, C_3, e_\ell)$, with C_3 being the clause at the line on ρ. At the corresponding line L_3, the strategies are at least as complex as those at L. Hence $\mathrm{var}(M_3^{x_k}) \cap \{d_k, e_k\} \neq \emptyset$. On the other hand, C_2 has the positive literal e_ℓ. By Lemma 3.20, for the corresponding line L_2, $\mathrm{UCI}(\Pi_{L_2}) = [\ell, c]$ for some $c \geq \ell$. Since $k < \ell$, by Lemma 3.20, $\{d_k, e_k\} \cap \mathrm{var}(M_2^{x_k}) = \emptyset$. However, the path from L_2 to L_1 and thence to $L_\Box$ along ρ witnesses that $L_2 \in \mathcal{S}$, so by Lemma 3.22, $M_2^{x_k} \neq *$. Thus $M_2^{x_k}$ and $M_3^{x_k}$ are non-trivial but not isomorphic, and this blocks the resolution on e_ℓ.

Thus our assumption that $\{d_k, e_k\} \cap \mathrm{var}(M^{x_k}) \neq \emptyset$ must be false. The lemma is proved. □

We will also use the following property of KBKF-lq formulas. It implies that in every countermodel, the strategy for every variable has self-dependence. This is used, towards the end of the proof of Theorem 3.17, to show that merge-maps for countermodels must be complex and large.

Proposition 3.24. *Let h be any countermodel for KBKF-lq$_n$. Let α be any assignment to D, and β be any assignment to E. For each $i \in [n]$, if $\alpha_j \neq \beta_j$ for all $1 \leq j \leq i$, then $h^{x_i}\big((\alpha, \beta)\restriction_{L_Q(x_i)}\big) = \alpha_i$. In particular, if $\alpha_j \neq \beta_j$ for all $j \in [n]$, then the countermodel computes $h(\alpha, \beta) = \alpha$.*

Proof. Let h be any countermodel for KBKF-lq$_n$. For $i \in [n]$, let α^i be an assignment to $\{d_1, \ldots, d_i\}$, and β^i be an assignment to $\{e_1, \ldots, e_i\}$. For $j \leq i$, let α^i_j (resp. β^i_j) be the assignment to d_j (resp. e_j) set by the assignment α^i_j (resp. β^i_j). We will show that for each $i \in [n]$, if $\alpha^i_j \neq \beta^i_j$ for all $1 \leq j \leq i$, then $h^{x_i}(\alpha^i, \beta^i) = \alpha^i_i$. This implies the claimed result.

Fix some $i \in [n]$. Assume to the contrary that $\alpha^i_j \neq \beta^i_j$ for all $1 \leq j \leq i$ and $h^{x_i}(\alpha^i, \beta^i) \neq \alpha^i_i$. We will give a winning strategy for the existential player. Note that all clauses in $\mathcal{A}[0, i-1]$ are satisfied by the partial assignment (α^i, β^i). The existential player sets $d_j = e_j = 1$ for all $j > i$ and sets $f_j = 1$ for all $j \in [n]$. This satisfies all the remaining clauses, irrespective of the strategy of the universal player. Therefore the existential player wins. This contradicts the assumption that h is a countermodel for KBKF-lq$_n$. □

Now we have all the required information; we put it together to obtain the lower bound.

Proof of Theorem 3.17. Let Π be a refutation of KBKF-lq$_n$ in MRes. Let $\mathcal{S}, \partial\mathcal{S}$ be as defined in the beginning of this section. Let the final line of Π be $L_\Box = (\Box, \{M_\Box^{x_i} \mid i \in [n]\})$, and for $i \in [n]$, let h_i be the functions computed by the merge map $M_\Box^{x_i}$. By soundness of MRes, the functions $\{h_i\}_{i \in [n]}$ form a countermodel for KBKF-lq$_n$.

For each $a \in \{0,1\}^n$, consider the assignment α to the variables of $D \cup E$ where $d_i = a_i$, $e_i = \overline{a_i}$. Call such an assignment an anti-symmetric assignment. Given such an assignment, walk from $L_\Box$ towards the leaves of Π as far as is possible while maintaining the following invariant at each line $L = (C, \{M^{x_i} \mid i \in [n]\})$ along the way:

(1) α falsifies C, and
(2) for each $i \in [n]$, $h_i(\alpha) = M^{x_i}(\alpha)$.

