

Adjoint functors and tree duality

Jan Foniok¹ and Claude Tardif^{2*}

¹ETH Zurich, Institute for Operations Research, Sälimstrasse 101, 8092 Zurich, Switzerland
foniok@math.ethz.ch

²Royal Military College of Canada, PO Box 17000, Stn Forces, Kingston, Ontario, Canada, K7K 7B4
Claude.Tardif@rmc.ca

received May 20, 2008, revised May 6, 2009, accepted July 14, 2009.

A family \mathcal{T} of digraphs is a *complete set of obstructions* for a digraph H if for an arbitrary digraph G the existence of a homomorphism from G to H is equivalent to the non-existence of a homomorphism from any member of \mathcal{T} to G . A digraph H is said to have *tree duality* if there exists a complete set of obstructions \mathcal{T} consisting of orientations of trees. We show that if H has tree duality, then its arc graph δH also has tree duality, and we derive a family of tree obstructions for δH from the obstructions for H .

Furthermore we generalise our result to right adjoint functors on categories of relational structures. We show that these functors always preserve tree duality, as well as polynomial CSPs and the existence of near-unanimity functions.

2000 Mathematics Subject Classification: 16B50, 68R10, 18A40, 05C15

Keywords: constraint satisfaction, tree duality, adjoint functor

1 Introduction

Our primary motivation is the *H-colouring problem* (which has become popular under the name *Constraint Satisfaction Problem—CSP*): for a fixed digraph H (a *template*) decide whether an input digraph G admits a homomorphism to H . The computational complexity of H -colouring depends on the template H . For some templates the problem is known to be NP-complete, for others it is tractable (a polynomial-time algorithm exists). Assuming that $P \neq NP$, infinitely many complexity classes lie strictly between P and NP [10], but it has been conjectured that H -colouring belongs to no such intermediate class for any template H [3]. This conjecture has indeed been proved for symmetric templates H [5].

In this paper the focus is on tractable cases. Several conditions are known to imply the existence of a polynomial-time algorithm for H -colouring (definitions follow in the next two paragraphs): it is the case if H has a near-unanimity function (nuf), if H has bounded-treewidth duality, if H has tree duality, if H has finite duality (see [2, 3, 7]). Some of the conditions are depicted in the diagram (Fig. 1).

A *near-unanimity function* is a homomorphism f from H^k to H with $k \geq 3$ such that for all $x, y \in V(H)$ we have $f(x, x, x, \dots, x) = f(y, x, x, \dots, x) = f(x, y, x, \dots, x) = \dots = f(x, x, x, \dots, y) = x$.

*Supported by grants from the Natural Sciences and Engineering Research Council of Canada and the Academic Research Program.




Figure 1: The structure of tractable templates

The power H^k is the k -fold product $H \times H \times \dots \times H$ in the category of digraphs and homomorphisms, see [6].

A digraph is a tree (has treewidth k) if its underlying undirected graph is a tree (has treewidth k , respectively). A set \mathcal{F} of digraphs is a *complete set of obstructions* for H if for an arbitrary digraph G there exists a homomorphism from G to H if and only if no $F \in \mathcal{F}$ admits a homomorphism to G . A template has *bounded-treewidth duality* if it has a complete set of obstructions with treewidth bounded by a constant; it has *tree duality* if it has a complete set of obstructions consisting of trees; and it has *finite duality* if it has a finite complete set of obstructions.

There is a fairly straightforward way to generate templates with finite duality. For an arbitrary tree T there exists a digraph $D(T)$ such that $\{T\}$ is a complete set of obstructions for $D(T)$. The digraph $D(T)$ is unique up to homomorphic equivalence⁽ⁱ⁾; it is called the *dual* of T . Several explicit constructions are known (see [4, 9, 15, 16]). If \mathcal{F} is a finite set of oriented trees, then the product $D = \prod_{T \in \mathcal{F}} D(T)$ is a template with finite duality and \mathcal{F} is a complete set of obstructions for D . This construction yields all digraphs with finite duality [15], thus also proving that finite duality implies tree duality.

Encouraged by the full description of finite dualities, we aim to provide a construction for some more digraphs with tree duality. To this end we use the arc-graph construction and consider the class $\delta_\pi \mathcal{C}$ of digraphs generated from finite duals by taking iterated arc graphs and finite Cartesian products. We show that all templates in this class have tree duality. We provide an explicit construction of the resulting tree obstructions, which allows us to show that all the digraphs in $\delta_\pi \mathcal{C}$ have in fact *bounded-height tree duality*, that is, they have a complete set of obstructions consisting of trees of bounded algebraic height (these are

⁽ⁱ⁾Two digraphs H and H' are *homomorphically equivalent* if there exists a homomorphism from H to H' as well as a homomorphism from H' to H . Clearly, if H and H' are homomorphically equivalent, then H -colouring and H' -colouring are equivalent problems, because H and H' admit homomorphisms from exactly the same digraphs.

tree obstructions that allow a homomorphism to a fixed directed path). In this context we also prove that the problem of existence of a complete set of obstructions consisting of trees with bounded algebraic height is decidable.

The arc-graph construction is a special case of a more general phenomenon: it is a right adjoint in the category of digraphs and homomorphisms. We show in the more general setting of the category of relational structures that right adjoints (characterised by Pultr [17] for all locally presentable categories) preserve tractability of templates and moreover they preserve tree duality and existence of a near-unanimity function. In this case, nevertheless, it remains open to provide a nice general description of complete sets of obstructions.

We use some notions and properties of graphs and homomorphisms which the reader can look up in [6], as well as some category-theory notions, for which, e.g. [1, 13] may be consulted.

2 Arc graphs and tree duality

Let $G = (V, A)$ be a digraph. The *arc graph* of G is the digraph $\delta G = (A, \delta A)$, where

$$\delta A = \{((u, v), (v, w)) : (u, v), (v, w) \in A\}.$$

Notice that δ is an endofunctor⁽ⁱ⁾ in the category of digraphs and homomorphisms. This implies in particular that if $G \rightarrow H$, then $\delta G \rightarrow \delta H$. (The notation $G \rightarrow H$ means that there exists a homomorphism from G to H .)

