Near Unanimity Constraints Have Bounded Pathwidth Duality

Libor Barto†
Mathematics & Statistics Department
McMaster University
Hamilton, Ontario L8S 4L8, Canada
and
Dept. of Algebra, Charles University
Praha 8, Czech Republic
Email: libor.barto@gmail.com

Marcin Kozik‡
Theoretical Computer Science Department
Faculty of Mathematics & Computer Science
Jagiellonian University
Kraków 30-348, Poland
Email: kozik@tcs.uj.edu.pl

Ross Willard§
Pure Mathematics Department
University of Waterloo
Waterloo, Ontario N2L 3G1, Canada
Email: rdwillar@uwaterloo.ca

***Abstract*—We show that if a finite relational structure has a near unanimity polymorphism, then the constraint satisfaction problem with that structure as its fixed template has bounded pathwidth duality (equivalently, the complement of the problem is definable in linear Datalog). As a consequence, the problem is in the complexity class NL. This generalizes the analogous result of Dalmau and Krokhin for majority polymorphisms and lends further support to a conjecture suggested by Larose and Tesson.**

I. Introduction

The *constraint satisfaction problem* (CSP) is a well-known and important protocol for declaring combinatorial problems arising from artificial intelligence [31]. It is also the source of deep research problems in theoretical computer science. In particular, Feder and Vardi [18] identified *fixed-template* versions of CSP as worthy of study and formulated their famous *CSP Dichotomy Conjecture*: for every template (i.e., finite relational structure) $\mathbf{B}$, the problem $\mathrm{CSP}(\mathbf{B})$ is either NP-complete or solvable in polynomial time. Considerable progress towards resolving this conjecture has been achieved during the last 12 years, in part because of the success of the so-called "algebraic approach" championed by Jeavons (e.g. [23], [11]). In this approach templates are classified according to their "polymorphisms," i.e., multi-variable functions that preserve the relations of the template; these functions connect CSP to universal algebra and its toolboxes and perspectives.

An important illustration of the power of the algebraic approach is the recent characterization by the first two authors of templates having *bounded width*. These are structures $\mathbf{B}$ for which $\mathrm{CSP}(\mathbf{B})$ can be solved in polynomial time by a standard local consistency checking algorithm. An "obvious" obstruction to having bounded width is the structure having relations which encode linear equations over some additive abelian group [18]. The algebraic perspective gives a precise, though technical, description ([29], [28]) of the class of templates which omit the obvious obstruction. In [30] a simple characterization of this class in terms of polymorphisms was given, and in [6] this characterization was used to show that every member of the class does indeed have bounded width.

Having bounded width can be characterized in many equivalent ways, including $\mathrm{CSP}(\mathbf{B})$ (here identified with the class of finite structures that admit a homomorphism to $\mathbf{B}$) having *bounded treewidth duality* (see [21], [18], [25], [10], [20], [29]), or the complement class $\neg\mathrm{CSP}(\mathbf{B})$ being definable in the logic Datalog [18]. (*Datalog* is a relational query language whose salient feature is its ability to formulate least-fixed-point recursive definitions [33], [1].)

A related property is that of $\mathrm{CSP}(\mathbf{B})$ having bounded *pathwidth* duality; this more restrictive property puts $\mathrm{CSP}(\mathbf{B})$ in the complexity class NL ([14], [15]) and has several equivalent formulations [15], including $\neg\mathrm{CSP}(\mathbf{B})$ being definable in *linear* Datalog (in which only non-branching recursion is permitted [1]). The "obvious" obstruction to bounded pathwidth duality, in addition to linear equations over an abelian group, is Horn 3-SAT ([2], [13]). Again, universal algebra gives a precise characterization of the class of templates which omit both obstructions [27], and in light of the available evidence (especially [26]) it is natural (as [27] noted) to conjecture that every template in this class has bounded pathwidth duality.

From the algebraic perspective, four reasonable intermediate steps on the journey to verifying this latter conjecture are:

1) Verify the conjecture on the 2-element domain.
2) Prove it for templates having a *majority* polymorphism.
3) Prove it for templates having a *near unanimity* polymorphism in $d+1$ variables for some $d \geq 2$.
4) Prove it for templates having *Jónsson* polymorphisms.

Step 1 was accomplished in [27], [2]. Step 2 was solved by Dalmau and Krokhin [16], who proved that if $\mathbf{B}$ has a majority polymorphism then $\mathrm{CSP}(\mathbf{B})$ has bounded pathwidth duality; they also posed Steps 3 and 4 as next steps. (In fact, Step 2 is the first case of Step 3, i.e. with $d = 2$.) The property stated in Step 3 has been called the *d-mapping property* by Feder

† Research supported by the Grant Agency of the Czech Republic, grant 201/09/P223 and by the Ministry of Education of the Czech Republic, grant MSM 0021620839.

‡ Research supported by the Foundation for Polish Science, grant No. HOM/2008/7 (supported by MF EOG) and National Science Center grant DEC-2011/01/B/ST6/01006.

§ Research supported by the NSERC of Canada.

and Vardi [18], who also showed that $\mathbf{B}$ having this property is equivalent to $\mathbf{B}$ having *bounded strict width* (implying that solutions to $\mathrm{CSP}(\mathbf{B})$ can be found by a greedy algorithm).

In this paper we verify Step 3. That is, we show (Theorem 7) that if a template $\mathbf{B}$ has a near unanimity polymorphism then $\mathrm{CSP}(\mathbf{B})$ has bounded pathwidth duality and hence is in NL. By a result of the first author [4], this also verifies Step 4.

Our proof is inspired by and follows to some extent the proof in [16] for the case $d = 2$. However the details are rather more complicated. In addition, we need (and establish) a surprising new algebraic fact about absorption (Theorem 6) which may be of independent interest to universal algebraists.

The plan of this paper is the following. In section II we summarize the background needed regarding constraint satisfaction problems and templates, bounded pathwidth duality, and algebra. The new algebraic result (Theorem 6) is stated at the end of subsection II-C but its proof is deferred until the end of section III. In section III the main result (Theorem 7) is stated and quickly reduced to the "binary" case; then in subsections III-B and III-C the binary case is proved using Theorem 6; finally in subsection III-D Theorem 6 is proved.

II. Basic Definitions and Tools

A. Structures and Constraint Satisfaction Problems

Everything in this subsection before Definition 1 is standard.

A (*relational*) *vocabulary* is any set of *relation symbols*, each of which is assigned an integer $n \geq 1$ called the *arity* of the symbol. In this paper all relational vocabularies are finite.

If τ is a vocabulary, a *τ-structure* is an object $\mathbf{B}$ consisting of a non-empty set B (the *universe* of $\mathbf{B}$) and, for each relation symbol $R \in \tau$ of arity n, an n-ary *relation* $R^{\mathbf{B}}$ on B, i.e., a subset $R^{\mathbf{B}} \subseteq B^n$. The relations $R^{\mathbf{B}}$ ($R \in \tau$) are the *basic relations* of $\mathbf{B}$. A structure is *finite* if its universe is finite, and is *binary* if every symbol in its vocabulary has arity 1 or 2.

Given two structures $\mathbf{A}, \mathbf{B}$ with the same vocabulary τ, a *homomorphism* from $\mathbf{A}$ to $\mathbf{B}$ is a function $h : A \to B$ which preserves basic relations; that is, for all $R \in \tau$ of arity n and for all $a_1, \ldots, a_n \in A$, if $(a_1, \ldots, a_n) \in R^{\mathbf{A}}$ then $(h(a_1), \ldots, h(a_n)) \in R^{\mathbf{B}}$. $\mathrm{Hom}(\mathbf{A}, \mathbf{B})$ denotes the set of homomorphisms from $\mathbf{A}$ to $\mathbf{B}$. We write $\mathbf{A} \to \mathbf{B}$ to assert $\mathrm{Hom}(\mathbf{A}, \mathbf{B}) \neq \varnothing$.

Given a finite τ-structure $\mathbf{B}$, the *constraint satisfaction problem with fixed template* $\mathbf{B}$ ("homomorphism version") is the decision problem, denoted $\mathrm{CSP}(\mathbf{B})$, which takes as input an arbitrary finite τ-structure $\mathbf{A}$ and asks whether $\mathbf{A} \to \mathbf{B}$.

Given a τ-structure $\mathbf{B}$ with universe B and a subset $X \subseteq B$, the *substructure induced by* $\mathbf{B}$ *on* X is the τ-structure $\mathbf{B}{\restriction}_X$ with universe X and relations defined by $R^{\mathbf{B}\restriction_X} = R^{\mathbf{B}} \cap X^n$ for each n-ary $R \in \tau$.

Given sets $A_1, \ldots, A_n$ and $X \subseteq A_1 \times \cdots \times A_n$, $\mathrm{proj}_i(X)$ denotes the projection of X onto coordinate i. We say that X is *subdirect* and write $X \subseteq_{sd} A_1 \times \cdots \times A_n$ if $\mathrm{proj}_i(X) = A_i$ for all $1 \leq i \leq n$.

Given a τ-structure $\mathbf{B}$, a relation $S \subseteq B^n$ is *$\wedge$-atomic definable over* $\mathbf{B}$ if there exists a formula $\varphi(x_1, \ldots, x_n)$ in the language of first-order logic with equality and vocabulary τ such that (i) φ is a conjunction of *atomic formulas* (assertions "$(x_{i_1}, \ldots, x_{i_n}) \in R$" or "$x_i = x_j$") and (ii) S is defined by φ, i.e., $\{\mathbf{b} \in B^n : \mathbf{B} \models \varphi(\mathbf{b})\} = S$. We say that S is *primitive positive* (or *pp*)-*definable over* $\mathbf{B}$ if, for some $m \geq 0$, S is the projection onto the first n coordinates of an $(n+m)$-ary $\wedge$-atomic-definable relation. If $\mathbf{B}$ is a binary structure, we say that the set of basic relations of $\mathbf{B}$ is *closed under $\wedge$-atomic definitions* provided every at-most-2-ary relation on B which is $\wedge$-atomic definable over $\mathbf{B}$ is already a basic relation of $\mathbf{B}$.

The remaining notions in this subsection are not standard.

Definition 1. Suppose $\mathbf{B}$ is a finite binary τ-structure and A is a finite non-empty set. A *potato system over* $\mathbf{B}$ *with domain* A is an indexed system $\mathcal{P} = (P_a, E_{a,b} : a, b \in A)$ satisfying the following. For all $a, b \in A$:

1) P_a is a 1-ary basic relation of $\mathbf{B}$.
2) $E_{a,b}$ is a 2-ary basic relation of $\mathbf{B}$.
3) $E_{a,b} \subseteq P_a \times P_b$.
4) $E_{a,a} = \{(x, x) : x \in P_a\}$.
5) $E_{b,a} = \{(y, x) : (x, y) \in E_{a,b}\}$.

Potato systems over $\mathbf{B}$ are similar to (1,2)-systems defined in [6]; they differ in that we do not require $E_{a,b} \subseteq_{sd} P_a \times P_b$.

Definition 2. Given a potato system $\mathcal{P} = (P_a, E_{a,b} : a, b \in A)$ over the finite τ-structure $\mathbf{B}$, the *structure associated to* $\mathcal{P}$ is the τ-structure $\mathbf{A}$ with universe A and basic relations defined as follows:

1) (For 1-ary $R \in \tau$): $R^{\mathbf{A}} := \{a \in A : P_a = R^{\mathbf{B}}\}$.
2) (For 2-ary $R \in \tau$): $R^{\mathbf{A}} := \{(a, b) \in A^2 : E_{a,b} = R^{\mathbf{B}}\}$.

A τ-structure is $\mathbf{B}$-*reduced* if it is the structure associated to some potato system over $\mathbf{B}$.

Lemma 1. *Suppose* $\mathbf{B}$ *is a finite binary* τ*-structure whose set of basic relations is closed under* $\wedge$*-atomic definitions. For every finite* τ*-structure* $\mathbf{A}$ *there exists a* $\mathbf{B}$*-reduced* τ*-structure* $\mathbf{A}^\circ$ *having the same domain as* $\mathbf{A}$ *and which satisfies the following: for all* $X \subseteq A$, $\mathrm{Hom}(\mathbf{A}{\restriction}_X, \mathbf{B}) = \mathrm{Hom}(\mathbf{A}^\circ{\restriction}_X, \mathbf{B})$.

B. Bounded Pathwidth Duality

In this section we present the facts we need about pathwidth duality. The following three definitions are from [14], [15].

Definition 3. Let $\mathbf{B}$ be a finite τ-structure. A set $\mathcal{O}$ of finite τ-structures is an *obstruction set for* $\mathrm{CSP}(\mathbf{B})$ if for all finite τ-structures $\mathbf{A}$, $\mathbf{A} \not\to \mathbf{B}$ if and only if there exists $\mathbf{C} \in \mathcal{O}$ with $\mathbf{C} \to \mathbf{A}$.

Obstruction sets are useful when their members are simple. One way they can be simple is by having bounded pathwidth.

