Electronic Colloquium on Computational Complexity, Report No. 152 (2015)

Are Short Proofs Narrow? QBF Resolution is *not* Simple.

Olaf Beyersdorff[1], Leroy Chew[1], Meena Mahajan[2], and Anil Shukla[2]

[1] School of Computing, University of Leeds, United Kingdom.
`{o.beyersdorff,mm12lnc}@leeds.ac.uk`
[2] The Institute of Mathematical Sciences, Chennai, India.
`{meena,anilsh}@imsc.res.in`

Abstract. The groundbreaking paper 'Short proofs are narrow – resolution made simple' by Ben-Sasson and Wigderson (J. ACM 2001) introduces what is today arguably *the* main technique to obtain resolution lower bounds: to show a lower bound for the width of proofs. Another important measure for resolution is space, and in their fundamental work, Atserias and Dalmau (J. Comput. Syst. Sci. 2008) show that lower bounds for space again can be obtained via lower bounds for width.
In this paper we assess whether similar techniques are effective for resolution calculi for quantified Boolean formulas (QBF). A mixed picture emerges. Our main results show that both the relations between size and width as well as between space and width drastically *fail* in Q-resolution, even in its weaker tree-like version. On the other hand, we obtain positive results for the expansion-based resolution systems $\forall$Exp+Res and IR-calc, however only in the weak tree-like models.
Technically, our negative results rely on showing width lower bounds together with simultaneous upper bounds for size and space. For our positive results we exhibit space and width-preserving simulations between QBF resolution calculi.

1 Introduction

The main objective in *proof complexity* is to obtain precise bounds on the size of proofs in various formal systems; and this objective is closely linked to and motivated by foundational questions in computational complexity (Cook's programme), first-order logic (separating theories of bounded arithmetic), and SAT solving. In particular, resolution is one of the best studied and most important propositional proof systems, as it forms the backbone of modern SAT solvers based on conflict-driven clause learning (CDCL). Complexity bounds for resolution proofs directly translate into bounds on the performance of SAT solvers.

What is arguably even more important than showing the actual bounds is to develop *general techniques* that can be applied to obtain lower bounds for important proof systems. A number of ingenious techniques have been designed to show lower bounds for the *size of resolution proofs*, among them feasible interpolation [21], which applies to many further systems. In their pioneering paper [6], Ben-Sasson and Wigderson showed that resolution size lower bounds can be elegantly obtained by showing lower bounds to the *width* of resolution proofs. Indeed, the discovery of this relation between width and size of resolution proofs was a milestone in our understanding of resolution, and today many if not most lower bounds for resolution are obtained via the size-width technique.

Another important measure for resolution is *space* [17], as it corresponds to memory requirements of solvers in the same way as resolution size relates to their running time. In their fundamental work [1], Atserias and Dalmau demonstrated that also space is tightly related to width. Indeed, showing lower bounds for width serves again as the primary method to obtain space lower bounds. Since these discoveries the relations between resolution size, width, and space have been subject to intense research (cf. [13]), and in particular sharp trade-off results between the measures have been obtained (cf. e.g. [4, 5, 24]).

ISSN 1433-8092

In this paper we initiate the study of width and space in resolution calculi for quantified Boolean formulas (QBF) and address the question whether similar relations between size, width and space as for classical resolution hold in QBF. Before explaining our results we sketch recent developments in QBF proof complexity.

QBF proof complexity is a relatively young field studying proof systems for quantified Boolean logic. Similarly as in the propositional case, one of the main motivations for the field comes via its intimate connection to solving. Although QBF solving is at an earlier state than SAT solving, due to its PSPACE completeness, QBF even applies to further fields such as formal verification or planning [7, 16, 25]. Each successful run of a solver on an unsatisfiable instance can be interpreted as a proof of unsatisfiability; and this connection turns proof complexity into the main theoretical tool to understand the performance of solving. As in SAT, QBF solvers are known to correspond to the resolution proof system and its variants.

However, compared to SAT, the QBF picture is more complex as there exist two main solving approaches utilising CDCL and expansion-based solving. To model the strength of these QBF solvers, a number of resolution-based QBF proof systems have been developed. Q-resolution (Q-Res) by Kleine Büning, Karpinski, and Flögel [20] forms the core of the CDCL-based systems. To capture further ideas from CDCL solving, Q-Res has been augmented to long-distance resolution [2], universal resolution [27], and their combinations [3]. Powerful proof systems for expansion-based solving were recently developed in the form of ∀Exp+Res [19], and the stronger IR-calc and IRM-calc [9].

In this paper we concentrate on the three QBF resolution systems Q-Res, ∀Exp+Res, and IR-calc. This choice is motivated by the fact that Q-Res and ∀Exp+Res form the base systems for CDCL and expansion-based solving, respectively, and IR-calc unifies both approaches in a natural way, as it simulates both Q-Res and ∀Exp+Res [9]. Recent findings show that CDCL and expansion are indeed orthogonal paradigms as Q-Res and ∀Exp+Res are incomparable with respect to simulations [10].

Understanding which lower bound techniques are effective in QBF proof complexity is of paramount importance for progress in the field. In [11] it was shown that the feasible interpolation technique applies to all QBF resolution systems. Another successful transfer of a classical technique was obtained in [12] for a game-theoretic characterisation of proof size in tree-like Q-Res.

Our contributions

The central question we address here is whether *lower bound techniques via width*, which have revolutionised classical proof complexity, are also effective for QBF resolution systems.

Though space and width have not been considered in QBF before, these notions straightforwardly apply to QBF resolution systems. However, due to the ∀-reduction rule in Q-Res handling universal variables, it is relatively easy to enforce that universal literals accumulate in clauses of Q-Res proofs, thus always leading to large width, irrespective of size and space requirements (Lemma 5). This prompts us to consider *existential width* — counting only existential literals — as an appropriate width measure in QBF. This definition aligns both with Q-Res, resolving only on existential variables, as well as with ∀Exp+Res and IR-calc, which like all expansion systems only operate on existential literals.

1. Negative results. Our main results show that the size-width relation of [6] as well as the space-width relation of [1] dramatically *fail* for Q-Res, even when considering the tighter existential width. We first notice that the proof establishing the size-width result in [6] almost fully goes through, except for some very inconspicuous step that fails in QBF (Proposition 6). But not only the technique fails: we prove that Tseitin transformations of formulas expressing a natural completion principle from [19] have small size and space, but require large existential width in tree-like Q-Res (Theorem 7). As the formulas for the completion principle have $O(n^2)$ variables, they do not rule out size-width relations in general dag-like Q-Res. However, we show that different formulas, hard for tree-like Q-Res [19], provide counterexamples for size-width and space-width relations in full Q-Res (Theorem 8).

Technically, our main contributions are width lower bounds for the above formulas, which we show by careful counting arguments. We complement these results by existential width lower bounds for parity-formulas from [10], providing an optimal width separation between Q-Res and ∀Exp+Res (Theorem 15).

2. Positive results and width-space-preserving simulations. Though the negative picture above prevails, we prove some positive results for size-width-space relations for tree-like versions of the expansion resolution systems ∀Exp+Res and IR-calc. Proofs in ∀Exp+Res can be decomposed into two clearly separated parts: an expansion phase followed by a classical resolution phase. This makes it easy to transfer almost the full spectrum of the classical relations to ∀Exp+Res (Theorem 16).

To lift these results to IR-calc (Theorem 17), we show a series of careful space and width-preserving simulations between tree-like Q-Res, ∀Exp+Res, and IR-calc. In particular, we show the surprising result that tree-like ∀Exp+Res and tree-like IR-calc are equivalent (Lemma 12), thus providing a rare example of two proof systems that coincide in the tree-like, but are separated in the dag-like model [10]. In addition, our simulations provide a simpler proof for the simulation of tree-like Q-Res by ∀Exp+Res (Corollary 14), shown in [19] via a more involved argument.

Our last positive result is a size-space relation in tree-like Q-Res (Theorem 17), which we show by a pebbling game analogous to the classical relation in [17]. Not surprisingly, this only positive result for Q-Res avoids any reference to the notion of width.

As the bottom line we can say that QBF proof complexity is not just a replication of classical proof complexity: it shows quite different and interesting effects as we demonstrate here. Especially for lower bounds it requires new ideas and techniques. We remark that in this direction, a new and 'genuine QBF technique' based on strategy extraction was recently developed, showing lower bounds for Q-Res [10] and indeed much stronger systems [8].

Organisation of the paper

The remainder of this paper is organised as follows. We start by reviewing background information on classical and QBF resolution systems (Section 2), including definitions of size, space, and width together with their main classical relations (Section 3). In Section 4 we prove our main negative results on the failure of the transfer of the classical size-width and space-width results to QBF. Section 5 contains the simulations between tree-like versions of Q-Res, ∀Exp+Res, and IR-calc, paying special attention to width and space. This enables us to show in Section 6 the positive results for relations between size, width, and space in these systems. We conclude in Section 7 with a discussion and directions for future research.

2 Notations and Preliminaries

Quantified Boolean Formulas. A (closed prenex) *Quantified Boolean Formula* (QBF) is a formula in quantified propositional logic where each variable is quantified at the beginning of the formula, using either an existential or universal quantifier. We denote such formulas as $\mathcal{Q} \,.\, \phi$, where ϕ is a propositional Boolean formula in Conjunctive Normal Form (CNF), called *matrix*, and $\mathcal{Q}$ is its *quantifier prefix*.

Given a variable y, either existentially quantified or universally quantified in $\mathcal{Q} \,.\, \phi$, the *quantification level* of y in $\mathcal{Q} \,.\, \phi$, $\mathrm{lv}(y)$, is the number of alternations of quantifiers y has on its left in the quantifier prefix of $\mathcal{Q} \,.\, \phi$. Given a variable y, we will sometimes refer to the variables with quantification level lower than $\mathrm{lv}(y)$ as variables *left* of y; analogously the variables with quantification lever higher than $\mathrm{lv}(y)$ will be *right* of y.

Resolution Calculi

Resolution (Res), introduced by Blake [14] and Robinson [26], is a refutational proof system manipulating unsatisfiable CNFs as sets of clauses. The only inference rule is

$$\frac{C \vee x \qquad D \vee \neg x}{C \cup D} \text{ (Res)}$$

where C, D denote clauses and x is a variable. A Res refutation derives the empty clause $\Box$. If we only allow proofs in form of a tree, i.e., each derived clause can be used at most once, we speak of *tree-like resolution*, denoted $\mathsf{Res_T}$.

