Electronic Colloquium on Computational Complexity, Report No. 113 (2013)

Revisiting Space in Proof Complexity: Treewidth and Pathwidth

Moritz Müller
Kurt Gödel Research Center
University of Vienna, Vienna, Austria

Stefan Szeider
Institute of Information Systems
Vienna University of Technology, Vienna, Austria

Abstract

So-called *ordered* variants of the classical notions of pathwidth and treewidth are introduced and proposed as proof theoretically meaningful complexity measures for the directed acyclic graphs underlying proofs. The ordered pathwidth of a proof is shown to be roughly the same as its formula space. Length-space lower bounds for $R(k)$-refutations are generalized to arbitrary infinity axioms and strengthened in that the space measure is relaxed to ordered treewidth.

1 Introduction

Proof complexity seeks to show that certain propositional contradictions do not admit short refutations in certain propositional refutation systems; here, short means polynomial in the size of the contradiction refuted. Of special interest are Resolution-based refutation systems and meaningful contradictions expressing combinatorial principles in some natural way. Common instances of the latter are given by propositional translations of first-order formulas, and in particular of *infinity axioms*[1]. By a Resolution-based system we mean one of the hierarchy $R(1), R(2), \ldots, R(\log)$ [15]; $R(1)$ is the same as Resolution and $R(k)$ is a straightforward generalization of Resolution operating with k-DNFs instead of clauses, i.e., cutting on conjunctions of k literals instead of single literals. From a practical perspective this special interest derives from the fact that SAT solvers are based on such systems. From a more theoretical perspective the special interest derives from the fact that lower bounds for these systems are prerequisite for understanding independence from bounded arithmetic (cf. [4]).

Besides *proof length* the most popular complexity measure of proofs is *proof space* (formula-space or clause-space) as introduced by Esteban and Toran [11]. Intuitively, a space 100 refutation of a set Γ of clauses, say in Resolution, is one that can be presented as follows.

> A teacher is in class equipped with a blackboard containing up to 100 clauses. The teacher starts from the empty blackboard and finally arrives at one containing the empty clause. The blackboard can be altered by either writing down a clause from Γ,

[1]Definitions are given in the following section.

ISSN 1433-8092

or by wiping out some clause, or by deriving a new clause from clauses currently written on the blackboard by means of the Resolution rule.

A motivation for studying the space of refutations is to understand memory requirements for SAT solvers [6]. The hierarchy of Resolution-based proof systems is not only strict with respect to length (see [23] for a survey) but it is also strict with respect to space [10, 6]. A sequence of work established lower bounds on space for $R(k)$ refutations, and especially for (translations of) infinity axioms [11, 1, 10]. Resolution lower bounds on space follow from lower bounds on width [11, 2] but not the other way around [19].

Ben-Sasson showed that size and space cannot in general be simultaneously optimized [5], laying the ground for so-called *length-space trade-offs*. An exponential length-space trade-off states that there exists a sequence of contradictions that have short Resolution refutations in small space while refutations in somewhat smaller space require exponential length (length-space lower bound). Ben-Sasson and Nordström found such sequences for various settings for the qualifications "small" and "somewhat smaller", e.g., for $O(n)$ versus $o(n/\log n)$. Moreover, they managed to extend the length-space lower bound to $R(k)$ for constant k when taking the $(k+1)$th root of the qualification "somewhat smaller." The contradictions are substitution instances of pebbling contradictions. What Ben-Sasson and Nordström showed is how to transfer trade-off results for pebbling games to Resolution proofs. We refer to the survey [6] for more information. The wording trade-off has to be taken with some care in that the upper bounds are claimed only for the very special contradictions constructed. In this paper we shall focus on the lower bound part of trade-offs.

This paper. We revisit refutation space by means of natural invariants of the refutation DAG, namely, we introduce so-called *ordered* variants of the notions of *pathwidth* and *treewidth*. These notions play an important role in Robertson and Seymour's graph minors project and have evolved as very successful and ubiquitously used complexity measures (see, for instance, Bodlaender's survey [7]). In contrast to earlier adaptions of the width notions to digraphs [3], the ordered width measures allow us to distinguish between DAGs. Our notions are well-motivated from a graph theoretic point of view; for example on DAGs, ordered pathwidth coincides with a straightforward variant of the vertex separation number [7] adapted to DAGs (Proposition 3.17). We show that the notions have proof theoretic sense: Resolution refutations of minimal ordered pathwidth are just Input Resolution refutations (Theorem 4.1), and those of minimal ordered treewidth are just the treelike ones. More importantly, we show that ordered pathwidth is roughly the same as refutation space (Theorem 5.1). Conceptually, these results allow to rethink space as a measure of how far a Resolution proof is from being an Input Resolution refutation.

This gives interest to *ordered treewidth*, a notion that relaxes ordered pathwidth in much the same way as treewidth relaxes pathwidth. Ordered treewidth of a refutation can be interpreted as measuring how far a refutation is from being treelike (Theorem 4.1). We also propose an interpretation of ordered treewidth in terms of space, using the following two player game, that continues the metaphor above.

A student visits the teacher in her office asking her to explain the proof. The teacher has a blackboard potentially containing up to 10 clauses and writes the empty clause on it. The student asks how to prove it. The teacher produces a length ≤ 10 proof from Γ plus some additional clauses. The student chooses one of these additional clauses and asks how to prove it. And so on. The game ends when the teacher comes up with a proof using no additional clauses.

The new graph invariants also provide the means for making progress with respect to the already mentioned *length-space lower bounds* from Ben-Sasson and Nordström [6]. Our main technical result (Theorem 6.1) is a lower bound on length and ordered treewidth for $R(k)$-refutations of infinity axioms in general. This makes progress with respect to the known length-space lower bounds in that it applies to infinity axioms in general, and thereby to a large class of formulas having a natural meaning. It relaxes the refutation space measure (i.e., ordered pathwidth) to ordered treewidth, and it gives nontrivial lower bounds for all $R(k)$ simultaneously, and for $R(\log)$. The latter feature overcomes a bottleneck in constructions from [6] which give good lower bounds for $R(k)$ with constant k but become trivial for $R(\log)$.

We also get strong lower bounds on space, i.e., ordered pathwidth, for infinity axioms in general (Corollary 6.3). While it is known that treelike $R(\log)$ refutations of infinity axioms need exponential length, there are some very few examples of infinity axioms known to have short DAG-like refutations, even in Resolution. We show that such short refutations need to be far from being treelike in the sense that they require large ordered treewidth.

Proof idea. The lower bound proof follows the adversary type argument of [16] against treelike $R(\log)$ refutations of infinity axioms. One uses restrictions that describe finite parts of some infinite model of the infinity axiom. Starting with the empty restriction, first choose a node as in Spira's theorem, namely one that splits the refutation tree into two subtrees of size at most 2/3 of total. Then distinguish two cases, namely whether no extension of the current restriction satisfies the formula at the chosen node or not. In the first case stick with the current restriction and recurse to the subtree rooted at the chosen node. In the second case delete this subtree and recurse with a "small" restriction satisfying the formula at the chosen node. The invariant maintained is a proof of a formula "forced false" from axioms plus some formulas "forced true." If the proof has length S, this process reaches a constant size proof after $O(\log S)$ steps. If S is not too large, it is argued that the final restriction can be further extended to force all remaining axioms true and a contradiction is reached. The proof of our lower bound proceeds similarly but by recursion on a tree decomposition of the refutation. To make sense of this idea we show that we can always find a tree decomposition whose underlying tree is binary (to find a Spira type split node) and whose size is linear in the size of the refutation (Lemma 3.12). Further care is needed to ensure that the partial tree decompositions during the recursion are decompositions of refutations with similar properties as the invariant described above (Lemma 3.8).

2 Preliminaries

2.1 Digraphs

We consider *directed graphs* (*digraphs*, for short) without self-loops and denote the set of vertices and the set of directed edges of a digraph D by $V(D)$ and $E(D)$, respectively. If $(u, v) \in E(D)$, then u is a *predecessor* of v and v a *successor* of u. An *ancestor* of $v \in V(D)$ is a vertex w such that there is a directed path from w to v in D; we understand that there is a directed path from any vertex to itself. The *in-degree (out-degree) of* v is the number of its predecessors (successors). The in-degree (out-degree) of D is the maximal in-degree (out-degree) over all vertices. Vertices of in-degree 0 are *sources*, vertices of out-degree 0 are *sinks*. An *(induced) subdigraph* of D is a digraph $D[X]$ induced on a nonempty $X \subseteq V(D)$; if $V(D) \setminus X$ is nonempty, we write $D - X$ for $D[V(D) \setminus X]$. The graph $\underline{D}$ *underlying* a digraph D has the same vertices as D and as edges the symmetric closure of $E(D)$. In general, a *graph* is a digraph D with symmetric $E(D)$. A *DAG* is a directed acyclic graph (i.e., a digraph without directed cycles), and a *tree* is a DAG T with a unique sink r_T called *root* such that for every $v \in V(T)$ there is exactly one directed path from v to r_T. We shall refer to vertices in a tree as *nodes*. The *subtree* T_t *rooted at* $t \in V(T)$ is the subtree of T induced on the set of ancestors of t in T; it has root $r_{T_t} = t$. The *height* of a tree is the maximal length (number of edges) in a *branch* (leaf-to-root path) in T. By the *perfect binary tree* B_h of height h we mean the tree where every node which is not a leaf has exactly two predecessors and all branches have length exactly h.

2.2 Propositional Logic

A *literal* is a propositional variable X or its negation $\neg X$; for a literal ℓ we let $\neg\ell$ denote $\neg X$, if $\ell = X$, and X, if $\ell = \neg X$. A *(k-)term* is a set of (at most k) literals. A *(k-)DNF* is a set of (k-)terms. The empty DNF is denoted by 0 and the empty term by 1. A *clause* is a 1-DNF. An *assignment* is a function from the propositional variables into $\{0, 1\}$. A *restriction* ρ is a finite partial assignment. For a restriction or assignment ρ and a term t we let $t \upharpoonright \rho$ be 0 if t contains a literal falsified by ρ (in the usual sense) and otherwise the subterm obtained by deleting all literals satisfied by ρ. For a DNF D we let $D \upharpoonright \rho$ be 1 if $t \upharpoonright \rho = 1$ for some $t \in D$; otherwise $D \upharpoonright \rho$ is obtained from D by deleting all $t \in D$ with $t \upharpoonright \rho = 0$ and deleting from every other $t \in D$ all literals satisfied by ρ. Note, if ρ is defined on all variables appearing in D then $D \upharpoonright \rho$ equals the truth value of D under ρ.

Definition 2.1. A *(k-)DNF proof* is a pair $(D, (F_v)_{v \in V(D)})$ where D is a DAG with a unique sink and in which every vertex has at most two predecessors, and F_v is a (k-)DNF for every $v \in V(D)$. The proof is said to be *of* F if $F = F_v$ for v the sink of D, and *from* Γ if $F_v \in \Gamma$ for all sources v of D. It is said to be *treelike* if D is a tree. Proofs of 0 are *refutations*. The *length* of the proof is $|V(D)|$. A *refutation system* is a set of refutations.

Usually one requires refutation systems to satisfy certain further properties like soundness or completeness or being polynomial time decidable (cf. [8]).

Definition 2.2. A DNF proof $(D, (F_v)_{v \in V(D)})$ is *sound* if for every inner vertex $v \in V(D)$ and every assignment ρ we have $F_v \restriction \rho = 1$ whenever $F_u \restriction \rho = 1$ for all predecessors u of v in D. It is *strongly sound* (cf. [24]) if for every inner vertex $v \in V(D)$ and every restriction ρ we have $F_v \restriction \rho = 1$ whenever $F_u \restriction \rho = 1$ for all predecessors u of v in D.

The next statement is obvious.

Lemma 2.3. *If there is a strongly sound proof of F from Γ and ρ is a restriction such that $G \restriction \rho = 1$ for all $G \in \Gamma$, then $F \restriction \rho = 1$.*

We consider the following rules of inference, namely *weakening*, *introduction of conjunction* and *cut*:

$$\frac{D}{D \cup \{t\}} \qquad \frac{D \cup \{t\} \quad D' \cup \{t'\}}{D \cup D' \cup \{t \cup t'\}} \qquad \frac{D \cup \{t\} \quad D' \cup D''}{D \cup D'},$$

where D, D', D'' are DNFs, t, t' are terms and in the cut rule we assume $\emptyset \neq D'' \subseteq \{\{\neg \ell\} \mid \ell \in t\}$. A k-DNF proof $(D, (F_v)_{v \in V(D)})$ is an $R(k)$-proof if for every inner vertex v with predecessors u, w the formula F_v is obtained from F_u and F_w by one of the three rules above. An $R(k)$-proof is an $R(\log)$*-proof* if its length is at least 2^k. An $R(1)$-proof is a *Resolution* proof. The refutation system consisting of all $R(k)$-refutations (resp. $R(\log)$-refutations) is denoted $R(k)$ (resp. $R(\log)$).