If G is a digraph and \sim is an equivalence relation on its vertex set $V(G)$, the *quotient* G/\sim is the digraph $(V(G)/\sim, A)$, where $V(G)/\sim$ is the set of all equivalence classes of \sim on $V(G)$, and for $X, Y \in V(G)/\sim$ we have $(X, Y) \in A$ if and only if there exist $x \in X$ and $y \in Y$ such that $(x, y) \in A(G)$.

Suppose still that $G = (V, A)$ is a digraph. Let $V' = \{o_u, t_u : u \in V\}$ and let $A' = \{(o_u, t_u) : u \in V\}$. Define the relation \sim_0 such that $t_u \sim_0 o_v$ if and only if $(u, v) \in A$. Let \sim be the minimal equivalence relation on V' containing \sim_0 . Set $\delta^{-1}G = (V', A')/\sim$. In the following, we use the notation $V'(G) = V'$, $A'(G) = A'$ and \sim_0 and \sim for the sets and relations appearing in the definition of δ^{-1} ; the precise meaning will be clear from the context. Now δ^{-1} is also an endofunctor in the category of digraphs. Strictly speaking, it is not an inverse of δ ; its name is chosen because of the following property.

Proposition 1 *For any digraphs G and H ,*

$$G \rightarrow \delta H \quad \text{if and only if} \quad \delta^{-1}G \rightarrow H.$$

Proof: Let $f : G \rightarrow \delta H$ be a homomorphism. Then there exist two homomorphisms $o, t : G \rightarrow H$ such that $f(u) = (o(u), t(u))$ for all $u \in V(G)$. Define the mapping $\hat{g} : V'(G) \rightarrow V(H)$ by $\hat{g}(o_u) = o(u)$ and $\hat{g}(t_u) = t(u)$. If $t_u \sim_0 o_v$, then $(u, v) \in A(G)$, whence $(f(u), f(v)) \in A(\delta H)$ and thus $t(u) = o(v)$. Therefore \hat{g} is constant on the equivalence classes of \sim , and it induces a homomorphism from $A'(G)/\sim = \delta^{-1}G$ to H .

Conversely, let $g : \delta^{-1}G \rightarrow H$ be a homomorphism. We define $f : V(G) \rightarrow V(\delta H)$ by $f(u) = (g(o_u/\sim), g(t_u/\sim))$. If $(u, v) \in A(G)$, then $t_u/\sim = o_v/\sim$, whence $(f(u), f(v)) \in A(\delta H)$. Therefore f is a homomorphism. \square

⁽ⁱ⁾An *endofunctor* is a functor from a category to itself.

Thus δ and δ^{-1} are Galois adjoints⁽ⁱ⁾ with respect to the ordering by existence of homomorphisms. They are in fact adjoint functors in the category of digraphs and homomorphisms. We return to this topic in Section 4. For the moment we aim to prove that δ preserves tree duality. More precisely, from the family \mathcal{T} of tree obstructions of H , we will derive the family $\text{Sproink}(\mathcal{T})$ of tree obstructions of δH .

The *algebraic height* of an oriented tree T is the minimum number of arcs of a directed path to which T maps homomorphically. The algebraic height of every finite oriented tree is well-defined and finite, since every such tree admits a homomorphism to some finite directed path. Thus a tree T is of *height at most one* if its vertex set can be split into two parts $0_T, 1_T$ in such a way that for every arc (x, y) of T we have $x \in 0_T$ and $y \in 1_T$. Note that if the tree T has no arcs, then it has only one vertex and thus one of the sets $0_T, 1_T$ is empty and the other one is a singleton.

Let T be a tree. For every vertex u of T , let $F(u)$ be a tree of height at most one. For each arc e of T incident with u , let there be a fixed vertex $v(e, F(u))$ in $F(u)$ such that if u is the initial vertex of e , then $v(e, F(u)) \in 1_{F(u)}$, and if u is the terminal vertex of e , then $v(e, F(u)) \in 0_{F(u)}$.⁽ⁱⁱ⁾ A tree S is now constructed by taking all the trees $F(u)$ for all vertices u of T , and by identifying the vertex $v(e, F(u))$ with $v(e, F(u'))$ whenever $e = (u, u')$ is an arc of T .

Any such tree S constructed from T by the above procedure is called a *sproink* of T . The set of all sproinks of a tree T is denoted by $\text{Sproink}(T)$. The following lemma asserts that sproinks of obstructions for a template H are indeed obstructions for its arc graph δH .

Lemma 2 *Let T be a tree and H a digraph such that $T \dashrightarrow H$. If $S \in \text{Sproink}(T)$, then $S \dashrightarrow \delta H$.*

Proof: We prove that $T \rightarrow \delta^{-1}S$. Consequently $\delta^{-1}S \dashrightarrow H$ because $T \dashrightarrow H$, and therefore $S \dashrightarrow \delta H$ by Proposition 1.

Thus let $S \in \text{Sproink}(T)$. For a vertex u of T , consider the tree $F(u)$, which is a subgraph of S . Since $F(u)$ has height at most one, its vertices are partitioned into the sets $0_{F(u)}$ and $1_{F(u)}$. The set $V'(S)$, which appears in the definition of $\delta^{-1}S$, contains $V'(F(u))$ as a subset. If (x, y) is an arc of $F(u)$, then $t_x \sim_0 o_y$. Thus whenever $x \in 0_{F(u)}$ and $y \in 1_{F(u)}$, then $t_x \sim o_y$. Hence for any vertex u of T there exists a unique vertex $f(u)$ of $\delta^{-1}S$ that is equal to t_x/\sim for all $x \in 0_{F(u)}$ and to o_y/\sim for all $y \in 1_{F(u)}$.

In this way, we have defined a mapping $f : V(T) \rightarrow V(\delta^{-1}S)$.