QBF resolution calculi. Q-resolution (Q-Res) [20] is a resolution-like calculus that operates on QBFs in prenex form where the matrix is a CNF. It uses the propositional resolution rule (Res) with the side conditions that variable x is existential and if $z \in C$, then $\neg z \notin D$. In addition Q-Res has a universal reduction rule

$$\frac{C \vee u}{C} \text{ ($\forall$-Red)}$$

where variable u is universal and all other existential variables $x \in C$ are left of u in the quantifier prefix.

In addition to Q-Res we consider two further QBF resolution calculi that have been introduced to model *expansion-based QBF solving.* These calculi are based on *instantiation* of universal variables: $\forall$Exp+Res [19], and IR-calc [9]. Both calculi operate on clauses that comprise only existential variables from the original QBF, which are additionally *annotated* by a substitution to some universal variables, e.g. $\neg x^{u/0, v/1}$. For any annotated literal l^σ, the substitution σ must not make assignments to variables right of l, i.e. if $u \in \mathsf{dom}(\sigma)$, then u is universal and $\mathrm{lv}(u) < \mathrm{lv}(l)$. To preserve this invariant, we use the *auxiliary notation* $l^{[\sigma]}$, which for an existential literal l and an assignment σ to the universal variables filters out all assignments that are not permitted, i.e. $l^{[\sigma]} = l^{\{u/c \in \sigma \mid \mathrm{lv}(u) < \mathrm{lv}(l)\}}$. We say that an assignment is complete if its domain is all universal variables. Likewise, we say that a literal x^τ is fully annotated if all universal variables u with $\mathrm{lv}(u) < \mathrm{lv}(x)$ in the QBF are in $\mathsf{dom}(\tau)$, and a clause is fully annotated if all its literals are fully annotated.

The calculus $\forall$Exp+Res from [19] works with fully annotated clauses on which resolution is performed. The rules of $\forall$Exp+Res are shown in Figure 1.

$$\overline{\left\{l^{[\tau]} \mid l \in C, l \text{ exist.}\right\} \cup \{\tau(l) \mid l \in C, l \text{ univ.}\}} \text{ (Ax)}$$

C is a clause from the matrix and τ is an assignment to all universal variables.

$$\frac{C_1 \vee x^\tau \qquad C_2 \vee \neg x^\tau}{C_1 \cup C_2} \text{ (Res)}$$

Fig. 1. The rules of $\forall\mathsf{Exp}{+}\mathsf{Res}$ [19]

In contrast, the system IR-calc from [9] is more flexible. It uses 'delayed' expansion and can mix instantiation with resolution steps. Formally, IR-calc works with partial assignments on which we use auxiliary operations of *completion* and *instantiation.* For assignments τ and μ, we write $\tau \veebar \mu$ for the assignment σ defined as $\sigma(x) = \tau(x)$ if $x \in \mathsf{dom}(\tau)$, otherwise $\sigma(x) = \mu(x)$ if $x \in \mathsf{dom}(\mu)$. The operation $\tau \veebar \mu$ is called *completion* as μ provides values for variables not defined in τ. For an assignment τ and an annotated clause C, the function $\mathsf{inst}(\tau, C)$ returns the annotated clause $\left\{l^{[\sigma \veebar \tau]} \mid l^\sigma \in C\right\}$. The system IR-calc uses the rules depicted in Figure 2.

$$\overline{\left\{x^{[\tau]} \mid x \in C, x \text{ is existential}\right\}} \text{ (Ax)}$$

C is a non-tautological clause from the matrix.
$\tau = \{u/0 \mid u \text{ is universal in } C\}$, where the notation $u/0$ for literals u is shorthand for $x/0$ if $u = x$ and $x/1$ if $u = \neg x$.

$$\frac{x^\tau \vee C_1 \qquad \neg x^\tau \vee C_2}{C_1 \cup C_2} \text{ (Res)} \qquad\qquad \frac{C}{\mathsf{inst}(\tau, C)} \text{ (Instantiation)}$$

τ is a (partial) assignment to universal variables.

Fig. 2. The rules of IR-calc [9]

Simulations. Given two proof systems P and Q for the same language (TAUT or QBF), P *p-simulates* Q (denoted $Q \leq_p P$) if each Q-proof can be transformed in polynomial time into a P-proof of the same formula. Two systems are called *p-equivalent* if they p-simulate each other. In [9] it was shown that IR-calc p-simulates both Q-Res and $\forall\mathsf{Exp}{+}\mathsf{Res}$, while [10] shows that Q-Res and $\forall\mathsf{Exp}{+}\mathsf{Res}$ are incomparable, i.e., IR-calc is exponentially stronger than both Q-Res and $\forall\mathsf{Exp}{+}\mathsf{Res}$. However, $\forall\mathsf{Exp}{+}\mathsf{Res}$ can p-simulate $\mathsf{Q\text{-}Res}_\mathsf{T}$ [19].

3 Size, width and space in resolution calculi

The purpose of the section is twofold: first to review the measures size, width, and space and their relations in classical resolution; and second to explain how to apply these measures to

QBF resolution systems. While this is straightforward for size and space, we need a more elaborate discussion on what constitutes a good notion of width for QBF resolution systems.

3.1 Defining size, width, and space for resolution

For a CNF F, $|F|$ denotes the number of clauses in it, and $w(F)$ denotes the maximum number of literals in any clause of F, and we extend the same notation to QBFs with a CNF matrix.

For S one of the resolution calculi Res, Q-Res, $\forall$Exp+Res, IR-calc, let $\pi \vdash_{\mathsf{S}} F$ (resp. $\pi \vdash_{\mathsf{S_T}} F$) denote that π is an S-proof (tree-like S-proof, respectively), of the formula F. For a proof π of F in system S, its size $|\pi|$ is defined as the number of clauses in π. The **size** complexity $S(\vdash_{\mathsf{S}} F)$ of deriving F in S is defined as $\min\{|\pi| \;:\; \pi \vdash_{\mathsf{S}} F\}$. The tree-like size complexity, denoted $S(\vdash_{\mathsf{S_T}} F)$, is $\min\{|\pi| \;:\; \pi \vdash_{\mathsf{S_T}} F\}$.

A second complexity measure is the minimal **width**. The width of a clause C is the number of literals in C, denoted $w(C)$. The width of a CNF F, denoted $w(F)$, is the maximum width of a clause in F; $w(F) = \max\{w(C) \;:\; C \in F\}$. The width $w(\pi)$ of a proof π is defined as the maximum width of any clause appearing in π, i.e, $w(\pi) = \max\{w(C) \;:\; C \in \pi\}$. The width $w(\vdash_{\mathsf{S}} F)$ of refuting a CNF F in S is defined as $\min\{w(\pi) \;:\; \pi \vdash_{\mathsf{S}} F\}$. Again the same notation extends to quantified CNFs.

Note that for width in any calculus, whether the proof is tree-like or not is immaterial, since a proof can always be made tree-like by duplication without increasing the width. We therefore drop the T subscript when talking about proof width.

The third complexity measure for resolution calculi is **space**. For classical resolution, this measure[3] was first defined by Esteban and Torán in [17]. Informally, it is the minimal number of clauses that must be kept simultaneously to refute a formula. Instead of viewing a proof as a DAG, we view it as a sequence of CNF formulas $F_0, F_1, \ldots, F_s$, where $F_0 = \emptyset$, $\Box \in F_s$, and each F_{i+1} is obtained from F_i by either erasing some clause, or by downloading an axiom, or by adding a resolvent of clauses in F_i. In the latter case, one of the clauses used in the resolution may also simultaneously be deleted. The space used by this proof is the maximum number of clauses in any F_i, $\max_{i \in [s]}\{|F_i|\}$. The space to refute F is the minimum, over all proofs σ, of the space used by σ. We can directly adapt this definition to QBF resolution calculi as well.

Definition 1 (Space-oriented proof sequences). *A false QBF sentence $\mathcal{F}$ can be refuted in system S within space k if there is a sequence σ of QBFs $\mathcal{F}_0, \mathcal{F}_1, \ldots, \mathcal{F}_s$, all having the same quantifier prefix as $\mathcal{F}$, and with matrix $F_0, F_1, \ldots, F_s$, respectively, such that $F_0 = \emptyset$, F_1 contains a subset of clauses obtained from the corresponding axiom download in the proof system S, $F_s = \{\Box\}$ (the empty clause), each F_i has at most k clauses, and for each $i < s$, F_{i+1} is obtained from F_i by one of the following rules:*

1. **Erase***: $F_{i+1} = F_i \setminus \{C\}$ for some clause $C \in F_i$.*
2. **Inference***: $F_{i+1} \subseteq F_i \cup \{C\}$ for C obtained by applying any inference rule of the proof system S. In this step, one of the hypotheses used in the inference rule may be erased.*
3. **Axiom Download***: $F_{i+1} = F_i \cup \{C\}$ for some clause C obtained by applying the axiom download rule of the proof system S.*

[3] Also called clause space, to distinguish it from variable space or total space. We consider only clause space in this paper, and so we call it just space.

For a proof written as a sequence σ as above, the clause space of σ, denoted by Space(σ), is $\max_{i\in[s]}\{|F_i|\}$. The clause space needed to refute a QBF $\mathcal{F}$ in the S-proof system, denoted by Space($\vdash_{\mathsf{S}} \mathcal{F}$), is the minimum Space($\sigma$) over all sequences σ refuting $\mathcal{F}$.

If we modify the inference step above so that the clause(s) used to obtain the inference are erased in the same step, then any clause D can be used at most once and we obtain a tree-like space-oriented $\mathcal{S}$-proof. Correspondingly we can define Space($\vdash_{\mathsf{S}_T} \mathcal{F}$) as the minimum space used by any tree-like proof sequence refuting $\mathcal{F}$.

3.2 Relations in classical resolution

We now state some of the main relations between size, width, and space for classical resolution. We start with the foundational size-width relation of Ben-Sasson and Wigderson [6].