Remark 2.4. $R(k)$ is strongly sound. We have completeness in the sense that for every k-DNF F implied by some set Γ of k-DNFs, there is an $R(k)$-proof of F from Γ plus some additional 'axioms' of the form $(X \vee \neg X)$, i.e., $\{\{X\}, \{\neg X\}\}$. $R(k)$ is refutation-complete in the sense that no such axioms are needed in case $F = 0$. If one adds a new rule allowing to infer such an axiom from any formula, then the system ceases to be strongly sound.

2.3 First-Order Logic and Propositional Translation

A *vocabulary* is a finite set τ of relation and function symbols, each with an associated *arity*; function symbols of arity 0 are *constants*. The *arity* of τ is the maximum arity of one of its symbols. τ*-terms* are first-order variables $x, y, z \dots$ or of the form $ft_1 \cdots t_r$ where $t_1, \dots, t_r$ are again τ-terms and $f \in \tau$ is a function symbol of arity r. τ*-atoms* are of the form $t_1 = t_2$ or $Rt_1 \cdots t_r$ where $t_1, t_2, \dots, t_r$ are τ-terms and $R \in \tau$ is a relation symbol of arity r. τ*-formulas* are built from τ-atoms using $\wedge, \vee, \neg$ and existential and universal quantifiers $\exists x, \forall x$. For a tuple of first-order variables $\bar{x}$ we write $\varphi(\bar{x})$ for a τ-formula φ to indicate that the free variables of φ are among the components of $\bar{x}$. A τ*-sentence* is a τ-formula without free variables. A τ*-structure* M consists of a nonempty set, its *universe*, that we also denote by M and for every, say, r-ary relation symbol $R \in \tau$ an *interpretation* $R^M \subseteq M^r$, and for every, say, r-ary function symbol $f \in \tau$ an *interpretation* $f^M : M^r \to M$; we identify the interpretation of a constant with its unique value. A τ-structure M is a *model of* a τ-sentence φ if φ is true in M.

The *spectrum* of a first-order sentence φ is the set of those naturals $n \geq 1$ such that φ has a model (with universe) of cardinality n. An *infinity axiom* is a satisfiable first-order sentence with empty spectrum, i.e., a sentence without a finite but with an infinite model. Skolemization and elementary formula manipulation allows to compute from every first-order sentence ψ a sentence φ with the same spectrum and of the form

$$\forall \bar{x} \bigwedge_{i \in I} C_i(\bar{x}), \tag{1}$$

where I is a nonempty finite set, the C_is are first-order clauses (disjunctions of atoms and negated atoms) whose atoms have the form $R\bar{y}$ or $f\bar{y} = z$ for some relation symbol R respectively function symbol f and variables $\bar{y}, z$. Moreover, φ has an infinite model if and only if ψ does.

Following Paris and Wilkie (cf. [20], see also [22]) we define for every natural $n \geq 1$ a set $\langle \varphi \rangle_n$ of clauses that is satisfied exactly by those assignments that describe a model of φ with universe $[n] := \{0, 1 \ldots, n-1\}$.

Let τ be the vocabulary of φ. We use as propositional variables $R\bar{a}, f\bar{a} = b$ where $r \in \mathbb{N}, \bar{a} \in [n]^r, b \in [n]$, R is an r-ary relation symbol in τ and f is an r-ary function symbol in τ. For $i \in I$ and $\bar{a} \in [n]^{|\bar{x}|}$ substitute $\bar{a}$ for $\bar{x}$ in $C_i(\bar{x})$; this transforms every literal into a propositional literal or into an expression of the form $a = a'$ or $\neg a = a'$ where a, a' are components of $\bar{a}$; the propositional clause $\langle C_i(\bar{a}) \rangle$ is $\{1\}$ if one of these expressions is "true" in the obvious sense; otherwise $\langle C_i(\bar{a}) \rangle$ is the clause whose terms are the singletons of the propositional literals (of the form $R\bar{a}, f\bar{a} = b$) obtained by the substitution. Then $\langle \varphi \rangle_n$ is the set of the clauses $\langle C_i(\bar{a}) \rangle$ obtained this way plus the *functionality clauses* $\{\{f\bar{a} = b\} \mid b \in [n]\}, \{\{\neg f\bar{a} = b\}, \{\neg f\bar{a} = b'\}\}$ for $f \in \tau$ an r-ary function symbol, $\bar{a} \in [n]^r$ and distinct $b, b' \in [n]$.

It should be clear that the assignments that satisfy the functional clauses bijectively correspond to τ-structures on $[n]$; moreover, such an assignment satisfies $\langle \varphi \rangle_n$ if and only if the corresponding τ-structure is a model of φ. Hence, $\langle \varphi \rangle_n$ is unsatisfiable if and only if n is not in the spectrum of φ; in particular, all $\langle \varphi \rangle_n, n \geq 1$, are contradictions if φ is an infinity axiom.

3 Width Notions for DAGs

3.1 Treewidth and Pathwidth

Let G be graph. A *tree decomposition* of G is a pair (T, χ) where T is a tree and χ is a function from $V(T)$ into the powerset of $V(G)$ such that:

(a) every vertex of G belongs to $\chi(t)$ for some $t \in V(T)$;

(b) for every edge $(v, w) \in E(G)$ of G there is some vertex t of T such that $v, w \in \chi(t)$;

(c) for every $v \in V(G)$ the set $\{t \in V(T) \mid v \in \chi(t)\}$ is connected in $\underline{T}$.

Recall, $\underline{T}$ is the graph underlying T. The *width* of a tree decomposition (T, χ) is the maximum $|\chi(t)| - 1$ over all $t \in V(T)$. The *treewidth* $tw(G)$ of G is the minimum width over all its tree decompositions. A *path decomposition* is a tree decomposition (T, χ) where T is a (directed) path. The *pathwidth* $pw(G)$ of a graph G is the minimum width over all its path decompositions.

Let (T, χ) be a tree decomposition of a graph G. We say that a vertex $v \in V(G)$ is *introduced* at $t \in V(T)$ if $v \in \chi(t)$ but $v \notin \chi(t')$ for any predecessor t' of t. Similarly, we say that v is *forgotten* at $t \in V(T)$ if $v \in \chi(t)$ and either $t = r_T$ or $v \notin \chi(t')$ for the successor t' of t. Note that every vertex $v \in V(G)$ is introduced at at least one tree node (by condition (a)) and forgotten at exactly one tree node (by condition (c)). In a path decomposition every vertex is introduced at exactly one tree node.

The same definitions apply literally to digraphs, so we can also speak of tree and path decompositions of digraphs. Consequently, the treewidth and pathwidth of a digraph equal the treewidth and pathwidth of the digraph's underlying graph, respectively. Thus the direction of edges is completely irrelevant for the treewidth or pathwidth of a digraph. For some considerations, however, one needs the direction of edges to be reflected in the decomposition and the associated width measure. For example [12] introduces the notion of *directed treewidth*, and it is known that every DAG has directed treewidth 1. We introduce new width measures that can distinguish between DAGs.

3.2 Ordered Treewidth and Ordered Pathwidth

Although we shall be mainly interested in DAGs, we give the definitions and some first observations generally for digraphs. All arguments in this section follow familiar lines hence we omit some of them.

Definition 3.1. A tree decomposition (T, χ) of a digraph D is *ordered* if the following condition holds:

(d) for every directed edge $(u, v) \in E(D)$ and every $t \in V(t)$ where v is introduced, $u \in \chi(t)$.

As above, we define the *ordered treewidth* $otw(D)$ of D as the minimum width over all ordered tree decompositions of D, and the *ordered pathwidth* $opw(D)$ of D as the minimum width over all ordered path decompositions of D.

We say that a class $\mathcal{C}$ of digraphs has *bounded ordered pathwidth* if there is a constant $w \in \mathbb{N}$ such that every digraph in $\mathcal{C}$ has ordered pathwidth at most w; we say $\mathcal{C}$ has *unbounded ordered pathwidth* if it does not have bounded ordered pathwidth. We use a similar mode of speech for the other width notions.

The ordered width measures are different from their classical counterparts:

Remark 3.2. For every digraph D, *otw*(D) is at least the in-degree of D.

Examples 3.3.

1. The ordered treewidth of a tree (with edges directed towards the root) is its in-degree.
2. A directed path with at least one edge has ordered pathwidth 1.
3. The class of perfect binary trees (with edges directed towards the root) has unbounded ordered pathwidth and bounded ordered treewidth.
4. The class of perfect binary trees with all edges reversed (edges directed away from the root) has unbounded ordered treewidth and bounded treewidth.

Proof of (1)–(3). (1) and (2). A tree (path) T has the ordered (path) tree decomposition (T, χ) where χ maps $t \in V(T)$ to the set containing t and its predecessors. It has minimal width by Remark 3.2.

(3). Recall B_h denotes the the perfect binary tree of height h (see Section 2.1). By (1) $otw(B_h)$ is 2 for $h > 0$ and 0 for $h = 0$. It is well-known that $pw(B_h) \geq \lceil h/2 \rceil$ (see, e.g., [7, Theorem 67]). This implies (3) noting *opw* $\geq$ *pw*. $\square$

We prove (4) after Lemma 3.10 below.

Lemma 3.4. *Let D be a digraph, (T, χ) an ordered tree decomposition of D and $X \subseteq V(D)$ be nonempty. Then (T, χ') is an ordered tree decomposition of $D[X]$ where χ' maps $t \in V(T)$ to $\chi(t) \cap X$.*

Lemma 3.5. *Let D be a DAG and (T, χ) an ordered tree decomposition of D. Assume $v \in V(D)$ has in-degree 1 and predecessor u and let D' be obtained by contracting the edge to v, i.e., by deleting v and adding edges from u to the successors of v. Then (T, χ') is an ordered tree decomposition of D', where for $t \in V(T)$*

$$\chi'(t) := \begin{cases} \chi(t) & \text{if } v \notin \chi(t); \\ (\chi(t) \setminus \{v\}) \cup \{u\} & \text{otherwise.} \end{cases}$$

These two lemmas are easy to prove. They show that ordered treewidth or pathwidth is not increased by taking "minors" in a certain sense (more restrictive than the one in [12, Section 5]).

Example 3.6. A star with n vertices and all edges directed towards the center can be obtained from B_h by contracting edges provided h is sufficiently large. Then $otw(B_h) = 2$ while the star has ordered treewidth $n - 1$.

Definition 3.7. A subtree T' of a tree T is *fully in T* if for every node of T' either all or none of its predecessors in T are in $V(T')$.

Lemma 3.8. *Let (T, χ) be an ordered tree decomposition of a digraph D, let T' be a subtree of T and set $\chi' := \chi \upharpoonright V(T')$. Assume $\bigcup_{t' \in V(T')} \chi(t') \neq \emptyset$ and set $D' := D[\bigcup_{t' \in V(T')} \chi(t')]$. Then*

1. *(T', χ') is an ordered tree decomposition of D';*
2. *if T' is fully in T, then there exists for every edge $(u, v) \in E(D)$ with $u \notin V(D')$ and $v \in V(D')$ a leaf t of T' which is not a leaf of T such that $v \in \chi(t)$.*

Proof. (1). That (T', χ') satisfies conditions (a) and (c) is easy to see. To verify (b), let $(u, v) \in E(D')$ and choose $t_u, t_v \in V(T')$ such that $u \in \chi(t_u)$ and $v \in \chi(t_v)$. By condition (b) for (T, χ) we find $t_{uv} \in V(T)$ such that $u, v \in \chi(t_{uv})$. Choose ℓ minimal such that there is a path $t_1 \cdots t_\ell$ in $\underline{T}$ with $t_1 = t_{uv}$ and $t_\ell \in V(T')$. Then every path in $\underline{T}$ from t_{uv} to some node in $V(T')$ contains t_ℓ. In particular, this holds for all paths in $\underline{T}$ connecting t_u and t_{uv}. Then $u \in \chi(t_\ell)$ since (T, χ) satisfies condition (c). Similarly $v \in \chi(t_\ell)$, and (b) for (T', χ') follows. Thus (T', χ') is a tree decomposition of D'.

