Diameters of duals are linear

Jaroslav Nešetřil[*]

Department of Applied Mathematics
and
Institute for Theoretical Computer Science (ITI)
Charles University
Malostranské nám. 25, 11800
Praha 1, Czech Republic
nesetril@kam.ms.mff.cuni.cz

Ida Švejdarová

Department of Applied Mathematics
Charles University
Malostranské nám. 25, 11800
Praha 1, Czech Republic
and
Department of Mathematics
University of Illinois at Urbana-Champaign
1409 W. Green St., Urbana, IL 61801, USA
isvejda2@math.uiuc.edu

Abstract

We prove that dual graphs and relational structures are connected. Moreover, these exponential structures have a linear diameter, which we determine up to a constant.

1 Introduction

How do local properties of graphs influence their global properties? The local-global phenomena were studied extensively and, in general, this is an area of negative results. See the seminal work of Erdős on high chromatic sparse graphs [1] (extended in this setting by [8]). However there are positive aspects of this local-global paradigm. For example for proper minor closed classes we can characterize optimal instances (see [4]) and for oriented graphs (and more generally for relational structures) one obtains a rich spectrum of

[*]Partially supported by the Projects LN00A056 and 1M0021620808 of the Czech Ministry of Education

global properties which are defined locally. The present paper is devoted to one such area — *homomorphism dualities.*

Our simplest model are oriented graphs. Recall that for oriented graphs $G = (V, E)$ and $G' = (V', E')$, a homomorphism $f : G \to G'$ is any mapping $f : V \to V'$ satisfying $(x, y) \in E \Rightarrow (f(x), f(y)) \in E'$. (See [2] for an introduction to graphs and their homomorphisms.) Let $G \to G'$ denote the existence of a homomorphism. A *homomorphism duality* (see [6], [2]) is any statement of the following type:

$$\text{for every graph } G \text{ holds: } F \not\to G \text{ iff } G \to H \tag{1}$$

(thus G is H-colorable iff G is F-free). The pair (F, H) is called a *dual pair* and H is the dual of F. This will be denoted by $H = D_F$. (The dual is uniquely determined up to homomorphism equivalence.) The following is a consequence of the main result of [6]:

Theorem 1. *The dual D_F exists iff F is homomorphically equivalent to an oriented tree.*

The original construction of duals was an indirect one, however recently [5] introduced a explicit and easy construction of dual D_F. Besides having the useful property (1), this construction (of size $2^{n \log n}$ for a tree with n vertices) is an interesting combinatorial structure in itself. However, due to its exponential size, not much is known about its properties. The construction is reviewed and analyzed in Section 2 where we prove the following:

Theorem 2. *After removing isolated vertices, D_F is a connected graph of diameter at most $|V(F)| + 3$ for every tree F.*

Although we can prove indirectly that the core of the dual is always a connected graph (see Theorem 5), for Theorem 2 we need a careful analysis of the explicit construction of D_F. For a fixed tree F, the vertices of D_F are (neighborly) mappings $V(F) \to V(F)$ with arcs defined by means of "switching". This result should be compared with a connectivity and diameter result for trees and their rotations (see [9]). Proof of Theorem 2 is given in Section 2.

It is important that the duality theorem holds for all finite relational structures. Let $\Delta = (\delta_i; i \in I)$ be a finite sequence of positive integers. A *relational structure S of type* Δ (shortly Δ-*structure*) is a pair $(X, (R_i; i \in I))$ when $R_i \subseteq X^{\delta_i}$ for all $i \in I$. Its base set X is sometimes denoted by $\underline{S}$.

We use $R_i(S)$ instead of R_i when necessary and call its elements *edges*. A homomorphism $S \to S'$ of Δ-structures is defined as mapping which preserves the relations R_i for all $i \in I$. We have

Theorem 3. *Δ-structure admits a dual iff it is homomorphically equivalent to a Δ-tree.*

(See [6] and Section 3 for definition of Δ-tree). We prove the connectivity even in this case:

Theorem 4. *For any Δ-tree F, its dual D_F is connected after removing isolated vertices.*

By isolated vertices we mean vertices which do not belong to any edge or those that belong only to edges in $R_i(A)$ with $\delta_i = 1$. This result will be proved in Section 3. Section 4 contains some remarks and open problems.

2 Oriented graphs

Let T be an arbitrary oriented tree. Although we can construct many dual graphs (graphs D_T such that $T \not\to G \Leftrightarrow G \to D_T$ holds for every G), any two duals D and D' are homomorphically equivalent, meaning that we have $D \to D'$ and $D' \to D$. Thus, up to an isomorphism, only one of the duals is a core (it has no proper retracts). This is why we often speak about *the* dual. In this section, D_T will denote a dual obtained by construction described in Definition 1, whereas $\mathrm{Core}(D_T)$ will be the dual that is a core. As a warm-up we prove that $\mathrm{Core}(D_T)$ is a connected graph.