Clearly this invariant is initially true at $L_\Box$, which is in $\mathcal{S}$. If we are currently at a line $L \in \mathcal{S}$ where the invariant is true, and if $L \notin \partial\mathcal{S}$, then L is obtained from lines L', L''. The resolution pivot in this step is not in F, since that would put L in $\partial\mathcal{S}$. So both L' and L'' are in $\mathcal{S}$, and the pivot is in $D \cup E$. Let the pivot be in $\{d_\ell, e_\ell\}$ for some $\ell \in [n]$. Depending on the pivot value, exactly one of C', C'' is falsified by α; say C' is falsified. By Lemma 3.22, for each $i \in [n]$, both $(M')^{x_i}$ and $(M'')^{x_i}$ are non-trivial. By definition of the MRes rule,

- For $i < \ell$, $(M')^{x_i}$ and $(M'')^{x_i}$ are isomorphic (otherwise the resolution is blocked), and $M^{x_i} = (M')^{x_i} = (M'')^{x_i}$.

- For $i \geq \ell$, there are two possibilities:
  (1) $(M')^{x_i}$ and $(M'')^{x_i}$ are isomorphic, and $M^{x_i} = (M')^{x_i}$.
  (2) M^{x_i} is a merge of $(M')^{x_i}$ and $(M'')^{x_i}$ with the pivot variable queried. By definition of the merge operation, since C' is falsified by α, $M^{x_i}(\alpha) = (M')^{x_i}(\alpha)$.

Thus in all cases, for each i, $h_i(\alpha) = M^{x_i}(\alpha) = (M')^{x_i}(\alpha)$. Hence L' satisfies the invariant.

We have shown that as long as we have not encountered a line in $\partial\mathcal{S}$, we can move further. We continue the walk until a line in $\partial\mathcal{S}$ is reached. We denote the line so reached by $P(\alpha)$. Thus P defines a map from anti-symmetric assignments to $\partial\mathcal{S}$.

We now show that the map P is one-to-one. Suppose, to the contrary, $P(\alpha) = P(\beta) = (C, \{M^{x_i} \mid i \in [n]\})$ for two distinct anti-symmetric assignments obtained from $a, b \in \{0,1\}^n$ respectively. Let j be the least index in $[n]$ where $a_j \neq b_j$. By Lemma 3.23, M^{x_j} depends only on $\{d_i, e_i \mid i < j\}$, and α, β agree on these variables. Thus we get the equalities $a_j = h_j(\alpha) = M^{x_j}(\alpha) = M^{x_j}(\beta) = h_j(\beta) = b_j$, where the first and last equalities follow from Proposition 3.24, the third equality from Lemma 3.23 and choice of j, and the second and fourth equalities by the invariant satisfied at $P(\alpha)$ and $P(\beta)$ respectively. This contradicts $a_j \neq b_j$.

We have established that the map P is one-to-one. Hence, $\partial\mathcal{S}$ has at least as many lines as anti-symmetric assignments, so $|\Pi| \geq |\partial\mathcal{S}| \geq 2^n$. □

4 RELATIONS AMONG PROOF SYSTEMS

In this section, we collect all the separations among proof systems which are implied by the lower bounds in Section 3.

Since any propositional formula is also a QBF formula and since MRes degenerates to Resolution on propositional formulas, it follows from propositional proof complexity that MRes strictly-simulates tree-like and regular MRes, and that tree-like MRes does not simulate regular or general MRes. Whether, regular MRes p-simulates tree-like MRes is unknown. Here we observe that the non-simulation of regular and general MRes by tree-like MRes is also witnessed by the QParity formulas (because the QParity formulas have polynomial-size refutations in regular and general MRes but require exponential-size refutations in tree-like MRes).

The following two theorems show that MRes and its restrictions are incomparable with some other resolution-based QBF proof systems. As observed in [7], one direction of the non-simulation follows from the Equality formulas: these formulas have polynomial-size refutations in tree-like, regular and general MRes but require exponential-size refutations in QRes, QURes, CP + ∀Red, ∀Exp + Res, and IR.

Theorem 4.1. *Tree-like and regular MRes are incomparable with the tree-like and general versions of QRes, QURes, CP + ∀Red, ∀Exp + Res, and IR.*

Proof. We showed in Theorem 3.12 that the Completion Principle CR_n requires exponential-size refutations in tree-like MRes. In Theorem 3.13, we showed that it requires exponential-size refutations in regular MRes. It has polynomial-size refutations in tree-like QRes [28] (and hence also in QURes and CP + ∀Red) and tree-like ∀Exp + Res [29] (and hence also in IR). (While [29] does not explicitly mention tree-like or regular refutations, the refutation provided there for CR_n is tree-like and regular.) Therefore, tree-like and regular MRes do not simulate the tree-like and general versions of QRes, QURes, CP + ∀Red, ∀Exp + Res, and IR.