Theorem 2 (Ben-Sasson, Wigderson [6]). *For all unsatisfiable CNFs F the following relations hold:* $S(\vdash_{\mathsf{Res}_T} F) \geq 2^{w\left(\vdash_{\mathsf{Res}} F\right)-w(F)}$ *and* $S(\vdash_{\mathsf{Res}} F) = \exp\left(\Omega\left(\frac{\left(w\left(\vdash_{\mathsf{Res}} F\right)-w(F)\right)^2}{n}\right)\right)$.

Space complexity was introduced in [17] and relations between space, size and width are explored (cf. also [13, 22]).

Theorem 3 (Esteban, Torán [17]). *For all unsatisfiable CNFs F the following relation holds:* $S(\vdash_{\mathsf{Res}_T} F) \geq 2^{Space\left(\vdash_{\mathsf{Res}_T} F\right)} - 1$.

The fundamental relation between space and width was obtained in [1]; a more direct proof was given recently in [18].

Theorem 4 (Atserias, Dalmau [1]). *For all unsatisfiable CNFs F the following relation holds:* $w(\vdash_{\mathsf{Res}} F) \leq Space(\vdash_{\mathsf{Res}} F) + w(F) - 1$.

3.3 Existential width: What is the right width notion for QBF?

We wish to explore the possibility of a similar approach as used in [6] to prove an analogue of Theorem 2 when dealing with QBFs. The following simple example shows that the relationships in Theorem 2 and Theorem 4 do not carry over for the system Q-Res.

Consider the following false QBF $\mathcal{F}_n$ over $2n$ variables:

$$\begin{aligned} \mathcal{F}_n =& \forall u_1 \dots u_n \exists e_0 \exists e_1 \dots e_n. \\ C_0 : \quad & (e_0) \wedge \\ \text{For } i \in [n], D_i : \quad & (\bar{e}_{i-1} \vee u_i \vee e_i) \wedge \\ D_{n+1} : \quad & (\bar{e}_n) \end{aligned}$$

Lemma 5. $S(\vdash_{\mathsf{Q\text{-}Res}_T} \mathcal{F}_n) \in O(n)$ *and* $Space(\vdash_{\mathsf{Q\text{-}Res}_T} \mathcal{F}_n) \in O(1)$, *but* $w(\vdash_{\mathsf{Q\text{-}Res}} \mathcal{F}_n) \in \Omega(n)$.

Proof. For the upper bounds consider the following proof. For $i \in [n]$, let $C_i = (u_1 \vee \cdots \vee u_i \vee e_i)$. For $i \in [n]$ in sequence, resolving C_{i-1} and D_i on variable e_{i+1} gives C_i. Resolving C_n and D_{n+1} on variable e_n gives the clause $U = (u_1 \vee \cdots \vee u_n)$. Finally, applying $\forall$-Red on the clause U yields the empty clause in n more steps.

This is a tree-like proof of size $O(n)$. Further, each resolution step involves an axiom clause, so at each step we need to hold just two clauses, and so the space requirement is $O(1)$.

Concerning the width lower bound, by the order of quantification in $\mathcal{F}_n$, every existential literal in $\mathcal{F}_n$ blocks any $\forall$-reduction. Therefore, in any refutation, when a $\forall$-reduction is first used, the clause C has only universal variables. At this point, the empty clause is derivable from C by a series of $\forall$ reductions. Note that if any clause is dropped from $\mathcal{F}_n$, the resulting sentence is no longer false. Thus any refutation must use all clauses. Hence C must have all universal variables in it; it must be $(u_1 \vee \cdots \vee u_n)$ as all u_i variables have been accumulated, without being reduced. Then clause C has width n. $\square$

Noting that $w(\mathcal{F}_n) = 3$, Lemma 5 implies that the relationships from Theorem 2 and Theorem 4 do not hold for Q-Res and Q-Res$_\mathsf{T}$.

As the above example illustrates, it is easy to enforce that universal variables are accumulated in a clause, thus leading to large width. Hence the following question naturally arises: can we obtain size-width or space-width relations by using the tighter measure of only counting existential variables?

This aligns with the situation in the expansion systems $\forall$Exp+Res and IR-calc, where clauses contain only existential variables. In this respect, it is worth noting that the above example indeed does not demonstrate the failure of the size-width relationship in expansion-based calculi. For instance, in $\forall$Exp+Res, a tree-like refutation could download the existential variables of axioms annotated with $u_i/0$ for $i \in [n]$, and generate the empty clause in $O(n)$ steps with width just 2 at the leaves and 1 at the internal nodes.

Thus, to get a consistent and interesting width measure for QBF calculi, we consider the notion of **existential width** that just counts the number of existential literals. This approach is justified also for Q-Res as the calculus can only resolve on existential variables, and rules out the easy counterexamples above. Formally, we define the existential width of a clause C to be the number of existential literals in C, and denote it by $w_\exists(C)$. Using $w_\exists$ instead of w everywhere, we obtain the existential width of a formula $w_\exists(F)$, of a proof $w_\exists(\pi)$, and of refuting a false sentence $w_\exists(\vdash_\mathsf{S} \mathcal{F})$.

For the expansion systems $\forall$Exp+Res and IR-calc the notions of existential width and width coincide. (In particular, distinct annotations of the same existential variable in a single clause are counted as distinct literals.) Hence we can drop the $\exists$ subscript in width of proofs in these systems. For the width of the sentence itself, there is still a difference between w and $w_\exists$.

4 Negative results: Size-width and space-width relations fail in Q-Res

In this section we show that in the Q-Res proof system, even replacing width by existential width, the relations to size or space as in classical resolution (Theorems 2 and 4) no longer hold for both tree-like and general proofs.

Firstly, we point out where the technique of [6] fails. A crucial ingredient of their proof is the following statement: if a clause A can be derived from $F|_{x=1}$ in width w, then the clause $A \vee \neg x$ can be derived from F in width $w + 1$ (possibly using a weakening rule at the end). We show that the statement no longer holds in Q-Res.

Proposition 6. *There are false sentences ψ_n, with an existential literal b quantified at the innermost level, such that the sentence $\psi_n|_{b=1}$ is false and has a small existential-width proof, but ψ_n itself needs large existential width to refute in Q-Res.*

Proof. The sentence ψ_n is constructed by taking the conjunction of two sentences with distinct variables. The first sentence is a very simple one: $\exists a \forall u \exists b \quad (a \vee u \vee \bar{b}) \wedge (\bar{a})$. It is a true sentence, but if b is set to 1, it becomes false. The second sentence is a false sentence of the form $\exists \boldsymbol{x} G_n(\boldsymbol{x})$, where G_n is any unsatisfiable CNF formula over the $\boldsymbol{x}$ variables, such that G_n needs large width in classical resolution. One such example is the CNF formula described by Bonet and Galesi [15], that we denote as BG_n. BG_n is an unsatisfiable 3-CNF formula over $O(n^2)$ variables with $w(\vdash BG_n) = \Omega(n)$. Now define ψ_n as:

$$\exists \boldsymbol{x} \exists a \forall u \exists b \ \ (a \vee u \vee \bar{b}) \wedge (\bar{a}) \wedge BG_n(\boldsymbol{x}).$$

Note that the clauses $(a \vee u \vee \bar{b}) \wedge (\bar{a})$ contain a contradiction if and only if $b = 1$. Thus $\psi_n|_{b=1}$ can be refuted with existential width 1 using just these two clauses: a $\forall$-Red on $(a \vee u)$ yields a which can be resolved with $\bar{a}$. On the other hand, to refute ψ_n, the contradiction in BG_n must be exposed. Since all the variables involved are existential, Q-Res degenerates to classical resolution, requiring (existential) width $\Omega(n)$. ◻

The example in the proof of Proposition 6 can be made 'less degenerate' by interleaving more existential and universal variables disjoint from $\boldsymbol{x}$ and putting them in the first sentence. All we need is that b is quantified existentially at the end, the first sentence is true as a whole but false if $b = 1$, and this latter sentence can be refuted in Q-Res with small existential width.

We now show that it is not just the technique of [6] that fails for Q-Res. No other technique will work either, because the relation from Theorem 2 between size and existential width itself fails to hold. The same example also shows that the relation from Theorem 4 between space and existential width also fails to hold.

We first give an example where the relation for tree-like proofs fails.

Theorem 7. *There is a family of false QBF sentences CR'_n over $O(n^2)$ variables, such that $S(\vdash_{\text{Q-Res}_T} CR'_n) = n^{O(1)}$, $w_\exists(CR'_n) = 3$, $Space(\vdash_{\text{Q-Res}_T} CR'_n) = O(1)$, and $w_\exists(\vdash_{\text{Q-Res}_T} CR'_n) = \Omega(n)$.*

Proof. Consider the following formulas CR_n, introduced by Janota and Marques-Silva [19]:

$$\begin{array}{rll} CR_n = & \exists x_{1,1} \dots x_{n,n} \, \forall z \, \exists a_1 \dots a_n \exists b_1 \dots b_n. & \\ (C_{i,j}) & (x_{i,j} \vee z \vee a_i), & i,j \in [n] \\ (D_{i,j}) & (\bar{x}_{i,j} \vee \bar{z} \vee b_j), & i,j \in [n] \\ (A) & \bigvee_{i \in [n]} \bar{a}_i & \\ (B) & \bigvee_{i \in [n]} \bar{b}_i. & \end{array}$$

CR_n is constructed from a principle called the *completion principle*. Consider two sets $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$, and depict their cross product $A \times B$ as in the table below.

a_1	a_1	$\dots$	a_1	a_2	a_2	$\dots$	a_2	$\dots$	$\dots$	a_n	a_n	$\dots$	a_n
b_1	b_2	$\dots$	b_n	b_1	b_2	$\dots$	b_n	$\dots$	$\dots$	b_1	b_2	$\dots$	b_n

The following two-player game is played on the above table. In the first round, player 1 deletes exactly one cell from each column. In the second round, player 2 chooses one of the two rows. Player 2 wins if the chosen row contains either the complete set A or the set B; otherwise player 1 wins. The completion principle states that player 2 has a winning strategy. The false sentence CR_n expresses the notion that player 1 has a winning strategy. For each column $\begin{bmatrix} a_i \\ b_j \end{bmatrix}$ of the table (denote this the $(i,j)^{th}$ column), there is a boolean variable $x_{i,j}$. Let $x_{i,j} = 0$ denote that player 1 'deletes b_j' (i.e, keeps a_i) from the $(i,j)^{th}$ column, and $x_{i,j} = 1$ denotes that player 1 keeps b_j in the $(i,j)^{th}$ column. There is a variable z to denote the choice of player 2: $z = 0$ means 'choose top row'. The Boolean variables a_i, b_j, for $i, j \in [n]$ encode that for the chosen values of all the $x_{k,\ell}$, and the row chosen via z, at least one copy of the element a_i and b_j respectively is kept. (eg. $(x_{i,j} \wedge z) \Rightarrow b_j$).