Definition 4. A finite τ-structure $\mathbf{C}$ has *pathwidth at most* (j, k) if there is a sequence $\mathcal{I} = (I_0, \ldots, I_N)$ of subsets of C such that:

1) $|I_t| \leq k$ for all t, and $|I_t \cap I_{t+1}| \leq j$ for all $t < N$.
2) $I_i \cap I_j \subseteq I_\ell$ for all $0 \leq i \leq \ell \leq j \leq N$.
3) For every $R \in \tau$ of arity n and every $(c_1, \ldots, c_n) \in R^{\mathbf{C}}$ there exists $t \leq N$ with $\{c_1, \ldots, c_n\} \subseteq I_t$.

The sequence $\mathcal{I}$ is called a (j, k)-*path decomposition* of $\mathbf{C}$.

Definition 5. Let $\mathbf{B}$ be a finite τ-structure, let $\mathcal{O}$ be a set of finite τ-structures, and let $0 \le j \le k$.

1) $\mathcal{O}$ has *pathwidth at most* (j,k) if every $\mathbf{C} \in \mathcal{O}$ has pathwidth at most (j,k).
2) $\mathrm{CSP}(\mathbf{B})$ has (j,k)*-pathwidth duality* if $\mathrm{CSP}(\mathbf{B})$ has an obstruction set of pathwidth at most (j,k).
3) $\mathrm{CSP}(\mathbf{B})$ has *bounded pathwidth duality* if $\mathrm{CSP}(\mathbf{B})$ has (j',k')-pathwidth duality for some $0 \le j' \le k'$.

The characterization of bounded pathwidth duality that will be most useful to us in this paper is one involving the following variation of Dalmau's "pebble relation games" [15].

Definition 6. Suppose $\mathbf{A},\mathbf{B}$ are finite τ-structures.

1) A *solo play of the* (j,k)*-PR game on* $(\mathbf{A},\mathbf{B})$ is a finite sequence $\mathcal{I} = (I_0, I_1, \dots, I_N)$ of subsets of A satisfying
   a) $|I_t| \le k$ for all $t \le N$.
   b) For all $t < N$, either $I_{t+1} \subseteq I_t$ or $I_t \subset I_{t+1}$. If the latter, then $|I_t| \le j$.
2) Given a solo play $\mathcal{I} = (I_0, \dots, I_N)$ of the (j,k)-PR game on $(\mathbf{A},\mathbf{B})$, the *resulting relations* $H_0, H_1, \dots, H_N$ are defined recursively as follows:
   a) $H_0 = \mathrm{Hom}(\mathbf{A}{\restriction}_{I_0}, \mathbf{B})$.
   b) If $t < N$ and $I_{t+1} \subseteq I_t$, then $H_{t+1} = H_t{\restriction}_{I_{t+1}}$.
   c) If $t < N$ and $I_t \subset I_{t+1}$, then $H_{t+1} = \{h \in \mathrm{Hom}(\mathbf{A}{\restriction}_{I_{t+1}}, \mathbf{B}) : h{\restriction}_{I_t} \in H_t\}$.
3) We write $\mathbf{A} \looparrowright_{j,k} \mathbf{B}$ to mean that for every solo play $\mathcal{I} = (I_0, \dots, I_N)$ of the (j,k)-PR game on $(\mathbf{A},\mathbf{B})$, the final resulting relation H_N is non-empty.

Solo plays and their resulting relations correspond to plays of Dalmau's two-player pebble relation game [15] where Spoiler chooses each set I_t and the resulting relation H_t is Duplicator's maximum allowable response. In particular, $\mathbf{A} \looparrowright_{j,k} \mathbf{B}$ if and only if Duplicator has a *strict* (in Dalmau's sense) winning strategy for the two-player pebble relation game played on $(\mathbf{A},\mathbf{B})$. Thus:

Proposition 2. ([15, Theorem 5 and Claim 1, p. 15]) *Let* $\mathbf{B}$ *be a finite* τ*-structure and* $j \le k$*. The following are equivalent:*

1) $\mathrm{CSP}(\mathbf{B})$ *has* (j,k)*-pathwidth duality.*
2) *For all finite* τ*-structures* $\mathbf{A}$*, if* $\mathbf{A} \looparrowright_{j,k} \mathbf{B}$ *then* $\mathbf{A} \to \mathbf{B}$.

Combining Proposition 2 with Lemma 1 we get:

Corollary 3. *Suppose* $\mathbf{B}$ *is a binary* τ*-structure whose set of basic relations is closed under* $\wedge$*-atomic definitions. For any* $0 \le j \le k$*, the following are equivalent:*

1) $\mathrm{CSP}(\mathbf{B})$ *has* (j,k)*-pathwidth duality.*
2) *For all finite* $\mathbf{B}$*-reduced* τ*-structures* $\mathbf{A}$*, if* $\mathbf{A} \looparrowright_{j,k} \mathbf{B}$ *then* $\mathbf{A} \to \mathbf{B}$.

C. Algebra

In this section we summarize the algebraic background needed in this paper. More in-depth treatments may be found in [12], [8]. Everything preceding Definition 8 is standard.

Given a non-empty set A, an *operation* on A is any function $\phi : A^n \to A$ for some $n \ge 1$; n is the *arity* of ϕ. An operation ϕ is *idempotent* if it satisfies the equation $\phi(x,x,\dots,x) = x$ for all $x \in A$. A 3-ary operation $\phi : A^3 \to A$ is a *majority operation on* A provided it is idempotent and satisfies the equations $\phi(y,x,x) = \phi(x,y,x) = \phi(x,x,y) = x$ for all $x,y \in A$. More generally, an n-ary operation $\phi : A^n \to A$ for $n \ge 3$ is a *near unanimity* (or *NU*) *operation on* A provided it is idempotent and for all $1 \le i \le n$ it satisfies

$$\phi(\underbrace{x,\dots,x}_{i-1}, y, \underbrace{x,\dots,x}_{n-i}) = x \qquad \text{for all } x,y \in A.$$

An *algebraic vocabulary* is any set (possibly infinite) of *operation symbols*, each of which has an assigned *arity* $n \ge 1$. If τ is an algebraic vocabulary, an *algebra of type* τ is an object $\mathbb{A}$ consisting of a non-empty set A (the universe) and, for each each operation symbol $\mathsf{f} \in \tau$ of arity n, an n-ary operation $\mathsf{f}^{\mathbb{A}}$ on A. The operations $\mathsf{f}^{\mathbb{A}}$ ($\mathsf{f} \in \tau$) are the *basic operations* of $\mathbb{A}$. An algebra is *finite* if its universe is finite, and is *idempotent* if each of its basic operations is idempotent.

Suppose $\mathbb{A}$ is an algebra of type τ and $X \subseteq A$. X is a *subuniverse* of $\mathbf{A}$ if X is closed under every basic operation of $\mathbf{A}$; that is, if $\mathsf{f}^{\mathbb{A}}(X^n) \subseteq X$ for all n-ary $\mathsf{f} \in \tau$. We denote this by $X \le \mathbb{A}$; if in addition $X \subseteq_{sd} A$ then we write $X \le_{sd} \mathbb{A}$. If $\varnothing \ne X \le \mathbb{A}$, the *subalgebra of* $\mathbb{A}$ *with universe* X is the algebra $\mathbb{X}$ of type τ whose operations are given by $\mathsf{f}^{\mathbb{X}} = \mathsf{f}^{\mathbb{A}}{\restriction}_X$.

For every $X \subseteq A$ there is a unique smallest subuniverse of $\mathbb{A}$ containing X, which is denoted $\mathrm{Sg}^{\mathbb{A}}(X)$.

We use the following device: any set $\mathcal{F}$ of operations on A can be considered as an algebraic vocabulary in the obvious way, making $(A, \mathcal{F})$ an algebra of type $\mathcal{F}$.

Two algebras are *similar* if they have the same vocabulary. The *product* of any number of similar algebras is defined naturally, that is, by defining operations coordinatewise. We sometimes use the same notation, e.g. ϕ, for both an operation symbol and its interpretations in similar algebras $\mathbb{A}, \mathbb{B}$, etc.

Definition 7. Given a structure $\mathbf{A}$ with universe A, an n-ary relation R on A, and an m-ary operation ϕ on A, we say that:

1) ϕ *preserves* R provided for all $\mathbf{a}_1, \dots, \mathbf{a}_m \in R$, if $\mathbf{c}_1, \dots, \mathbf{c}_n \in A^m$ are the columns of the matrix whose rows are $\mathbf{a}_1, \dots, \mathbf{a}_m$, then $(\phi(\mathbf{c}_1), \dots, \phi(\mathbf{c}_n)) \in R$.
2) ϕ is a *polymorphism* of $\mathbf{A}$ if ϕ preserves every basic relation of $\mathbf{A}$.

Let $\mathrm{Pol}(\mathbf{A})$ denote the set of all polymorphisms of $\mathbf{A}$. The *polymorphism algebra* of $\mathbf{A}$ is the algebra $\mathbb{A} = (A, \mathrm{Pol}(\mathbf{A}))$ (of type $\mathrm{Pol}(\mathbf{A})$). It is a well-known fact (e.g., [19], [9], [24]) that if R is an arbitrary nonempty n-ary relation on A, then R is pp-definable over $\mathbf{A}$ if and only if $R \le \mathbb{A}^n$.

The next definition slightly extends a notion from [7], [6].

Definition 8. Suppose $\mathbb{A}$ is an algebra, $B, C \le \mathbb{A}$, and ϕ is an n-ary operation of $\mathbb{A}$. We say that B *absorbs* C *with respect to* ϕ provided the following condition holds: for all $1 \le i \le n$, all $b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n \in B$ and all $c \in C$, $\phi(b_1, \dots, b_{i-1}, c, b_{i+1}, \dots, b_n) \in B$. We write $B \lhd_\phi C$ to mean $B \subseteq C$ and B absorbs C with respect to ϕ.

Here are two easy facts about absorption.

Lemma 4. *Suppose ϕ is an operation of the algebra $\mathbb{A}$.*

1) *If $B \lhd_\phi C \leq \mathbb{A}$ then $B \cap D \lhd_\phi C \cap D$ for any $D \leq \mathbb{A}$.*
2) *If $\mathbb{A}$ is idempotent, then ϕ is an NU operation if and only if $\{a\} \lhd_\phi A$ for all $a \in A$.*

The following claim is a good exercise, or can be extracted from [7, Lemma 2.5].

Lemma 5. *Suppose $\mathbb{B}_1, \mathbb{C}_1$ are similar algebras, ϕ is an operation symbol in their common vocabulary, $S \leq \mathbb{B}_1 \times \mathbb{C}_1$ with $\mathrm{proj}_1(S) = B_1$, and $C_0 \lhd_\phi C_1$. Define $B_0 = \{b \in B_1 : \exists c \in C_0 \text{ with } (b,c) \in S\}$. Then $B_0 \lhd_\phi B_1$.*

We need one new result about absorption, whose proof will be postponed until subsection III-D. Given an algebra $\mathbb{D}$ and an integer $n \geq 2$, define $0_D^{(n)} = \{(b, b, \ldots, b) : b \in D\} \subseteq D^n$, the set of constant n-tuples over D. Note that $0_D^{(n)} \leq \mathbb{D}^n$.

Definition 9. Let $\mathbb{D}$ be an algebra, ϕ an operation of $\mathbb{D}$, and $b \in D$. We call b an *absorption constant for $\mathbb{D}$ with respect to ϕ* provided, for all $n \geq 2$ and every $R \leq_{sd} \mathbb{D}^n$, if R absorbs $0_D^{(n)}$ with respect to ϕ then $(b, b, \ldots, b) \in R$.

Theorem 6. *Let $\mathbb{D}$ be a finite algebra and ϕ an idempotent operation of $\mathbb{D}$. There exists an absorption constant for $\mathbb{D}$ with respect to ϕ.*

III. Main Result

A. Statement and Reduction to the Binary Case

The main result of this paper is the following.

Theorem 7. *Suppose the finite τ-structure $\mathbf{B}$ has a $(d+1)$-ary NU polymorphism for some $d \geq 2$. Then $\mathrm{CSP}(\mathbf{B})$ has bounded pathwidth duality and hence is in NL.*

The rest of the paper is devoted to proving this theorem. In this subsection we will reduce it to the case of binary structures. The first reduction is a variant of one step in the proof of the well-known "CSP reduction to digraphs" [18, Theorem 11] (see also e.g. the proof of [5, Theorem 4.4]). It applies to any structure.

Lemma 8. *Suppose τ is a vocabulary, $n \geq 1$ is an integer such that every relation symbol in τ has arity at most $2n$, $\mathbf{B}$ is a finite τ-structure, and $\mathbb{B}$ is its polymorphism algebra. There exists a binary structure $\mathbf{B}^{(n)}$ with universe B^n such that:*

1) *The polymorphism algebra of $\mathbf{B}^{(n)}$ is $\mathbb{B}^n$.*
2) *For any $0 \leq j \leq k$, if $\mathrm{CSP}(\mathbf{B}^{(n)})$ has (j,k)-pathwidth duality, then $\mathrm{CSP}(\mathbf{B})$ has (jn, kn)-pathwidth duality.*

The next reduction applies only to structures with an NU polymorphism. It is the obvious and straightforward generalization of [16, Lemma 2] for majority polymorphisms.