Now assume that $e = (u, v)$ is an arbitrary arc of T . Then the vertex $v(e, F(u))$, which belongs to $1_{F(u)}$, has been identified with $v(e, F(v))$, which belongs to $0_{F(v)}$. Let this identified vertex be x ; it is a vertex of S . By definition, $f(u) = o_x/\sim$ because $x \in 1_{F(u)}$, and $f(v) = t_x/\sim$ because $x \in 0_{F(v)}$. Of course $(o_x/\sim, t_x/\sim) \in A(\delta^{-1}S)$. Therefore $f : T \rightarrow \delta^{-1}S$ is a homomorphism, as we have promised to prove. \square

For a set \mathcal{F} of trees, let $\text{Sproink}(\mathcal{F}) = \bigcup_{T \in \mathcal{F}} \text{Sproink}(T)$.

Theorem 3 *Let \mathcal{F} be a set of trees which is a complete set of obstructions for a template H . Then $\text{Sproink}(\mathcal{F})$ is a complete set of obstructions for δH .*

⁽ⁱ⁾Let X and Y be partially ordered sets. Mappings $\phi : X \rightarrow Y$ and $\psi : Y \rightarrow X$ are *Galois adjoints* if $\phi(x) \leq_Y y \Leftrightarrow x \leq_X \psi(y)$ for all elements $x \in X$ and $y \in Y$.

⁽ⁱⁱ⁾It follows that if u is neither a source nor a sink of T , then both $0_{F(u)}$ and $1_{F(u)}$ are non-empty, and so in this case $F(u)$ is not a single vertex. If u is a source or a sink of T , then $F(u)$ may be an arbitrary tree of height at most one.




Figure 2: A thunderbolt

Proof: Lemma 2 implies that $\text{Sproink}(\mathcal{F})$ is a set of obstructions for δH . It remains to prove that it is complete, that is whenever $G \dashrightarrow \delta H$, then there exists some $S \in \text{Sproink}(\mathcal{F})$ such that $S \rightarrow G$.

So let $G \dashrightarrow \delta H$. Thus by Proposition 1 we have $\delta^{-1}G \dashrightarrow H$. Hence there exists a tree $T \in \mathcal{F}$ such that $T \rightarrow \delta^{-1}G$, because \mathcal{F} is a complete set of obstructions for H . Consequently it suffices to prove that if $T \rightarrow \delta^{-1}G$ then there exists $S \in \text{Sproink}(T)$ such that $S \rightarrow G$.

Thus assume that $f : T \rightarrow \delta^{-1}G$ is a homomorphism. For every $u \in V(T)$, the image $f(u)$ is a \sim -equivalence class; put

$$1_u = \{y \in V(G) : o_y \in f(u)\},$$

$$0_u = \{x \in V(G) : t_x \in f(u)\}.$$

Then $f(u) = 1_u \cup 0_u$, and by the definition of \sim as the least equivalence containing \sim_0 , there exists a tree $F(u)$ of height at most one and a homomorphism $g_u : F(u) \rightarrow G$ such that $g_u(0_{F(u)}) = 0_u$ and $g_u(1_{F(u)}) = 1_u$. For every arc (u, v) of T , we have $(f(u), f(v)) \in A(\delta^{-1}G)$ so there exists $x \in V(G)$ such that $o_x \in f(u)$ and $t_x \in f(v)$.

We then select $y \in 1_{F(u)}$ and $z \in 0_{F(v)}$ such that $g_u(y) = g_v(z) = x$, and identify them. Proceeding with all such identifications, we construct a tree $S \in \text{Sproink}(T)$ such that $g = \bigcup_{u \in V(T)} g_u : S \rightarrow G$ is a well-defined homomorphism. \square

Corollary 4 *If a digraph H has tree duality, then its arc graph δH also has tree duality.* \square

Example 1 Consider $T = \vec{P}_4$, the directed path with four arcs, and its dual $D = \vec{T}_4$, the transitive tournament on four vertices. Here δD has six vertices, but its core⁽ⁱ⁾ is the directed path \vec{P}_2 with two arcs. It is well known that a directed graph G admits a homomorphism to \vec{P}_2 if and only if it does not admit a homomorphism from a “thunderbolt”, that is, an oriented path with two forward arcs at the beginning and at the end, and with an odd-length alternating path between them (see Fig. 2). Thus the family of all thunderbolts is a complete set of tree obstructions for \vec{P}_2 .

Our construction $\text{Sproink}(T)$ gives all obstructions obtained by stacking five trees L_0, L_1, L_2, L_3, L_4 of height at most one, with one top vertex of L_i identified with one bottom vertex of L_{i+1} for $i = 0, 1, 2, 3$.

⁽ⁱ⁾The *core* of a digraph is any of its smallest subgraphs to which it admits a homomorphism. Every digraph H has a unique core C (up to isomorphism), which is moreover the only core homomorphically equivalent to it. In fact, the core C of H is a *retract* of H , which means that there exists a homomorphism $\rho : H \rightarrow C$ whose restriction on C is the identity mapping (such a homomorphism is called a *retraction*).

The example of thunderbolts shows that in fact L_0 can be restricted to a single (top) vertex, and L_4 can be restricted to a single (bottom) vertex. The same holds for leaves of general trees. Also, L_1, L_2, L_3 can be restricted to *paths* of height one, and it is also true in general that it is sufficient to consider sproinks obtained by replacing vertices by *paths* of height at most one. In fact the name ‘‘sproink’’ is inspired by picturing such a path springing out of every non-leaf of T .

The results of this section show that we can construct an interesting class of templates with tree duality by repeatedly applying the arc-graph construction to digraphs with finite duality. Moreover, if templates H_1, H_2, \dots, H_k all have tree duality, then also their product $H_1 \times H_2 \times \dots \times H_k$ has tree duality as the union of the respective complete sets of obstructions of the factors is a complete set of obstructions for the product. The resulting class of templates is subject to examination in the next section.