Theorem 5. *$\mathrm{Core}(D_T)$ is a connected graph.*

Proof. For contradiction, suppose that there exist two graphs, D_1 and D_2, such that $\mathrm{Core}(D_T) = D_1 + D_2$ (there is no edge uv for $u \in V(D_1)$ and $v \in V(D_2)$). Each of the two graphs contains at least one edge, otherwise $\mathrm{Core}(D_T)$ would not be a core. Choose arbitrarily $u_1u_2 \in E(D_1)$ and $v_1v_2 \in E(D_2)$ and pick some odd k such that $k > |V(T)|$. Next, build a new graph D' from $\mathrm{Core}(D_T)$ by inserting new vertices $w_1, \ldots, w_k$ and edges w_jw_{j-1} and w_jw_{j+1} for all even j as well as u_1w_1 and v_1w_k. This D', contrary to $\mathrm{Core}(D_T)$, contains a path with alternating directions of edges with endpoints in D_1 and D_2. Clearly $D_1 \not\to D_2$ and $D_2 \not\to D_1$ (as $D_1 + D_2$ is a core). It follows that $D' \not\to D_1 + D_2$. On the other hand $T \not\to D'$ as if $\phi : T \to D'$ is a homomorphism, then $\phi[V(T)] \cap V(D_i) = \emptyset$ for some $i = 1, 2$, because T is connected and the length of the path which connects D_1 and D_2 is greater than $|V(T)|$. Without loss of generality $i = 2$. The subgraph induced by vertices $\phi[V(T)]$ consists of some vertices of D_1 and some vertices that belong to the path between D_1 and D_2. Formally, $\phi[V(T)] \subseteq (V(D_1) \cup \{w_1, \dots, w_k\})$. However, the subgraph induced by vertices $V(D_1) \cup \{w_1, \dots, w_k\}$ is homomorphically equivalent to D_1: consider homomorphism ψ such that $\psi \upharpoonright V(D_1)$ is the identity, $\psi(w_j) = u_1$ for j even and $\psi(w_j) = u_2$ for j odd. Then $\psi\phi$ is a homomorphism mapping T to D_1, which is a contradiction with $T \not\to D_1 + D_2$. Thus we indeed have $T \not\to D'$ which together with $D' \not\to D_1 + D_2$ contradicts the assumption that $D_1 + D_2$ is a dual of T. $\square$

As a corollary of Theorem 8 we obtain that Core(D_T) has diameter at most $n+3$ for T with n vertices. Let us remark that [3] contains example of trees T for which Core(D_T) has diameter at least $\lfloor \frac{n-1}{2} \rfloor$.

In [5], Nešetřil and Tardif introduced the following explicit construction of D_T:

Definition 1. D_T is the graph defined the following way: $V(D_T) = \{f : V(T) \to V(T);$ for all $u \in V(T)$ we have $(u, f(u)) \in E(T)$ or $(f(u), u) \in E(T)\}$,

$E(D_T) = \{(f, g);$ for all $(u, v) \in E(T)$ we have $f(u) \neq v$ or $g(v) \neq u\}$.

Theorem 6. *D_T defined above is a dual of T.*

We have shown (Theorem 5) that Core(D_T) is a connected graph, i.e. for every $u, v \in V(\text{Core}(D_T))$ there exists an oriented path starting with u and finishing with v. But the above proof of Theorem 5 does not construct such a path and does not provide an information about its length. In particular, we would like to estimate the diameter of D_T. We will prove a stronger statement: not only the core of D_T is connected, but D_T itself is connected after removing isolated vertices. Moreover, its diameter is linear in the number of vertices of T, which is perhaps surprising considering that the number of vertices of D_T can be exponential in $|V(T)|$ (see [7]).

To prove this, we first characterize the isolated vertices of D_T.

Definition 2. Vertex $u \in V(T)$ is a source if its indegree is zero. It is a problematic source for $f \in V(D_T)$ if for all its neighbors w we have $f(w) = u$.

Similarly, u is a sink if its outdegree is zero and it is a problematic sink for $f \in V(D_T)$ if $f(w) = u$ for all vertices w adjacent to u.

The proof of the next lemma follows directly from Definition 1.

Lemma 7 (Characterization of isolated vertices of dual)**.** *Outdegree of f in D_T is zero if and only if there exists a problematic sink for f in T. Indegree of f in D_T is zero if and only if there exists a problematic source for f in T.*

Theorem 8. *Let T be oriented tree with n vertices and D_T its dual constructed in Definition 1. Let f and g be two vertices of D_T which are not isolated. Then there exists an oriented path between f and g of length at most $n + 3$.*

Proof. First, let f and g be vertices with outdegree greater than zero. Let Z be the set of sources of T and let S be its sinks. Define a mapping $f^* : V(T) \to V(T)$ as follows: if $w \in V(T)$ is a source, then we put $f^*(w) = f(w)$, and if is not a source, then we pick (an arbitrary) u such that uw is an edge and we put $f^*(w) = u$. The mapping f^* has the following property: for every edge (f, h) there exists an edge (f^*, h). We will define g^* analogously, only now we require also that g^* coincide with f^* on the set $V(T) \setminus Z$. The set of the sources for which f^* and g^* differ will be denoted by Y.

Claim 9. *There exists a path of length at most $2|Y|$ which connects f^* and g^*.*

Proof. The proof of Claim proceeds by induction on $|Y|$.