Lemma 9. *Suppose $\mathbf{B}$ is a finite τ-structure with universe B and a $(d+1)$-ary NU polymorphism for some $d \geq 2$. Let $s = \max(\{arity(R) : R \in \tau\} \cup \{d\})$. There exists a vocabulary τ_d and a τ_d-structure $\mathbf{B}_d$ with universe B satisfying:*

1) *Every relation symbol in τ_d has arity at most d.*
2) *$\mathbf{B}$ and $\mathbf{B}_d$ have the same polymorphisms.*
3) *If $\mathrm{CSP}(\mathbf{B}_d)$ has (j,k)-pathwidth duality, then $\mathrm{CSP}(\mathbf{B})$ has $(k, k+s-d)$-pathwidth duality.*

Lemmas 8 and 9 reduce the task of proving Theorem 7 to proving it for the special case of binary structures. In subsections III-C and III-D we will verify this special case by proving the following.

Proposition 10. *Suppose $d \geq 2$ and $\mathbf{B}$ is a binary structure having a $(d+1)$-ary NU polymorphism. Let $k = |B|$, $c = \lfloor \log_3(2d-3) \rfloor + 2$, and $p = 2c^k - k - 1$. Then $\mathrm{CSP}(\mathbf{B})$ has $(p, p+1)$-pathwidth duality. If $d = 2$ then $\mathrm{CSP}(\mathbf{B})$ has $(2k, 2k+1)$-pathwidth duality.*

Hence we get the following sharpening of Theorem 7.

Corollary 11. *Suppose the finite τ-structure $\mathbf{B}$ has a $(d+1)$-ary NU polymorphism for some $d \geq 2$. Let k, p be defined as in Proposition 10, let $q = \lceil d/2 \rceil (p+1)$, and let s be defined as in Lemma 9. Then $\mathrm{CSP}(\mathbf{B})$ has $(q, q+s-d)$-pathwidth duality.*

Proof of Theorem 7 and Corollary 11: Given $\mathbf{B}$, Let $\mathbf{B}_d$ be the structure defined in Lemma 9, let $e = \lceil d/2 \rceil$, and let $(\mathbf{B}_d)^{(e)}$ be the structure obtained from $\mathbf{B}_d$ via Lemma 8. Both $\mathbf{B}_d$ and $(\mathbf{B}_d)^{(e)}$ inherit $(d+1)$-ary NU polymorphisms from $\mathbf{B}$. As $(\mathbf{B}_d)^{(e)}$ is binary, $(\mathbf{B}_d)^{(e)}$ has $(p, p+1)$-pathwidth duality by Proposition 10; hence $\mathbf{B}$ has $(q, q+s-d)$-pathwidth duality by Lemmas 8 and 9. ■

B. A-Trees

Proposition 10 will be proved via an intricate analysis of realizations of certain trees in τ-structures. In this subsection we define these trees and state some facts about them that will be needed in subsection III-C.

Following [15], [16], if $G = (V, E)$ is an undirected graph, we denote by $\mathrm{pw}(G)$ the least k for which G has pathwidth (j, k) for some $j \leq k$, and call $\mathrm{pw}(G)$ the *pathwidth* of G. (Note that $\mathrm{pw}(G)$ is 1 greater than the usual graph-theoretic measure of the pathwidth of G as defined in [32].)

Definition 10. Let A be a non-empty set. An *A-tree* is a pair (T, χ) where $T = (V, E)$ is a tree (i.e., a connected undirected graph with no cycles) and χ is a coloring of the vertices of T by elements of A (i.e., $\chi : V \to A$).

Definition 11. Let $T_0, \ldots, T_n$ be trees on disjoint vertex sets. A *tree composition* of $T_0, \ldots, T_n$ is any tree T that can be constructed from the union of $T_0, \ldots, T_n$ by identifying some vertices among the leaves of $T_0, \ldots, T_n$ (enough to connect the graph, but not so many as to introduce cycles).

If T is a tree composition of $T_0, \ldots, T_n$, then the vertices of T which were formed by identifying leaves of the original trees are called the *composition vertices* of T. The subtrees of T corresponding to the original trees $T_0, \ldots, T_n$ are called the *components* of T. Among the components of T we distinguish the *leaf components*, which are those which had at most one leaf identified in the construction of T. (Note that, strictly speaking, the notions in this paragraph are not invariants of T itself, but of a particular composition that produces T. Whenever discussing tree compositions we shall assume that such trees come equipped with a record of a particular composition that produced it.)

Definition 12. Suppose $T = (V, E)$ is a tree composition and $L \subseteq A$. An A-coloring $\chi : V \to A$ is said to *lie over* L if χ sends all composition vertices of T into L.

Definition 13. Suppose $d \geq 2$.

1) $\mathcal{T}_d^0$ is the set of all trees consisting of exactly one edge.
2) For $i \geq 1$, $\mathcal{T}_d^{i+1}$ is the set of all tree compositions having components from $\mathcal{T}_d^i$ and at most d leaf components.

Note that $\mathcal{T}_d^0 \subseteq \mathcal{T}_d^1 \subseteq \mathcal{T}_d^2 \subseteq \cdots$. Note also that for $d = 2$ we have $\mathcal{T}_2^1 = \mathcal{T}_2^2 = \cdots = \{\text{all paths}\}$; hence every tree in $\bigcup_i \mathcal{T}_2^i$ has pathwidth at most 2. For $d > 2$ we have the following:

Lemma 12. *Suppose $d \geq 3$ and $i \geq 1$. The trees in $\mathcal{T}_d^i$ have bounded pathwidth. More precisely, every tree $T \in \mathcal{T}_d^i$ satisfies $\mathrm{pw}(T) \leq c^i + c^{i-1} - 1$ where $c = \lfloor \log_3(2d-3) \rfloor + 2$.*

Definition 14. Suppose A is a non-empty set, $L \subseteq A$, $d \geq 2$, and $i \geq 0$. $\mathcal{T}_d^i(A)$ denotes the set of all A-trees (T, χ) where $T \in \mathcal{T}_d^i$. If $i > 0$ we also define $\mathcal{T}_d^i(A, L) = \{(T, \chi) \in \mathcal{T}_d^i(A) : \chi \text{ lies over } L\}$.

Note that for $i > 0$ we have $\mathcal{T}_d^i(A, A) = \mathcal{T}_d^i(A)$, and $L \subseteq L'$ implies $\mathcal{T}_d^i(A, L) \subseteq \mathcal{T}_d^i(A, L')$. Also note that $\mathcal{T}_d^i(A) = \mathcal{T}_d^{i+1}(A, \varnothing)$ for all i.

We will use A-trees in the following way. Suppose $\mathbf{B}$ is a binary τ-structure whose set of basic relations is closed under $\wedge$-atomic definitions. Let $\mathcal{P} = (P_a, E_{a,b} : a, b \in A)$ be a potato system over $\mathbf{B}$ indexed by A. Given an A-tree (T, χ) with $T = (V, E)$, a *realization of* (T, χ) *in* $\mathcal{P}$ is a function $f : V \to B$ satisfying the following:

1) $f(u) \in P_{\chi(u)}$ for all $u \in V$.
2) $(f(u), f(v)) \in E_{\chi(u), \chi(v)}$ for all $\{u, v\} \in E$.

The proof of the next lemma is straightforward.

Lemma 13. *Suppose $\mathbf{B}$ and $\mathcal{P}$ are as above, (T, χ) is an A-tree with $T = (V, E)$, $U = \{u_1, \dots, u_n\}$ is a non-empty subset of V, and for each $v \in V$ we have an associated subset B_v of B which is pp-definable over $\mathbf{B}$. Then the following relations are pp-definable over $\mathbf{B}$:*

1) $R := \{(f(u_1), \dots, f(u_n)) \in B^n : f$ *is a realization of* (T, χ) *in* $\mathcal{P}$ *and* $f(v) \in B_v$ *for all* $v \in V\}$.
2) $S := \{b \in B : (b, \dots, b) \in R\}$.

C. Proof of Proposition 10

In this subsection we prove Proposition 10, modulo the proof of Theorem 6. Much of the argument in this subsection is inspired by the proof of [16, Theorem 1].

Let $\mathbf{B}$ be a binary τ-structure having a $(d+1)$-ary NU polymorphism ϕ. Because adding (arbitrary) relations to the vocabulary of $\mathbf{B}$ cannot decrease the pathwidth of obstruction sets for its constraint satisfaction problem, we may assume that the set of basic relations of $\mathbf{B}$ contains every ϕ-invariant unary and binary relation on B. In particular, the set of basic relations of $\mathbf{B}$ contains all 1-element subsets of B and is closed under $\wedge$-atomic definitions. Let k, c, p be defined as in the statement of Proposition 10, and let $r = p$ if $d \geq 3$ while $r = 2k$ if $d = 2$. Let $\mathbf{A}$ be a $\mathbf{B}$-reduced τ-structure such that $\mathbf{A} \looparrowright_{r,r+1} \mathbf{B}$. By Corollary 3, to prove Proposition 10 it suffices to show that $\mathbf{A} \to \mathbf{B}$.

Let $\mathcal{P} = (P_a, E_{a,b} : a, b \in A)$ be the potato-system to which $\mathbf{A}$ is associated. For each $a \in A$, define a sequence of "levels" $P_a^0 \supseteq P_a^1 \supseteq P_a^2 \supseteq \cdots$ within P_a as follows:

1) $P_a^0 = P_a$.
2) If $i \geq 0$ and $(T, \chi) \in \mathcal{T}_d^{i+1}(A)$ with $T = (V, E)$, then $P_a^{i+1}(T, \chi)$ is the set of $b \in P_a^i$ for which there exists a realization f of (T, χ) in $\mathcal{P}$ such that
   a) f maps V into level i; i.e., $f(u) \in P_{\chi(u)}^i$ for all $u \in V$.
   b) f maps each a-labelled vertex to b; i.e., $f(u) = b$ for all vertices $u \in V$ such that $\chi(u) = a$.
3) $P_a^{i+1} = \bigcap \{P_a^{i+1}(T, \chi) : (T, \chi) \in \mathcal{T}_d^{i+1}(A)\}$.

Lemma 13 implies that each set P_a^i is pp-definable over $\mathbf{B}$.

Claim 1. *$P_a^k \neq \varnothing$ for all $a \in A$.*

Proof sketch: For $i \geq 0$ define $g(i) = 2i$ if $d = 2$ and $g(i) = (c+1)(c-1)^{-1}(c^i - 1) - i$ otherwise. We claim that for all $i \geq 0$ and all $a \in A$, there exists a solo play $\mathcal{I}(i, a) = (I_0, I_1, \dots, I_n)$ of the $(g(i), g(i)+1)$-PR game on $(\mathbf{A}, \mathbf{B})$ such that $a \in I_j$ for all $j \leq n$, $I_n = \{a\}$, and the final resulting relation of $\mathcal{I}(i, a)$ is contained in P_a^i. This will suffice, since $\mathbf{A} \looparrowright_{r,r+1} \mathbf{B}$ and $g(k) \leq r$ imply that the final resulting relation of each $\mathcal{I}(k, a)$ is non-empty.

The proof of the claim is by induction on i. If $i = 0$, then we can choose $\mathcal{I}(0, a) = (I_0)$ where $I_0 = \{a\}$. Now assume that $i \geq 0$ and the claim has been verified for i. For each $a' \in A$ inductively choose and fix a solo play $\mathcal{I}(a')$ of the $(g(i), g(i)+1)$-PR game each of whose sets contains a', whose final set is $\{a'\}$, and whose final resulting relation is contained in $P_{a'}^i$. Fix $a \in A$.

Call a solo play $\mathcal{I} = (I_0, \dots, I_n)$ of the $(g(i+1), g(i+1)+1)$-PR game on $(\mathbf{A}, \mathbf{B})$ an *a-play* if $a \in I_j$ for all $j \leq n$ and $I_n = \{a\}$. If $\mathcal{I}, \mathcal{J}$ are a-plays with final resulting relations R, S respectively, then the concatenation of $\mathcal{I}$ and $\mathcal{J}$ is also an a-play and its final resulting relation is contained in $R \cap S$. Thus it suffices to show that for every $(T, \chi) \in \mathcal{T}_d^{i+1}(A)$ there exists an a-play $\mathcal{I}_a^{T,\chi}$ whose final resulting relation is contained in $P_a^{i+1}(T, \chi)$. Fix $(T, \chi) \in \mathcal{T}_d^{i+1}(A)$ with $T = (V, E)$.

1) Using Lemma 12 and the comment preceding it, let $(J_0, \dots, J_m)$ be a $(t, t+1)$-path decomposition of T where $t = 1$ if $d = 2$ and $t = c^{i+1} + c^i - 2$ otherwise.
2) Let $\{a_1^0, \dots, a_{k_0}^0\}$ be an enumeration of $\chi(J_0)$. For each $1 \leq j \leq m$, let $\{a_1^j, \dots, a_{k_j}^j\}$ be an enumeration of $\chi(J_j) \setminus \chi(J_{j-1})$.
3) For each $0 \leq j \leq m$ and $1 \leq \ell \leq k_j$, let $\mathcal{I}_a^*(j, \ell)$ be the play defined as follows: if $\mathcal{I}(a_\ell^j) = (I_0, I_1, \dots, I_t)$ then $\mathcal{I}_a^*(j, \ell) = (I_0^*, I_1^*, \dots, I_t^*)$ where $I_u^* = I_u \cup \chi(J_j) \cup \{a\}$. (Note in particular that the final set is $I_t^* = \chi(J_j) \cup \{a\}$.) Also define $\mathcal{I}_a^*(j, 0)$ to be the 1-step play $(\chi(J_j) \cup \{a\})$.