It is known that CR_n has $O(n^2)$ proofs in $\mathsf{Q\text{-}Res}$, and even in $\mathsf{Q\text{-}Res_T}$ [23]. However CR_n has large existential width, and in order to prove Theorem 7, we need a formula with constant initial existential width. To achieve this we proceed similarly as in the Tseitin transformations, i.e., we introduce $2n+2$ new existential variables (i.e, $\boldsymbol{y}, \boldsymbol{p}$) at the innermost level in CR_n, and replace the two large clauses in CR_n by any CNF formula which preserves their satisfiability. Let CR'_n denote the modified formula

$$CR'_n = \exists x_{1,1} \dots x_{n,n} \, \forall z \, \exists a_1 \dots a_n \exists b_1 \dots b_n \exists y_0 \dots y_n \exists p_0 \dots p_n.$$

$$(C_{i,j}) \qquad (x_{i,j} \vee z \vee a_i), \qquad i, j \in [n] \tag{1}$$

$$(D_{i,j}) \qquad (\bar{x}_{i,j} \vee \bar{z} \vee b_j), \qquad i, j \in [n] \tag{2}$$

$$\bar{y}_0 \wedge \bigwedge_{i \in [n]} (y_{i-1} \vee \bar{a}_i \vee \bar{y}_i) \wedge y_n \tag{3}$$

$$\bar{p}_0 \wedge \bigwedge_{i \in [n]} (p_{i-1} \vee \bar{b}_i \vee \bar{p}_i) \wedge p_n. \tag{4}$$

Note that $w_\exists(CR'_n) = 3$.

It is clear that from the type-(3) clauses of CR'_n, we can derive the large clause $\bigwedge_{i \in [n]} \bar{a}_i$ of CR_n in $n+1$ resolution steps. Similarly we can derive the large clause $\bigwedge_{i \in [n]} \bar{b}_i$ of CR_n from the type (4) clauses in $n+1$ steps. The proof refuting CR_n uses each of these large clauses n times; see below. Thus $S(\vdash_{\mathsf{Q\text{-}Res_T}} CR'_n) \leq S(\vdash_{\mathsf{Q\text{-}Res_T}} CR_n) + O(n^2) = O(n^2)$.

We briefly sketch the refutation of CR_n from [23] to analyse its space requirement. The fragment W_j starts with clause A, successively resolves it with clauses from $C_{*,j}$ to get $z \vee x_{1,j} \vee \dots \vee x_{n,j}$, eliminates z through a $\forall$-reduction, then successively resolves it with clauses from $D_{*,j}$ to get $W_j = \bar{z} \vee b_j$. It is easy to see that $O(1)$ space suffices to construct this fragment. The overall proof starts with the clause B, successively resolves it with $W_1, W_2, \dots, W_n$ (reusing the space to construct successive W_j's), and finally gets $\bar{z}$ which is eliminated through a $\forall$-reduction. Again $O(1)$ space suffices.

Finally, we show that CR'_n needs large existential width.

Let π be a proof in $\mathsf{Q\text{-}Res}$, $\pi \vdash_{\mathsf{Q\text{-}Res}} CR'_n$. List the clauses of π in sequence, $\pi = \{D_0, D_1, \dots, D_s = \Box\}$, where each clause in the sequence is either a clause from CR'_n, or is derived from clause(s) preceding it in the sequence using resolution or $\forall$-Red. There must be at least one universal reduction step in π, since all the initial clauses are necessary for refuting CR'_n, some of them contain universal variables, and the only way to remove a universal variable in $\mathsf{Q\text{-}Res}$ is by $\forall$-Red. Let t be the least index such that in the clause D_t, a $\forall$-Red step has been performed on the only universal variable. Without loss of generality, let the universal literal be the positive literal z; the argument for $\bar{z}$ is identical. As the existential variables, $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{y}$, and $\boldsymbol{p}$ all block the universal variable z, none of them is present in the clause D_t. We use this fact to show that $w_\exists(D_t) = \Omega(n)$. Our strategy is to associate some set with each clause in π in a specific way, and use the set size to bound existential width.

We associate the following sets with the literals of CR'_n and the clauses of π.

$$\begin{array}{lrl}
 & \sigma(z) &= \emptyset = \sigma(\bar{z}) \\
\forall i \in [n] & \sigma(a_i) &= [n] \setminus \{i\} = \{1, \dots, n\} \setminus \{i\} \\
\forall i \in [n] & \sigma(x_{i,j}) = \sigma(\bar{a}_i) &= \{i\} \\
\forall i \in [n] & \sigma(\bar{y}_i) &= [n] \setminus [i] = \{i+1, \dots, n\} \\
\forall i \in [n] & \sigma(y_i) &= [i] = \{1, \dots, i\} \\
\forall j \in [n] & \sigma(b_j) &= [n] \setminus \{j\} = \{1, \dots, n\} \setminus \{j\} \\
\forall j \in [n] & \sigma(\bar{x}_{i,j}) = \sigma(\bar{b}_j) &= \{j\} \\
\forall j \in [n] & \sigma(\bar{p}_j) &= [n] \setminus [j] = \{j+1, \dots, n\} \\
\forall j \in [n] & \sigma(p_j) &= [j] = \{1, \dots, j\} \\
\forall D \in \pi & \sigma(D) &= \bigcup_{l \in D} \sigma(l).
\end{array}$$

Note that for variables v in $\boldsymbol{a}$, $\boldsymbol{b}$, $\boldsymbol{p}$, $\boldsymbol{y}$, the sets $\sigma(v)$ and $\sigma(\bar{v})$ form a partition of $[n]$.

For $D \in \pi$, let π_D be the sub-DAG of π, rooted at D. Consider the sub-DAG π_{D_t} of π. We have the following observations:

Observation 1. π_{D_t} contains at least one type-(1) clause as a source; this is because $z \in D_t$, and the only initial clauses containing z are the type-(1) clauses.

Observation 2. π_{D_t} does not contain any clause of type (2): as $z \in D_t$, we know that $\bar{z} \notin D_t$. Therefore if some type-(2) clause is present in this sub-DAG, the only way to remove $\bar{z}$ is via $\forall$-Red. This reduction will take place before the reduction on D_t, contradicting our choice of index t. We also conclude that the literal $\bar{z}$ cannot appear anywhere in π_{D_t}.

Observation 3. π_{D_t} does not contain any type-(4) clause: we know that D_t does not contain $\boldsymbol{p}$ and $\boldsymbol{b}$ variables (because they block z). Any use of type (4) clauses introduces $\boldsymbol{p}$ variables and possibly $\bar{b}$ literals. Removing $\boldsymbol{p}$ variables introduces $\bar{b}$ literals. But $\bar{b}$ can be removed only by resolving with b, which is only in type-(2) clauses. We have already seen that type-(2) clauses are not present in π_{D_t}.

Observation 4. No clause in π_{D_t} contains a literal $\bar{x}_{i,j}$, since $\bar{x}_{i,j}$ are introduced only in type (2) clauses.

Observation 5. For any clause C derived solely from type (3) clauses, $\sigma(C) = [n]$. This is true for type-(3) clauses by definition of σ. Using only these clauses, the only resolution step possible is with a y variable as pivot. The claim can be verified by induction on depth: Since $\sigma(y_i)$ and $\sigma(\bar{y}_i)$ partition $[n]$, $[n] \setminus \sigma(y_i)$ and $[n] \setminus \sigma(\bar{y}_i)$ also partition $[n]$.

We show that all clauses in π_{D_t} that are descendants of some type-(1) clause have large sets associated with them. In particular, we show:

Claim. Every clause D in π_{D_t} such that π_D contains a type-(1) clause has $\sigma(D) = [n]$.

Deferring the proof briefly, we continue with our argument. From the Claim we conclude that $\sigma(D_t) = [n]$. Recall that the variables $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{y}, \boldsymbol{p}$ and the literals $\bar{x}_{i,j}$'s are not present in D_t. The only literals left are positive $x_{i,j}$'s. These literals are associated with singleton sets, and the variables $x_{i,j}$ for different values of j give the same singleton set. So we conclude that for each $i \in [n]$, there must be some $x_{i,j} \in D_t$. Hence $w_\exists(D_t) = \Omega(n)$.

It remains to establish the claimed set size.

Proof (of claim). We proceed by induction on the depth of descendants of type-(1) clauses in π_{D_t}. The base case is a type-(1) clause itself and follows from the definition of σ.

For the inductive step, let D be obtained by resolving $(E \vee r)$ and $(F \vee \bar{r})$. There are two cases to consider: both are descendants of some type-(1) clauses, or only one of them, say $(E \vee r)$, is a descendant of a type-(1) clause. In the former case, by the induction hypothesis, $\sigma(E \vee r) = [n]$ and $\sigma(F \vee \bar{r}) = [n]$. In the latter case, $\sigma(E \vee r) = [n]$ by induction hypothesis, and $\sigma(F \vee \bar{r}) = [n]$ from the observations above. ($(F \vee \bar{r})$ is not a descendant of any type (1) clause. But it belongs to π_{D_t} which has only type-(1) and type-(3) clauses. So it must be a descendant of only type (3) clauses, and hence has $[n]$ associated with it.)