3 Finite duality

Following [15], every tree T admits a dual $D(T)$ such that for every digraph G , we have $G \rightarrow D(T)$ if and only if $T \rightarrow G$. A digraph H has finite duality if and only if it is homomorphically equivalent to a finite product of duals of trees.

In this section, we consider the class $\delta_\pi\mathcal{C}$, the smallest class of digraphs that contains all duals of trees and is closed under taking arc graphs, finite products and homomorphically equivalent digraphs. It follows from Corollary 4 that all elements of $\delta_\pi\mathcal{C}$ have tree duality. Moreover we know how to construct a complete set of obstructions for each of these templates, using iterated Sproink constructions and unions. The question then arises as to how significant the class $\delta_\pi\mathcal{C}$ is within the class of digraphs with tree duality. It turns out that the digraphs in $\delta_\pi\mathcal{C}$ have properties that are not shared by all digraphs with tree duality.

A digraph H has *bounded-height tree duality* provided there exists a constant m such that H admits a complete set of obstructions consisting of trees of algebraic height at most m .

Proposition 5 (i) *Every core in $\delta_\pi\mathcal{C}$ admits a near-unanimity function.*

(ii) *Every member of $\delta_\pi\mathcal{C}$ has bounded-height tree duality.*

Proof: (i): By Corollary 4.5 of [11], every structure with finite duality admits a near-unanimity function. Therefore it suffices to show that the class of structures admitting a near-unanimity function is closed under taking cores, finite products and the arc-graph construction.

Let C be the core of H , $\rho : H \rightarrow C$ a retraction and $f : H^k \rightarrow H$ a near-unanimity function. Since C is an induced subgraph of H , the restriction $\rho \circ f \upharpoonright C^k$ is a near-unanimity function on C .

Suppose $f_i : H_i^{k_i} \rightarrow H_i, i = 1, \dots, m$ are near-unanimity functions. For $k = \max\{k_i : i = 1, \dots, m\}$, we define k -ary near-unanimity functions $g_i : H_i^k \rightarrow H_i$ by $g_i(x_1, \dots, x_k) = f_i(x_1, \dots, x_{k_i})$. For $H = \prod_{i=1}^m H_i$ we then define a near-unanimity function $g : H^k \rightarrow H$ coordinate-wise, by putting

$$g((x_{1,1}, \dots, x_{m,1}), \dots, (x_{1,k}, \dots, x_{m,k})) = (g_1(x_{1,1}, \dots, x_{1,k}), \dots, g_m(x_{m,1}, \dots, x_{m,k})).$$

Now suppose that $f : H^k \rightarrow H$ is a near-unanimity function. Then $(\delta H)^k$ is naturally isomorphic to $\delta(H^k)$, and we define $g : (\delta H)^k \rightarrow \delta H$ by

$$g((u_1, v_1), \dots, (u_k, v_k)) = (f(u_1, \dots, u_k), f(v_1, \dots, v_k)).$$

The fact that f is a homomorphism implies that g is well defined, and g is a homomorphism by the definition of adjacency in δH . Also, g clearly satisfies the near-unanimity identities, so it is a near-unanimity function on δH .

(ii): The class of digraphs with bounded-height tree duality is obviously preserved by taking cores and finite products. By Theorem 3, if H has a complete set of obstructions consisting of trees of algebraic height at most k , then δH has a complete set of obstructions consisting of trees of algebraic height at most $k + 1$, so the class of digraphs with bounded-height tree duality is also preserved by the arc-graph construction. \square

We know a core digraph with tree duality but no near-unanimity function and no bounded-height tree duality. (The example is complicated and out of the scope of this paper, therefore we omit it.) Thus the class $\delta_\pi \mathcal{C}$ does not capture all core digraphs with tree duality. The problem of generating all structures with tree duality by means of suitable functors applied to structures with finite duality remains nevertheless interesting.

Membership in $\delta_\pi \mathcal{C}$ is not known to be decidable. In the remainder of this section, we show that bounded-height tree duality is decidable.

Given a digraph H , the n -th crushed cylinder H_n^* is the quotient $(H^2 \times P_n)/\simeq_n$, where P_n is the path with arcs $(0, 0), (0, 1), (1, 2), \dots, (n-1, n), (n, n)$, and \simeq_n is the equivalence defined by

$$(u, v, i) \simeq_n (u', v', j) \Leftrightarrow \begin{cases} i = j = 0 \text{ and } u = u', \\ \text{or } i = j = n \text{ and } v = v', \\ \text{or } (u, v, i) = (u', v', j). \end{cases}$$

Theorem 6 *For a core digraph H with tree duality, the following are equivalent:*

- (1) H has bounded-height tree duality,
- (2) For some n we have $H_n^* \rightarrow H$.
- (3) There exists a directed (upward) path from the first projection to the second in H^{H^2} .

Proof: (1) \Rightarrow (2): The two subgraphs obtained from H_n^* by removing the two ends both admit homomorphisms to H . Therefore, if a tree obstruction of H admits a homomorphism to H_n^* , its image must intersect the two ends hence its algebraic length must be at least n .

$\neg(1) \Rightarrow \neg(2)$: Let T be a critical obstruction of H of algebraic length $n + 2$. Let T_0, T_n be the subgraphs of T obtained by removing the vertices of height 0 and $n + 2$ respectively. Then there exists homomorphisms $f_0 : T_0 \rightarrow H$ and $f_n : T_n \rightarrow H$. Let $h : T \rightarrow P_{n+2}$ be the height function of T . We define a map $f : T \rightarrow H_n^*$ by

$$f(u) = \begin{cases} (f_n(u), f_0(u), h(u) - 1)/\simeq_n & \text{if } h(u) \notin \{0, n + 2\}, \\ (f_n(u), f_n(u), 0)/\simeq_n & \text{if } h(u) = 0, \\ (f_0(u), f_0(u), n)/\simeq_n & \text{if } h(u) = n + 2. \end{cases}$$

Let (u, v) be an arc of T . Then $h(v) = h(u) + 1$. If $\{h(u), h(v)\} \cap \{0, n + 2\} = \emptyset$, we clearly have $(f(u), f(v)) \in A(H_n^*)$. If $h(u) = 0$, then $f(u) = (f_n(u), f_n(u), 0)/\simeq_n$ is an in-neighbour of $(f_n(v), f_n(v), 0)/\simeq_n = (f_n(v), f_0(v), 0)/\simeq_n = f(v)$, and if $h(v) = n + 2$, $f(v) = (f_0(v), f_0(v), n)/\simeq_n$ is an out-neighbour of $f(u)$ because

$$(f_0(u), f_0(u), n)/\simeq_n = (f_n(u), f_0(u), n)/\simeq_n = f(u).$$

Therefore f is a homomorphism.