4) For each $0 \leq j \leq m$ let $\mathcal{I}^*_a(j)$ be the concatenation of the plays $\mathcal{I}^*_a(j,0), \mathcal{I}^*_a(j,1), \dots, \mathcal{I}^*_a(j,k_j)$.
5) Let $\mathcal{I}^{T,\chi}_a$ be the concatenation of $\mathcal{I}(a), \mathcal{I}^*_a(0), \dots, \mathcal{I}^*_a(m)$, and the 1-step play $(\{a\})$.

Suppose $0 \leq j \leq m$ and $0 \leq \ell \leq k_j$. Let $\mathcal{I}^\star$ denote the initial segment of $\mathcal{I}^{T,\chi}_a$ consisting of $\mathcal{I}(a)$, $\mathcal{I}^{T,\chi}_a(0), \dots, \mathcal{I}^{T,\chi}_a(j-1)$ and that portion of $\mathcal{I}^{T,\chi}_a(j)$ up to and including $\mathcal{I}^*_a(j,\ell)$. Note that the last set of $\mathcal{I}^*$ is $\chi(J_j) \cup \{a\}$. Let $V^* = J_0 \cup \cdots \cup J_j$, let $C = \{a^j_1, \dots, a^j_\ell\}$, and let T^* be the subtree of T with vertex set V^*. It can be shown inductively that if h is in the final resulting relation of $\mathcal{I}^*$ then $h(a) \in P^i_a$ and there exists a realization f of the A-tree $(T^*, \chi\restriction_{T^*})$ such that

1) f maps $J_0 \cup \cdots \cup J_{j-1}$ into level i;
2) f maps $J_j \cap \chi^{-1}(\{a^j_1, \dots, a^j_\ell\})$ into level i;
3) $f(u) = h(\chi(u))$ for all $u \in J_j$.
4) f maps all a-labelled vertices of T^* to $h(a)$.

In particular, if h is in the penultimate resulting relation of $\mathcal{I}^{T,\chi}_a$ then $h(a) \in P^{i+1}_a(T,\chi)$; hence the final resulting relation of $\mathcal{I}^{T,\chi}_a$ is contained in $P^{i+1}_a(T,\chi)$.

It remains to check that every set in $\mathcal{I}^{T,\chi}_a$ has size at most $g(i+1)+1$. Note that any set in $\mathcal{I}^{T,\chi}_a$ of maximal size is of the form $I_u \cup \chi(J_j) \cup \{a\}$ where I_u is a set in $\mathcal{I}(a')$ for some $a' \in \chi(J_j)$. We have $|J_j| \leq t+1$ and $|I_u| \leq g(i)+1$ and $|I_u \cap \chi(J_j)| \geq 1$. Thus $|I_u \cup \chi(J_j) \cup \{a\}| \leq g(i)+t+2 = g(i+1)+1$ as required. ∎

Thus for each $a \in A$ we have a chain $P^0_a \supseteq P^1_a \supseteq \cdots \supseteq P^k_a$ of $k+1$ non-empty subsets of B, where $|B| = k$. Hence there exists $r < k$ such that $P^r_a = P^{r+1}_a$. Let r_a be the least r with this property. Enumerate A as $\{a_0, a_1, \dots, a_N\}$ so that $r_{a_0} \geq r_{a_1} \geq \cdots$ and for each $j \leq N$ define $\mathrm{rank}(j) = r_{a_j}$.

Our aim is to construct a homomorphism $h : \mathbf{A} \to \mathbf{B}$, and we will do this by inductively defining $h(a_0), h(a_1)$, etc. At stage i we will have defined $h(a_0), h(a_1), \dots, h(a_i)$. In this context we will consider certain A-trees and their realizations in $\mathcal{P}$. Suppose (T,χ) is an A-tree and f is a realization of (T,χ) in $\mathcal{P}$. If Λ_T is the set of leaves of T and $U \subseteq \Lambda_T$, then we say that f is *fixed on U up to i* if for all $u \in U$, if $\chi(u) \in \{a_0, \dots, a_i\}$ then $f(u) = h(\chi(u))$. We say that f is *fixed up to i* if it is fixed on Λ_T up to i.

The inductive property that we will establish at stage i is the following:

1) $h(a_j) \in P^{\mathrm{rank}(j)}_{a_j}$ for all $j \leq i$.
2) Let $r = \mathrm{rank}(i)$ and $L = \{a_j : j \leq N, \mathrm{rank}(j) = r\}$. Then for every A-tree $(T,\chi) \in \mathcal{T}^{r+1}_d(A,L)$ with $T = (V,E)$ there exists a realization of (T,χ) in $\mathcal{P}$ which sends V into level r and is fixed up to i.

This will suffice since at stage N we will have fully defined a function $h : A \to B$ satisfying property 1) above; moreover, as $\mathcal{T}^{r+1}_d(A,L) \supseteq \mathcal{T}^0_d(A)$, property 2) above can be applied as follows: for any $a_i, a_j \in A$, if we let (T,χ) be the 2-vertex A-tree $\{u,v\}$ with $\chi(u) = a_i$ and $\chi(v) = a_j$, and if f is a realization of (T,χ) in $\mathcal{P}$ which is fixed up to N, then $(h(a_i), h(a_j)) = (f(u), f(v)) \in E_{\chi(u),\chi(v)} = E_{a_i,a_j}$, proving h is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$.

At stage 0 we can define $h(a_0)$ to be any element of $P^r_{a_0}$ where $r = \mathrm{rank}(0)$. This works as $P^r_{a_0} = P^{r+1}_{a_0}$, so the definition of $P^{r+1}_{a_0}$ and the fact that $\mathcal{T}^{r+1}_d(A,L) \subseteq \mathcal{T}^{r+1}_d(A)$ imply property 2) above.

Assume that we have finished stage $i-1$ and want to define $h(a_i)$. Define $r^* = \mathrm{rank}(i-1)$, $L^* = \{a_j : \mathrm{rank}(j) = r^*\}$, $r = \mathrm{rank}(i)$, and $L = \{a_j : \mathrm{rank}(j) = r\}$.

Claim 2. *For every A-tree $(T,\chi) \in \mathcal{T}^{r+1}_d(A,L)$ with $T = (V,E)$ there exists a realization of (T,χ) in $\mathcal{P}$ which maps V into level r and is fixed up to $i-1$.*

Proof: If $r^* = r$, then $L^* = L$ and the claim follows directly from the inductive property at stage $i-1$. If on the other hand $r^* > r$, then use the fact that $\mathcal{T}^{r+1}_d(A,L) \subseteq \mathcal{T}^{r^*+1}_d(A,\varnothing) \subseteq \mathcal{T}^{r^*+1}_d(A,L^*)$ and the inductive property. ∎

Definition 15. Let $\mathcal{T}^*$ denote the set of all A-trees (T,χ) where T is an (arbitrary) tree composition of component trees from $\mathcal{T}^r_d$ and χ lies over L. (Thus $\mathcal{T}^{r+1}_d(A,L) \subseteq \mathcal{T}^*$.)

Claim 3. *Suppose $(T,\chi) \in \mathcal{T}^*$ with $T = (V,E)$.*

1) *For every leaf u of T such that $\chi(u) = a \in L$, and for every $b \in P^r_a$, there exists a realization f of (T,χ) in $\mathcal{P}$ which maps V into level r and satisfies $f(u) = b$.*
2) *There exists a realization of (T,χ) in $\mathcal{P}$ which maps V into level r and is fixed up to $i-1$.*

Proof: (1) Let $T_0 = (V_0, E_0)$ be the component of T containing u, let $\chi_0 = \chi\restriction_{V_0}$, and note that $(T_0,\chi_0) \in \mathcal{T}^{r+1}_d(A)$. Because $a \in L$ we have $P^r_a = P^{r+1}_a$. Thus if $b \in P^r_a$ then by definition of P^{r+1}_a there exists a realization f_0 of (T_0,χ_0) in $\mathcal{P}$ which sends V_0 into level r and maps u to b. Consider a composition vertex u' of T which lies in V_0 and let $a' = \chi(u')$ and $b' = f(u')$. We have $a' \in L$ (as χ lies over L) and $b' \in P^r_{a'}$ by our choice of f; hence $b' \in P^{r+1}_{a'}$. If we now let $T_1 = (V_1, E_1)$ be another component of T containing u' and put $\chi_1 = \chi\restriction_{V_1}$, then we can repeat the previous argument to get a realization f_1 of (T_1,χ_1) in $\mathcal{P}$ which sends V_1 into level r and maps u' to b'; in particular, $f_1(u') = f_0(u')$. Repeating this process one component of T at a time, we can eventually construct the required realization f of (T,χ).

(2) Let Λ_T denote the set of leaves of T. It will suffice to prove that, for every $U \subseteq \Lambda_T$, there exists a realization of (T,χ) in $\mathcal{P}$ which maps V into level r and is fixed on U up to $i-1$. We do this by induction on $|U|$. Assume first that $|U| \leq d$. Let $T' = (V',E')$ be the smallest subtree of T which is a tree composition of (some of) the components of T and contains U, and let $\chi' = \chi\restriction_{T'}$. Also let $C_1, \dots, C_m$ be the connected components of $T \setminus T'$. Observe that for each C_j there exists a composition vertex u_j of T such that

1) u_j is either a composition vertex or a leaf of T'.
2) The subtree of T with vertex set $C_j \cup \{u_j\}$ is a tree composition of (some of the) components of T which are not components of T'. Call this subtree T_j and let $\chi_j = \chi\restriction_{T_j}$. Note that $(T_j,\chi_j) \in \mathcal{T}^*$.

T' has at most d leaf components and so $(T',\chi') \in \mathcal{T}^{r+1}_d(A,L)$. By Claim 2 there exists a realization f' of

$(T', \chi\restriction_{T'})$ which maps V' into level r and is fixed up to $i-1$; in particular, f' is fixed on U up to $i-1$. By applying part (1) to the A-trees (T_j, χ_j) we can extend f' to the desired realization f of (T, χ).

Assume next that $|U| > d$. Choose distinct leaves $u_0, u_1, \dots, u_d \in U$ and for each $j \leq d$ let $U_j = U \setminus \{u_j\}$. By induction, there exist realizations f_j of (T, χ) in $\mathcal{P}$ which map V into level r and are such that f_j is fixed on U_j up to $i-1$. Let $f : V \to B$ be defined by

$$f(v) = \phi(f_0(v), f_1(v), \dots, f_d(v)).$$

Because ϕ is a polymorphism of $\mathbf{B}$, f is a realization of (T, χ) which sends V into level r. It remains to check that f is fixed on U up to $i-1$; this follows from the NU property of ϕ. ■

Claim 4. *There exists a non-empty set $D \subseteq P^r_{a_i}$ such that (i) D is pp-definable over $\mathbf{B}$, and (ii) for every $(T, \chi) \in \mathcal{T}^*$ with $T = (V, E)$, if Δ is the set of all leaves u of T with $\chi(u) = a_i$, then for every $u \in \Delta$ and every $b \in D$ there exists a realization f of (T, χ) in $\mathcal{P}$ satisfying:*

1) *f sends V into level r and is fixed up to $i-1$.*
2) *$f(v) \in D$ for every $v \in \Delta$.*
3) *$f(u) = b$.*

Proof: Suppose no such set D exists. Let $D_0 = P^r_{a_i}$. As D_0 is a non-empty, is pp-definable over $\mathbf{B}$, but does not satisfy the statement of the Claim, there must exist an A-tree $(T_0, \chi_0) \in \mathcal{T}^*$ with $T_0 = (V_0, E_0)$, whose set of leaves u with $\chi_0(u) = a_i$ is Δ_0, and there must exist $u_0 \in \Delta_0$ and $b_0 \in D_0$, such that for all realizations f of (T_0, χ_0) in $\mathcal{P}$, if f sends V_0 into level r and is fixed up to $i-1$, then $f(u_0) \neq b_0$. Define

$$\begin{aligned} D_1 = \{f(u_0) : \, & f \text{ is realization of } (T_0, \chi_0) \text{ in } \mathcal{P}, \\ & \text{sends } V_0 \text{ into level } r \text{, and is fixed up to } i-1\}. \end{aligned}$$

D_1 is non-empty by Claim 3(2), is pp-definable over $\mathbf{B}$ by Lemma 13, satisfies $D_1 \subseteq D_0$ because f maps V_0 into level r, and satisfies $D_1 \neq D_0$ because $b_0 \in D_0 \setminus D_1$.

Again as D_1 does not satisfy the statement of the Claim, there must exist an A-tree $(T_1, \chi_1) \in \mathcal{T}^*$ with $T_1 = (V_1, E_1)$, whose set of leaves u with $\chi_1(u) = a_i$ is Δ_1, and there must exist $u_1 \in \Delta_1$ and $b_1 \in D_1$, such that for all realizations f of (T_1, χ_1) in $\mathcal{P}$, if f sends V_1 into level r, sends Δ_1 into D_1, and is fixed up to $i-1$, then $f(u_1) \neq b_1$. Let $(T_1^\circ, \chi_1^\circ)$ be the A-tree with $T_1^\circ = (V_1^\circ, E_1^\circ)$ obtained by

1) starting with (T_1, χ_1);
2) gluing to every leaf $u \in \Delta_1$ a copy (T_0^u, χ_0^u) of (T_0, χ_0) at the vertex u_0; that is, u and the image of u_0 in T_0^u are identified;
3) adding a new leaf u_1° with an edge to u_1, and defining $\chi_1^\circ(u_1^\circ) = a_i$.