The other direction of the non-simulation follows from the Equality formulas, as mentioned in Example 2.7. □

Theorem 4.2. *MRes is incomparable with QURes and CP + ∀Red.*

Proof. Theorem 3.17 shows that the KBKF-lq$_n$ formula requires exponential-size refutations in MRes. It has polynomial-size refutations in QURes [2], and also in CP + ∀Red (since CP + ∀Red simulates QURes [15]). Therefore MRes does not simulate QURes and CP + ∀Red. The other direction of the non-simulation follows from the Equality formulas, as mentioned in Example 2.7. □

5 CONCLUSIONS AND FUTURE WORK

The proof system MRes was introduced in [7], using the novel idea of building strategies directly into the proof and using them to enable additional sound applications of resolution. In [7], the strengths of the proof system were demonstrated. In this paper, we complement that study by exposing some limitations of MRes. We obtain hardness for tree-like MRes by transferring computational hardness of the countermodels in decision trees, and for regular and general MRes by ad hoc combinatorial arguments.

Several questions still remain.

(1) One of the driving goals behind the definition of MRes was overcoming a perceived weakness of LD-QRes: its criterion for blocking unsound applications of resolution also blocks several sound applications. However, whether MRes actually overcomes this weakness is not demonstrated, neither in [7] nor here. In [7], MRes is shown to be more powerful than the reductionless variant of LD-QRes (introduced in [19] and further investigated in [7, 38]). Very recently, in [35], this question has been resolved; another variant of KBKF has been shown to be easy in MRes but exponentially hard for LD-QRes and even for systems more powerful than LD-QRes. The other direction, whether there is a formula easy for LD-QRes but hard for MRes, is still open. One possible candidate for this separation might appear to be the original KBKF formula, which is easy for LD-QRes [26] (that paper uses the name φ_t). However the KBKF formulas can be shown to have short refutations in MRes as well, and hence cannot be used for this purpose. Perhaps the completion principle formulas CR$_n$ may demonstrate this separation.

(2) In the propositional case, regular resolution simulates tree-like resolution. This relation may not hold in the case of MRes, and even if it does, it will need a different proof. The trick used in the propositional case — (i) interpret the proof tree as a decision tree for search, (ii) make the decision tree read-once, (iii) then return from the search tree to a refutation, — does not work here because when we prune away parts of the decision tree to get a read-once tree, we may end up destroying isomorphism of strategies of blocking variables. Perhaps modifying the proof system itself to require not isomorphism but only semantic equivalence, as was done in [21], could lead to a simulation more easily, but that would be in the context of the modified proof system, not MRes itself.

6 ACKNOWLEDGEMENTS

Olaf Beyersdorff's research was supported by grants from the John Templeton Foundation (grant no. 60842), the DFG (BE 4209/3-1), and the Carl Zeiss Foundation. Tomáš Peitl's research was supported by Grant J-4361 of the Austrian Science Fund FWF. Olaf Beyersdorff and Meena Mahajan were supported by a DAAD/DST grant. Part of this work was done during the Dagstuhl Seminar 'SAT and Interactions' (Seminar 20061).

REFERENCES

[1] Valeriy Balabanov and Jie-Hong R. Jiang. 2012. Unified QBF Certification and Its Applications. *Form. Methods Syst. Des.* 41, 1 (Aug. 2012), 45–65.

[2] Valeriy Balabanov, Magdalena Widl, and Jie-Hong R. Jiang. 2014. QBF Resolution Systems and Their Proof Complexities. In *Theory and Applications of Satisfiability Testing – SAT 2014*, Carsten Sinz and Uwe Egly (Eds.). Springer International Publishing, Cham, 154–169.

[3] Olaf Beyersdorff. 2022. Proof Complexity of Quantified Boolean Logic – a Survey. In *Mathematics for Computation (M4C)*, Marco Benini, Olaf Beyersdorff, Michael Rathjen, and Peter Schuster (Eds.). World Scientific, Singapore, 353–391.

[4] Olaf Beyersdorff and Joshua Blinkhorn. 2017. Formulas with Large Weight: a New Technique for Genuine QBF Lower Bounds. *Electron. Colloquium Comput. Complex.* 24 (2017), 32.