Thus in both cases, we have $\sigma(E \vee r) = \sigma(F \vee \bar{r}) = [n]$. So we have $\sigma(E) \supseteq [n] \setminus \sigma(r)$ and $\sigma(F) \supseteq [n] \setminus \sigma(\bar{r})$. Observe that the pivot variable r can only be either an $\boldsymbol{a}$ or a $\boldsymbol{y}$ variable. Thus $\sigma(r)$ and $\sigma(\bar{r})$ are disjoint, and hence $\sigma(E) \cup \sigma(F) = [n]$. Thus $\sigma(D) = \sigma(E) \cup \sigma(F) = [n]$ as claimed. ☐

This completes the proof of the Theorem. ☐

Since tree-like space is at least as large as space, Theorem 7 also rules out the space-width relation for general dag-like Q-Res proofs. However, observe that Theorem 7 cannot be used to show that the size-existential-width relationship for general dag-like proofs fails in Q-Res, because the sentences CR'_n have $O(n^2)$ variables. However, we show via another example that the relation fails to hold in Q-Res as well. This example cannot be used for proving Theorem 7 because it is known to be hard for Q-Res_{T} [19]. (In [19] the hardness for ∀Exp+Res is shown, which implies hardness for Q-Res_{T}, as ∀Exp+Res p-simulates Q-Res_{T}.)

Theorem 8. *There is a family of false QBFs ϕ'_n in $O(n)$ variables such that $S(\vdash_{\text{Q-Res}} \phi'_n) = n^{O(1)}$, $w_\exists(\phi'_n) = 3$, and $w_\exists(\vdash_{\text{Q-Res}} \phi'_n) = \Omega(n)$.*

Proof. Consider the following formulas ϕ_n, also introduced by Janota and Marques-Silva [19]:

$$\begin{aligned}\phi_n = {} & \exists e_1 \forall u_1 \exists c_1 c_2 \dots \exists e_n \forall u_n \exists c_{2n-1} c_{2n} \\ & \bigwedge_{i \in [n]} \big((\bar{e}_i \vee c_{2i-1}) \wedge (\bar{u}_i \vee c_{2i-1}) \wedge (e_i \vee c_{2i}) \wedge (u_i \vee c_{2i}) \big) \wedge \\ & \Big(\bigvee_{i \in [2n]} \bar{c}_i \Big).\end{aligned}$$

We know from [19], that ϕ_n have polynomial-size proofs in Q-Res (but require exponential-size proofs in Q-Res_{T}). However, in order to prove Theorem 8, we need a formula with constant initial width. To achieve this we consider quantified Tseitin transformations of ϕ_n, i.e. we introduce $2n+1$ new existential variables x_i at the innermost quantification level in ϕ_n, and replace the only large clause in ϕ_n by any CNF formula that preserves satisfiability. Let ϕ'_n denote the modified formula:

$$\begin{aligned}\phi'_n = {} & \exists e_1 \forall u_1 \exists c_1 c_2 \dots \exists e_n \forall u_n \exists c_{2n-1} c_{2n} \exists x_0 \dots x_{2n} \\ & \bigwedge_{i \in [n]} \big((\bar{e}_i \vee c_{2i-1}) \wedge (\bar{u}_i \vee c_{2i-1}) \wedge (e_i \vee c_{2i}) \wedge (u_i \vee c_{2i}) \big) \wedge \quad (5) \\ & \bar{x}_0 \wedge \bigwedge_{i \in [2n]} (x_{i-1} \vee \bar{c}_i \vee \bar{x}_i) \wedge x_{2n}. \quad (6)\end{aligned}$$

Note that $w_\exists(\phi'_n) = 3$.

We refer to the clauses in (6) as x-clauses. It is clear that from the x-clauses, we can derive the large clause of ϕ_n in $2n + 1$ resolution steps and get back ϕ_n. Thus $S(\vdash_{\mathsf{Q\text{-}Res}} \phi'_n) \leq S(\vdash_{\mathsf{Q\text{-}Res}} \phi_n) + 2n + 1 \in n^{O(1)}$.

We now show that ϕ'_n needs large existential width. We follow the same strategy used in proving Theorem 7.

Let π be a proof in Q-Res, $\pi \vdash_{\mathsf{Q\text{-}Res}} \phi'_n$. List the clauses of π in sequence, $\pi = \{D_0, D_1, \ldots, D_s = \Box\}$, where each clause in the sequence is either a clause from ϕ'_n, or is derived from clause(s) preceding it in the sequence using resolution or $\forall$-Red. There must be at least one universal reduction step in π, since all the initial clauses are necessary for refuting ϕ'_n, some of them contain universal variables, and the only way to remove a universal variable in Q-Res is by $\forall$-Red. Let i be the least index such that the clause D_i is obtained by $\forall$-Red on D_j for some $0 < i$. Since all x variables block all u variables, D_j and D_i cannot contain any x variables. We use this fact to show that $w_\exists(D_i) = \Omega(n)$. Our strategy is to associate some set with each clause in π in a specific way, and use the set size to bound existential width.

We associate the following sets with the literals of ϕ'_n and the clauses of π.

$$\begin{array}{rr}
 & \sigma(x_0) = \emptyset \\
\forall i \in [2n] & \sigma(x_i) = [i] = \{1, 2, \ldots, i\} \\
 & \sigma(\bar{x}_0) = [2n] \\
\forall i \in [2n] & \sigma(\bar{x}_i) = [2n] \setminus [i] = \{i + 1, \ldots, 2n\} \\
\forall i \in [n] & \sigma(e_i) = \sigma(u_i) = \sigma(\bar{c}_{2i}) = \sigma(c_{2i-1}) = \{2i\} \\
\forall i \in [n] & \sigma(\bar{e}_i) = \sigma(\bar{u}_i) = \sigma(\bar{c}_{2i-1}) = \sigma(c_{2i}) = \{2i - 1\} \\
\forall D \in \pi & \sigma(D) = \bigcup_{l \in D} \sigma(l).
\end{array}$$

Note that for any literal ℓ, $\sigma(\ell)$ and $\sigma(\bar{\ell})$ are disjoint.

For $D \in \pi$, let π_D be the sub-DAG of π, rooted at D.

Claim. π_{D_i} contains at least one x-clause (axiom clause of type (6)).

Proof. The child D_j of node D_i contains a universal variable which is then removed through $\forall$-Red to get D_i. The universal variables appear only in clauses of type (5), but are blocked by the c-variables in every clause where they appear. Thus, before a reduction is permitted, a c variable must be eliminated by resolution. Since all c variables appear only positively in type (5) clauses, some x-clause must be used in the resolution. $\Box$

We show that all clauses in π_{D_i} that are descendants of some x-clause have large sets associated with them. In particular, we show:

Claim. Every clause D in π_{D_i} such that π_D contains an x-clause has $\sigma(D) = [2n]$.

Deferring the proof briefly, we continue with our argument. From the Claim we conclude that $\sigma(D_i) = [2n]$. Recall that none of the x variables belongs to D_i. All other literals are associated with singleton sets, so D_i must contains at least $2n$ literals in order to be associated with the complete set $[2n]$. Since Q-Res proofs prohibit a variable and its negation in the same clause, at most n of the literals in D_i can be universal variables. Thus D_i has at least n existential literals, hence $w_\exists(D_i) = \Omega(n)$.

It remains to establish the claimed set size.

Proof (of claim). We proceed by induction on the depth of descendants of x-clauses in π_{D_i}. The base case is an x-clause itself and follows from the definition of σ.

For the inductive step, let D be obtained by resolving $(E \vee z)$ and $(F \vee \bar{z})$. There are two cases to consider:

Case 1: Both $(E \vee z)$ and $(F \vee \bar{z})$ are descendants of x-clauses (not necessarily the same x-clause). Then by induction, $\sigma(E \vee z) = \sigma(F \vee \bar{z}) = [2n]$. So $\sigma(E) \supseteq [2n] \setminus \sigma(z)$ and $\sigma(F) \supseteq [2n] \setminus \sigma(\bar{z})$. Since $\sigma(z)$ and $\sigma(\bar{z})$ are disjoint, $\sigma(E) \cup \sigma(F) = [2n]$. Thus $\sigma(D) = \sigma(E) \cup \sigma(F) = [2n]$ as claimed.

Case 2: Exactly one of $(E \vee z)$ and $(F \vee \bar{z})$ is a descendant of an x-clause. Without loss of generality, let $F \vee \bar{z}$ be the descendant. Then $E \vee z$ is either a type-(5) clause or is derived solely from type-(5) clauses using resolution. However, observe that the only clauses derivable solely from type-(5) clauses via resolution, without creating tautologies as mandated in Q-Res, are of the form $(c_{2i-1} \vee c_{2i})$ for some i. It follows that z is not an x variable. Hence $\sigma(z)$ and $\sigma(\bar{z})$ are distinct singleton sets. Further, z cannot be a u variable either, since resolution on universal variables is not permitted in Q-Res.

Now note that for any type-(5) clause C, $\sigma(C) = \{2i-1, 2i\}$ for the appropriate i. Similarly, $\sigma(c_{2i-1} \vee c_{2i}) = \{2i-1, 2i\}$. So if $E \vee z$ is one of these clauses, then $\sigma(E \vee z) = \sigma(z) \cup \sigma(\bar{z})$ and $\sigma(E) = \sigma(\bar{z})$. Further, as in Case 1, by induction we know that $\sigma(F \vee \bar{z}) = [2n]$ and $\sigma(F) \supseteq [2n] \setminus \sigma(\bar{z})$. Hence, $\sigma(E \vee F) = [2n]$ as claimed.

This completes the proof of the claim, and of the theorem. □

The above counterexamples are provided by formulas that require small size, but large existential width. We will now illustrate via another example that also *large size and large width* can occur. These examples are very natural formulas based on the parity function, which have recently been used in [10] to show exponential size lower bounds for Q-Res, and indeed a separation between Q-Res and ∀Exp+Res. We will later use these formulas in Section 5 to also show a separation for width between Q-Res and ∀Exp+Res.

Let $\mathrm{xor}(o_1, o_2, o)$ be the set of clauses expressing $o \equiv o_1 \oplus o_2$; that is, $\{\neg o_1 \vee \neg o_2 \vee \neg o,\ o_1 \vee o_2 \vee \neg o,\ \neg o_1 \vee o_2 \vee o,\ o_1 \vee \neg o_2 \vee o\}$. In [10], the sentence $\mathrm{QP{\scriptstyle ARITY}}_n$ is defined as follows:

$$\exists x_1, \ldots, x_n \,\forall z\, \exists t_2, \ldots, t_n.\ \mathrm{xor}(x_1, x_2, t_2) \cup \bigcup_{i=3}^{n} \mathrm{xor}(t_{i-1}, x_i, t_i) \cup \{z \vee t_n, \neg z \vee \neg t_n\}.$$

The x_i variables act as the input for the parity function, and the t_i variables are defined inductively to calculate $\mathrm{P{\scriptstyle ARITY}}(x_1, \ldots, x_i)$.