We verify condition (d), i.e., that (T', χ') is ordered. Consider an edge $(u, v) \in E(D')$ and choose a node t_1 of T' where v is introduced in (T', χ'). We have to show $u \in \chi(t_1)$. In (T, χ), the vertex v must be introduced at some ancestor t_2 of t_1, that is, at some $t_2 \in V(T_{t_1})$. Since (T, χ) is ordered, $u \in \chi(t_2)$. We already verified (b) for (T', χ'), so there must be a node $t_3 \in V(T')$ with $u, v \in \chi(t_3)$. If $t_3 = t_1$ we are done, so assume $t_3 \neq t_1$. Then t_3 cannot be an ancestor of t_1: otherwise, by condition (c) for (T, χ), v is contained in the bag at the predecessor of t_1 in T on the path from t_3 to t_1; since T' is a subtree containing t_3 and t_2 it also contains this predecessor, contradicting that v is introduced at t_1 in T'. Hence $t_3 \in V(T) \setminus V(T_{t_1})$. Then every path between t_2 and t_3 in $\underline{T}$ contains t_1. By condition (c) for (T, χ) then $u \in \chi(t_1)$.

(2). Assume T' is fully in T and let $(u, v) \in E(D)$ with $u \notin V(D')$ and $v \in V(D')$. Choose $t' \in V(T')$ such that $v \in \chi(t')$. In (T, χ), v is introduced at some ancestor t of t'. Then $u \in \chi(t)$ because (T, χ) is ordered. Since $u \notin V(D')$, we have $t \notin V(T')$. In (T, χ), v is contained in every bag on the directed path from $t \notin V(T')$ to $t' \in V(T')$, and in particular, in the bag of the first node $t'' \in V(T')$ that we reach on this path. Then t'' is not a leaf of T (since t'' has ancestor $t \neq t''$). It also has some predecessor outside $V(T')$, namely its predecessor on the mentioned path. Since T' is fully in T, all predecessors of t'' are outside $V(T')$, i.e. t'' is a leaf of T'. □

Definition 3.9. A tree decomposition (T, χ) is *succinct* if every node forgets some vertex.

Lemma 3.10. *Every digraph D has a succinct ordered tree decomposition of width $otw(D)$, and a succinct ordered path decomposition of width $opw(D)$.*

Proof. We only prove the first statement. Let (T, χ) be a width $otw(D)$ ordered tree decomposition of D with the smallest number of nodes. We claim (T, χ) is succinct. Assume there is a node $s \in V(T)$ that does not forget some vertex. It suffices to construct a new tree decomposition (T', χ') with $\chi' := \chi \upharpoonright V(T')$ where $V(T') = V(T) \setminus \{s\}$.

If $s = r_T$, then $\chi(r_T) = \emptyset$. In this case, r_T has predecessors $t_1, \ldots, t_r$ for some $r > 0$. We define T' by (declaring t_1 to be the new root and) adding edges (t_i, t_1) for $1 < i \leq r$.

If $s \neq r_T$, then s has a successor t in T with $\chi(s) \subseteq \chi(t)$. In this case we define T' by adding all edges (t', t) for $(t', s) \in E(T)$. □

Proof of Examples 3.3 (4). Write B_h^{-1} for B_h with all edges reversed. Clearly, $tw(B_h^{-1})$ is 1 for $h > 0$ and 0 for $h = 0$. For $h > 0$ we show that

$$otw(B_h^{-1}) \geq \tfrac{1}{2} \log h. \tag{2}$$

By Lemma 3.10 there exists a succinct ordered tree decomposition (T, χ) of B_h^{-1} of minimal width $w := otw(B_h^{-1})$.

Claim 1. T has at most $2^{w+1} - 1$ many leaves.

Proof of Claim 1. For every leaf t of T let $N_t \subseteq \chi(t)$ be the set of vertices forgotten at t; this set is nonempty by succinctness. Consider a leaf t of T and a vertex $v \in N_t$. Because t is a leaf and the decomposition is ordered, $\chi(t)$ contains all ancestors of v in B_h^{-1}. Since $|\chi(t)| \leq w + 1$, it follows that v has at most $w + 1$ ancestors in B_h^{-1}, hence v is of distance at most w from the root r_{B_h} of B_h.

Now, B_h^{-1} has exactly $2^{w+1} - 1$ vertices that are of distance at most w from the root. Since each such vertex can occur in at most one set N_t for a leaf t, the claim follows. $\dashv$

Claim 2. $pw(B_h^{-1}) < (2^{w+1} - 1)(w + 1)$.

Proof of Claim 2. Let P be a longest branch in T, and let $t_1, \ldots, t_m = r_T$ be its nodes in order. For $i \in [m]$ let $\chi'(t_i)$ be the union of all sets $\chi(s)$ where $s \in V(T)$ is a node in T of distance exactly $m - i$ from the root $t_m = r_T$. By Claim 1 there are at most $2^{w+1} - 1$ such nodes s. Then (P, χ') is a path decomposition (in fact, even an ordered one) of B_h^{-1} and has width at most $(2^{w+1} - 1)(w + 1) - 1$. $\dashv$

As mentioned above $pw(B_h^{-1}) = pw(B_h) \geq \lceil h/2 \rceil$, so $h < 2(2^{w+1} - 1)(w + 1) < 2^{2w+2}$ by Claim 2. This implies (2). $\square$

Lemma 3.11. *A succinct ordered tree decomposition of a digraph D has at most $|V(D)|$ many nodes.*

Proof. We proceed by induction on $n = |V(D)|$. The statement is trivial for $n = 1$. Hence let $n > 1$ and assume the statement of the lemma holds for all digraphs D' with $|V(D')| < n$. Let (T, χ) be a succinct tree decomposition of a digraph D with $n = |V(D)|$. If $|V(T)| = 1$ there is nothing to show, so suppose $|V(T)| > 1$. Then T has a leaf s with a successor t. For $N := \chi(s) \setminus \chi(t)$ consider the digraph $D' := D - N$. By Lemma 3.8, $(T - \{s\}, \chi \upharpoonright V(T - \{s\}))$ is an ordered tree decomposition of D', and it is succinct because (T, χ) is. Since (T, χ) is succinct, $N \neq \emptyset$, and so $V(D') < n$. By induction, $|V(T - \{s\})| \leq n - 1$, and hence $|V(T)| \leq n$. $\square$

Lemma 3.12. *For every digraph D there exists an ordered tree decomposition (T, χ) of width $otw(D)$ where T has in-degree at most 2 and $|V(T)| < 2|V(D)|$.*

Proof. Let D be a DAG with $|V(D)| = n$. By Lemmas 3.10 and 3.11, D has a an ordered tree decomposition (T, χ) of width $otw(D)$ and $|V(T)| \leq n$. Let $\ell(T)$ be the sum of all in-degrees of those vertices in T that have in-degree bigger than 2.

Clearly, $\ell(T) \leq |E(T)| = |V(T)| - 1 = n - 1$. If $\ell(T) = 0$ then all nodes have in-degree at most 2 and we are done. Hence assume $\ell(T) > 0$ and consider a node $t_0 \in V(T)$ of in-degree $d > 2$ and let $t_1, \ldots, t_d$ be its predecessors. We transform (T, χ) into a new tree decomposition (T', χ') as follows. From T we remove the edges $(t_1, t_0), \ldots, (t_d, t_0)$ and add instead a binary tree with root r and leaves $t_1, \ldots, t_d$; we also add the edge (r, t_0). Let T' denote the tree obtained. Observe that $\ell(T') = \ell(T) - d$ and that $|V(T')| = |V(T)| + d - 1$ (a binary tree with d leaves has $d - 1$ inner nodes). Let N be the vertices introduced at t_0 let χ' be the extension of χ giving every new node $t' \in V(T') \setminus V(T)$ the bag $\chi(t_0) \setminus N$. It is easy to check that (T', χ') is a tree decomposition of D of the same width as (T, χ). To see that it is ordered, note that no new node introduces some vertex.

By repeating the above replacement we can successively replace all nodes whose in-degree exceeds 2, and we end up with an ordered tree decomposition (T^*, χ^*) of D of width $otw(D)$ where all nodes have in-degree at most 2. In total we have added $< \ell(T) \leq n - 1$ many nodes to T, so $|V(T^*)| < n + n - 1$. $\square$

Proposition 3.13. *Let $w, \ell \geq 1$ and (T, χ) a width w ordered tree decomposition of a digraph D such that T has height ℓ. Then $opw(D) < (w+1) \cdot (\ell+1)$.*

Proof. By adding if necessary some nodes with empty bags we can assume that all branches of T have the same length, say, ℓ. If we order $V(T)$ in an arbitrary way, then every branch naturally corresponds to a tuple from $[d]^\ell$ where d is the in-degree of T. Then branches are ordered via the lexicographical order on $[d]^\ell$. Use the path of branches according to this order as the path underlying a path decomposition. The bag at the ith path node is the union of the $\ell + 1$ many bags $\chi(t)$ for t ranging over the ith branch in T. It is straightforward to verify that this defines an ordered path decomposition of D. The size of bags is bounded by $(w+1) \cdot (\ell+1)$. □

3.3 Vertex Separation Numbers

A *linear layout* (or *linear arrangement*) of a graph G with n vertices is a bijection $\phi : V(G) \to [n]$. For every $i \in [n]$ we define four sets of vertices:

$$\begin{aligned}
L_G(i, \phi) &:= \{u \in V(G) \mid \phi(u) \leq i\}, \\
R_G(i, \phi) &:= \{u \in V(G) \mid \phi(u) > i\}, \\
L^*_G(i, \phi) &:= \{u \in L_G(i, \phi) \mid \exists v \in R_G(i, \phi) : (u, v) \in E(G)\}, \\
R^*_G(i, \phi) &:= \{v \in R_G(i, \phi) \mid \exists u \in L_G(i, \phi) : (u, v) \in E(G)\}.
\end{aligned}$$

The *in-degree* and the *out-degree* of ϕ is defined as $\max_{i \in [n-1]} |R^*_G(i, \phi)|$ and $\max_{i \in [n-1]} |L^*_G(i, \phi)|$, respectively. The *vertex separation number* $vsn(G)$ of G is defined as the smallest out-degree over all linear layouts of G (which equals the smallest in-degree over all linear layouts of G).

Proposition 3.14 ([13]). $pw(G) = vsn(G)$ *for every graph G.*

Note that the definition of in-degree and out-degree of a linear layout makes sense for digraphs. Recalling that a digraph is a DAG if and only if there exist linear layouts such that all (directed) edges run from left to right, it is natural to consider the following variant of the vertex separation number (for DAGs):

Definition 3.15. A linear layout ϕ of a DAG D is *ordered* if for every $(u, v) \in E(D)$ we have $\phi(u) < \phi(v)$. The *ordered vertex separation number* $ovsn(D)$ of a DAG D is the smallest out-degree over all ordered linear layouts of D.

We note that this definition is not symmetric in the sense that, in general, if we replace "smallest out-degree" by "smallest in-degree" we get a different number.

Example 3.16. Let D be a star with n vertices and all edges directed towards the center v. Then every ordered linear layout ϕ satisfies $\phi(v) = n$, has in-degree 1 and out-degree $n - 1$.

We prove an ordered analogue of Proposition 3.14.

Proposition 3.17. $opw(D) = ovsn(D)$ *for every DAG D.*

Proof. Let D be a DAG with n vertices. First we show that $opw(D) \geq ovsn(D)$. Let (T, χ) be an ordered path decomposition of D of width w. Let $t_1, \ldots, t_m$ be the vertices of T given in the ordering as we visit them when traversing T from the leaf to the root. Recall that every vertex of D is introduced at exactly one node in T. Let $\psi : V(D) \to [m]$ be the function such that a vertex v of D is introduced at node $t_{\psi(v)}$. We define the inverse ϕ^{-1} of a linear layout ϕ of D recursively as follows. To determine $\phi^{-1}(j)$, let r be minimal such that $N_r := \chi(t_r) \setminus \{\phi^{-1}(i) \mid i < j\} \neq \emptyset$; choose a source v (say, the smallest according to some fixed order of $V(G)$) of the DAG $D[N_r]$ induced on N_r, and set $\phi^{-1}(j) := v$.

To see that the linear layout ϕ is ordered, consider a directed edge $(u, v) \in E(D)$. Let $\phi(v) = j$ and $\psi(v) = r$. We have $u \in \chi(t_{\psi(v)})$ since (T, χ) is ordered, so $\psi(u) \leq \psi(v)$. If $\psi(u) < \psi(v)$, then clearly $\phi(u) < \phi(v)$. If $\psi(u) = \psi(v)$, then in the above process we assign $\phi(u)$ before we assign $\phi(v)$, hence $\phi(u) < \phi(v)$ as well.