(2) \Leftrightarrow (3): This equivalence follows easily from the definition. \square

Corollary 7 *The problem whether an input digraph has bounded-height tree duality is decidable.*

Proof: It is decidable whether a digraph has tree duality [3] (see Theorem 11 below). For a digraph with tree duality, bounded height of the obstructions (the condition (1) of Theorem 6) is equivalent to the condition (3), which involves directed reachability in a finite graph. Hence bounded-height tree duality is decidable. \square

4 Adjoint functors and generation of tractable templates

The correspondence of Proposition 1 can be extended to a wide class of functors presented in this section. To illustrate this extension, we first redefine δ in terms of patterns. Let P be the digraph with vertices 0, 1 and arc $(0, 1)$, and Q the digraph with vertices 0, 1, 2 and arcs $(0, 1)$, $(1, 2)$. Furthermore let $q_1, q_2 : P \rightarrow Q$ be the homomorphisms mapping the arc $(0, 1)$ to $(0, 1)$ and $(1, 2)$ respectively. For an arbitrary digraph G , its arc graph δG can be described as follows: The vertices of δG are the arcs of G , that is, the homomorphisms $f : P \rightarrow G$. The arcs of δG are the couples of consecutive arcs in G , that is, the couples (f_1, f_2) such that there exists a homomorphism $g : Q \rightarrow G$ satisfying $g \circ q_1 = f_1$ and $g \circ q_2 = f_2$. Thus the functor δ is generated by the pattern $\{P, (Q, q_1, q_2)\}$ in a way that generalises quite naturally.

The rest of this section deals with relational structures. A *vocabulary* is a finite set $\sigma = \{R_1, \dots, R_m\}$ of relation symbols, each with an arity r_i assigned to it. A σ -*structure* is a relational structure $A = \langle \underline{A}; R_1(A), \dots, R_m(A) \rangle$ where \underline{A} is a non-empty set called the *universe* of A , and $R_i(A)$ is an r_i -ary relation on \underline{A} for each i . Homomorphisms of relational structures are relation-preserving mappings between universes; a homomorphism is defined only between structures with the same vocabulary. Cores, trees, quotient structures, etc. can also be defined in the context of relational structures, consult [12] for the details (see also [8, 11]). The notions of the constraint satisfaction problem, template, and tree duality also carry over naturally from the setting of digraphs.

Let σ and τ be two vocabularies. Let P be a σ -structure, and for every relation R of τ of arity $r = a(R)$, let Q_R be a σ -structure with r fixed homomorphisms $q_{R,i} : P \rightarrow Q_R$ for $i = 1, \dots, r$. Then the family $\{P\} \cup \{(Q_R, q_{R,1}, \dots, q_{R,a(R)}) : R \in \tau\}$ defines a functor Ψ from the category \mathcal{A} of σ -structures to the category \mathcal{B} of τ -structures as follows.

- For a σ -structure A , let $B = \Psi A$ be a τ -structure whose universe is the set of all homomorphisms $f : P \rightarrow A$.
- For every relation R of τ of arity $r = a(R)$, let $R(B)$ be the set of r -tuples (f_1, \dots, f_r) such that there exists a homomorphism $g : Q_R \rightarrow A$ such that for $i = 1, \dots, r$ we have $g \circ q_{R,i} = f_i$.

It was shown by Pultr [17] that functors Ψ defined by means of patterns are right adjoints into a category of relational structures characterised by axioms of a specific type. We exhibit their corresponding left adjoints Ψ^{-1} in the case when both the domain and the range of Ψ is the category of all relational structures with a given vocabulary.

For every τ -structure B , we define a σ -structure $\Psi^{-1}B = A/\sim$, where

- A is a disjoint union of σ -structures; for every element x of the universe of B , A contains a copy P_x of P , and for every $R \in \tau$ and $(x_1, \dots, x_r) \in R(B)$, A contains a copy $Q_{R,(x_1, \dots, x_r)}$ of Q_R .
- \sim is the least equivalence which identifies every element u of P_{x_i} with its image $q_{R,i}(u)$ in $Q_{R,(x_1, \dots, x_r)}$, for every $R \in \tau$, every $(x_1, \dots, x_r) \in R(B)$ and every $i \in \{1, \dots, r\}$.

Proposition 8 ([17]) *For any τ -structure B and σ -structure A ,*

$$B \rightarrow \Psi A \quad \text{if and only if} \quad \Psi^{-1}B \rightarrow A.$$

Proof: Let $h : B \rightarrow \Psi A$ be a homomorphism, and put $h(b) = f_b : P \rightarrow A$. Then for every $b \in B$, the mapping f_b corresponds to a well-defined homomorphism to A from a copy P_b of P . Also, for every $R \in \tau$ and $(b_1, \dots, b_r) \in R(B)$, we have $(h(b_1), \dots, h(b_r)) \in R(\Psi A)$, so there exists a homomorphism $g_{(b_1, \dots, b_r)} : Q_R \rightarrow A$ such that $f_{b_i} = g_{(b_1, \dots, b_r)} \circ q_{R,i}$ for $i = 1, \dots, r$; the mapping $g_{(b_1, \dots, b_r)}$ corresponds to a well-defined homomorphism from a copy $Q_{R,(b_1, \dots, b_r)}$ of Q_R to A . Therefore $\bigcup_{b \in B} f_b \cup \bigcup_{R \in \tau} \bigcup_{(b_1, \dots, b_r) \in R(B)} g_{(b_1, \dots, b_r)}$ corresponds to a well-defined homomorphism $\hat{h} : \bigcup_{b \in B} P_b \cup \bigcup_{R \in \tau} \bigcup_{(b_1, \dots, b_r) \in R(B)} Q_{R,(b_1, \dots, b_r)} \rightarrow A$, such that if $x \sim y$, then $\hat{h}(x) = \hat{h}(y)$. Therefore \hat{h} induces a homomorphism from the quotient structure $\Psi^{-1}B$ to A .