[5] Olaf Beyersdorff and Joshua Blinkhorn. 2020. Lower Bound Techniques for QBF Expansion. *Theory of Computing Systems* 64, 3 (2020), 400–421. https://doi.org/10.1007/s00224-019-09940-0

[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* Volume 15, Issue 1 (Feb. 2019). https://doi.org/10.23638/LMCS-15(1:13)2019

[7] Olaf Beyersdorff, Joshua Blinkhorn, and Meena Mahajan. 2021. Building Strategies into QBF Proofs. *J. Autom. Reason.* 65, 1 (2021), 125–154. https://doi.org/10.1007/s10817-020-09560-1

[8] Olaf Beyersdorff, Joshua Blinkhorn, Meena Mahajan, and Tomás Peitl. 2023. Hardness Characterisations and Size-width Lower Bounds for QBF Resolution. *ACM Trans. Comput. Log.* 24, 2 (2023), 10:1–10:30. https://doi.org/10.1145/3565286

[9] Olaf Beyersdorff, Joshua Blinkhorn, Meena Mahajan, Tomáš Peitl, and Gaurav Sood. 2020. Hard QBFs for Merge Resolution. In *40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 182)*. 12:1–12:15. https://doi.org/10.4230/LIPIcs.FSTTCS.2020.12

[10] Olaf Beyersdorff and Benjamin Böhm. 2023. Understanding the Relative Strength of QBF CDCL Solvers and QBF Resolution. *Log. Methods Comput. Sci.* 19, 2 (2023).

[11] 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)*. 249–260.

[12] Olaf Beyersdorff, Ilario Bonacina, Leroy Chew, and Jan Pich. 2020. Frege Systems for Quantified Boolean Logic. *J. ACM* 67, 2, Article 9 (2020), 36 pages.

[13] Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. 2019. New Resolution-Based QBF Calculi and Their Proof Complexity. *ACM Trans. Comput. Theory* 11, 4, Article 26 (Sept. 2019), 42 pages. https://doi.org/10.1145/3352155

[14] 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.

[15] Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. 2018. Understanding Cutting Planes for QBFs. *Information and Computation* 262 (2018), 141–161.

[16] 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.

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

[18] Olaf Beyersdorff, Mikolás Janota, Florian Lonsing, and Martina Seidl. 2021. Quantified Boolean Formulas. In *Handbook of Satisfiability*, Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh (Eds.). IOS Press, 1177–1221.

[19] Nikolaj Bjørner, Mikolás Janota, and William Klieber. 2015. On Conflicts and Strategies in QBF. In *20th International Conferences on Logic for Programming, Artificial Intelligence and Reasoning LPAR 2015 (EPiC Series in Computing, Vol. 35)*, Ansgar Fehnker, Annabelle McIver, Geoff Sutcliffe, and Andrei Voronkov (Eds.). EasyChair, 28–41.

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

[21] Joshua Blinkhorn, Tomás Peitl, and Friedrich Slivovsky. 2021. Davis and Putnam Meet Henkin: Solving DQBF with Resolution. In *Theory and Applications of Satisfiability Testing - SAT 2021 - 24th International Conference, Barcelona, Spain, July 5-9, 2021, Proceedings (Lecture Notes in Computer Science, Vol. 12831)*, Chu-Min Li and Felip Manyà (Eds.). Springer, 30–46. https://doi.org/10.1007/978-3-030-80223-3_4

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

[23] Sravanthi Chede and Anil Shukla. 2023. Extending Merge Resolution to a Family of QBF-Proof Systems. In *40th International Symposium on Theoretical Aspects of Computer Science, STACS 2023, March 7-9, 2023, Hamburg, Germany (LIPIcs, Vol. 254)*, Petra Berenbrink, Patricia Bouyer, Anuj Dawar, and Mamadou Moustapha Kanté (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 21:1–21:20. https://doi.org/10.4230/LIPICS.STACS.2023.21

[24] Leroy Chew and Friedrich Slivovsky. 2022. Towards Uniform Certification in QBF. In *39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, March 15-18, 2022, Marseille, France (Virtual Conference) (LIPIcs, Vol. 219)*, Petra Berenbrink and Benjamin Monmege (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 22:1–22:23. https://doi.org/10.4230/LIPICS.STACS.2022.22

[25] Stephen A. Cook and Robert A. Reckhow. 1979. The relative efficiency of propositional proof systems. *The Journal of Symbolic Logic* 44, 1 (1979), 36–50.

[26] Uwe Egly, Florian Lonsing, and Magdalena Widl. 2013. Long-Distance Resolution: Proof Generation and Strategy Extraction in Search-Based QBF Solving. In *Logic for Programming, Artificial Intelligence, and Reasoning - 19th International Conference, LPAR-19*. 291–308.