To see that the out-degree of ϕ is at most w, consider $L^*_D(i, \phi)$ for some $i \in [n-1]$. Let $v = \phi^{-1}(i+1)$ and consider a vertex $u \in L^*_D(i, \phi)$. By definition, u has a successor $u' \in R_D(\phi, i)$, and clearly $\phi(u) \leq \phi(v) \leq \phi(u')$. This implies $\psi(u) \leq \psi(v) \leq \psi(u')$. Since ϕ is ordered and $(u, u') \in E(D)$ it follows that $u \in \chi(t_{\psi(u')})$, and by definition, $u \in \chi(t_{\psi(u)})$. By condition (c) of a tree decomposition, it follows that $u \in \chi(t_j)$ for all $\psi(u) \leq j \leq \psi(u')$, and in particular $u \in \chi(t_{\psi(v)})$. Thus $L^*_D(i, \phi) \subseteq \chi(t_{\psi(v)})$. Moreover, we have $v \in \chi(t_{\psi(v)}) \setminus L^*_D(i, \phi)$. Thus $L^*_D(i, \phi) \subseteq \chi(t_{\psi(v)}) \setminus \{v\}$, and hence $|L^*_D(i, \phi)| \leq |\chi(t_{\psi(v)}) \setminus \{v\}| \leq w$.

Next we show that $opw(D) \leq ovsn(D)$. Let ϕ be an ordered layout of D with out-degree w. We define a path decomposition (T, χ) letting T be the directed path $([n], \{(i, i+1) \mid i \in [n-1]\})$ and setting $\chi(i) := L^*_D(i-1, \phi) \cup \{\phi^{-1}(i)\}$ for $i \in [n]$; here, we understand $L^*_D(i-1, \phi) = \emptyset$ for $i = 0$. It is easy to verify that (T, χ) is a path decomposition of D and each bag has size at most $w + 1$. To see it is ordered, let $(u, v) \in E(D)$ and note that v is introduced at node $\phi(v)$. We claim $u \in \chi(\phi(v))$. But since ϕ is ordered, $\phi(u) < \phi(v)$ and in particular $\phi(v) \neq 0$. Then $u \in L^*_D(\phi(v) - 1, \phi) \subseteq \chi(\phi(v))$ as claimed. $\square$

4 Resolution Proofs of Minimal Width

Recall, the ordered treewidth of a proof containing an application of the cut rule is at least 2 (Remark 3.2). Clearly, when talking about the ordered pathwidth or ordered treewidth of a proof we mean the ordered pathwidth or ordered treewidth of its underlying DAG.

4.1 Minimal Ordered Pathwidth

A Resolution refutation of Γ is in *Input Resolution* if it contains only applications of the cut rule and each such application has at least one premiss (i.e., label of a predecessor) in Γ (see, e.g., [14]).

Theorem 4.1. *Let ℓ be a natural and Γ a set of clauses. There is a Resolution refutation of Γ of ordered pathwidth at most 2 and length at most ℓ if and only if there is an Input Resolution refutation of Γ of length at most ℓ.*

This allows us to think of ordered pathwidth as a measure of how far a Resolution refutation is from being in Input Resolution.

To prove this we need some preparations. A digraph is *triangle-free* if so is its underlying graph. A clause is *tautological* if it contains (as a term) a variable and its negation.

Lemma 4.2. *Let $w \in \mathbb{N}$ and $(D, (F_v)_{v \in V(D)})$ be a Resolution refutation of a set Γ of clauses such that D has a width w ordered tree decomposition with underlying tree T. Then there is a Resolution refutation $(D', (F'_v)_{v \in V(D)})$ of Γ such that*

1. *$V(D') \subseteq V(D)$ contains the sink of D;*
2. *D' has an ordered tree decomposition with underlying tree T of width at most w;*
3. *no vertex in $V(D')$ has in-degree 1 in D';*
4. *no $v \in V(D')$ has a tautological label F'_v;*
5. *D' is triangle-free.*

Proof. Let Γ be a set of clauses, $(D, (F_v)_{v \in V(D)})$ a Resolution refutation of Γ, and (T, χ) an ordered tree decomposition of D of width at most w. Let v^* denote the sink of D.

In a *first step* we transform $(D, (F_v)_{v \in V(D)})$ into a Resolution refutation $(D_1, (F_v)_{v \in V(D_1)})$ where no F_v for $v \in V(D_1)$ is tautological, and where D_1 is a sub-DAG of D with unique sink v^*. If there is some $v \in V(D)$ with tautological F_v, then there is such a v having a successor w with non-tautological F_w (the sink label is not tautological). Clearly F_w must be obtained by a cut from F_v and $F_{w'}$, where w' is the other predecessor of w. Then F_w is a weakening of $F_{w'}$ and we delete the edge (v, w). The deletion of (v, w) may have caused that v has become a sink. Then we repeatedly delete sinks different from v^*. Iterating this leads to a refutation $(D_1, (F_v)_{v \in V(D_1)})$ as desired.

In a *second step* we transform $(D_1, (F_v)_{v \in V(D)})$ into a Resolution refutation $(D_2, (F'_v)_{v \in V(D_2)})$ such that $F'_v \subseteq F_v$ for all $v \in V(D_2)$ and all weakenings are *improper* in the sense that if F'_v is obtained from F'_u by weakening, then $F'_v = F'_u$. Let $\phi : V(D_1) \to [|V(D_1)|]$ be a linear layout of D_1 and write $v_i = \phi^{-1}(i)$. Define F'_{v_i} recursively for each $i \in [|V(D_1)|]$ as follows. If v_i is a source in D_1, we set $F'_{v_i} := F_{v_i}$ (in particular this is the case for $i = 0$). If F_{v_i} is obtained by weakening from F_{v_j} for $j < i$, we set $F'_{v_i} := F'_{v_j}$. If F_{v_i} is obtained by a cut from F_{v_j} and F_{v_k} for $j, k < i$, then

- either F_{v_i} is a weakening of F'_{v_j} or of F'_{v_k} and we set $F'_{v_i} := F'_{v_j}$ resp. $F'_{v_i} := F'_{v_k}$;
- or otherwise F_{v_i} is a weakening of a clause F obtainable by cut on F'_{v_j} and F'_{v_k} and we set $F'_{v_i} := F$.

The digraph D'_2 is obtained from D_1 by deleting edges (v_j, v_i) resp. (v_k, v_i) in the first case above. Then D_2 is the digraph induced in D'_2 on the ancestors of the sink v^*.

Finally, in a *third step* we obtain a DAG D' from D_2 by contracting all edges (u, v) such that F'_v is obtained by weakening from F'_v. As weakenings are improper, such contractions preserve the property of being a refutation. In fact, $(D', (F'_v)_{v \in V(D')})$ is as desired: (1), (3) and (4) are easy to see and (2) follows from Lemmas 3.4 and 3.5. We verify (5). For contradiction, assume $\underline{D'}$ contains a triangle. As D' is a acyclic, this means there are $u, w, v \in V(D')$ such that $(u, v), (w, v), (u, w) \in E(D)$. Choose literals ℓ, ℓ' from F_u such that F_v is obtained cutting F_u with F_w on ℓ and F_w is obtained cutting F_u with $F_{w'}$ on ℓ', where w' is the second predecessor of w. In particular, $\neg\ell'$ is in $F_{w'}$ and $\neg\ell$ in F_w. First note that $\neg\ell'$ is not in F_w as it is cut from $F_{w'}$, so would have to appear in F_u and then F_u would be tautological. That F_u is not tautological, clearly implies $\ell \neq \neg\ell'$. Further $\ell \neq \ell'$ because otherwise $\neg\ell' = \neg\ell$ would be in F_w. Hence, ℓ and ℓ' are literals over distinct variables. But then ℓ comes into F_w from the premiss F_u. As $\neg\ell$ is in F_w, this clause is tautological, a contradiction. $\square$

Proof of Theorem 4.1. To see the backward direction, let $(D, (F_v)_{v \in V(D)})$ be an Input Resolution refutation of Γ. We can assume that D has vertices $V(D) = \{v_i \mid i \leq n\} \cup \{v'_i \mid i < n\}$ and edges $E(D) = \{(v_i, v_{i+1}), (v'_i, v_{i+1}) \mid i \leq n - 1\}$ for some suitable natural n. Then we have an ordered path decomposition (P, χ) of D where $V(P) := [n], E(P) := \{(i, i + 1) \mid i < n\}$ and $\chi(i) := \{v_i, v'_i, v_{i+1}\}$ for $i \in [n]$.

To verify the forward direction, let $(D, (F_v)_{v \in V(D)})$ be a refutation of Γ and let (T, χ) be a path decomposition of D of width at most 2. By Lemma 4.2 we can assume that D satisfies (3) and (5). We claim that every non-source has at least one source as a predecessor.

Assume otherwise, say $v \in V(D)$ has predecessors u_1, u_2 in D which are not sources of D. By Proposition 3.17 there exists an ordered linear layout ϕ of D of out-degree 2. As the layout is ordered $\phi(u_1), \phi(u_2) < \phi(v)$. Assume $\phi(u_1) < \phi(u_2)$ (the case $\phi(u_2) < \phi(u_1)$ is symmetrical) and consider the predecessors w_1, w_2 of u_2 in D'. We can assume $\phi(w_1) < \phi(w_2) < \phi(u_2)$. Further we have that w_1, w_2, u_1 are pairwise distinct because otherwise u_1, u_2, v would form a triangle in $\underline{D}$, a contradiction. Hence $\phi(w_2) < \phi(u_1)$ or $\phi(w_2) > \phi(u_1)$. In the first case, $w_1, w_2, u_1 \in L^*_D(\phi(u_1), \phi)$, so $|L^*_D(\phi(u_1), \phi)| > 2$, a contradiction. In the second case, $\phi(w_2) > \phi(u_1)$ and $w_1, w_2, u_1 \in L^*_D(\phi(w_2), \phi)$, again a contradiction. $\square$

4.2 Minimal Ordered Treewidth

Recall that treelike refutations have ordered treewidth 2 (Example 3.3 (1)). We prove a weak converse to this observation. This allows us to think of ordered treewidth as a measure of how far a Resolution refutation is from being treelike.

Theorem 4.3. *Let ℓ be a natural and Γ a set of clauses. If there is a Resolution refutation of Γ of ordered treewidth at most 2 and length at most ℓ, then there is a treelike Resolution refutation of Γ of length at most 3ℓ.*

Proof. Let $(D, (F_v)_{v \in V(D)})$ be a refutation of Γ with $otw(D) = 2$. By Lemma 4.2 we can assume that D satisfies (3) and (5). Let (T, χ) be a width 2 ordered tree decomposition of D. We claim that D is *almost treelike* in the sense that all its vertices of out-degree ≥ 2 are sources. This implies the theorem: for each source v with $\ell \geq 2$ successors $w_1, \ldots, w_\ell$ replace the edges $(v, w_2), \ldots, (v, w_\ell)$ by edges $(v_2, w_2), \ldots, (v_\ell, w_\ell)$ for $\ell - 1$ new vertices $v_2, \ldots, v_\ell$; this transforms D in a treelike refutation D' and adds at most two new vertices per successor of some source, so $|V(D')|$ is at most $3|V(D)|$.

To prove our claim, we show that D_v is almost treelike for every $v \in V(D)$; here, D_v is the sub-DAG of D induced on the ancestors of v in D.

This is clear for sources v. If v is not a source, but has predecessors u_1, u_2 we assume that both D_{u_1} and D_{u_2} are almost treelike, and show that also D_v is almost treelike.

Assume $u \in V(D_v)$ has out-degree ≥ 2 in D_v. We have to show that u is a source.

Claim 1. $u \in V(D_{u_1}) \cap V(D_{u_2})$.

Proof of Claim 1. Since $V(D_v) = V(D_{u_1}) \cup V(D_{u_2}) \cup \{v\}$ and $u \neq v$ we can assume that u is in $V(D_{u_1})$. For the sake of contradiction assume $u \notin V(D_{u_2})$. As D_{u_2} is closed under predecessors, no successor of u is in $V(D_{u_2})$. But u has at least two successors w, w' and these cannot be both in $V(D_{u_1})$ because D_{u_1} is almost treelike. Hence one of them, say w, equals v, and $w' \in V(D_{u_1})$. Then u is a predecessor of v outside $V(D_{u_2})$, so $u = u_1$. It follows that w' is both a successor and an ancestor of u and this contradicts acyclicity. $\dashv$

If one of u_1, u_2 is a source, say u_1, then $V(D_{u_1}) = \{u_1\}$. By Claim 1 then $u = u_1$ and we are done. Hence, assume that none of u_1, u_2 is a source. Choose $t \in V(T)$ where v is introduced. Then $u_1, u_2 \in \chi(t)$ so we find ancestors t_1, t_2 of t in T where u_1, u_2 are introduced respectively.

Claim 2. t_1, t_2 are incomparable in the sense that none is an ancestor of the other.