Conversely, if $h : \Psi^{-1}B \rightarrow A$ is a homomorphism, we define a homomorphism $\hat{h} : B \rightarrow \Psi A$ by $\hat{h}(b) = f_b$, where f_b corresponds to the restriction of h to the quotient of P_b in $\Psi^{-1}B$. Indeed, if $R \in \tau$ and $(b_1, \dots, b_r) \in R(B)$, then the restriction of h to the quotient of $Q_{R,(b_1, \dots, b_r)}$ in $\Psi^{-1}B$ corresponds to a homomorphism $g : Q_R \rightarrow A$ such that $f_{b_i} = g \circ q_{R,i}$ for $i = 1, \dots, r$, whence $(\hat{h}(b_1), \dots, \hat{h}(b_r)) \in R(\Psi A)$. \square

Corollary 9 *If a σ -structure A has polynomial CSP, then the τ -structure ΨA also has polynomial CSP.* \square

In fact, Corollary 4 generalises as follows.

Theorem 10 *If a σ -structure A has tree duality, then the τ -structure ΨA also has tree duality.*

We prove Theorem 10 using Feder and Vardi's characterisation of structures with tree duality. For a σ -structure A , let $\mathcal{U}A$ be the σ -structure defined as follows. The universe of $\mathcal{U}A$ is the set of all nonempty subsets of A , and for $R \in \sigma$ of arity r , $R(\mathcal{U}A)$ is the set of all r -tuples (X_1, \dots, X_r) such that for all $j \in \{1, \dots, r\}$ and $x_j \in X_j$ there exist $x_k \in X_k, k \in \{1, \dots, r\} \setminus \{j\}$ such that $(x_1, \dots, x_r) \in R(A)$.

Theorem 11 ([3]) *A structure A has tree duality if and only if there exists a homomorphism from $\mathcal{U}A$ to A .* \square

Proof of Theorem 10: Suppose A has tree duality. Then there is a homomorphism $f : \mathcal{U}A \rightarrow A$. Let $U = \mathcal{P}(\Psi A) \setminus \{\emptyset\}$ be the universe of $\mathcal{U}\Psi A$ and let $S \in U$. For $p \in P$, define $S_p = \{f(p) : f \in S\} \in \mathcal{U}A$, and $f_S(p) = f(S_p)$. We claim that $f_S : P \rightarrow A$ is a homomorphism. Indeed, for $R \in \sigma$ and $(p_1, \dots, p_r) \in R(P)$, the r -tuples $(f(p_1), \dots, f(p_r)) \in R(A)$ for all $f \in S$ prove that $(S_{p_1}, \dots, S_{p_r}) \in R(\mathcal{U}A)$, whence $(f_S(p_1), \dots, f_S(p_r)) = (f(S_{p_1}), \dots, f(S_{p_r})) \in R(A)$.

Thus we define a map $\hat{f} : \mathcal{U}\Psi A \rightarrow \Psi A$ by $\hat{f}(S) = f_S$. We show that it is a homomorphism. For $R \in \tau$ and $(S_1, \dots, S_r) \in R(\mathcal{U}\Psi A)$, every $f_i \in S_i$, $1 \leq i \leq r$ is contained in an r -tuple $(h_1, \dots, h_r) \in R(\Psi A)$ with $f_j \in S_j$ for $1 \leq j \leq r$ and $h_i = f_i$, whence there exists a homomorphism $g_{(h_1, \dots, h_r)} : Q_R \rightarrow A$ such that $h_j = g_{(h_1, \dots, h_r)} \circ q_{R,j}$ for $j = 1, \dots, r$. For $x \in Q$, let T_x be the set of all images $g_{(h_1, \dots, h_r)}(x) \in A$ (with (S_1, \dots, S_r) fixed), and $g_{(S_1, \dots, S_r)}(x) = f(T_x)$. Then $g_{(S_1, \dots, S_r)} : Q_R \rightarrow A$ is a homomorphism, and for $x \in q_{R,j}(P)$ we have $T_x = S_x$ (because they are images of x under restrictions of the same homomorphisms), whence $g_{(S_1, \dots, S_r)}(x) = f_{S_j}(x)$. Thus $f_{S_j} = g_{(S_1, \dots, S_r)} \circ q_{R,j}$ for $j = 1, \dots, r$. Consequently $(f_{S_1}, \dots, f_{S_r}) = (\hat{f}(S_1), \dots, \hat{f}(S_r)) \in R(\Psi A)$. This shows that \hat{f} is a homomorphism. \square

Unlike the case of the arc-graph construction, we are unable to provide an explicit description of the tree obstructions of ΨA in terms of those of A for a general right adjoint Ψ . However, in isolated cases we can do it, as the following example shows.

Example 2 The endofunctor Ψ on the category of digraphs is defined via the pattern $\{P, (Q, q_1, q_2)\}$, where $P = \vec{P}_1$ is the one-arc path $u \rightarrow v$, $Q = \vec{P}_3$ is the directed path $0 \rightarrow 1 \rightarrow 2 \rightarrow 3$, the homomorphism $q_1 : u \mapsto 0$, $v \mapsto 1$, and finally $q_2 : u \mapsto 2$, $v \mapsto 3$.