[27] Marijn Heule, Martina Seidl, and Armin Biere. 2014. A Unified Proof System for QBF Preprocessing. In *IJCAR*. 91–106.

[28] Mikoláš Janota. 2016. On Q-Resolution and CDCL QBF Solving. In *Theory and Applications of Satisfiability Testing – SAT 2016*, Nadia Creignou and Daniel Le Berre (Eds.). Springer International Publishing, Cham, 402–418.

[29] Mikoláš Janota and Joao Marques-Silva. 2015. Expansion-based QBF solving versus Q-resolution. *Theoretical Computer Science* 577 (2015), 25 – 42. https://doi.org/10.1016/j.tcs.2015.01.048

[30] Manuel Kauers and Martina Seidl. 2018. Short proofs for some symmetric Quantified Boolean Formulas. *Inf. Process. Lett.* 140 (2018), 4–7.

[31] Hans Kleine Büning, Marek Karpinski, and Andreas Flögel. 1995. Resolution for Quantified Boolean Formulas. *Inf. Comput.* 117, 1 (1995), 12–18.

[32] Jan Krajíček. 1995. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*. Encyclopedia of Mathematics and Its Applications, Vol. 60. Cambridge University Press, Cambridge.

[33] Florian Lonsing, Uwe Egly, and Martina Seidl. 2016. Q-Resolution with Generalized Axioms. In *Theory and Applications of Satisfiability Testing - SAT 2016 - 19th International Conference, Proceedings (Lecture Notes in Computer Science, Vol. 9710)*, Nadia Creignou and Daniel Le Berre (Eds.). Springer, 435–452.

[34] Meena Mahajan and Anil Shukla. 2016. Level-ordered Q-resolution and tree-like Q-resolution are incomparable. *Inform. Process. Lett.* 116, 3 (2016), 256–258. https://doi.org/10.1016/j.ipl.2015.11.017

[35] Meena Mahajan and Gaurav Sood. 2022. QBF Merge Resolution Is Powerful but Unnatural. In *25th International Conference on Theory and Applications of Satisfiability Testing, SAT 2022, August 2-5, 2022, Haifa, Israel (LIPIcs, Vol. 236)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 22:1–22:19. https://doi.org/10.4230/LIPIcs.SAT.2022.22

[36] Jakob Nordström. 2015. On the interplay between proof complexity and SAT solving. *SIGLOG News* 2, 3 (2015), 19–44.

[37] Tomáš Peitl, Friedrich Slivovsky, and Stefan Szeider. 2019. Long-Distance Q-Resolution with Dependency Schemes. *J. Autom. Reasoning* 63, 1 (2019), 127–155.

[38] Tomáš Peitl, Friedrich Slivovsky, and Stefan Szeider. 2019. Proof Complexity of Fragments of Long-Distance Q-Resolution. In *Theory and Applications of Satisfiability Testing - SAT 2019 - 22nd International Conference, SAT, Proceedings (Lecture Notes in Computer Science, Vol. 11628)*, Mikolás Janota and Inês Lynce (Eds.). Springer, 319–335.

[39] Luca Pulina and Martina Seidl. 2019. The 2016 and 2017 QBF solvers evaluations (QBFEVAL'16 and QBFEVAL'17). *Artif. Intell.* 274 (2019), 224–248.

[40] John Alan Robinson. 1965. A machine-oriented logic based on the resolution principle. *J. ACM* 12 (1965), 23–41.

[41] Ankit Shukla, Armin Biere, Luca Pulina, and Martina Seidl. 2019. A Survey on Applications of Quantified Boolean Formulas. In *31st IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2019*. 78–84.

[42] Friedrich Slivovsky and Stefan Szeider. 2016. Soundness of Q-resolution with dependency schemes. *Theoretical Computer Science* 612 (2016), 83–101.

[43] G. S. Tseitin. 1983. On the Complexity of Derivation in Propositional Calculus. In *Automation of Reasoning: 2: Classical Papers on Computational Logic 1967–1970*, Jörg H. Siekmann and Graham Wrightson (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 466–483. https://doi.org/10.1007/978-3-642-81955-1_28

[44] Allen Van Gelder. 2012. Contributions to the Theory of Practical Quantified Boolean Formula Solving. In *Proc. Principles and Practice of Constraint Programming (CP'12)*. 647–663.

[45] Lintao Zhang and Sharad Malik. 2002. Conflict driven learning in a quantified Boolean Satisfiability solver. In *IEEE/ACM International Conference on Computer-aided Design, ICCAD 2002*. 442–449.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il