Note that (i) $(T_1^\circ, \chi_1^\circ) \in \mathcal{T}^*$, and (ii) if f is a realization of $(T_1^\circ, \chi_1^\circ)$ in $\mathcal{P}$ which sends V_1° into level r and is fixed up to $i-1$, then $f\restriction_{T_1}$ is a realization of (T_1, χ_1) which sends V_1 into level r, is fixed up to $i-1$, and sends Δ_1 into D_1; hence $f(u_1) \neq b_1$. As $(f(u_1^\circ), f(u_1)) \in E_{a_i, a_i} = \{(x, x) : x \in P_{a_i}\}$, this implies $f(u_1^\circ) = f(u_1) \neq b_1$. Thus if we define

$$\begin{aligned} D_2 = \{f(u_1^\circ) : \, & f \text{ is realization of } (T_1^\circ, \chi_1^\circ) \text{ in } \mathcal{P}, \\ & \text{sends } V_1^\circ \text{ into level } r \text{, and is fixed up to } i-1\}, \end{aligned}$$

then again D_2 is non-empty, is pp-definable over $\mathbf{B}$, and is a proper subset of D_1. By repeating this process we get an infinite strictly decreasing sequence $D_0 \supsetneq D_1 \supsetneq D_2 \supsetneq \cdots$ which is impossible. ■

We are now ready to define $h(a_i)$. Let $\mathbb{B}$ be the algebra (B, ϕ) and let $\mathbb{D}$ be its subalgebra $(D, \phi^{\mathbb{D}})$ where D is given by Claim 4. Choose any absorption constant β for $\mathbb{D}$ with respect to $\phi^{\mathbb{D}}$ (at least one exists by Theorem 6) and define $h(a_i) = \beta$. We claim that this choice achieves stage i. What needs to be shown is that for every A-tree $(T, \chi) \in \mathcal{T}_d^{r+1}(A, L)$ with $T = (V, E)$ there exists a realization of (T, χ) in $\mathcal{P}$ which:

$$(\dagger) \quad \begin{cases} 1. & \text{sends } V \text{ into level } r \text{, and} \\ 2. & \text{is fixed up to } i \text{; that is, is fixed up to } i-1 \\ & \quad \text{and sends every } a_i\text{-labelled leaf to } \beta. \end{cases}$$

Let (T, χ) be such an A-tree. Let $\Delta = \{u_1, \dots, u_n\}$ be an enumeration of the set of leaves u of T satisfying $\chi(u) = a_i$, and let $\Gamma = \{v_1, \dots, v_m\}$ be an enumeration of the set of leaves u of T satisfying $\chi(u) \in \{a_0, a_1, \dots, a_{i-1}\}$. Define two subsets of B^{n+m}:

$$\begin{aligned} S^+ = \{(f(u_1), \dots, f(u_n), f(v_1, \dots, f(v_m)) : \, & f \text{ is a} \\ & \text{realization of } (T, \chi) \text{ in } \mathcal{P} \text{ sending } V \text{ into level } r\}, \\ R^+ = \{(f(u_1), \dots, f(u_n), f(v_1, \dots, f(v_m)) : \, & f \text{ is a} \\ & \text{realization of } (T, \chi) \text{ in } \mathcal{P} \text{ sending } V \text{ into level } r \\ & \text{and which is fixed up to } i-1\}. \end{aligned}$$

Let B_0 and B_1 be the projections onto the first n coordinates of R^+ and S^+ respectively, and let $R = B_0 \cap D^n$ and $S = B_1 \cap D^n$. R^+ and S^+ are pp-definable over $\mathbf{B}$ by Lemma 13 and so are subuniverses of $\mathbb{B}^{n+m}$; we consider them as subuniverses of $\mathbb{B}^n \times \mathbb{B}^m$. Define $\mathbb{C}_1 = \mathbb{B}^m$ and $C_0 = \{\mathbf{c}\} \subseteq C_1$ where $\mathbf{c} = (h(\chi(v_1)), \dots, h(\chi(v_m)))$; then $S^+ \leq \mathbb{B}_1 \times \mathbb{C}_1$ and $B_0 = \{\mathbf{x} \in B_1 : \exists \mathbf{y} \in C_0 \text{ with } (\mathbf{x}, \mathbf{y}) \in S^+\}$. Note that $C_0 \lhd_\phi C_1$ by Lemma 4(2) (because $\mathbb{C}_1$ is idempotent and $\phi^{\mathbb{C}_1}$ is NU); hence $B_0 \lhd_\phi B_1$ by Lemma 5 and so $R \lhd_\phi S$ by Lemma 4(1). R is subdirect in D^n for the following reason: for any $1 \leq j \leq n$ and $b \in D$, Claim 4 gives a realization f of (T, χ) which sends V into level r, is fixed up to $i-1$, and satisfies $f(u_\ell) \in D$ for all $1 \leq \ell \leq n$ and $f(u_j) = b$. This f puts b into $\mathrm{proj}_j(R)$, proving $R \subseteq_{sd} D^n$. And S contains all constant tuples (over D), for the following reason: if $b \in D$, then because $D \subseteq P^r_{a_i} = P^{r+1}_{a_i}$ and $(T, \chi) \in \mathcal{T}_d^{r+1}(A)$, there exists a realization f of (T, χ) which maps V into level r and maps every a_i-labelled leaf to b. This f puts $(b, b, \dots, b)$ into S, as claimed.

In summary, R and S are subuniverses of $\mathbb{D}^n$; R is subdirect in D^n; S contains all constant tuples; and $R \lhd_\phi S$. As β is an absorption constant for $\mathbb{D}$ with respect to $\phi^{\mathbb{D}}$, it follows that the constant tuple $(\beta, \dots, \beta)$ is in R. This witnesses the existence of the desired realization of (T, χ) satisfying $(\dagger)$, completing the proof that stage i can be achieved in the construction of $h : \mathbf{A} \to \mathbf{B}$, and thus completes the proof of Proposition 10.

D. Proof of Theorem 6

Theorem 6 (restated). *Let $\mathbb{D}$ be a finite algebra and ϕ an idempotent operation of $\mathbb{D}$. There exists an absorption constant for $\mathbb{D}$ with respect to ϕ.*

Proof: We may assume that $\mathbb{D} = (D, \phi)$. We argue by induction on $|D|$. The claim is clearly true if $|D| = 1$, so we may assume $|D| \geq 2$. For $n \geq 2$ let $\mathcal{A}_n = \{R \leq_{sd} \mathbb{D}^n : R \text{ absorbs } 0_D^{(n)} \text{ with respect to } \phi\}$ and let $\mathcal{A} = \bigcup_{n=2}^{\infty} \mathcal{A}_n$. We must prove the existence of $b \in D$ such that $(b, \ldots, b) \in R$ for all $R \in \mathcal{A}$.

Let $\mathcal{P}$ be the set of all subuniverses of $\mathbb{D}$ (including $\varnothing$ and D) and let $\mathcal{P}^+ = \mathcal{P} \setminus \{\varnothing, D\}$. Observe that $\mathcal{P}^+ \neq \varnothing$, as $\mathcal{P}^+$ contains all the 1-element subsets of D. Suppose there exists $S \in \mathcal{P}^+$ such that, for all $n \geq 2$, we have $R \cap S^n \subseteq_{sd} S^n$ for all $R \in \mathcal{A}_n$. Let $\mathbb{S}$ be the subalgebra of $\mathbb{D}$ with universe S and note that $\phi\restriction_S$ is an idempotent operation of $\mathbb{S}$. Hence by the inductive hypothesis, $\mathbb{S}$ has an absorption constant b with respect to $\phi\restriction_S$. For any $n \geq 2$ and $R \in \mathcal{A}_n$, let $\overline{R} = R \cap S^n$; then $\overline{R} \leq_{sd} \mathbb{S}^n$ by our choice of S and $\overline{R}$ absorbs $0_S^{(n)}$ with respect to $\phi\restriction_S$ by Lemma 4(1); hence $(b, \ldots, b) \in \overline{R} \subseteq R$. Thus b is an absorption constant for $\mathbb{D}$, completing the proof.

Hence it suffices to prove that there exists $S \in \mathcal{P}^+$ such that, for all $n \geq 2$, we have $R \cap S^n \subseteq_{sd} S^n$ for all $R \in \mathcal{A}_n$. Let τ be the following algebraic vocabulary: for each $n \geq 2$, $R \in \mathcal{A}_n$ and $1 \leq i \leq n$, let $\mathsf{f}_{R,i}$ be an $(n-1)$-ary operation symbol and let $\tau = \{\mathsf{f}_{R,i} : n \geq 2, R \in \mathcal{A}_n, 1 \leq i \leq n\}$. Now define $\mathbb{P}$ to be the algebra with universe $\mathcal{P}$ and vocabulary τ where, for $n \geq 2$, $R \in \mathcal{A}_n$, $1 \leq i \leq n$ and $C_1, \ldots, C_{n-1} \in \mathcal{P}$,

$$\mathsf{f}_{R,i}^{\mathbb{P}}(C_1, \ldots, C_{n-1}) = \{x \in D : \exists c_1 \in C_1, \ldots, \exists c_{n-1} \in C_{n-1} \text{ such that } (c_1, \ldots, c_{i-1}, x, c_i, \ldots, c_{n-1}) \in R\}.$$

For convenience, we also let $\mathsf{f}_R^{\mathbb{P}}$ denote $\mathsf{f}_{R,1}^{\mathbb{P}}$ whenever $R \in \mathcal{A}$. Note that

1) $\{\mathsf{f}_{R,i}^{\mathbb{P}} : n \geq 2, R \in \mathcal{A}_n, 1 \leq i \leq n\} = \{\mathsf{f}_R^{\mathbb{P}} : R \in \mathcal{A}\}$. That is, every basic operation of $\mathbb{P}$ can be written as $\mathsf{f}_R^{\mathbb{P}}$ for some $R \in \mathcal{A}$.
2) Every operation of $\mathbb{P}$ is monotone with respect to $\subseteq$; that is, if $R \in \mathcal{A}_n$ and $B_i, C_i \in \mathcal{P}$ with $B_i \subseteq C_i$ for $1 \leq i \leq n-1$, then $\mathsf{f}_R^{\mathbb{P}}(B_1, \ldots, B_{n-1}) \subseteq \mathsf{f}_R^{\mathbb{P}}(C_1, \ldots, C_{n-1})$.
3) If $R = 0_D^{(3)}$, then $R \in \mathcal{A}_3$ and $\mathsf{f}_R^{\mathbb{P}}(C_1, C_2) = C_1 \cap C_2$. That is, $\cap$ is a basic operation of $\mathbb{P}$.
4) $\varnothing$ is a "zero" element for $\mathbb{P}$; that is, for all $n \geq 2$, $R \in \mathcal{A}_n$, and $C_1, \ldots, C_{n-1} \in \mathcal{P}$, if $\varnothing \in \{C_1, \ldots, C_{n-1}\}$ then $\mathsf{f}_R^{\mathbb{P}}(C_1, \ldots, C_{n-1}) = \varnothing$.
5) $\mathsf{f}_R^{\mathbb{P}}(D, \ldots, D) = D$ for all $R \in \mathcal{A}$.

Define a quasi-ordering $\preccurlyeq$ on $\mathcal{P}$ as follows: $C_1 \preccurlyeq C_2$ if and only if C_1 is an element of the subuniverse of $\mathbb{P}$ generated by $\{C_2, D\}$, i.e., $C_1 \in \mathrm{Sg}^{\mathbb{P}}(\{C_2, D\})$. Also define $C_1 \sim C_2$ to mean $C_1 \preccurlyeq C_2$ and $C_2 \preccurlyeq C_1$. Thus $\sim$ is an equivalence relation on $\mathcal{P}$ and $\preccurlyeq$ naturally induces a partial ordering $\leq$ of the set $\mathcal{P}/{\sim}$ of $\sim$-equivalence classes. Note that $\{\varnothing\}$ and $\{D\}$ are $\sim$-classes; furthermore, $\{D\}$ is the unique minimum element of $(\mathcal{P}/{\sim}, \leq)$, and $\{\varnothing\}$ is a minimal element of the poset obtained from $(\mathcal{P}/{\sim}, \leq)$ by deleting $\{D\}$. If we delete both $\{D\}$ and $\{\varnothing\}$, we obtain the poset $\mathcal{Q} := (\mathcal{P}^+/{\sim}, \leq)$.

Choose and fix a minimal element M of $\mathcal{Q}$. Let $M_{\min}$ denote the set of members of M which are minimal with respect to $\subseteq$. Also let $X = \bigcup\{C : C \in M_{\min}\}$.

Claim 5.