Let T be a tree of algebraic height h and consider the unique homomorphism t from T to the directed path \vec{P}_h . The arcs of T are of two kinds: *blue arcs* $A_b(T) = \{(x, y) : t(x) = 2k, t(y) = 2k + 1 \text{ for some integer } k\}$ and *red arcs* $A_r(T) = \{(x, y) : t(x) = 2k + 1, t(y) = 2k + 2 \text{ for some integer } k\}$. We define two equivalence relations on the vertices of T : $x \sim_b y$ if the (not necessarily directed) path from x to y in T has only blue arcs, and $x \sim_r y$ if the path from x to y in T has only red arcs. Then T has two Ψ -Sproinks, namely T/\sim_b and T/\sim_r with loops removed.

For a collection \mathcal{T} of trees, let $\Psi\text{-Sproink}(\mathcal{T})$ be the set of all Ψ -Sproinks of the trees contained in \mathcal{T} . We claim that if \mathcal{T} is a complete set of obstructions for a template H , then $\Psi\text{-Sproink}(\mathcal{T})$ is a complete set of obstructions for ΨH . To prove it, we follow the idea of the proofs of Lemma 2 and Theorem 3.

First we prove that $T \rightarrow \Psi^{-1}(T/\sim_b)$. This is not difficult: every blue arc of T was contracted to a vertex of T/\sim_b and this vertex was blown up to an arc in $T \rightarrow \Psi^{-1}(T/\sim_b)$. Thus we can map blue arcs to the corresponding blown-up arcs. Red arcs of T are also arcs of T/\sim_b , and hence we can map each red arc to the arc $(1, 2)$ of the corresponding copy of Q in $\Psi^{-1}(T/\sim_b)$. Clearly such a mapping is a homomorphism.

Analogously we show that $T \rightarrow \Psi^{-1}T/\sim_r$.

Finally we want to prove that if $T \rightarrow \Psi^{-1}G$, then either $T/\sim_b \rightarrow G$ or $T/\sim_r \rightarrow G$. Suppose that $f : T \rightarrow \Psi^{-1}G$. Then some arcs of T are mapped by f to arcs corresponding to vertices of G (arcs of copies of P), and others are mapped to arcs corresponding to arcs of G (arcs $(1, 2)$ of copies of Q). Let us call the former v-arcs and the latter a-arcs. It follows from the definition of Ψ^{-1} that either all blue arcs of T are v-arcs and all red arcs of T are a-arcs, or all blue arcs of T are a-arcs and all red arcs of T are v-arcs. In the former case $T/\sim_b \rightarrow G$, while in the latter case $T/\sim_r \rightarrow G$.

It is notable that in the above example each tree obstruction for H generates finitely many obstructions for ΨH . This is no accident.

Theorem 12 *Let Ψ be a functor generated by a pattern $\{P\} \cup \{(Q_R, q_{R,1}, \dots, q_{R,a(R)}) : R \in \tau\}$, where for every $R \in \tau$ and $1 \leq i < j \leq a(R)$, the image $q_{R,i}(P)$ is vertex-disjoint from $q_{R,j}(P)$. If a σ -structure A has finite duality, then the τ -structure ΨA also has finite duality.*

The proof uses the characterisation of structures with finite duality of [11]. The *square* of a σ -structure B is the structure $B \times B$. It contains the *diagonal* $\Delta_{B \times B} = \{(b, b) : b \in B\}$. An element a of B is *dominated* by an element b of B if for every $R \in \sigma$, for every i and every $(x_1, \dots, x_{a(R)}) \in R(B)$ with $x_i = a$, we have $(y_1, \dots, y_{a(R)}) \in R(B)$ with $y_i = b$ and $y_j = x_j$ for $j \neq i$. A structure B *dismantles* to its induced substructure C if there exists a sequence x_1, \dots, x_k of distinct elements of B such that $B \setminus C = \{x_1, \dots, x_k\}$ and for each $1 \leq i \leq k$ the element x_i is dominated in the structure induced by $C \cup \{x_i, \dots, x_k\}$. The sequence x_1, \dots, x_k is then called a *dismantling sequence* of B on C .

Theorem 13 ([11]) *A structure has finite duality if and only if it has a retract whose square dismantles to its diagonal.* \square

Proof Proof of Theorem 12: Let A be a σ -structure with finite duality. Without loss of generality, we assume that A is a core, so that A has no proper retracts; thus the square of A dismantles to its diagonal. Let $(x_1, y_1), \dots, (x_k, y_k)$ be a dismantling sequence of $A \times A$ on $\Delta_{A \times A}$. Then $\Psi(A \times A) \cong \Psi A \times \Psi A$; we want to prove that it dismantles to $\Delta_{\Psi A \times \Psi A} \cong \Psi \Delta_{A \times A}$.

For $i \in \{1, \dots, k\}$, define X_i to be the substructure of $A \times A$ induced by the set

$$\Delta_{A \times A} \cup \{(x_i, y_i), \dots, (x_k, y_k)\},$$

and let $X_{k+1} = \Delta_{A \times A}$. We will show that ΨX_i can be dismantled to ΨX_{i+1} , $i = 1, \dots, k$.

Let $b = (b_1, b_2)$ be an element dominating $a = (x_i, y_i)$ in X_i . Let $f \in \Psi X_i \setminus \Psi X_{i+1}$, and assume that $f = (f_1, f_2) : P \rightarrow A \times A$. Then there exists (at least one) $p_0 \in P$ such that $f(p_0) = a$. We define $g = (g_1, g_2) : P \rightarrow A \times A$ by $g(p_0) = b$ and $g(p) = f(p)$ if $p \neq p_0$. Since b dominates a , g is a homomorphism, and obviously $g \in \Psi X_i$. We claim that g dominates f . Indeed, for $R \in \tau$ and $(f_1, \dots, f_{a(R)}) \in R(\Psi X_i)$ such that $f = f_j$, there exists a homomorphism $h : Q_R \rightarrow X_i$ such that $f = h \circ q_{R,j}$. Define $h' : Q_R \rightarrow X_i$ by $h'(q_{R,j}(p_0)) = b$ and $h'(z) = h(z)$ for $z \neq q_{R,j}(p_0)$. Since b dominates $a = h(q_{R,j}(p_0))$, the mapping h' is a homomorphism. By hypothesis, for $\ell \neq j$, the image $q_{R,\ell}(P)$ is disjoint from $q_{R,j}(P)$, whence $f_\ell = h' \circ q_{R,\ell}$, while $h' \circ q_{R,j} = g$. Therefore $R(\Psi X_i)$ contains all the $a(R)$ -tuples needed to establish the domination of f by g .