Proof of Claim 2. Assume otherwise, say, t_2 is an ancestor of t_1. Then t_1 lies on the path in T from t_2 to t and hence $u_2 \in \chi(t_1)$. As u_1 is not a source and introduced at t_1, the bag $\chi(t_1)$ also contains the two predecessors w_1, w_2 of u_1. Then $u_1, u_2, w_1, w_2 \in \chi(t_1)$ so these vertices cannot be pairwise distinct. Then $u_2 \in \{w_1, w_2\}$. It follows that $\{u_1, u_2, v\}$ induces a triangle in $\underline{D}$, a contradiction. $\dashv$

So we know t_1, t_2 are incomparable, say with t_0 as least upper bound, i.e., t_0 has both t_1, t_2 as ancestors but no predecessor of t_0 has this property. This t_0 lies on all paths in $\underline{T}$ from t_1 or t_2 to t, so $u_1, u_2 \in \chi(t_0)$.

It is not hard to show that all ancestors of u_1 in D are introduced at an ancestor of t_1 in T; similarly for u_2 and t_2. In particular u is introduced at ancestors s_1, s_2 of t_1, t_2 respectively. All paths in $\underline{T}$ from s_1 to s_2 contain a path from t_1 to t_2, and hence contain t_0. We finally show that u is a source. Otherwise its two predecessors are different from u_1 and from u_2. As they are in $\chi(s_1)$ and well as in $\chi(s_2)$, it follows they are in $\chi(t_0)$. As $\chi(t_0)$ also contains u_1, u_2 it has cardinality at least 4, a contradiction. $\square$

5 Proof Space

Let $k, w, \ell > 0$ be naturals, F a k-DNF and Γ a set of k-DNFs.

5.1 Ordered Pathwidth is Proof Space

In the Introduction we informally explained a bounded space proof by a sequence of blackboards. Formally, we follow [11] and define a *space w $R(k)$-proof of F from* Γ to be a finite sequence

$(\mathbb{B}_0, \dots, \mathbb{B}_{\ell-1})$ of sets $\mathbb{B}_i$ of k-DNFs, called *blackboards*, each of cardinality at most w such that $\mathbb{B}_0 = \emptyset$ and $F \in \mathbb{B}_{\ell-1}$ and for all $0 < i < \ell$ there is a formula G such that

(B1) $\mathbb{B}_i = \mathbb{B}_{i-1} \cup \{G\}$ and $G \in \Gamma$, or

(B2) $\mathbb{B}_i = \mathbb{B}_{i-1} \cup \{G\}$ and G is derived from at most two formulas in $\mathbb{B}_{i-1}$ by one application of some inference rule of $R(k)$, or

(B3) $\mathbb{B}_i = \mathbb{B}_{i-1} \setminus \{G\}$.

The space measure above is known as "formula space" or, in case $k = 1$, as "clause space." This is roughly the same as ordered pathwidth:

Theorem 5.1.

1. *If there is a space w $R(k)$-proof of F from Γ of length ℓ, then there is an $R(k)$-proof of F from Γ of length $< \ell$ and ordered pathwidth $< w$.*
2. *If there is a $R(k)$-proof of F from Γ of length ℓ and ordered pathwidth w, then there is a space $(w+1)$ $R(k)$-proof of F from Γ of length at most $(2w\ell + w + 1)$.*

Proof. (1). Let $(\mathbb{B}_0, \dots, \mathbb{B}_{\ell-1})$ be a space w $R(k)$-proof of F from Γ. We can assume that $\mathbb{B}_1 \neq \emptyset$. It suffices to show that for every $0 < i < \ell$ there are a positive $n_i \in \mathbb{N}$, a map vert_i from $[n_i]$ into $\bigcup_{j \leq i} \mathbb{B}_j$, an irreflexive set $E_i \subseteq [n_i]^2$ and a set $X_i \subseteq [n_i]$ such that

(a) the identity on $[n_i]$ is an ordered linear layout of the digraph $([n_i], E_i)$ (cf. Definition 3.15); in particular, this is a DAG.

(b) vert_i labels every source of $([n_i], E_i)$ with an element from Γ and every inner vertex with a k-DNF that can be obtained by one application of an inference rule of $R(k)$ from the labels of its predecessors;

(c) $\mathit{vert}_i \upharpoonright X_i$ is injective and has image $\mathbb{B}_i$;

(d) $(P_i, (X_j)_{1<j<\ell})$ is an ordered path decomposition of $([n_i], E_i)$;

(e) $n_i \leq i$.

Here, P_i denotes the path $(\{1, \dots, i\}, \{(j, j+1) \mid 0 < j < i\})$. We prove this by induction on i.

For $i = 1$, note $\mathbb{B}_1 = \{G\}$ for some $G \in \Gamma$, and take $n_1 := 1, E_1 := \emptyset, \mathit{vert}_1 := \{(0, G)\}$ and $X_1 := \{0\}$. Assume you found the desired objects up to i. If $\mathbb{B}_{i+1} = \mathbb{B}_i$ we take the the same objects for $i+1$, so assume $\mathbb{B}_i \neq \mathbb{B}_{i+1}$.

If $\mathbb{B}_{i+1} \subsetneq \mathbb{B}_i$ then $\mathbb{B}_{i+1} = \mathbb{B}_i \setminus \{G\}$ for some $G \in \mathbb{B}_i$; in this case we set $X_{i+1} := X_i \setminus \mathit{vert}_i^{-1}(G)$ and keep the other objects unchanged.

If $\mathbb{B}_{i+1} \supsetneq \mathbb{B}_i$, then $\mathbb{B}_{i+1} = \mathbb{B}_i \cup \{G\}$ for some $G \notin \mathbb{B}_i$. We set $n_{i+1} := n_i + 1, \mathit{vert}_{i+1} := \mathit{vert}_i \cup \{(n_i, G)\}$ and $X_{i+1} := X_i \cup \{n_i\}$. If $G \in \Gamma$, we set $E_{i+1} := E_i$. If $G \notin \Gamma$, then there are $G', G'' \in \mathbb{B}_i$ such that G is obtained from G', G'' by an $R(k)$-rule. Using (c) for i there are $m', m'' \in X_i$ with $\mathit{vert}_i(m') = G'$ and $\mathit{vert}_i(m'') = G''$. We set $E_{i+1} := E_i \cup \{(m', n_i), (m'', n_i)\}$.

In all cases it is easy to verify (a)–(e) for $i+1$.

(2). Let $(D, (F_v)_{v \in V(D)})$ be a length ℓ proof of F from Γ of ordered pathwidth w. By Lemma 3.10 we find a succinct ordered path decomposition (T, χ) of D of width w. By Lemma 3.11

we have $|V(T)| \leq \ell$. Using the notation from above we can assume $T = P_\ell$. Define $(\ell+1)$ many blackboards $\mathbb{B}_0 = \emptyset$ and $\mathbb{B}_i := \{F_v \mid v \in \chi(i)\}$ for $1 \leq i \leq \ell$. Then each $\mathbb{B}_i$ has cardinality at most $w+1$.

Between $\mathbb{B}_i$ and $\mathbb{B}_{i+1}$ add blackboards deleting all labels $F_v \in \mathbb{B}_i \setminus \mathbb{B}_{i+1}$ one by one. This gives inferences of type (B3) and ends with blackboard $\mathbb{B}^0 := \mathbb{B}_i \cap \mathbb{B}_{i+1}$. Note that $\mathbb{B}^0$ contains the set $\{F_v \mid v \in \chi(i+1) \cap \chi(i)\}$. We remark that $\mathbb{B}^0 \neq \emptyset$: otherwise the bags $\chi(i)$ and $\chi(i+1)$ are disjoint, and it is straightforward to see that then there is no edge in D between a vertex in $\bigcup_{1 \leq i' \leq i} \chi(i')$ and a vertex in $\bigcup_{i+1 \leq i' \leq \ell} \chi(i')$; since the decomposition is succinct both sets are nonempty, so D is not connected; but then D has at least two sinks, a contradiction.

If $\mathbb{B}^0 \neq \mathbb{B}_{i+1}$ add further blackboards as follows. Every formula in $\mathbb{B}_{i+1} \setminus \mathbb{B}^0$ is a label of some vertex in $\chi(i+1) \setminus \chi(i)$. Choose a linear layout ϕ of $D[\chi(i+1) \setminus \chi(i)]$ and let $m := |\chi(i+1) \setminus \chi(i)|$. Note $m \leq w$ since $\chi(i+1) \cap \chi(i) \neq \emptyset$ as seen above. Set $\mathbb{B}^j := \mathbb{B}^0 \cup \{F_{\phi^{-1}(j')} \mid j' < j\}$ for $1 \leq j \leq m$. Then $\mathbb{B}^m = \mathbb{B}_{i+1}$ and we add blackboards $\mathbb{B}^1, \ldots, \mathbb{B}^{m-1}$ between $\mathbb{B}^0$ and $\mathbb{B}_{i+1}$. We verify that each $\mathbb{B}^j, 1 \leq j \leq m$, can be obtained from $\mathbb{B}^{j-1}$ by an inference of type (B1) or (B2).

Note $\mathbb{B}^j = \mathbb{B}^{j-1} \cup \{F_v\}$ for $v := \phi^{-1}(j-1)$. If v is a source of D, then $F_v \in \Gamma$ and $\mathbb{B}^j$ can be obtained by an inference of type (B1). If v is not a source of D, then $\mathbb{B}^j$ can be obtained by an inference of type (B2). To see this we show that every predecessor u of v in D has label F_u in $\mathbb{B}^{j-1}$. But $u \in \chi(i+1)$ since v is introduced at $i+1$ and the decomposition is ordered; so either $u \in \chi(i)$, and then $F_u \in \mathbb{B}^0 \subseteq \mathbb{B}^{j-1}$, or $\phi(u) < j-1$, and then too $F_u \in \mathbb{B}^{j-1}$.

This adds at most $|\mathbb{B}_i \setminus \mathbb{B}^0| + (m-1) \leq 2w-1$ many blackboards between $\mathbb{B}_i$ and $\mathbb{B}_{i+1}$, each contained in $\mathbb{B}_i$ or in $\mathbb{B}_{i+1}$. Further, between $\mathbb{B}_0 = \emptyset$ and $\mathbb{B}_1$ at most w many blackboards are added. In total, this gives $(\ell+1) + (2w-1)\ell + w$ many blackboards each of cardinality at most $w+1$. $\square$

Combining with Theorem 4.1 this result allows to think of the space of a Resolution refutation as a measure of how far it is from being in Input Resolution.

5.2 Ordered Treewidth as Interactive Proof Space

The conversation of a teacher with her student described informally in the Introduction is described more formally by a game $\Pi^k_w(\Gamma, F)$ between two players called *Student* and *Teacher* on the following game graph.

Its vertices are partitioned into *Student positions* and *Teacher positions*, the former are $R(k)$-proofs of length at most w and the latter are k-DNFs. Its directed edges run from each k-DNF to all length $\leq w$ proofs of it, and from each proof to all labels of its sources that are outside of Γ. In particular, precisely the proofs from Γ are sinks. The *initial* position is the Teacher position F. Paths starting at the initial position are *plays*.

A *strategy for Teacher* (in $\Pi^k_w(\Gamma, F)$) is a function that maps plays ending in a Teacher position to a successor of this position; it is *positional* in case this value depends only on the Teacher position reached by the play. A play is *conform* to the strategy if every Student position in it is the value of the strategy on the initial segment of the play up to it. The strategy is *winning* if all plays conform to it are finite, and *ℓ-winning* if all plays conform to it have length at most $2\ell - 1$, i.e., Teacher wins making at most ℓ moves.

Remark 5.2. The game $\Pi_w^k(\Gamma, F)$ can be seen as a parity game, so it is memory-less determined; in particular, if a winning strategy for Teacher exists, then so does a positional one [18].

By a standard argument we get the following result.