1) *Suppose $n \geq 2$, $R \in \mathcal{A}_n$, $C_1, \ldots, C_{n-1} \in M \cup \{D\}$, and let $A = \mathsf{f}_R^{\mathbb{P}}(C_1, \ldots, C_{n-1})$. If $A \neq \varnothing$, then $A \in M \cup \{D\}$; hence there exists $C \in M_{\min}$ with $C \subseteq A$.*
2) *For all $C_1 \in M_{\min}$ and $C_2 \in M \cup \{D\}$, if $C_1 \cap C_2 \neq \varnothing$ then $C_1 \subseteq C_2$.*

Proof: (1) With $n, R, C_1, \ldots, C_{n-1}, A$ as in the statement of the Claim, assume $A \notin \{\varnothing, D\}$, i.e., $A \in \mathcal{P}^+$, and pick any $C \in M$. Since $C_i \preccurlyeq C$ for each i, we get $A \preccurlyeq C$, so $A \in M$ as $C/{\sim} = M$ is a minimal member of $\mathcal{Q}$.

(2) This follows from (1) and the fact that $\cap$ is a basic operation of $\mathbb{P}$. ∎

Claim 6. *For all $n \geq 2$ and $R \in \mathcal{A}_n$, $R \cap X^n \subseteq_{sd} X^n$.*

Proof: We will show $\mathrm{proj}_1(R \cap X^n) = X$, the argument for the other coordinates being similar. Pick any $C \in M_{\min}$. We will show, by induction on $1 \leq i \leq n$, that there exist $C_j \in M_{\min}$ for $2 \leq j \leq i$ such that $C \subseteq \mathsf{f}_R^{\mathbb{P}}(C_2, \ldots, C_i, D, \ldots, D)$. When $i = 1$ the claim is simply that $C \subseteq \mathsf{f}_R^{\mathbb{P}}(D, \ldots, D)$, which is true by a previous observation. Assume now that $1 \leq i < n$ and the claim has been verifed for i and must be proved for $i+1$. Thus we have $C_2, \ldots, C_i \in M_{\min}$ such that

$$C \subseteq \mathsf{f}_R^{\mathbb{P}}(C_2, \ldots, C_i, D, \ldots, D). \quad (1)$$

Equation (1) implies $\mathsf{f}_{R,i+1}^{\mathbb{P}}(C, C_2, \ldots, C_i, D, \ldots, D) \neq \varnothing$; hence by Claim 5 we can choose $C_{i+1} \in M_{\min}$ with

$$C_{i+1} \subseteq \mathsf{f}_{R,i+1}^{\mathbb{P}}(C, C_2, \ldots, C_i, D, \ldots, D). \quad (2)$$

Similarly, equation (2) implies

$$C \cap \mathsf{f}_R^{\mathbb{P}}(C_2, \ldots, C_i, C_{i+1}, D, \ldots, D) \neq \varnothing$$

and hence $C \subseteq \mathsf{f}_R^{\mathbb{P}}(C_2, \ldots, C_{i+1}, D, \ldots, D)$ by Claim 5, as desired, which completes the inductive argument. When $i = n$ this gives $C \subseteq \mathsf{f}_R^{\mathbb{P}}(C_2, \ldots, C_n) \subseteq \mathsf{f}_R^{\mathbb{P}}(X, \ldots, X)$, which implies $C \subseteq \mathrm{proj}_1(R \cap X^n)$. As $C \in M_{\min}$ was arbitrary, this proves $\mathrm{proj}_1(R \cap X^n) = X$. ∎

Corollary 7. *Suppose $R \in \mathcal{A}_n$ and $B_1, \ldots, B_{n-1} \in \mathcal{P}$. Let $B = \mathsf{f}_R^{\mathbb{P}}(B_1, \ldots, B_{n-1})$.*

1) *For all $(b_1, \ldots, b_{n-1}) \in B_1 \times \cdots \times B_{n-1}$ and $b \in D$, if $(b, b_1, \ldots, b_{n-1}) \in R$ then $b \in B$.*
2) *For all $b \in B$ there exists $(b_1, \ldots, b_{n-1}) \in B_1 \times \cdots \times B_{n-1}$ such that $(b, b_1, \ldots, b_{n-1}) \in R$.*
3) *For all $x \in X$ there exists $(x_1, \ldots, x_{n-1}) \in X^{n-1}$ such that $(x, x_1, \ldots, x_{n-1}) \in R$.*

Proof: (1) and (2) follow from the definition of $\mathsf{f}_R^{\mathbb{P}}$, while (3) follows from Claim 6. ∎

Observe that if $M_{\min}$ contains just one member C, then $X = C \in \mathcal{P}^+$ and so by Claim 6 we could choose $S = C$ and the proof of Theorem 6 would be complete. For the remainder of the proof, let $M_{\min} = \{C_1, \dots, C_k\}$ be an enumeration of $M_{\min}$ and assume for the sake of contradiction that $k \geq 2$.

Definition 16. An *evaluated term tree* for $\mathbb{P}$ consists of a finite ordered directed tree $T = (V, E)$, an assignment to each non-leaf node v of an operation symbol $\mathsf{f}_R \in \tau$ whose arity equals the number of children of v, and a map $A : V \to \mathcal{P}$ satisfying the following recursive condition: if v is a non-leaf node, $v_1, \dots, v_t$ are its children (listed in increasing order), and f_R is the operation symbol assigned to v, then

$$A(v) = \mathsf{f}_R^{\mathbb{P}}(A(v_1), \dots, A(v_t)).$$

We call $A(v)$ the *value* of the evaluated term tree at v.

Note that if T_1, T_2 are evaluated term trees for $\mathbb{P}$, v is a leaf of T_1, and the value of T_1 at v is equal to the value of T_2 at its root, then T_2 may be glued to T_1 by identifying the root of T_2 with v.

Evaluated term trees witness subuniverse generation in $\mathbb{P}$; that is, if $B_1, \dots, B_n, C \in \mathbb{P}$, then $C \in \mathrm{Sg}^{\mathbb{P}}(B_1, \dots, B_n)$ if and only if there exists an evaluated term tree whose root has value C and whose leaves have values in $\{B_1, \dots, B_n, D\}$.

Claim 8. *For any $i, j \in \{1, \dots, k\}$ with $i \neq j$ there exists an evaluated term tree $T_{i,j}$ for $\mathbb{P}$ satisfying the following:*

1) *The value at each leaf is in $\{C_i, D\}$. At least one leaf has value C_i.*
2) *The value at the root is C_j.*

Proof: The existence of a tree satisfying (2) and the first sentence of (1) follows from the fact that $C_i \preccurlyeq C_j$. That at least one leaf has value C_i follows from the fact that $\{D\}$ is a subuniverse of $\mathbb{P}$. ■

We now construct a special evaluated term tree T for $\mathbb{P}$ as follows. We start with $T_2 := T_{2,1}$. To each leaf of T_2 whose value is C_2 we glue a copy of $T_{3,2}$; call the resulting evaluated term tree T_3. Then to each leaf of T_3 whose value is C_3 we glue a copy of $T_{4,3}$, to get T_4, etc. After we have constructed T_k, we glue to each leaf whose value is C_k a copy of $T_{1,k}$, to get T. Here are the salient properties of T:

1) The value at each leaf is in $\{C_1, D\}$.
2) The value at the root is C_1.
3) For every $1 \leq i \leq k$ and every path from the root to a leaf whose value is C_1, some node on the path has value in C_i.

Let V denote the set of nodes of T, r the root, Λ the set of leaves, and Λ_1 the set of leaves whose value is C_1. Also let $A : V \to \mathbb{P}$ be the value map.

Definition 17. A map $\alpha : V \to D$ is called a *selection map* (for T). Suppose α is a selection map, u is an arbitrary node, and v is a non-leaf node of T. Let f_R be the operation symbol in τ assigned to v with $R \in \mathcal{A}_n$, let $v_1, \dots, v_{n-1}$ be the children of v listed in increasing order, and let $\alpha(\mathbf{v}) = (\alpha(v), \alpha(v_1), \dots, \alpha(v_{n-1}))$. We say that:

1) α *respects values at* u if $\alpha(u) \in A(u)$.
2) α *quasi-respects values at* u if $\alpha(u) \in X$.
3) α *respects relations at* v if $\alpha(\mathbf{v}) \in R$.
4) α *quasi-respects relations at* v if $\alpha(\mathbf{v}) \in R \cup 0_D^{(n)}$.

Note that, by Corollary 7(1), if α is a selection map which respects values at all C_1-valued leaves and respects relations at all non-leaf nodes, then α respects values at all nodes. In particular, $\alpha(r) \in C_1$. Conversely:

Claim 9.

1) *For all $a \in C_1$ there exists a selection map α_a which respects values at all nodes, respects relations at all non-leaf nodes, and satisfies $\alpha_a(r) = a$.*
2) *For all $a \in X$ there exists a selection map β_a which quasi-respects values at all nodes, respects relations at all non-leaf nodes, and satisfies $\beta_a(r) = a$.*
3) *For all $a \in X$ there exists a selection map γ_a which respects values at all C_1-valued leaves, quasi-respects relations at all non-leaf nodes, and satisfies $\gamma_a(r) = a$.*

Proof: (1) We inductively define α_a, starting at the root. Of course, $\alpha_a(r) = a$. Suppose now that v is a non-leaf node and $\alpha_a(v) \in A(v)$. Let f_R be the operation symbol assigned to v with $R \in \mathcal{A}_n$, and let $v_1, \dots, v_{n-1}$ be the children of v in increasing order. As $A(v) = \mathsf{f}_R^{\mathbb{P}}(A(v_1), \dots, A(v_{n-1}))$, by Corollary 7(2) there exists $(c_1, \dots, c_{n-1}) \in A(v_1) \times \dots \times A(v_{n-1})$ such that $(\alpha_a(v), c_1, \dots, c_{n-1}) \in R$. We can thus define $\alpha_a(v_i) = c_i$ and continue inductively.

(2) is proved similarly, using Corollary 7(3).

(3) Suppose $a \in C_i$. Let V_i denote the set of nodes of T such that the path from v to the root includes a node having value C_i. We will construct γ_a so that it will inductively satisfy the following additional property: γ_a has constant value a on $V \setminus V_i$ and respects values on V_i. By one of the salient properties of T, this will imply that γ_a respects values at C_1-valued leaves.

We start by defining $\gamma_a(r) = a$. Suppose now that v is a non-leaf node and $\gamma_a(v)$ has already been defined. Let f_R be the operation assigned to v with $R \in \mathcal{A}_n$, and let $v_1, \dots, v_{n-1}$ be the children of v in increasing order. If $v \in V_i$ then inductively we have $\gamma_a(v) \in A(v)$, so we define γ_a at the children of v exactly as in the definition of α_a; thus γ_a respects relations at v. If instead $v \in V \setminus V_i$, then inductively we have $\gamma_a(v) = a$, and we define $\gamma_a(v') = a$ for all children v' of v; thus γ_a quasi-respects relations at v. It remains to check that the inductive property is maintained by this construction. The only problem to consider is if $v \in V \setminus V_i$ but a child v' of v is in V_i. If this is the case, then it must be that v' has value C_i. As the construction assigns $\gamma_a(v') = a$ and as $a \in C_i$, there is no problem. ■

Recall that ϕ is the basic operation of $\mathbb{D}$ referenced in the statement of Theorem 6. Let m be its arity.

Claim 10. *For all $0 \leq j \leq m$,*

$$\phi(\underbrace{C_1, \dots, C_1}_{m-j}, \underbrace{X, \dots, X}_{j}) \subseteq C_1.$$

Proof: By induction on j. When $j = 0$ the claim is simply that $\phi(C_1, \ldots, C_1) \subseteq C_1$, which is true because C_1 is a subuniverse of $\mathbb{D}$. Suppose $j < m$ and the claim is true for j. To prove that it is true for $j + 1$, let $\ell = m - j - 1$ and assume $a_1, \ldots, a_\ell \in C_1$ and $x_0, x_1 \ldots, x_j \in X$. It suffices to prove $\phi(a_1, \ldots, a_\ell, x_0, \ldots, x_j) \in C_1$. Let $\alpha_{a_1}, \ldots, \alpha_{a_\ell}, \beta_{x_1}, \ldots, \beta_{x_j}, \gamma_{x_0}$ be selection maps for T constructed according to Claim 9. Define $\delta : V \to D$ by

$$\delta(v) = \phi(\alpha_{a_1}(v), \ldots, \alpha_{a_\ell}(v), \gamma_{x_0}(v), \beta_{x_1}(v), \ldots, \beta_{x_j}(v)).$$

δ is a selection map. Suppose v is a C_1-valued leaf of T. Then $\alpha_{a_i}(v) \in C_1$ for all $1 \leq i \leq \ell$, $\gamma_{x_0}(v) \in C_1$, and $\beta_{x_i}(v) \in X$ for all $1 \leq i \leq j$. Thus $\delta(v) \in \phi(\underbrace{C, \ldots, C, C}_{\ell+1}, \underbrace{X, \ldots, X}_{j}) \subseteq C_1$ by the inductive assumption. This proves that δ respects values at C_1-valued leaves.