Let p_1, p_2, \dots, p_m be an enumeration of the elements of P . We dismantle ΨX_i to ΨX_{i+1} by successively removing the functions f such that $f(p_j) = (x_i, y_i)$ for $j = 1, \dots, m$. Proceeding in this way for $i = 1, \dots, k$, we get a dismantling of $\Psi A \times \Psi A \cong \Psi X_1$ to $\Psi X_{k+1} \cong \Delta_{\Psi A \times \Psi A}$. Therefore ΨA has finite duality. \square

Perhaps the lack of knowledge of a general construction is natural since there is no restriction on the pattern $\{P\} \cup \{(Q_R, q_{R,1}, \dots, q_{R,a(R)}) : R \in \tau\}$. On the other hand, there are many possible transformations T' on a family T of tree obstructions, in the style of $\text{Sproink}(T)$. Any such transformation gives rise to a complete set of obstructions to homomorphisms into a structure $H' = \prod_{T \in T'} D_T$; however in general there is no way of guaranteeing that such structure H' is finite, even when T is a complete set of obstructions for a finite structure H .

5 Concluding comments

In this paper we tried to shed more light on the structure of tractable templates with tree duality. Let us turn our attention one more time to Fig. 1. The grey areas in the diagram are areas that need a closer look in future research.

Currently we do not know any digraph with a near-unanimity function and with bounded-height tree duality that could not be generated using right adjoints and products, starting from digraphs with finite duality; it is not clear whether any such “reasonable” class of structures with tree duality can be generated from structures with finite duality with a “reasonable” set of adjoint functors.

We have shown here that possession of bounded-height tree duality is decidable. It is natural to ask what its complexity is; in particular, whether it is complete for some class of problems.

Equally interesting is the decidability of membership in other classes depicted in Fig. 1. Tree duality is known to be decidable [3], but not known to be in PSPACE. Our decision procedure for bounded-height duality is in PSPACE for graphs with tree duality; this suggests that checking tree duality may be harder than checking bounded height of the obstructions.

Furthermore, finite duality is NP-complete [11]. The decidability of bounded-treewidth duality is unknown⁽ⁱ⁾, and so is the decidability of a near-unanimity function (see [14] for a related result).

The properties of near-unanimity functions proved in the proof of Proposition 5, (i) in the context of digraphs and the arc-graph construction, also hold in the context of general structures and right adjoints. The proofs carry over naturally.

⁽ⁱ⁾Added in Proof: The decidability of bounded-treewidth duality has recently been proved in L. Barto, M. Kozik, Constraint Satisfaction Problems of Bounded Width, preprint, 2009, 18 pages ms.

References

- [1] M. Barr and C. Wells. *Category Theory for Computing Science*. Les Publications CRM, Montréal, 3rd edition, 1999.
- [2] D. Cohen and P. Jeavons. The complexity of constraint languages. In F. Rossi, P. van Beek, and T. Walsh, editors, *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*, chapter 8. Elsevier, 2006.
- [3] T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, 1999.
- [4] J. Foniok. *Homomorphisms and Structural Properties of Relational Systems*. PhD thesis, Charles University, Prague, 2007.
- [5] P. Hell and J. Nešetřil. On the complexity of H -coloring. *J. Combin. Theory Ser. B*, 48(1):92–110, 1990.
- [6] P. Hell and J. Nešetřil. *Graphs and Homomorphisms*, volume 28 of *Oxford Lecture Series in Mathematics and Its Applications*. Oxford University Press, 2004.
- [7] P. Hell, J. Nešetřil, and X. Zhu. Duality and polynomial testing of tree homomorphisms. *Trans. Amer. Math. Soc.*, 348(4):1281–1297, 1996.
- [8] W. Hodges. *A shorter model theory*. Cambridge University Press, 1997.
- [9] P. Komárek. *Good characterisations in the class of oriented graphs*. PhD thesis, Czechoslovak Academy of Sciences, Prague, 1987. In Czech (Dobré charakteristiky ve třídě orientovaných grafů).
- [10] R. E. Ladner. On the structure of polynomial time reducibility. *J. Assoc. Comput. Mach.*, 22(1):155–171, 1975.
- [11] B. Larose, C. Loten, and C. Tardif. A characterisation of first-order constraint satisfaction problems. In *Proceedings of the 21st IEEE Symposium on Logic in Computer Science (LICS'06)*, pages 201–210. IEEE Computer Society, 2006.
- [12] C. Loten and C. Tardif. Majority functions on structures with finite duality. *European J. Combin.*, 29(4):979–986, 2008.
- [13] S. Mac Lane. *Categories for the Working Mathematician*, volume 5 of *Graduate texts in mathematics*. Springer-Verlag, New York Berlin Heidelberg, 2nd edition, 1998.
- [14] M. Maróti. The existence of a near-unanimity term in a finite algebra is decidable. Manuscript, 2005.
- [15] J. Nešetřil and C. Tardif. Duality theorems for finite structures (characterising gaps and good characterisations). *J. Combin. Theory Ser. B*, 80(1):80–97, 2000.
- [16] J. Nešetřil and C. Tardif. Short answers to exponentially long questions: Extremal aspects of homomorphism duality. *SIAM J. Discrete Math.*, 19(4):914–920, 2005.
- [17] A. Pultr. The right adjoints into the categories of relational systems. In *Reports of the Midwest Category Seminar, IV*, volume 137 of *Lecture Notes in Mathematics*, pages 100–113, Berlin, 1970. Springer.