Proposition 5.3. *If there is an ℓ-winning strategy for Teacher in $\Pi_w^k(\Gamma, F)$, then there is also a positional one.*

Proof. Assume there is an ℓ-winning strategy for Teacher in $\Pi_w^k(\Gamma, F)$. Let W_i be the set of Teacher positions G such that an i-winning strategy for teacher in $\Pi_w^k(\Gamma, G)$ exists. Then W_1 is the set of predecessors of sinks, i.e., formulas that have length $\leq w$ $R(k)$-proofs from Γ. Recursively, W_{i+1} equals W_i plus those G that have an *i-good* successor, namely one all of whose successors are in W_i; in other words, W_{i+1} is the set of formulas with length $\leq w$ $R(k)$-proofs from $W_i \cup \Gamma$. If an i-winning strategy for Teacher exists in $\Pi_w^k(\Gamma, G)$, then $G \in W_i$. An "attractor strategy" maps every $G \in W_1$ to a sink and every $G \in W_{i+1} \setminus W_i$ to an i-good successor of it. Such a strategy is positional and ℓ-winning in $\Pi_w^k(\Gamma, F)$; note $F \in W_\ell$ by assumption. □

Theorem 5.4. *There is an ℓ-winning strategy for Teacher in $\Pi_w^k(\Gamma, F)$ if and only if there is an $R(k)$-proof of F from Γ with an ordered tree-decomposition of width $< w$ and height $< \ell$.*

Sketch of proof. Assume there is an ℓ-winning strategy for Teacher in $\Pi_w^k(\Gamma, F)$. By the previous proposition, we can assume the strategy is positional. Consider the following tree T with nodes t labeled with Student positions $\pi(t)$. The label of the root is the value the winning strategy gives the initial position. Every node t has exactly one predecessor (in T) for each of the at most w many successors of $\pi(t)$ in the game graph; the label of such a node is the value given by the strategy to the corresponding successor. Note that, the leafs of T are labeled with sinks in the game graph, i.e., with length $\leq w$ proofs from Γ. The labels on branches of this tree correspond to the sequence of Student positions in a play conform to the strategy. Since the strategy is ℓ-winning, T has height at most $\ell - 1$.

First assume that the sets of vertices (of the DAG underlying) the proofs $\pi(t), t \in V(T)$, are pairwise disjoint. Then for each t with successor t' in T identify the sink of $\pi(t)$, say labeled F, and all sources of $\pi(t')$ with label F. This ensures that the union of the $\pi(t), t \in V(T)$, is an $R(k)$-proof. Letting $\chi(t)$ denote the vertices of $\pi(t)$, then (T, χ) witnesses that this proof has of ordered treewidth at most $w - 1$.

Conversely, let $(D, (F_v)_{v \in V(D)})$ be an $R(k)$-proof of F from Γ and (T, χ) an ordered tree decomposition of D of height $< \ell$ and width $< w$. We can assume that $\chi(t) \neq \emptyset$ for all $t \in V(T)$. For $t \in V(T)$ and $v \in \chi(t)$ let $\pi(t, v)$ be the proof induced on the ancestors of v in $D[\chi(t)]$. We informally describe a winning strategy for Teacher.

On the initial position F, Teacher chooses $v_1 \in V(D)$ and $t_1 \in V(T)$ such that $F_{v_1} = F$ and v_1 is introduced at t_1, and moves to $\pi(t_1, v_1)$. If Student moves to source label G, Teacher chooses a source v_2 of $\pi(t_1, v_1)$ such that $G = F_{v_2}$, chooses an ancestor t_2 of t_1 where v_2 is introduced and answers $\pi(t_2, v_2)$. And so on. Note that a strategy implementing such moves is not positional. Namely, for her ith move Teacher remembers a vertex $v_i \in V(D)$ and a node $t_i \in V(T)$ and these satisfy

(a) F_{v_i} is a Teacher position in the play,

(b) v_i is introduced at t_i,

(c) v_{i+1} is a source in $\pi(t_i, v_i)$,

(d) t_{i+1} is an ancestor of t_i.

Note, no Student position in the play is a formula in Γ (here we assume $F \notin \Gamma$; otherwise there is a 1-winning strategy). In particular, $F_{v_{i+1}} \notin \Gamma$ (by (a)), so v_{i+1} has predecessors in D. By (c), these predecessors are not in $\pi(t_i, v_i)$ and, by definition of $\pi(t_i, v_i)$, also not in $\chi(t_i)$. As the tree-decomposition is ordered, v_{i+1} is not introduced at t_i. As v_{i+1} is introduced at t_{i+1} (by (b)) we have $t_{i+1} \neq t_i$. By (d) then the sequence $t_1, t_2, \ldots$ has length $\leq \ell$. This implies that the strategy is ℓ-winning. □

Remark 5.5. Assume Γ is a set of clauses. If Teacher wins $\Pi^k_w(\Gamma, 0)$ then Γ is contradictory and hence has a treelike Resolution refutation; by Theorems 4.1 and 5.4, Teacher wins $\Pi^1_3(\Gamma, 0)$. It follows that for $w \geq 3$, Teacher wins $\Pi^k_w(\Gamma, 0)$ if and only if Teacher wins $\Pi^1_3(\Gamma, 0)$. Thus, the parameters k and w only matter when taking into account how fast Teacher can win, that is, when considering ℓ-winning strategies.

As a side remark we observe that if the teacher knows how to convince visiting students very quickly then she also does not need a large blackboard in class.

Corollary 5.6. *If Teacher has an ℓ-winning strategy in $\Pi^k_w(\Gamma, F)$, then there is a space ℓw $R(k)$-proof of F from Γ.*

Proof. Teacher has an ℓ-winning strategy in $\Pi^k_w(\Gamma, F)$. By the previous result there is a $R(k)$-proof of F from Γ with an ordered tree decomposition (T, χ) of width $\leq w - 1$ and height $\leq \ell - 1$. By Proposition 3.13 this proof has ordered pathwidth $< w \cdot \ell$. Now apply Theorem 5.1 (2). □

Remark 5.7. It is well-known and easy to see that every contradictory set of clauses Γ has a treelike Resolution refutation of height at most the number n of variables in Γ. By the previous remark this gives an n-winning strategy of Teacher in $\Pi^1_3(\Gamma, 0)$ and the last corollary thus shows that Γ has a Resolution refutation in space $3n$. Even $n + 1$ is known and proved in [11, Theorem 12].

6 Lower Bounds

Theorem 6.1. *Let φ be a first-order τ-sentence of the form* (1) *that has an infinite model. Let r be the maximal arity of some function symbol in τ and assume $r \geq 1$. Then there exists a real $c_\varphi > 0$ such that for every natural $n \geq 1$ and every natural $k \geq 1$, every strongly sound k-DNF refutation $(D, (F_v)_{v \in V(D)})$ of $\langle\varphi\rangle_n$ satisfies*

$$k \cdot otw(D) \cdot \log |V(D)| > c_\varphi \cdot n^{1/r}.$$

Remark 6.2. The assumption that $r \geq 1$ does not exclude interesting cases. If $r = 0$, all function symbols of τ are constants. In an infinite model of φ every nonempty set containing the interpretations of these constants carries a submodel which too is a model of φ (being universal). Hence, the spectrum of φ is co-finite, so all but finitely many translations $\langle\varphi\rangle_n$ are satisfiable and have no sound refutations at all.

Proof. Let a τ-sentence φ and a natural r accord the assumption of the theorem, and let M be an infinite model of φ. Let $m_0, \ldots, m_{\ell-1}$ be a list without repetitions of the interpretations of constants of τ in M. It suffices to find $c_\varphi > 0$ satisfying our claim for every positive $n \geq \ell$.

So let $n \geq \ell$ and k be positive naturals, and let $(D, (F_v)_{v \in V(D)})$ be a strongly sound k-DNF refutation of $\langle\varphi\rangle_n$. Write $w := \mathit{otw}(D)$. By Lemma 3.12, D has an ordered tree decomposition with a tree of in-degree at most 2 and at most $2|V(D)|$ many nodes. Add the sink of D to all bags on nodes on the path from the node where it is forgotten to the root. The resulting tree decomposition (T_0, χ) has width at most $w + 1$ with the sink of D contained in $\chi(r_{T_0})$, the bag at the root.

For $X \subseteq V(T_0)$ we write
$$\chi(X) := \bigcup_{t \in X} \chi(t).$$

Conditions For $N \subseteq M$ let $\partial N := \bigcup_f \mathrm{im}(f^M \upharpoonright N)$, where f ranges over the function symbols of τ, i.e., ∂N contains the values which M's functions take on N. Note $m_0, \ldots, m_{\ell-1} \in \partial N$ for every $N \subseteq M$ and
$$|\partial N| \leq |\tau| \cdot |N|^r. \tag{3}$$

We define a *condition* to be a pair (κ, λ) of partial bijections from $[n]$ to M such that $\kappa \subseteq \lambda$ and $\mathrm{im}(\lambda) = \mathrm{im}(\kappa) \cup \partial\mathrm{im}(\kappa)$. We say a condition (κ^*, λ^*) *extends* another (κ, λ) if $\kappa \subseteq \kappa^*$ and $\lambda \subseteq \lambda^*$. With a condition (κ, λ) we associate the restriction $\rho(\kappa, \lambda)$ which is defined on a propositional atom of the form $R\bar{a}$ or $f\bar{a} = b$ if and only if κ is defined on all components of $\bar{a}$; in this case it maps

- $R\bar{a}$ to 1 if $\kappa(\bar{a}) \in R^M$, and to 0 otherwise;
- $f\bar{a} = b$ to 1 if $\lambda^{-1}(f^M(\kappa(\bar{a}))) = b$, and to 0 otherwise;

note that λ^{-1} is defined on $f^M(\kappa(\bar{a})) \in \partial\mathrm{im}(\kappa)$. Naturally, here $\kappa(\bar{a})$ for a tuple $\bar{a} = a_1 \cdots a_r$ stands for the tuple $\kappa(a_1) \cdots \kappa(a_r)$.

Observe that, if (κ^*, λ^*) extends (κ, λ) in the sense above, then $\rho(\kappa^*, \lambda^*)$ extends $\rho(\kappa, \lambda)$ as a partial function. The *rank* of (κ, λ) is $|\mathrm{dom}(\kappa)|$. For example, $(\emptyset, \lambda_0)$ is a condition of rank 0, where λ_0 is the function that maps $i < \ell$ to m_i.

Claim 1. If (κ, λ) is a condition and C a clause in $\langle\varphi\rangle_n$, then $C \upharpoonright \rho(\kappa, \lambda) \neq 0$.

Proof of Claim 1. We assume $\rho(\kappa, \lambda)$ is defined on all variables appearing in C (otherwise there is nothing to show). If C is a functionality clause $\bigvee_b f\bar{a} = b$ or $\neg f\bar{a} = b \vee \neg f\bar{a} = b'$, then $\rho(\kappa, \lambda)$ is defined on all variables of the form $f\bar{a} = c$ for $c \in [n]$. By definition it evaluates exactly one of them, namely the one for $c := \lambda^{-1}(f^M(\kappa(\bar{a})))$, to 1 and all others to 0. This implies $C \upharpoonright \rho(\kappa, \lambda) = 1$.

Suppose C is $\langle C_i(\bar{a})\rangle$ and choose an injection $\lambda' : [n] \to M$ extending λ (we only need it to be defined on all components of $\bar{a}$). As $\langle C_i(\bar{a})\rangle \neq 1$, all pure equality literals in $C_i(\bar{x})$ become "false" under the replacement of $\bar{a}$ for $\bar{x}$. Since λ' is injective, the tuple $\lambda'(\bar{a})$ falsifies all these literals in M. But since M is a model of φ, the tuple $\lambda'(\bar{a})$ satisfies $C_i(\bar{x})$ in M, so satisfies some literal mentioning a symbol from τ. Writing $\bar{x} = x_0 \cdots x_{s-1}$ and $\bar{a} = a_0 \cdots a_{s-1}$ we can write our literal as $(\neg)fx_{i_0}\cdots x_{i_{r-1}} = x_{i_r}$ or $(\neg)Rx_{i_0}\cdots x_{i_{r-1}}$ where $f, R \in \tau$ are r-ary symbols for some $r \in \mathbb{N}$ and $i_0, \ldots, i_r \in [s]$. Assume our literal is $fx_{i_0}\cdots x_{i_{r-1}} = x_{i_r}$, the other cases are treated analogously. Then $f^M(\lambda'(a_{i_0}), \ldots, \lambda'(a_{i_{r-1}})) = \lambda'(a_{i_r})$ and $fa_{i_0}\cdots a_{i_{r-1}} = a_{i_r}$ is a propositional literal in $\langle C_i(\bar{a})\rangle$. Since we assumed $\rho(\kappa, \lambda)$ to be defined on all atoms in $\langle C_i(\bar{a})\rangle$, we have $a_{i_0}, \ldots, a_{i_{r-1}} \in \mathrm{dom}(\kappa)$ and $f^M(\kappa(a_{i_0}), \ldots, \kappa(a_{i_{r-1}})) = \lambda'(a_{i_r})$. Hence $\lambda'(a_{i_r}) \in \mathrm{im}(\lambda)$, and as $\lambda' \supseteq \lambda$ is an injection, $\lambda^{-1}(\lambda'(a_{i_r})) = a_{i_r}$. By definition then the restriction $\rho(\kappa, \lambda)$ evaluates $fa_{i_0}\cdots a_{i_{r-1}} = a_{i_r}$ to 1 and $C \upharpoonright \rho(\kappa, \lambda) = 1$ follows. $\dashv$

Claim 2. Let $B \subseteq [n]$ and assume (κ, λ) is a condition of rank at most d. If

$$n \geq 3|\tau| \cdot (d + |B|)^r, \tag{4}$$

then there exists a condition (κ', λ') extending (κ, λ) such that $B \subseteq \mathrm{dom}(\kappa')$.