Suppose next that v is a non-leaf node of T. Let f_R be the operation symbol assigned to v with $R \in \mathcal{A}_n$, let $v_1, \ldots, v_{n-1}$ be the children of v listed in increasing order, and for any selection map η let $\eta(\mathbf{v}) = (\eta(v), \eta(v_1), \ldots, \eta(v_{n-1}))$. Then $\alpha_{a_i}(\mathbf{v}) \in R$ for $1 \leq i \leq \ell$, $\beta_{x_i}(\mathbf{v}) \in R$ for $1 \leq i \leq j$, $\gamma_{x_0}(\mathbf{v}) \in R \cup 0_D^{(n)}$, and

$$\begin{aligned}\delta(\mathbf{v}) &= \phi^{\mathbb{D}^n}(\alpha_{a_1}(\mathbf{v}), \ldots, \alpha_{a_\ell}(\mathbf{v}), \gamma_{x_0}(\mathbf{v}), \beta_{x_1}(\mathbf{v}), \ldots, \beta_{x_j}(\mathbf{v})) \\ &\in \phi^{\mathbb{D}^n}(R, \ldots, R, R \cup 0_D^{(n)}, R, \ldots, R) \quad \subseteq \quad R,\end{aligned}$$

where the last inclusion follows because R absorbs $0_D^{(n)}$ with respect to ϕ. This proves that δ respects relations at all non-leaf nodes.

It follows from the observation preceding Claim 9 that δ respects values at all nodes. In particular, δ respects values at the root, i.e., $\delta(r) = \phi(a_1, \ldots, a_\ell, x_0, \ldots, x_j) \in C_1$, which finishes the proof of the Claim. ■

In particular, Claim 10 yields $\phi(X, \ldots, X) \subseteq C_1$, while $X \subseteq \phi(X, \ldots, X)$ because ϕ is idempotent. Thus $X \subseteq C_1$, implying $X = C_1$. This contradicts our assumption that $|M_{\min}| \geq 2$ and thus completes the proof of Theorem 6. ■

IV. Conclusion

We have proved that for every finite relational structure $\mathbf{B}$ having a near-unanimity polymorphism, the corresponding constraint satisfaction problem $\mathrm{CSP}(\mathbf{B})$ has bounded pathwidth duality; equivalently, $\neg\mathrm{CSP}(\mathbf{B})$ is definable in linear Datalog, and as a consequence, $\mathrm{CSP}(\mathbf{B})$ is in the complexity class NL. This answers a question from [16].

The natural algebraic conjecture alluded to in the introduction, suggested by Larose and Tesson [27], is the following: if $\mathbf{B}$ is core and $\mathbb{B}_e$ is the idempotent reduct of the polymorphism algebra of $\mathbf{B}$, then $\mathrm{CSP}(\mathbf{B})$ has bounded pathwidth duality if and only if the variety generated by $\mathbb{B}_e$ "omits types 1, 2 and 5" in the sense of tame congruence theory. A characterization of this property in terms of idempotent polymorphisms of $\mathbf{B}$ is given by [22, Theorem 9.11]. Our result lends further support to this conjecture which, however, remains open.

References

[1] S. Abiteboul, R. Hull and V. Vianu, *Foundations of Databases*, Addison-Wesley, 1995.

[2] E. Allender, M. Bauland, N. Immerman, H. Schnoor, and H. Vollmer, The complexity of satifiability problems: Refining Schaefer's theorem, *Proc. 30th Math. Found. Comp. Sci.* (2005), 71–82.

[3] K. Baker and A. Pixley, Polynomial interpolation and the Chinese remainder theorem for algebraic systems, *Math. Z.* **143** (1975), 165–174.

[4] L. Barto, Finitely related algebras in congruence distributive varieties have near unanimity terms, *Canad. J. Math.*, published electronically Dec. 2011, doi:10.4153/CJM-2011-087-3, 19 pp.

[5] L. Barto and M. Kozik, Congruence distributivity implies bounded width, *SIAM J. Comput.* **39** (2009), 1531–1542.

[6] L. Barto and M. Kozik, Constraint satisfaction problems of bounded width, *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2009)*, 595–603, IEEE Computer Soc., 2009.

[7] L. Barto and M. Kozik, Absorbing subalgebras, cyclic terms and the constraint satisfaction problem, *Logical Methods in Computer Science* **8** (1:07) (2012), 1–26.

[8] C. Bergman, *Universal Algebra*, CRC Press, Boca Raton, 2012.

[9] V.G. Bodnarčuk, L.A. Kalužnin, V.N. Kotov, and B.A. Romov, Galois theory for Post algebras. I, *Cybernetics and Systems Analysis* **5** (1969), 243–252.

[10] A. Bulatov, A graph of a relational structure and constraint satisfaction problems, *LICS '04* (2004), 448–457.

[11] A. Bulatov, A. Krokhin, and P. Jeavons, Classifying the complexity of constraints using finite algebras, *SIAM J. Comput.* **34** (2005), 720–742.

[12] S. Burris and H.P. Sankappanavar, *A Course in Universal Algebra*, Springer-Verlag, New York, 1981.

[13] S. Cook and R. Sethi, Storage requirements for deterministic polynomial time recognizable languages, *J. Compu. System Sci.* **13** (1976), 25–37.

[14] V. Dalmau, Constraint satisfaction problems in non-deteministic logarithmic space, *LNCS* **2380** (2002), 414–425.

[15] V. Dalmau, Linear datalog and bounded path duality of relational structures, *Logical Methods in Computer Science* **1** (2005), 1–32.

[16] V. Dalmau and A. Krokhin, Majority constraints have bounded pathwidth duality, *Euro. J. Combin.* **29** (2008), 821–837.

[17] J. A. Ellis, I. H. Sudborough, and J. S. Turner, The vertex separation and search number of a graph, *Inform. Comput.* **113** (1994), 50-79.

[18] T. Feder and M. Vardi, The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory, *SIAM J. Comput.* **28** (1998), 57–104.

[19] D. Geiger, Closed systems of functions and predicates, *Pacific J. Math.* **27** (1968), 95–100.

[20] P. Hell and J. Nešetřil, *Graphs and Homomorphisms*, Oxford University Press, 2004.

[21] P. Hell, J. Nešetřil, and X. Zhu, Duality and polynomial testing of tree homomorphisms, *Trans. Amer. Math. Soc.* **348** (1996), 147–156.

[22] D. Hobby and R. McKenzie, *The Structure of Finite Algebras*, American Mathematical Society, Providence, 1988.

[23] P. Jeavons, On the algebraic structure of combinatorial problems, *Theoret. Comput. Sci.* **200** (1998), 185—204.

[24] P. Jeavons, D. Cohen, and M. Gyssens, How to determine the expressive power of constraints, *Constraints* **4** (1999), 113–131.

[25] Ph. Kolaitis and M.Vardi, Conjunctive-query containment and constraint satisfaction, *J. Comput. System Sci.* **61** (2000), 302–332.

[26] A. Krokhin and B. Larose, Solving order constraints in logarithmic space, *STACS'03*, *LNCS* **2607** (2003), 379–390.

[27] B. Larose and P. Tesson, Universal algebra and hardness problems for constraint satisfaction problems, *Theoret. Comput. Sci.* **410** (2009), 1629–1647.

[28] B. Larose, M. Valeriote, and L. Zádori, Omitting types, bounded width and the ability to count, *Internat. J. Algebra Comput.* **19** (2009), 647–668.

[29] B. Larose and L. Zádori, Bounded width problems and algebras, *Algebra Universalis* **56** (2007), 439–466.

[30] M. Maróti and R. McKenzie, Existence theorems for weakly symmetric operations, *Algebra Universalis* **59** (2008), 463–489.

[31] U. Montanari, Networks of constraints: Fundamental properties and applications to picture processing, *Information Science* **7** (1974), 95–132.

[32] N. Robertson and P. Seymour, Graph minors. I. Excluding a forest, *J. Combin. Theory Ser. B* **35** (1983), 39–61.

[33] J. Ullman, *Principles of Database and Knowledge-base Systems*, Vol. 1, Computer Science Press, 1988.

Appendix

Lemma 1. *Suppose $\mathbf{B}$ is a finite binary τ-structure whose set of basic relations is closed under $\wedge$-atomic definitions. For every finite τ-structure $\mathbf{A}$ there exists a $\mathbf{B}$-reduced τ-structure $\mathbf{A}^\circ$ having the same domain as $\mathbf{A}$ and which satisfies the following: for all $X \subseteq A$, $\operatorname{Hom}(\mathbf{A}{\restriction_X}, \mathbf{B}) = \operatorname{Hom}(\mathbf{A}^\circ{\restriction_X}, \mathbf{B})$.*

Proof: Let $\mathcal{R}_1$ and $\mathcal{R}_2$ denote the sets of 1-ary and 2-ary basic relations of $\mathbf{B}$ respectively. By hypothesis, (i) $B \in \mathcal{R}_1$; (ii) the equality relation $\{(x, x) : x \in B\}$ is in $\mathcal{R}_2$; (iii) if $U, V \in \mathcal{R}_1$ then $U \times V \in \mathcal{R}_2$; (iv) if $R \in \mathcal{R}_2$ then $R^{-1} := \{(y, x) : (x, y) \in R\} \in \mathcal{R}_2$ and $R_\Delta := \{x : (x, x) \in R\} \in \mathcal{R}_1$; (v) each of $\mathcal{R}_1$ and $\mathcal{R}_2$ is closed under intersections.

Define a potato system $\mathcal{P} = (P_a, E_{a,b} : a, b \in A)$ over $\mathbf{B}$ as follows: For $a, b \in A$,

$$\begin{aligned} P_a &= \bigcap\{R^{\mathbf{B}} : R \in \tau \text{ is 1-ary and } a \in R^{\mathbf{A}}\} \cap \\ &\quad \bigcap\{(R^{\mathbf{B}})_\Delta : R \in \tau \text{ is 2-ary and } (a, a) \in R^{\mathbf{A}}\}, \\ E_{a,b} &= (P_a \times P_b) \cap \\ &\quad \bigcap\{R^{\mathbf{B}} : R \in \tau \text{ is 2-ary and } (a, b) \in R^{\mathbf{A}}\} \cap \\ &\quad \bigcap\{(R^{\mathbf{B}})^{-1} : R \in \tau \text{ is 2-ary and } (b, a) \in R^{\mathbf{A}}\}. \end{aligned}$$

Then let $\mathbf{A}^\circ$ be the $\mathbf{B}$-reduced structure associated with $\mathcal{P}$. ∎

Lemma 8. *Suppose τ is a vocabulary, $n \geq 1$ is an integer such that every relation symbol in τ has arity at most $2n$, $\mathbf{B}$ is a τ-structure, and $\mathbb{B}$ is its polymorphism algebra. There exists a binary structure $\mathbf{B}^{(n)}$ with universe B^n such that:*

1) *The polymorphism algebra of $\mathbf{B}^{(n)}$ is $\mathbb{B}^n$.*
2) *For any $0 \leq j \leq k$, if $\operatorname{CSP}(\mathbf{B}^{(n)})$ has (j, k)-pathwidth duality, then $\operatorname{CSP}(\mathbf{B})$ has (jn, kn)-pathwidth duality.*

Proof: Define $\tau^{(n)}$ to be the signature consisting of a binary relation symbol $\widehat{R}$ for each $R \in \tau$, and binary relation symbols $E_{i,j}$ for $1 \leq i, j \leq n$. For any τ-structure $\mathbf{A}$ with domain A we define $\mathbf{A}^{(n)}$ to be the $\tau^{(n)}$-structure with domain A^n and relations given by: if $R \in \tau$ has arity m, then

$$\begin{aligned} \widehat{R}^{\mathbf{A}^{(n)}} = \{((a_1, \ldots, a_n), (a_{n+1}, \ldots, a_{2n})) \in (A^n)^2 : \\ (a_1, \ldots, a_m) \in R^{\mathbf{A}}\}, \end{aligned}$$

while $E_{i,j}^{\mathbf{A}^{(n)}} = \{(\mathbf{a}, \mathbf{b}) \in (A^n)^2 : a_i = b_j\}$.

The map $\mathbf{A} \mapsto \mathbf{A}^{(n)}$ is a functor of the obvious categories. In particular, if $\mathbf{A}_1, \mathbf{A}_2$ are τ structures and $h : \mathbf{A}_1 \to \mathbf{A}_2$ is a homomorphism, then $h^{(n)} : \mathbf{A}_1^{(n)} \to \mathbf{A}_2^{(n)}$ is a homomorphism where $h^{(n)}((a_1, \ldots, a_n)) = (h(a_1), \ldots, h(a_n))$. Conversely, it can be shown that every homomorphism $f : \mathbf{A}_1^{(n)} \to \mathbf{A}_2^{(n)}$ is of the form $f = h^{(n)}$ for some $h : \mathbf{A}_1 \to \mathbf{A}_2$. Finally, $(\mathbf{A}^{(n)})^k$ and $(\mathbf{A}^k)^{(n)}$ are isomorphic in a natural way for all $k \geq 1$. These remarks suffice to prove item (1).