We now complement the exponential size lower bound from [10] by a width lower bound.

Theorem 9. *$w_\exists(\vdash_{\textsf{Q-Res}} \mathrm{QP{\scriptstyle ARITY}}_n) \geq n$.*

Proof. Observe that the contradiction occurs semantically because of $z \vee t_n$, $\neg z \vee \neg t_n$. In response to the existential player's choice of $x_1, \ldots, x_n$ the universal player has to play z as $x_1 \oplus \cdots \oplus x_n$ in order to win. In Q-Res we cannot reduce z while any of the t variables are present; and due to the restrictions in Q-Res we cannot resolve the descendants of $z \vee t_n$ with descendants of $\neg z \vee \neg t_n$ until there is at least one ∀-reduction.

We will assume without loss of generality that this happens on the positive literal z. Therefore before this ∀-reduction step we have essentially a resolution proof π from $\Gamma = \mathrm{xor}(x_1, x_2, t_2) \cup \bigcup_{i=3}^{n} \mathrm{xor}(t_{i-1}, x_i, t_i) \cup \{t_n\}$, where we can ignore the z literal in $z \vee t_n$ because it does not restrict any resolution steps in this part of the proof.

The clause D that occurs in π immediately before the $\forall$-reduction must only contain variables from $\{x_1, \ldots, x_n\}$, else the reduction is blocked.

We now use the following observation.

Claim. Suppose $x_1 \oplus \cdots \oplus x_n \models C$ with C a clause, then C is either a tautology or C contains all variables $x_1, \ldots, x_n$.

Proof. Suppose the clause C is not a tautology, but the variables x_i, $i \in I \neq \emptyset$, do not appear in C. Since C is a non-tautological clause, there is exactly one partial assignment α falsifying C. By setting the variables x_i, $i \in I$, appropriately, we can increase α to an assignment satisfying $x_1 \oplus \cdots \oplus x_n$, but still falsifying C. Hence $x_1 \oplus \cdots \oplus x_n \not\models C$. ☐

Γ introduces new variables, but these variables are definitions: given an assignment to the x variables they have exactly one satisfying assignment. Furthermore, the theory of Γ is a conservative extension of the theory of $x_1 \oplus \cdots \oplus x_n$. This and the previous lemma mean that $\forall i \in [n], x_i \in \mathrm{var}(D)$, and therefore D has existential width n. ☐

5 Simulations: Preserving size, width, and space across calculi

After these strong negative results, ruling out size-width and space-width relations in Q-Res and Q-Res$_\mathsf{T}$, we aim to determine whether any positive results hold in the expansion systems $\forall$Exp+Res and IR-calc. Before we can do this we need to relate the measures of size, width, and space across the three calculi Q-Res, $\forall$Exp+Res, IR-calc. Of course, such a comparison in terms of refined simulations is also interesting in its own as it determines the relative strength of the different proof systems. As size corresponds to running time, and space to memory consumption of QBF solvers, such a comparison yields interesting insights into the power of QBF solvers using CDCL vs. expansion techniques.

It is known that IR-calc p-simulates $\forall$Exp+Res and Q-Res [9], and that $\forall$Exp+Res p-simulates Q-Res$_\mathsf{T}$ [19]. We revisit these proofs, with special attention to the width parameter, and also obtain simulating proofs that are tree-like if the original proof is tree-like. The relationships we establish are stated in the following theorem:

Theorem 10. *For all false QBFs $\mathcal{F}$, the following relations hold:*

1. $\frac{1}{2}S(\vdash_{\overline{\mathsf{IR}_T\text{-}\mathsf{calc}}} \mathcal{F}) \leq S(\vdash_{\overline{\forall\mathsf{Exp+Res}_T}} \mathcal{F}) \leq S(\vdash_{\overline{\mathsf{IR}_T\text{-}\mathsf{calc}}} \mathcal{F}) \leq 3S(\vdash_{\overline{\mathsf{Q\text{-}Res}_T}} \mathcal{F})$.
2. $w(\vdash_{\overline{\mathsf{IR\text{-}calc}}} \mathcal{F}) = w(\vdash_{\overline{\forall\mathsf{Exp+Res}}} \mathcal{F}) \leq w_\exists(\vdash_{\overline{\mathsf{Q\text{-}Res}}} \mathcal{F})$.
3. $Space(\vdash_{\overline{\forall\mathsf{Exp+Res}_T}} \mathcal{F}) = Space(\vdash_{\overline{\mathsf{IR}_T\text{-}\mathsf{calc}}} \mathcal{F}) \leq Space(\vdash_{\overline{\mathsf{Q\text{-}Res}_T}} \mathcal{F})$.

These results follow from Proposition 11 and Lemmas 12, 13 that are stated and established below.

Proposition 11 ([9]). *Any proof in $\forall$Exp+Res of size S, width W, and space C can be efficiently converted into a proof in IR-calc of size at most $2S$, width W, and space C. If the proof in $\forall$Exp+Res is tree-like, so is the resulting IR-calc proof.*

Proof. In IR-calc, when an axiom is downloaded, the existential literals in it are annotated partially. However in $\forall$Exp+Res, the annotations are *complete*; all universal variables at a lower level than a literal appear in its annotation. To convert a proof π in $\forall$Exp+Res to one in IR-calc, all that is needed is to follow up each axiom-download with an instantiation that completes the annotations as in π. This introduces at most one extra step per leaf but does not increase width. Also observe that the space required has not changed: to instantiate a clause we can reuse the same space. $\square$

Lemma 12. *$\forall$Exp+Res$_\mathsf{T}$ p-simulates IR$_\mathsf{T}$-calc while preserving its width, size, and space.*

Proof. Recall the main reason why IR$_\mathsf{T}$-calc proofs differ from those in $\forall$Exp+Res$_\mathsf{T}$: axioms are downloaded with partial rather than complete annotations, and annotations can be extended at any stage by the inst operation.

The idea is to systematically transform an IR$_\mathsf{T}$-calc proof, proceeding downwards from the top where we have the empty clause, and modifying annotations as we go down, so that when all leaves have been modified the resulting proof is in fact an $\forall$Exp+Res$_\mathsf{T}$ proof. This crucially requires that we start with a tree-like proof; if the underlying graph is not a tree, we cannot always find a way of modifying the annotations that will work for all descendants.

Let π be an IR$_\mathsf{T}$-calc proof of a false QBF $\mathcal{F}$. Without loss of generality, we can assume that every resolution node has, as parent, an instantiation node. (If it does not, we introduce the dummy $\mathsf{inst}(\emptyset, *)$ node between it and its parent.) Since the proof is tree-like, we can also collapse contiguous instantiation nodes into a single instantiation node. Thus, as we move down a path from the root, nodes are alternately instantiation and resolution nodes. We consider each resolution node and its parent instantiation node to be at the same level.

Starting from the top, which we call level zero, we transform π to another proof π' in IR$_\mathsf{T}$-calc maintaining the following invariants: after the i^{th} step, all the instantiated clauses up to level i are fully annotated and the instantiating assignments are complete. Thus the instantiation steps become redundant. This further implies that after the last level (when we reach the axiom farthest from the top), the resulting proof is in fact a $\forall$Exp+Res$_\mathsf{T}$ proof.

- **At level** 0**:** The node at this level must be a resolution producing the empty clause, followed by a dummy instantiation with the empty assignment. Thus the clauses at this level are already fully annotated, but the instantiating assignment is far from complete. Pick an arbitrary complete assignment, say σ, and instantiate the empty clause with σ. Clearly the invariants hold now.
- Assume that the invariants holds after processing all nodes at level $i-1$.
- **At level** i**:** Let D be an instantiated clause at level $i-1$, obtained by instantiating some clause C by an assignment σ. That is, $D = \mathsf{inst}(C, \sigma)$. By the induction hypothesis, D is fully annotated and σ is complete. Let C be obtained by resolving $E = (G \vee x^\tau)$ and $F = (H \vee \neg x^\tau)$. We need to make E and F fully annotated. Let $E = \mathsf{inst}(I, \beta_1)$ and $F = \mathsf{inst}(J, \beta_2)$ in π. Replace E by $E' = \mathsf{inst}(I, \beta_1 \veebar \sigma)$ and F by $F' = \mathsf{inst}(J, \beta_2 \veebar \sigma)$. As σ is complete, both $\beta_1 \veebar \sigma$ and $\beta_2 \veebar \sigma$ are complete, and hence both E' and F' are fully annotated. The resolution step is now performed on $x^{\tau'}$, where $\tau' = \tau \veebar \sigma$ is the resulting annotation on x. It is easy to see that the resolvent of E' and F' is D, so the intermediate instantiation step going from C to D becomes redundant.

It is clear that the simulation preserves width. It also does not increase size: we may introduce dummy instantiation nodes to make the proof 'alternating', but after the transformation, all instantiations — dummy and actual — are eliminated completely. It is also clear that the simulation preserves the space needed, since the structure of the proof is preserved. $\square$

The simulation in Lemma 12 exhibits an interesting phenomenon: while it shows that the tree-like versions of $\forall$Exp+Res and IR-calc are p-equivalent, it was shown in [10] that in the dag-like versions, IR-calc is exponentially stronger than $\forall$Exp+Res. Thus $\forall$Exp+Res and IR-calc provide a rare example in proof complexity of two systems that coincide in the tree-like model, but are separated in the dag-like model.

Lemma 13. *$\mathsf{IR_T\text{-}calc}$ p-simulates $\mathsf{Q\text{-}Res_T}$ while preserving space and existential width exactly and size upto a factor of 3. That is, $S(\vdash_{\mathsf{IR_T\text{-}calc}} \mathcal{F}) \leq 3S(\vdash_{\mathsf{Q\text{-}Res_T}} \mathcal{F})$, $Space(\vdash_{\mathsf{IR_T\text{-}calc}} \mathcal{F}) \leq Space(\vdash_{\mathsf{Q\text{-}Res_T}} \mathcal{F})$, and $w(\vdash_{\mathsf{IR\text{-}calc}} \mathcal{F}) \leq w_\exists(\vdash_{\mathsf{Q\text{-}Res}} \mathcal{F})$.*

Proof. We use the same simulation as given in [9]. This simulation was originally for dag-like proof systems, but here we check that it also works for tree-like systems, and we observe that space and existential width are preserved.