Proof of Claim 2. By (3) we have $|\mathrm{dom}(\lambda)| \leq |\tau| \cdot d^r + d$. Choose some minimal injective extension κ' of κ such that $B \subseteq \mathrm{dom}(\kappa)$ and then a minimal injective extension λ' of λ such that $\mathrm{im}(\lambda') \supseteq \partial\mathrm{im}(\kappa') \cup \mathrm{im}(\kappa' \upharpoonright B)$. The choice of κ' is possible if $n - d \geq |B|$ – and this follows from (4) (note $|\tau| > 0$ as $r \geq 1$). By (3) we see that the choice of λ' is possible if $n - |\mathrm{dom}(\lambda)| \geq |\tau| \cdot |\mathrm{dom}(\kappa')|^r + |B|$, and hence if $n \geq |\tau| \cdot (d + |B|)^r + (|\tau| \cdot d^r + d) + |B|$. This is implied by (4). $\dashv$

Adversary positions Recall Definition 3.7. An *adversary position* is a tuple (T, L, κ, λ) such that

(A1) T is a subtree of T_0 which is fully in T_0;

(A2) $L \subseteq V(T_0)$ contains every leaf of T which is not a leaf of T_0;

(A3) (κ, λ) is a condition such that

- (A3a) $F_v \upharpoonright \rho(\kappa, \lambda) = 1$ for all $v \in \chi(L)$, and
- (A3b) for every condition (κ^*, λ^*) such that (κ^*, λ^*) extends (κ, λ) there exists $v \in \chi(r_T)$ such that $F_v \upharpoonright \rho(\kappa^*, \lambda^*) \neq 1$.

Adversary positions exist: for example, $(T_0, \emptyset, \emptyset, \lambda_0)$ is one; property (A3b) holds because $\chi(r_{T_0})$ contains the sink v of D and $F_v = 0$.

Claim 3. Suppose (T, L, κ, λ) is an adversary position and $v \in \chi(r_T)$. Let

$$\Gamma_T := \{F_u \mid u \in \chi(V(T)) \text{ is a source in } D\}.$$

Then there exists a strongly sound k-DNF proof of F_v from $\Gamma_T \cup \{F_u \mid u \in \chi(L)\}$.

Proof of Claim 3. Let $D' := D[\chi(T)]$. It suffices to show that for every $v \in V(D')$ either all predecessors of v in D are in $V(D')$ or $v \in \chi(L)$. But if $(u, v) \in E(D)$ and $u \notin V(D')$, then Lemma 3.8 (2) and (A1) imply that $v \in \chi(t)$ for some leaf t of T which is not a leaf in T_0; by (A2) then $v \in \chi(L)$. $\dashv$

Recall that T_t denotes the subtree of a tree T rooted at t. An adversary position (T, L, κ, λ) has as *successor* any tuple $(T', L', \kappa', \lambda')$ that can be obtained as follows.

> Choose $t \in V(T)$ such that both $V(T_t)$ and $V(T) \setminus (V(T_t) \setminus \{t\})$ have cardinality at most $\lfloor 2|V(T)|/3 \rfloor + 1$. Such a t exists because T has in-degree at most 2 (as a subtree of T_0).
>
> *Case 1.* Property (A3b) holds for t, i.e. for every extension (κ^*, λ^*) of (κ, λ) there exists $v \in \chi(t)$ such that $F_v \restriction \rho(\kappa^*, \lambda^*) \neq 1$.
>
> Set $T' := T_t, L' := L, \kappa' := \kappa$ and $\lambda' := \lambda$.
>
> *Case 2.* Otherwise, choose an extension (κ^*, λ^*) of (κ, λ) of minimal rank among those satisfying $F_v \restriction \rho(\kappa^*, \lambda^*) = 1$ for all $v \in \chi(t)$.
>
> Set $T' := T - (V(T_t) \setminus \{t\}), L' := L \cup \{t\}, \kappa' := \kappa^*$ and $\lambda' := \lambda^*$.

Claim 4. If (T, L, κ, λ) is an adversary position with successor $(T', L', \kappa', \lambda')$, then $(T', L', \kappa', \lambda')$ too is an adversary position. The rank of (κ', λ') is at most $(w + 2) \cdot k \cdot r_\varphi$ bigger than the rank of (κ, λ) where r_φ denotes the maximal arity of some symbol in τ.

For the proof, we use a mode of speech from [9] and say that the propositional variables $R\bar{a}$ and $f\bar{a} = b$ *mention* an element $a \in [n]$ if a is a component of $\bar{a}$; in particular $f\bar{a} = b$ does not necessarily mention b. A formula mentions an element if so does some variable appearing in it.

Proof of Claim 4. Let $t \in V(T)$ be the node chosen to compute $(T', L', \kappa', \lambda')$ from (T, L, κ, λ). Both subtrees T_t and $T - (V(T_t) \setminus \{t\})$ are fully in T. Since T is a fully in T_0, so is T' and $(T', L', \kappa', \lambda')$ satisfies (A1). Properties (A2) and (A3a) are clear. Property (A3b) follows in Case 1 because $r_{T'} = t$, and in Case 2 because $r_{T'} = r_T$, (κ', λ') extends (κ, λ) and (T, L, κ, λ) satisfies (A3b).

To see the second statement, assume $(T', L', \kappa', \lambda')$ is obtained according to Case 2 (in Case 1 there is nothing to show). Choose a condition $(\tilde{\kappa}, \tilde{\lambda})$ extending (κ, λ) such that $F_v \restriction \rho(\tilde{\kappa}, \tilde{\lambda}) = 1$ for every $v \in \chi(t)$. For $v \in \chi(t)$ choose a k-term t_v in the k-DNF F_v such that $t_v \restriction \rho(\tilde{\kappa}, \tilde{\lambda}) = 1$. Any restriction ρ that agrees with $\rho(\tilde{\kappa}, \tilde{\lambda})$ on the atoms appearing in these k-terms is such that $F_v \restriction \rho = 1$ for every $v \in \chi(t)$. In particular, this is the case for $\rho(\tilde{\kappa} \upharpoonright (A \cup \mathrm{dom}(\kappa)), \tilde{\lambda}')$ where A is the set of elements mentioned by $\bigwedge_{v \in \chi(t)} t_v$ and $\tilde{\lambda}'$ is a suitable restriction of $\tilde{\lambda}$ such that $(\tilde{\kappa} \upharpoonright (A \cup \mathrm{dom}(\kappa)), \tilde{\lambda}')$ is a condition. Every k-term t_v mentions at most $k \cdot r_\varphi$ many elements, and there are at most $|\chi(t)| \leq w + 2$ many terms t_v. Thus, the rank of $(\tilde{\kappa} \upharpoonright (A \cup \mathrm{dom}(\kappa)), \tilde{\lambda}')$ and hence of (κ', λ') is at most $|A| \leq (w + 2) \cdot k \cdot r_\varphi$ bigger than the rank of (κ, λ). $\dashv$

Wrapping up Let $((T_i, L_i, \kappa_i, \lambda_i))_{i\in\mathbb{N}}$ be a sequence such that for all $i \in \mathbb{N}$, $(T_{i+1}, L_{i+1}, \kappa_{i+1}, \lambda_{i+1})$ is a successor of $(T_i, L_i, \kappa_i, \lambda_i)$ and $(T_0, L_0, \kappa_0, \lambda_0)$ is $(T_0, \emptyset, \emptyset, \lambda_0)$; we already noted that this is an adversary position. By Claim 4 all tuples $(T_i, L_i, \kappa_i, \lambda_i)$ are adversary positions. Further, $|V(T_{i+1})| \leq \lfloor 2|V(T_i)|/3 \rfloor + 1$, so $|V(T_m)| \leq 3$ for $m := \lceil \log_{3/2} |V(T_0)| \rceil$. Recalling that $|V(T_0)| \leq 2|V(D)|$ the theorem follows once we show

$$n < 3|\tau| \cdot (m \cdot (w+2) \cdot k \cdot r_\varphi + 3(w+2) \cdot w_\varphi \cdot r_\varphi)^r, \tag{5}$$

where w_φ is the maximal number of literals in some first order clause $C_i(\bar{x})$ of φ. We now verify (5).

Since (κ_0, λ_0) has rank 0, Claim 4 implies that (κ_m, λ_m) has rank at most

$$d_m := m \cdot (w+2) \cdot k \cdot r_\varphi.$$

Recall the notation Γ_{T_m} from Claim 3 and let $B \subseteq [n]$ denote the set of elements mentioned by formulas in Γ_{T_m}. Note $\Gamma_{T_m} \subseteq \langle\varphi\rangle_n$. Since $|V(T_m)| \leq 3$ we have $|\Gamma_{T_m}| \leq 3(w+2)$ and hence

$$|B| \leq 3(w+2) \cdot w_\varphi \cdot r_\varphi.$$

Assume for contradiction that (5) fails. Then $n \geq 3|\tau| \cdot (d_m + |B|)^r$. By Claim 2 there exists a condition (κ, λ) extending (κ_m, λ_m) such that $B \subseteq \mathrm{dom}(\kappa)$. By (A3b) there exists $v_m \in \chi(r_{T_m})$ such that $F_{v_m} \restriction \rho(\kappa, \lambda) \neq 1$. By Claim 3 there is a strongly sound k-DNF proof of F_{v_m} from $\Gamma_{T_m} \cup \{F_u \mid u \in \chi(L_m)\}$. For every clause $C \in \Gamma_{T_m} \subseteq \langle\varphi\rangle_n$, we have that $\mathrm{dom}(\kappa)$ contains all elements mentioned by C. Hence $\rho(\kappa, \lambda)$ evaluates every atom in C, so $C \restriction \rho(\kappa, \lambda) \in \{0, 1\}$, and hence $C \restriction \rho(\kappa, \lambda) = 1$ by Claim 1. Further we have $F_u \restriction \rho(\kappa, \lambda) = 1$ for every $u \in \chi(L_m)$ since $\rho(\kappa, \lambda)$ extends $\rho(\kappa_m, \lambda_m)$ and (A3a). In summary, F_{v_m} does not restrict to 1 under $\rho(\kappa, \lambda)$ and there is a strongly sound k-DNF proof of F_{v_m} from formulas that do restrict to 1 under $\rho(\kappa, \lambda)$. This contradicts Lemma 2.3. $\square$

This proof has the following corollary.

Corollary 6.3. *Let φ be a first-order τ-sentence of the form* (1) *that has an infinite model. Let r be the maximal arity of some function symbol in τ and assume $r \geq 1$. Then there exists a real $c_\varphi > 0$ such that for every natural $n \geq 1$ and every natural $k \geq 1$, every strongly sound k-DNF refutation $(D, (F_v)_{v\in V(D)})$ of $\langle\varphi\rangle_n$ satisfies*

$$k \cdot \mathit{opw}(D) > c_\varphi \cdot n^{1/r}.$$

Proof. Let $(D, (F_v)_{v\in V(D)})$ be a strongly sound k-DNF refutation of $\langle\varphi\rangle_n$ with path decomposition (T, χ) of width $w := \mathit{opw}(D)$. Assume the nodes of the path T are $0, 1, 2 \ldots, \ell$ in order. We can assume that for each i there is v_i such that $\chi(i) \setminus \{v_i\} = \chi(i+1)$ or $\chi(i) \cup \{v_i\} = \chi(i+1)$. Furthermore we assume $\chi(0) = \emptyset$.

Call a condition (κ, λ) *good for i* if for every $v \in \chi(i)$

(a) if F_v is a clause from $\langle\varphi\rangle_n$, then dom(κ) contains all elements mentioned by F_v;

(b) otherwise there is a term t_v of F_v such that $\mathrm{dom}(\kappa)$ contains all elements mentioned by t_v and $t_v \upharpoonright \rho(\kappa, \lambda) = 1$.

Recall the constants r_φ, w_φ from the previous proof.