Assume that $\operatorname{CSP}(\mathbf{B}^{(n)})$ has (j, k)-pathwidth duality and let $\mathbf{B}$ be a finite τ-structure with $\mathbf{A} \not\to \mathbf{B}$; hence $\mathbf{A}^{(n)} \not\to \mathbf{B}^{(n)}$. As $\operatorname{CSP}(\mathbf{B}^{(n)})$ has (j, k)-pathwidth duality, we get $\mathbf{A}^{(n)} \not\leadsto_{j,k} \mathbf{B}^{(n)}$ by Proposition 2. Hence there exists a solo play $\mathcal{I} = (I_0, I_1, \ldots, I_N)$ of the (j, k)-PR game on $(\mathbf{A}^{(n)}, \mathbf{B}^{(n)})$ whose final relation is empty. For each $0 \leq t \leq N$ let J_t be the smallest subset of A such that $(a_1, \ldots, a_n) \in I_t$ implies $\{a_1, \ldots, a_n\} \subseteq J_t$. Thus $|J_t| \leq n \cdot |I_t|$ and the sequence $\mathcal{J} = (J_0, J_1, \ldots, J_N)$ is a solo play of the (jn, kn)-PR game on $(\mathbf{A}, \mathbf{B})$. Let $S_0, S_1, \ldots, S_N$ be the relations resulting from playing $\mathcal{I}$ on $(\mathbf{A}^{(n)}, \mathbf{B}^{(n)})$, and let $T_0, T_1, \ldots, T_N$ be the relations resulting from playing $\mathcal{J}$ on $(\mathbf{A}, \mathbf{B})$. One can show inductively that for all $t \leq N$, if $h \in T_t$ then $h^{(n)}{\restriction_{I_t}} \in S_t$. Since $S_N = \varnothing$, it follows that $T_N = \varnothing$. Thus $\mathbf{A} \not\leadsto_{jn,kn} \mathbf{B}$.

In summary, if $\mathbf{A} \not\to \mathbf{B}$ then $\mathbf{A} \not\leadsto_{jn,kn} \mathbf{B}$. By Proposition 2 we get that $\operatorname{CSP}(\mathbf{B})$ has (jn, kn)-pathwidth duality. ∎

Lemma 9. *Suppose $\mathbf{B}$ is a finite τ-structure with universe B and a $(d+1)$-ary NU polymorphism for some $d \geq 2$. Let $s = \max(\{\mathit{arity}(R) : R \in \tau\} \cup \{d\})$. There exists a vocabulary τ_d and a τ_d-structure $\mathbf{B}_d$ with universe B satisfying:*

1) *Every relation symbol in τ_d has arity at most d.*
2) *$\mathbf{B}$ and $\mathbf{B}_d$ have the same polymorphisms.*
3) *If $\operatorname{CSP}(\mathbf{B}_d)$ has (j, k)-pathwidth duality, then $\operatorname{CSP}(\mathbf{B})$ has $(k, k + s - d)$-pathwidth duality.*

Proof: τ_d is the vocabulary consisting of the following relation symbols: for each $R \in \tau$ of arity n and for each subset $U \subseteq \{1, 2, \ldots, n\}$ satisfying $|U| = \min(n, d)$, a $|U|$-ary relation symbol R_U. (1) is then obvious.

For each τ-structure $\mathbf{A}$, $\mathbf{A}_d$ denotes the τ_d-structure having the same universe as $\mathbf{A}$ and whose interpretation of each symbol $R_U \in \tau_d$ is the projection onto U of the interpretation of R in $\mathbf{A}$.

(2) is a well-known consequence of [3, Theorem 2.1 (1) $\Rightarrow$ (3)]. As in the proof of [16, Lemma 2], for any τ-structure $\mathbf{A}$ we have $\mathbf{A} \to \mathbf{B}$ if and only if $\mathbf{A}_d \to \mathbf{B}_d$. (Both assertions make use of the NU polymorphism of $\mathbf{B}$.)

Suppose $\operatorname{CSP}(\mathbf{B}_d)$ has (j, k)-pathwidth duality and let $\mathbf{A}$ be a finite τ-structure with $\mathbf{A} \not\to \mathbf{B}$. Then $\mathbf{A}_d \not\to \mathbf{B}_d$ by a comment above. As $\operatorname{CSP}(\mathbf{B}_d)$ has (j, k)-pathwidth duality, we get $\mathbf{A}_d \not\leadsto_{j,k} \mathbf{B}_d$ by Proposition 2. Hence there exists a solo play $\mathcal{I} = (I_0, I_1, \ldots, I_N)$ of the (j, k)-PR game on $(\mathbf{A}_d, \mathbf{B}_d)$ whose final relation is empty. From $\mathcal{I}$ we construct a solo play $\mathcal{I}^*$ for $(\mathbf{A}, \mathbf{B})$ as follows. If $s = d$ then $\mathcal{I}^* = \mathcal{I}$. If $s > d$, then let $p = \binom{|A|}{s-d}$ and let $V_1, \ldots, V_p$ be an enumeration of the subsets of A of size $s - d$. Put $m = 2p$ and for $0 \leq t \leq N$ put

$$\mathcal{I}_t = (I_t, I_t \cup V_1, I_t, I_t \cup V_2, \ldots, I_t, I_t \cup V_p, I_t)$$

and let $\mathcal{I}^*$ be the concatenation of $\mathcal{I}_0, \mathcal{I}_1, \ldots, \mathcal{I}_N$. $\mathcal{I}^*$ is a solo play of the $(k, k + s - d)$-PR game on $(\mathbf{A}, \mathbf{B})$. Furthermore, if $T_0, T_1, \ldots, T_N$ are the relations resulting from $\mathcal{I}$ played on $(\mathbf{A}_d, \mathbf{B}_d)$ and $T_0[0], \ldots, T_0[m], \ldots, T_N[0], \ldots, T_N[m]$ are the relations resulting from $\mathcal{I}^*$ played on $(\mathbf{A}, \mathbf{B})$, then it can be shown that $T_t[m] \subseteq T_t$ for each $t \leq N$. In particular, since $T_N = \varnothing$ we get $T_N[m] = \varnothing$, i.e., the final relation of $\mathcal{I}^*$ is empty. Thus $\mathbf{A} \not\leadsto_{k,k+s-d} \mathbf{B}$. In summary, if $\mathbf{A} \not\to \mathbf{B}$ then $\mathbf{A} \not\leadsto_{k,k+s-d} \mathbf{B}$. Hence $\operatorname{CSP}(\mathbf{B})$ has $(k, k + s - d)$-pathwidth duality by Proposition 2. ∎

Lemma 12. *Suppose $d \geq 3$ and $i \geq 1$. The trees in $\mathcal{T}_d^i$ have bounded pathwidth. More precisely, every tree $T \in \mathcal{T}_d^i$ satisfies $\operatorname{pw}(T) \leq c^i + c^{i-1} - 1$ where $c = \lfloor \log_3(2d - 3) \rfloor + 2$.*

Proof: We prove this in three steps.

Step 1. *If a tree T has $k \geq 2$ leaves, then* $\mathrm{pw}(T) \leq \log_3(2k-3)+2$.

Proof: Equivalently, we show that $k \geq (3^{\mathrm{pw}(T)-2}+3)/2$. We do this by induction on $(\mathrm{pw}(T), |T|)$. The claim is obvious if $\mathrm{pw}(T) = 2$. Assume $\mathrm{pw}(T) = p > 2$. By [17, Corollary 3.1], there exists a vertex $a \in T$ such that the forest induced on $T \setminus \{a\}$ contains at least three trees T_1, T_2, T_3 with $\mathrm{pw}(T_i) \geq p-1$. Let k_i be the number of leaves in T_i. By induction, $k_i \geq (3^{p-3}+3)/2$ for each $i = 1, 2, 3$. Clearly the leaves of T_1, T_2, T_3 are distinct and for each i all the leaves of T_i except one are leaves of T. Thus $k \geq (k_1+k_2+k_3)-3 \geq 3(3^{p-3}+3)/2-3 = (3^{p-2}-3)/2$ as required. ∎

Let T be any composition tree, and let T_0 be the tree formed as follows:

- The vertices of T_0 are the components and composition vertices of T.
- For each composition vertex v and component C, we have an edge $\{v, C\}$ if and only if $v \in C$.

Note that the leaves of T_0 are just the leaf components of T.

Step 2. *If* $\mathrm{pw}(T_0) \leq s$ *and* $\mathrm{pw}(C) \leq c$ *for every component C of T, then* $\mathrm{pw}(T) \leq s(c+1)-1$.

Proof: We argue by induction on $\mathrm{pw}(T_0)$. If $\mathrm{pw}(T_0) = 1$, i.e., T_0 consists of a single vertex, then T equals its unique component and the inequality holds. Now assume $\mathrm{pw}(T_0) = k \geq 2$. By [17, pp. 64–65], there exists a path $P = (0, 1, 2, \ldots, n)$ in T_0 such that every connected component of the forest $T_0 \setminus P$ has pathwidth at most $k-1$. We may assume that P is a maximal path in T_0, so begins and ends at leaves of T_0. Thus P looks like

$$P = (C_0, v_0, C_1, v_1, \ldots, C_{n-1}, v_{n-1}, C_n)$$

where $C_0, C_1, \ldots, C_n$ are components of T and each v_i is the unique composition node joining C_i to C_{i+1}. Let X be the set of all composition nodes which belong to some C_i in P and let $P' = P \cup X$. It remains true that every connected component of $T_0 \setminus P'$ has pathwidth at most $k-1$.

We construct a path decomposition of T. To start, we pick a path decomposition $\mathcal{I}^0 = (I^0_0, \ldots, I^0_N)$ of C_0 with each set I^0_i of size at most c. For each node v of C_0 let $[\![v]\!]$ denote the interval $\{i : v \in I_i\}$. We may assume that $|[\![a]\!]| = 1$ for each leaf a of C_0. We may also assume that if a, b are distinct leaves of C_0 then $[\![a]\!] \cap [\![b]\!] = \varnothing$, is if $[\![a]\!] = [\![b]\!] = \{i\}$, then we can simply replace I^0_i with the two sets $I^0_i \setminus \{a\}, I^0_i \setminus \{b\}$.

We may assume that $I^0_N = \varnothing$. Now modify $\mathcal{I}^0$ as follows: if $[\![v_0]\!] = \{i\}$, we add v_0 to the sets $I^0_{i+1}, \ldots, I^0_N$ so that $[\![v_0]\!] = \{i, \ldots, N\}$. This increases the size of the sets in $\mathcal{I}^0$ by at most 1.

Let V_0 be the set of composition nodes of T that belong to C_0. Define $i(v_0) = N$ and, if $v \in V_0 \setminus \{v_0\}$, let $i(v)$ be the unique element of $[\![v]\!]$. The important fact is that $i(v) \neq i(v')$ for $v, v' \in V_0$ with $v \neq v'$. Now for each $v \in V_0$ we modify $\mathcal{I}^0$ as follows. Let $C^0_1, \ldots, C^0_k$ denote the components of T other than C_0 or C_1 which have an edge in T_0 to v, and for each $1 \leq \ell \leq k$ let S_ℓ denote the connected component of $T_0 \setminus P'$ containing C^0_ℓ. Note that $S_\ell = (T_\ell)_0$ for an appropriate composition tree T_ℓ whose components are among those of T. Furthermore, $\mathrm{pw}(S_\ell) \leq k-1$; hence by induction, $\mathrm{pw}(T_\ell) \leq (s-1)(c+1)-1$. Let $\mathcal{I}_\ell$ denote a path decomposition of T_ℓ where each set has size at most $(s-1)(c+1)-1$. We now replace the set $I^0_{i(v)}$ in $\mathcal{I}^0$ with the sequence formed by taking the union of $I^0_{i(v)}$ with each member of $\mathcal{I}_1$, then with each member of $\mathcal{I}_2$, etc., and finally with each member of $\mathcal{I}_k$.

When this has been done for each $v \in V_0$, append the final set $\{v_0\}$ to the augmented $\mathcal{I}^0$ and call the resulting sequence $\mathcal{I}^0_+$. Note that each set in $\mathcal{I}^0_+$ has size at most $(c+1)+[(s-1)(c+1)-1] = s(c+1)-1$.

We proceed to do the same thing to C_1, with one small adjustment. Given a path decomposition $\mathcal{I}^1$ of C_1, if $[\![v_0]\!] = \{i\}$ and $[\![v_1]\!] = \{j\}$ and $i > j$, then we replace $\mathcal{I}^1$ with its "reverse," so that we may assume $i < j$. We then increase the sets in $\mathcal{I}^1$ to make $[\![v_0]\!]$ an initial interval and $[\![v_1]\!]$ a terminal interval. We then proceed as before, letting V_1 be the set of of composition nodes of T other than v_0 that belong to C_1, producing $\mathcal{I}^1_+$ whose final set if $\{v_1\}$. In the end we take the concatenation of $\mathcal{I}^0_+, \mathcal{I}^1_+, \ldots, \mathcal{I}^n_+$; it is a path decomposition of T in which every set has size at most $s(c+1)+1$. ∎

Step 3. We now prove the claim of the Lemma by induction on i. For $i = 1$ the claim follows from Step 1. For $i > 1$, suppose $T \in \mathcal{T}^i_d$ and define T_0 as above. Since the leaves of T_0 are the leaf components of T, Step 1 gives $\mathrm{pw}(T_0) \leq c$. By induction, each component of T has pathwidth at most $c^{i-1}+c^{i-2}-1$. Hence by Step 2, T has pathwidth at most $c((c^{i-1}+c^{i-2}-1)+1)-1 = c^i+c^{i-1}-1$. ∎