Let $C_1, \ldots, C_k$ be a Q-Res$_\mathsf{T}$ proof. We translate the clauses into clauses $D_1, \ldots, D_k$, which will form the skeleton of a proof in IR-calc.

- For an axiom C_i in Q-Res$_\mathsf{T}$ we introduce the same clause D_i by the axiom rule of IR-calc, i.e., we remove all universal variables and add annotations.
- If C_i is obtained via $\forall$-reduction from C_j, then $D_i = D_j$; we make no change.
- Consider now the case that C_i is derived by resolving C_j and C_k with pivot variable x. Then $D_j = x^\tau \vee K_j$ and $D_k = \bar{x}^\sigma \vee K_k$. It is shown in [9] that the annotations τ and σ are not contradictory; in fact, no annotations in the two clauses are contradictory. So if we define $D'_j = \mathsf{inst}(\sigma, D_j)$ and $D'_k = \mathsf{inst}(\tau, D_k)$, then the annotations of x in D'_j and $\bar{x}$ in D'_k match, and we can resolve on this literal. Define D'_i as the resolvent of D'_j and D'_k. We can perform a further instantiation to obtain $D_i = \mathsf{inst}(\eta, D_i)$, where η is the set of all assignments to universal variables appearing anywhere in D'_i. D_i has no more literals than C_i. For details, see [9].

Note that to complete this skeleton into a proof, we only add instantiation rules. Thus, if the original proof was tree-like, so is the new proof. If the original proof has size S, the new proof has size at most $4S$, since each resolution may now be preceded by two instantiations and followed by one instantiation. However, this is an overcount, since we are counting two instantiations per edge, one from the parent and one from the child, and contiguous instantiations can be collapsed. That is, every instantiation following a resolution step can be merged with the instantiation preceding the next resolution and need not be counted separately. The only exception is at the root, where there is nothing to collapse it with. However, at the root, the instantiation itself is redundant and can be discarded. Thus we obtain a new proof of size at most $3S$.

Further, if the original proof had existential width w, then the new proof has width w since each D_i has at most (annotated versions of) the existential literals of C_i.

Regarding space, observe that simulating axiom download and $\forall$-Red do not require additional space. At the resolution step, the simulation first performs additional instantiations. But instantiation does not need additional space. So the space bound remains the same. ☐

As a by-product, these simulations enable us to give an easy and elementary proof of the simulation of Q-Res$_\mathsf{T}$ by $\forall$Exp+Res, shown in [19] via a more involved argument. In fact, our result improves upon the simulation from [19] as we show that even *tree-like* $\forall$Exp+Res suffices to p-simulate Q-Res$_\mathsf{T}$.

Corollary 14 (Janota, Marques-Silva [19]). *$\forall$Exp+Res$_\mathsf{T}$ p-simulates Q-Res$_\mathsf{T}$.*

Proof. By Lemma 12, $\forall$Exp+Res$_\mathsf{T}$ p-simulates IR$_\mathsf{T}$-calc, which in turn p-simulates Q-Res$_\mathsf{T}$ by Lemma 13. $\square$

Using again the width lower bound for $\mathrm{QP{\scriptsize ARITY}}_n$ (Theorem 9) we can show that item 2 of Theorem 10 cannot be improved, i.e. we obtain an optimal width separation between Q-Res and $\forall$Exp+Res.

Theorem 15. *There exist false QBFs ψ_n with $w_\exists(\vdash_{\textsf{Q-Res}} \psi_n) = \Omega(n)$, but $w(\vdash_{\forall\textsf{Exp+Res}} \psi_n) = O(1)$.*

Proof. We use the $\mathrm{QP{\scriptsize ARITY}}_n$ formulas, which by Theorem 9 require existential width n in Q-Res. To get the separation it remains to show $w(\vdash_{\forall\textsf{Exp+Res}} \mathrm{QP{\scriptsize ARITY}}_n) = O(1)$. For this we use the following $\forall$Exp+Res proofs of $\mathrm{QP{\scriptsize ARITY}}_n$ from [10]: the formulas $\mathrm{QP{\scriptsize ARITY}}_n$ have exactly one universal variable z, which we expand in both polarities 0 and 1. This does not affect the x_i variables, but creates different copies $t_i^{z/0}$ and $t_i^{z/1}$ of the existential variables right of z. Using the clauses of $\mathrm{xor}(t_{i-1}, x_i, t_i)$, we can inductively derive clauses representing $t_i^{z/0} = t_i^{z/1}$. This lets us derive a contradiction using the clauses $t_n^{z/0}$ and $\neg t_n^{z/1}$.

Clearly, this proof only contains clauses of constant width, giving the result. $\square$

6 Positive results: Size, width, and space in tree-like QBF calculi

We are now in a position to show some positive results on size-width and size-space relations for QBF resolution calculi. However, most of these results only apply to rather weak tree-like proof systems.

6.1 Relations in the expansion calculi $\forall$Exp+Res and IR-calc

We first observe that for $\forall$Exp+Res almost the full spectrum of relations from classical resolution remains valid.

Theorem 16. *For all false QBFs $\mathcal{F}$, the following relations hold:*

1. $S(\vdash_{\forall\textsf{Exp+Res}_\mathsf{T}} \mathcal{F}) \geq 2^{(w(\vdash_{\forall\textsf{Exp+Res}} \mathcal{F})) - w_\exists(\mathcal{F}))}$.
2. $S(\vdash_{\forall\textsf{Exp+Res}_\mathsf{T}} \mathcal{F}) \geq 2^{Space(\vdash_{\forall\textsf{Exp+Res}_\mathsf{T}} \mathcal{F})} - 1$.
3. $Space(\vdash_{\forall\textsf{Exp+Res}_\mathsf{T}} \mathcal{F}) \geq Space(\vdash_{\forall\textsf{Exp+Res}} \mathcal{F}) \geq w(\vdash_{\forall\textsf{Exp+Res}} \mathcal{F}) - w_\exists(\mathcal{F}) + 1$.

Proof. This theorem follows from the analogous statements for classical resolution. We just describe how to apply those results to $\forall$Exp+Res.

We know that in $\forall$Exp+Res$_\mathsf{T}$ proofs, leaves corresponds to the expanded clauses from $\mathcal{F}$. The expanded clauses contain only existential (annotated) literals and no universal literals. Let $\mathcal{G}$ be the QBF obtained after expanding $\mathcal{F}$ based on all possible assignments of universal variables. Clearly, $\mathcal{G}$ contains no universal variables and hence can be treated as a propositional CNF formula (all variables are only existentially quantified). That is, if G is the matrix of clauses in $\mathcal{G}$, then $\mathcal{G}$ asserts that G is satisfiable. Also, $w(G) = w(\mathcal{G}) = w_\exists(\mathcal{F})$.

Refutations of $\mathcal{F}$ in $\forall$Exp+Res (respectively, $\forall$Exp+Res$_\mathsf{T}$) are precisely refutations (resp. tree-like refutations) of G in classical resolution; the size, space and width are exactly the same, by definition. That is, $S(\vdash_{\textsf{Res}_\mathsf{T}} G) = S(\vdash_{\forall\textsf{Exp+Res}_\mathsf{T}} \mathcal{F})$, $w(\vdash_{\textsf{Res}} G) = w(\vdash_{\forall\textsf{Exp+Res}} \mathcal{F})$, $Space(\vdash_{\textsf{Res}} G) = Space(\vdash_{\forall\textsf{Exp+Res}} \mathcal{F})$, and $Space(\vdash_{\textsf{Res}_\mathsf{T}} G) = Space(\vdash_{\forall\textsf{Exp+Res}_\mathsf{T}} \mathcal{F})$. Now the Theorem follows by applying Theorems 2, 3, and 4, on G. $\square$

By the equivalence of $\forall$Exp+Res$_\mathsf{T}$ and IR$_\mathsf{T}$-calc with respect to all the three measures size, width, and space (Theorem 10) we can immediately transfer all results from Theorem 16 to IR$_\mathsf{T}$-calc.

Theorem 17. *For all false QBFs $\mathcal{F}$, the following relations hold:*

1. $S(\vdash_{\textsf{IR}_\mathsf{T}\textsf{-calc}} \mathcal{F}) \geq 2^{(w(\vdash_{\textsf{IR-calc}} \mathcal{F})) - w_\exists(\mathcal{F}))}$.
2. $S(\vdash_{\textsf{IR}_\mathsf{T}\textsf{-calc}} \mathcal{F}) \geq 2^{Space(\vdash_{\textsf{IR}_\mathsf{T}\textsf{-calc}} \mathcal{F})} - 1$.
3. $Space(\vdash_{\textsf{IR}_\mathsf{T}\textsf{-calc}} \mathcal{F}) \geq w(\vdash_{\textsf{IR-calc}} \mathcal{F}) - w_\exists(\mathcal{F}) + 1$.

6.2 The size-space relation in tree-like Q-resolution

We finally return to Q-Res. Most relations were already ruled out in Section 4 for both Q-Res and Q-Res$_\mathsf{T}$. The only relation that we can still show to hold is the classical size-space relation (Theorem 3), which we transfer from Res$_\mathsf{T}$ to Q-Res$_\mathsf{T}$.

In classical resolution, this relationship was obtained using pebbling games [17]. We observe that the same holds for Q-Res$_\mathsf{T}$ as well, giving the analogous relationship. That is, we show:

Theorem 18. *For a false QBF sentence $\mathcal{F}$,*

$$S(\vdash_{\textsf{Q-Res}_T} \mathcal{F}) \geq 2^{Space\left(\vdash_{\textsf{Q-Res}_T} \mathcal{F}\right)} - 1.$$

Before getting into the proof, we describe the pebbling game.