Claim 5. Let $i \leq \ell$. If there is a condition good for i, then there is one of rank at most

$$(w + 1) \cdot k \cdot w_\varphi \cdot r_\varphi.$$

Proof of Claim 5. Let (κ, λ) be good for i. For $v \in \chi(i)$ let $B_v \subseteq [n]$ be the set of elements mentioned by F_v if $F_v \in \langle\varphi\rangle_n$, and otherwise the set of elements mentioned by t_v (chosen according (b) above). In the first case $|B_v| \leq w_\varphi \cdot r_\varphi$ and in the second $|B_v| \leq k \cdot r_\varphi$. Set $B := \bigcup_{v \in \chi(i)} B_v$ and note $|B| \leq (w + 1) \cdot k \cdot r_\varphi \cdot w_\varphi$. Define $\kappa' := \kappa \upharpoonright B$ and $\lambda' := \lambda \upharpoonright (\mathrm{im}(\kappa') \cup \partial\mathrm{im}(\kappa'))$. Then (κ', λ') is good for i and has rank at most $|B|$. ⊣

Observe that if (κ, λ) is good for i, then $F_v \upharpoonright \rho(\kappa, \lambda) = 1$ for all $v \in \chi(i)$ (in case (a) this follows from Claim 1). In particular, there is a condition good for 0 (namely $(\emptyset, \lambda_0)$ from the previous proof) but there is no condition good for i^* where $i^* \leq \ell$ is the node where the sink of D is introduced. Hence there exists $i_* < i^*$ such that there exists a condition (κ_*, λ_*) good for i_* and such that there does not exist a condition good for $i_* + 1 \leq \ell$. In particular, (κ_*, λ_*) is not good for $i_* + 1$. It follows that $\chi(i_* + 1) = \chi(i_*) \cup \{v_{i_*}\}$ and v_{i_*} is introduced at $i_* + 1$.

By Claim 5 we can assume that (κ_*, λ_*) has rank at most $d_* := (w + 1) \cdot k \cdot w_\varphi \cdot r_\varphi$.

Claim 6. v_{i_*} is a source of D.

Proof of Claim 6. Otherwise, because (T, χ) is ordered, the predecessors u, w of v_{i_*} in D are present in $\chi(i_* + 1) = \chi(i_*) \cup \{v_{i_*}\}$, so $u, w \in \chi(v_{i_*})$. Since (κ_*, λ_*) is good for i_* we have $F_u \upharpoonright \rho(\kappa_*, \lambda_*) = F_w \upharpoonright \rho(\kappa_*, \lambda_*) = 1$. By strong soundness $F_{v_{i_*}} \upharpoonright \rho(\kappa_*, \lambda_*) = 1$, so (b) is satisfied for v_{i_*}. Hence, (κ_*, λ_*) is good for $i_* + 1$, a contradiction. ⊣

By Claim 6 $F_{v_{i_*}}$ is a clause from $\langle\varphi\rangle_n$. Let B_* denote the set of elements mentioned by $F_{v_{i_*}}$. Any condition (κ, λ) extending (κ_*, λ_*) with $B_* \subseteq \mathrm{dom}(\kappa)$ would satisfy $F_{v_{i_*}} \upharpoonright \rho(\kappa^*, \lambda^*) = 1$ by Claim 1 and would thus be good for $i_* + 1$. That such a condition does not exist, implies by Claim 2 that $n < 3|\tau| \cdot (d_* + |B_*|)^r$. Noting $|B_*| \leq w_\varphi \cdot r_\varphi$, the corollary follows. □

Remark 6.4. It is well-known that Input Resolution is not refutation-complete (cf. [14]). Indeed, if φ is as above, then for sufficiently large n there is no Input Resolution refutation of $\langle\varphi\rangle_n$. This follows from the above corollary and Theorem 4.1.

The above corollary generalizes bounds on space (recall Theorem 5.1) known for particular infinity axioms (cf. Introduction). Concerning the more peculiar notion of space from Section 5.2 we find it worthwhile to explicitly note the following rather direct corollary.

Corollary 6.5. *Let φ be a first-order τ-sentence of the form* (1) *that has an infinite model. Let r be the maximal arity of some function symbol in τ and assume $r \geq 1$. Then there exists a real $c_\varphi > 0$ such that for all naturals $n, k, w, \ell \geq 1$ the following holds. If there exists an ℓ-winning strategy for Teacher in $\Pi^k_w(\langle\varphi\rangle_n)$, then*

$$k \cdot w \cdot \ell > c_\varphi \cdot n^{1/r}.$$

Proof. Let φ and r accord the assumptions. Choose c_φ according Corollary 6.3, let $n, k, w, \ell \geq 1$ be given and assume there exists an ℓ-winning strategy for Teacher in $\Pi_w^k(\langle\varphi\rangle_n)$. We claim that $k \cdot w \cdot \ell > c_\varphi \cdot n^{1/r}$. But by Corollary 5.6 there exists a space $\ell \cdot w$ $R(k)$-refutation of $\langle\varphi\rangle_n$. By Theorem 5.1 (1) there exists a $R(k)$-refutation of $\langle\varphi\rangle_n$ of ordered pathwidth $< \ell \cdot w$. Then our claim follows from Corollary 6.3. □

7 Infinity Axioms

An *infinity axiom* is a first-order sentence φ of the form (1) that does not have finite models but does have an infinite model. Note that in this case all propositional translations $\langle\varphi\rangle_n, n \geq 1$, are contradictory. Strong lower bounds on the length of refutations of these principles are known for the treelike systems [16, 22, 9]. One also knows, however, some few short DAG-like refutations:

Example 7.1. The *least number principle* is formulated using a unary function symbol f and a binary relation symbol $<$:

$$\textit{lnp} := \forall xyz(\neg x < x \wedge (\neg x < y \vee \neg y < z \vee x < z) \wedge fx < x).$$

This sentence states that $<$ is a strict linear order and f maps every element to a smaller one. Stålmarck [25] gave polynomial length Resolution refutations of $\langle \textit{lnp}\rangle_n$.

Example 7.2. The *very weak pigeonhole principle* states that n^2 pigeons cannot fly injectively into n holes. This principle can be formulated as a first-order infinity axiom using a binary function symbol f:

$$\textit{wphp} := \forall xx'yy'z\big((\neg fxx' = z \vee \neg fyy' = z \vee x = y) \wedge (\neg fxx' = z \vee \neg fyy' = z \vee x' = y')\big).$$

Razborov [21] showed that DAG-like Resolution refutations of $\langle \textit{wphp}\rangle_n$ need length $2^{\Omega(n/(\log n)^2)}$. Maciel, Pitassi and Woods [17] showed that there exist quasipolynomial length $R(\log)$-refutations of $\langle \textit{wphp}\rangle_n$.

It is not understood what kind (say in some model-theoretic sense) of infinity axioms do have short DAG-like refutations in Resolution or in $R(k)$ for small k. We note that short DAG-like refutations of translations of infinity axioms need to be far from being treelike in that they require unbounded ordered treewidth. This is the first statement in the corollary below. The second can be seen as a generalization of known lower bounds for treelike $R(\log)$ [16].

Corollary 7.3. *Let φ be as in Theorem 6.1.*

1. *Every $R(\log)$-refutation of $\langle\varphi\rangle_n$ of length at most $2^{n^{o(1)}}$ has ordered treewidth at least $n^{\Omega(1)}$.*
2. *Every $R(\log)$-refutation of $\langle\varphi\rangle_n$ of ordered treewidth at most $n^{o(1)}$ has length at least $2^{n^{\Omega(1)}}$.*

Specifically for the above two examples we can say the following.

Corollary 7.4.

1. *Polynomial length $R(100)$-refutations of $\langle \textit{lnp}\rangle_n$ have ordered treewidth at least $\Omega(n/\log n)$.*
2. *Quasipolynomial length $R(\log)$-refutations of $\langle \textit{wphp}\rangle_n$ have ordered treewidth at least $\Omega(n^{0.4})$.*

8 Conclusion

In this paper we have revisited proof complexity using the graph invariants ordered treewidth and ordered pathwidth. Whereas the first corresponds to ordinary proof space, the latter gives rise to a notion of interactive proof space, which can be described in terms of a student-teacher game. These graph invariants provide the means for length-space lower bounds for $R(k)$-refutations that apply to a large class of formulas having a natural meaning (infinity axioms). It relaxes the refutation space measure (i.e., ordered pathwidth) to ordered treewidth and applies to $R(\log)$.

Acknowledgements.

The restrictions used in the proof of Theorem 6.1 come from unpublished work of Albert Atserias, Sergi Oliva and the first author. We thank Albert Atserias and Sergi Oliva for their kind allowance to use them here. The first author thanks the FWF (Austrian Science Fund) for its support through Project P 24654 N25. The second author thanks the ERC (European Research Council) for its support through Project COMPLEX REASON 239962.

References

[1] M. Alekhnovich, E. Ben-Sasson, A. A. Razborov, and A. Wigderson. Space complexity in propositional calculus. *SIAM J. Comput.*, 31(4):1184–1211, 2002.

[2] A. Atserias and V. Dalmau. A combinatorial characterization of resolution width. *J. of Computer and System Sciences*, 74(3):323–334, 2008.

[3] J. Bang-Jensen and G. Gutin. *Digraphs*. Springer Monographs in Mathematics. Springer-Verlag London Ltd., London, second edition, 2009.

[4] A. Beckmann and J. Johannsen. *Bounded Arithmetic and Resolution-Based Proof Systems*, volume 7 of *Collegium Logicum*. Kurt Gödel Society, 2004.

[5] E. Ben-Sasson. Size-space tradeoffs for resolution. *SIAM J. Comput.*, 38(6):2511–2525, 2009. Journal version of STOC'02 paper.

[6] E. Ben-Sasson and J. Nordström. Understanding space in proof complexity: Separations and trade-offs via substitutions. Electronic Colloquium on Computational Complexity *(ECCC)*, 17:125, 2010.

[7] H. L. Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theoretical Computer Science*, 209(1-2):1–45, 1998.

[8] S. A. Cook and R. A. Reckhow. The relative efficiency of propositional proof systems. *J. Symbolic Logic*, 44(1):36–50, Mar. 1979.

[9] S. Dantchev and S. Riis. On relativisation and complexity gap for resolution-based proof systems. In *Computer science logic*, volume 2803 of *Lecture Notes in Computer Science*, pages 142–154. Springer Verlag, 2003.

[10] J. L. Esteban, N. Galesi, and J. Messner. On the complexity of resolution with bounded conjunctions. *Theoretical Computer Science*, 321(2-3):347–370, 2004.

[11] J. L. Esteban and J. Torán. Space bounds for resolution. *Information and Computation*, 171(1):84–97, 2001.

[12] T. Johnson, N. Robertson, P. D. Seymour, and R. Thomas. Directed tree-width. *J. Combin. Theory Ser. B*, 82(1):138–154, 2001.

[13] N. G. Kinnersley. The vertex separation number of a graph equals its path-width. *Information Processing Letters*, 42(6):345–350, 1992.

[14] H. Kleine Büning and T. Lettman. *Propositional logic: deduction and algorithms*. Cambridge University Press, Cambridge, 1999.

[15] J. Krajíček. On the weak pigeonhole principle. *Fund. Math.*, 170(1-2):123–140, 2001. Dedicated to the memory of Jerzy Łoś.

[16] J. Krajíček. Combinatorics of first order structures and propositional proof systems. *Archive for Mathematical Logic*, 43(4):427–441, 2004.

[17] A. Maciel, T. Pitassi, and A. R. Woods. A new proof of the weak pigeonhole principle. *J. of Computer and System Sciences*, 64(4):843–872, 2002. Special issue on STOC 2000 (Portland, OR).

[18] R. Mazala. Infinite games. In E. Grädel and W. Thomas, editors, *Automata, Logics, and Infinite Games*, volume 2500 of *Lecture Notes in Computer Science*, chapter 2. Springer Verlag, 2002.

[19] J. Nordström. Narrow proofs may be spacious: separating space and width in resolution. *SIAM J. Comput.*, 39(1):59–121, 2009.

[20] J. Paris and A. Wilkie. Counting problems in bounded arithmetic. In *Methods in mathematical logic (Caracas, 1983)*, volume 1130 of *Lecture Notes in Math.*, pages 317–340. Springer Verlag, 1985.

[21] A. Razborov. Resolution lower bounds for the weak functional pigeonhole principle. *Theoretical Computer Science*, 303(1):233–243, 2003. Logic and complexity in computer science (Créteil, 2001).

[22] S. Riis. A complexity gap for tree resolution. *Computational Complexity*, 10(3):179–209, 2001.

[23] N. Segerlind. The complexity of propositional proofs. *Bull. of Symbolic Logic*, 13(4):417–481, 2007.

[24] N. Segerlind, S. Buss, and R. Impagliazzo. A switching lemma for small restrictions and lower bounds for k-DNF resolution. *SIAM J. Comput.*, 33(5):1171–1200, 2004.

[25] G. Stålmarck. Short resolution proofs for a sequence of tricky formulas. *Acta Informatica*, 33(3):277–280, 1996.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de