Definition 19. *(Pebbling Game) Let $G = (V, E)$ be a connected directed acyclic graph with a unique sink s, where every vertex of G has fan-in at most 2. The aim of the game is to put a pebble on the sink of the graph following this set of rules:*

1. *A pebble can be placed on any source vertex, that is, on a vertex with no predecessors.*
2. *A pebble can be removed from any vertex.*
3. *A pebble can be placed on an internal vertex provided all of its children are pebbled. In this case, instead of placing a new pebble on it, one can shift a pebble from a child to the vertex.*

The minimum number of pebbles needed to pebble the unique sink following the above rules is said to be the **pebbling number** *of G.*

Consider the proof graph G_π corresponding to a Q-Res proof π of a false QBF $\mathcal{F}$. In G_π clauses are the vertices and edges go from the hypotheses to the conclusion of inference rules (i.e, $\forall$-Red, resolution steps). Clearly G_π is a DAG with initial clauses as sources and the empty clause as the unique sink. Also the in-degree of each vertex in G_π is at most 2. Hence the pebbling game is well defined in G_π.

We now show that the space required to refute a false QBF sentence $\mathcal{F}$ (as per Definition 1) coincides with the minimum number of pebbles needed to play the pebble game on the graph of a Q-Res proof of $\mathcal{F}$. The relation holds for tree-like proofs as well.

Lemma 20. *Let $\mathcal{F}$ be a false QBF in prenex form. Then the following holds:*

1. $Space(\vdash_{\textsf{Q-Res}} \mathcal{F}) = \min\{k : \exists$ *Q-Res proof π of $\mathcal{F}$, G_π can be pebbled with k pebbles*$\}$;

2. Space$(\vdash_{\text{Q-Res}_T} \mathcal{F}) = \min\{k : \exists$ *Q-Res$_T$ proof π of $\mathcal{F}$, G_π can be pebbled with k pebbles*$\}$.

Proof (Sketch). The proof is exactly the same as in classical resolution.

Let π be a Q-Res proof whose proof graph G_π can be pebbled with k pebbles. (If π is treelike, then G_π is a tree.) Note that the vertices of G_π are clauses in the proof. The space-oriented Q-Res (respectively Q-Res$_T$) proof sequence with clause space k is constructed by maintaining at each stage exactly the pebbled clauses. By the rules of the pebbling game, adding a clause to the current set is valid because the added clause is either at a source node and hence an axiom, or it has all predecessors pebbled and hence can be inferred. Further, if π is tree-like, then it can be shown that there is a k-pebble sequence where no node is pebbled more than once (once a node is pebbled, no predecessor of the node need be pebbled again). So the above construction will yield a tree-like space-k proof sequence.

In the other direction, given a space-k proof as a sequence σ, we can construct a corresponding DAG G with nodes for each clause appearing anywhere in σ, and edges reflecting how the clauses are used for inference in σ. Thus we obtain a proof π with $G_\pi = G$ (it is the same proof as σ, just represented differently). We can pebble G with k pebbles by maintaining the invariant that at each stage, pebbles are placed on exactly the clauses present in the corresponding formula in the sequence σ. If σ is a tree-like space-k proof, we construct a corresponding tree with a distinct node for every copy of a clause introduced at some stage in σ, and then pebble it as above. We omit the details. ◻

We can now prove Theorem 18.

Proof (of Theorem 18). This proof too is almost identical to the proof for classical resolution [17]. We give a brief sketch.

Let $S(\vdash_{\text{Q-Res}_T} \mathcal{F}) = s$. Consider a tree-like Q-Res$_T$ proof π of $\mathcal{F}$ (i.e, $\pi \vdash_{\text{Q-Res}_T} \mathcal{F}$), of size s, and let T be the underlying proof-tree.

In contrast to classical resolution, a proof graph in Q-Res may have unary nodes corresponding to $\forall$-reductions. In particular, for a proof in Q-Res$_T$, there may be paths corresponding to series of $\forall$-reductions. Once the lower end of such a path is pebbled, the same pebble can be slid up to the top of the path; no additional pebbles are needed. So without loss of generality we work with the tree T' obtained by shortcutting all paths containing unary nodes.

Let $d_c(T)$ be the depth of the biggest complete binary tree that can be embedded in T' or in T. (We say that a graph G_1 is embeddable in a graph G_2 if a graph isomorphic to G_2 can be obtained from G_1 by adding vertices and edges or subdividing edges of G_1.) Clearly, $2^{d_c(T)+1} - 1 \leq s$.

By induction on $|T'|$, we can show that $d_c(T) + 1$ pebbles suffice to pebble T'. Hence, by the argument given above, $d_c(T)+1$ pebbles suffice to pebble T as well. Now, using Lemma 20, we obtain $\mathit{Space}(\vdash_{\text{Q-Res}_T} \mathcal{F}) \leq d_c(T) + 1$. Hence

$$2^{\mathit{Space}(\vdash_{\text{Q-Res}_T} \mathcal{F})} - 1 \leq 2^{d_c(T)+1} - 1 \leq s = S(\vdash_{\text{Q-Res}_T} \mathcal{F})$$

as claimed. ◻

7 Conclusion

Our results show that the success story of width in resolution needs to be rethought when moving to QBF. Indeed, the question arises: is width a central parameter in QBF resolution?

Is there another parameter that plays a similar role as classical width for understanding QBF resolution size and space?

Our findings almost completely uncover the picture for size, space, and width for the most basic and arguably most important QBF resolution systems Q-Res, ∀Exp+Res, and IR-calc. The most immediate open question arising from our investigation is whether size-width relations hold for general dag-like ∀Exp+Res or IR-calc proofs. The issue here is that in the classical size-width relation of [6] the number of variables enters the formula in a crucial way. For the instantiation calculi it is not clear what should qualify as the right count for this as different annotations of the same existential variable are formally treated as distinct variables (which is also clearly justified by the semantic meaning of expansions).

For further research it will also be interesting whether size-width or space-width relations apply to any of the stronger QBF resolution systems QU-Res [27], LD-Q-Res [2], or IRM-calc [9]. However, we conjecture that the negative picture also prevails for these systems.

Acknowledgements

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

References

1. Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, 2008.
2. Valeriy Balabanov and Jie-Hong R. Jiang. Unified QBF certification and its applications. *Form. Methods Syst. Des.*, 41(1):45–65, August 2012.
3. Valeriy Balabanov, Magdalena Widl, and Jie-Hong R. Jiang. QBF resolution systems and their proof complexities. In *SAT*, pages 154–169, 2014.
4. Paul Beame, Christopher Beck, and Russell Impagliazzo. Time-space tradeoffs in resolution: superpolynomial lower bounds for superlinear space. In *STOC*, pages 213–232, 2012.
5. Eli Ben-Sasson and Jakob Nordström. Understanding space in proof complexity: Separations and trade-offs via substitutions. In *ICS*, pages 401–416, 2011.
6. Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow - resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001.
7. Marco Benedetti and Hratch Mangassarian. QBF-based formal verification: Experience and perspectives. *JSAT*, 5(1-4):133–191, 2008.
8. Olaf Beyersdorff, Ilario Bonacina, and Leroy Chew. Lower bounds: from circuits to QBF proof systems. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:133, 2015.
9. Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. On unification of QBF resolution-based calculi. In *MFCS, II*, pages 81–93, 2014.
10. Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. Proof complexity of resolution-based QBF calculi. In *Proc. Symposium on Theoretical Aspects of Computer Science*, pages 76–89. LIPIcs series, 2015.
11. Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Feasible interpolation for QBF resolution calculi. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015*, pages 180–192, 2015.
12. Olaf Beyersdorff, Leroy Chew, and Karteek Sreenivasaiah. A game characterisation of tree-like Q-resolution size. In *Proc. 9th International Conference on Language and Automata Theory and Applications*, volume 8977 of *Lecture Notes in Computer Science*, pages 486–498. Springer-Verlag, Berlin Heidelberg, 2015.
13. Olaf Beyersdorff and Oliver Kullmann. Unified characterisations of resolution hardness measures. In *SAT*, pages 170–187, 2014.
14. Archie Blake. *Canonical expressions in boolean algebra*. PhD thesis, University of Chicago, 1937.
15. Maria Luisa Bonet and Nicola Galesi. A study of proof search algorithms for resolution and polynomial calculus. *In Proceedings 40th Annual Symposium on Foundations of Computer Science, IEEE*, 1999.

16. Uwe Egly, Martin Kronegger, Florian Lonsing, and Andreas Pfandler. Conformant planning as a case study of incremental QBF solving. In *Artificial Intelligence and Symbolic Computation AISC 2014*, pages 120–131, 2014.
17. Juan Luis Esteban and Jacobo Torán. Space bounds for resolution. *Information and Computation*, 171(1):84–97, 2001.
18. Yuval Filmus, Massimo Lauria, Mladen Miksa, Jakob Nordström, and Marc Vinyals. From small space to small width in resolution. In *31st International Symposium on Theoretical Aspects of Computer Science, STACS 2014*, pages 300–311, 2014.
19. Mikolás Janota and Joao Marques-Silva. Expansion-based QBF solving versus Q-resolution. *Theor. Comput. Sci.*, 577:25–42, 2015.
20. Hans Kleine Büning, Marek Karpinski, and Andreas Flögel. Resolution for quantified Boolean formulas. *Inf. Comput.*, 117(1):12–18, 1995.
21. Jan Krajíček. Interpolation theorems, lower bounds for proof systems and independence results for bounded arithmetic. *The Journal of Symbolic Logic*, 62(2):457–486, 1997.
22. Oliver Kullmann. Investigating a general hierarchy of polynomially decidable classes of CNF's based on short tree-like resolution proofs. *Electronic Colloquium on Computational Complexity (ECCC)*, 99(41), 1999.
23. M. Mahajan and A. Shukla. Level-ordered Q-resolution and tree-like Q-resolution are incomparable. submitted manuscript, 2015.
24. Jakob Nordström. Pebble games, proof complexity, and time-space trade-offs. *Logical Methods in Computer Science*, 9(3), 2013.
25. Jussi Rintanen. Asymptotically optimal encodings of conformant planning in QBF. In *AAAI*, pages 1045–1050. AAAI Press, 2007.
26. John Alan Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12:23–41, 1965.
27. Allen Van Gelder. Contributions to the theory of practical quantified Boolean formula solving. In *CP*, pages 647–663, 2012.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de