If $|Y| = 0$, then we have $f^* = g^*$ and the statement is true.

Let $|Y| = 1$. We have $Y = \{y\}$, $f^*(y) = z_1$ and $g^*(y) = z_2 \neq z_1$. We want to find some h such that (f^*, h) and (g^*, h) are edges. Choose h_f such that (f, h_f) is an edge and define $h(x) = h_f(x)$ for all $x \in S \setminus \{z_1, z_2\}$. For every u that is not a sink, pick an arbitrary w such that uw is an edge and put $h(u) = w$. Now we are left with at most two vertices for which h is not defined yet, namely z_1 and z_2 (if one or both of them are sinks - otherwise h is already defined). If z_1 is a sink, then (since the outdegree of f^* is not zero, see Lemma 7), it has a neighbor $u \neq y$ such that $f^*(u) \neq z_1$ (thus also $g^*(u) \neq z_1$) and put $h(z_1) = u$. Similarly, if z_2 is a sink, then it has a neighbor $v \neq y$ such that $g^*(v) \neq z_2$ and we may define $h(z_2) = v$.

Finally, let $|Y| = m > 1$. Let V^* be the set of all vertices f with nonzero outdegree such that the corresponding mapping goes against the directions of edges whenever possible, i.e. we have $f(w) = z$ for an edge wz only if w is a source. Suppose that for each pair $f^*, g^* \in V^*$ of mappings that differ on $m-1$ sources and coincide on $V(T) \setminus Z$ there exists a path of length $2(m-1)$ connecting f^* and g^* . Let $f^*, g^* \in V^*$ be mappings that differ on the set $Y \subseteq Z$ such that $|Y| = m$ (and again, they coincide on the rest of vertices). Choose an arbitrary $u \in Y$.

(i) If we can redefine g^* on u according to f^* and no problematic sink arise in this process, then we do it. This way we get a mapping $\overline{g^*}$ such that $\overline{g^*}(u) = f^*(u)$, $\overline{g^*}(v) = g^*(v)$ for $v \in V(D_T) \setminus \{u\}$.

(ii) Suppose (i) is not possible: by redefining g^* on the vertex u we would get a problematic sink v_1. But this sink is not problematic for f^*. So v_1 has a neighbor $w_1 \neq u$ which is a source and for which $g^*(w_1) = v_1$ but $f^*(w_1) = v_2 \neq v_1$ and w_1 also belongs to Y. If we can redefine g^* on w_1 without obtaining problematic sinks, then we do it. If not, find $w_2 \in Y$ in a similar manner. The sequence $\{w_i, i \in \mathbb{N}\}$ is not infinite, since vertices never repeat in it. Thus there exists at least one vertex $y \in Y$ (the last element of sequence $\{w_i, i \in \mathbb{N}\}$) such that we can redefine $g^*(y)$ according to $f^*(y)$ without creating problematic sinks; again, we call this redefined mapping $\overline{g^*}$.

Now we have $\overline{g^*}$ which differs from f^* on the set of $m-1$ sources. By induction hypothesis there exists a path at most $2(m-1)$ long which connects f^* and $\overline{g^*}$. Meanwhile, $\overline{g^*}$ differs from g^* only on one vertex in Z and thus these two vertices are connected by a path of length at most 2. Consequently, there exists a path of length at most $2m = 2|Y|$ between f^* and g^*. This finishes the proof of the Claim. $\square$

Now we need to compare $|Y|$ with the number of vertices of T. Let $Y = \{y_1, \ldots, y_m\}$. Let R be the subgraph of T induced by the set $\{y_1, f^*(y_1), g^*(y_1), y_2, f^*(y_2), g^*(y_2), \ldots, y_m, f^*(y_m), g^*(y_m)\}$. This is either a tree or a forest, in any case we have $|V(R)| \geq |E(R)| + 1$. Each of the vertices $y_1, \ldots, y_m$ has degree at least 2 (there are edges from y_i to $f^*(y_i)$ and $g^*(y_i)$) and no two of these vertices share an edge (all of them are sources), so $|E(R)| \geq 2m$. We get:

$$n = |V(T)| \geq |V(R)| \geq |E(R)| + 1 \geq 2m + 1$$

So $m \leq (n-1)/2$; and thus there is a path connecting f^* and g^* of length at most $2m \leq n-1$. The distance of f and f^* is at most 2, as well as the distance of g and g^*. Thus any two vertices f and g with outdegree greater than zero are at most $n-1+4=n+3$ apart.

If f' is a sink (but is not isolated), then we consider its neighbor f and find the corresponding f^*. The distance of f' and f^* is 1, so in this case the distance of f' from other vertices is even less than $n+3$. Consequently, any two non-isolated vertices of D_T are at most $n+3$ apart.

□

This bound can be improved to $n+2$. We do not include the proof because it is merely a tedious case analysis based on the ideas presented in the above proof.

3 Relational structures

Let A be a relational structure of type $\Delta = (\delta_i; i \in I)$. Its *incidence graph* $\mathrm{Inc}(A)$ is the bipartite graph with parts $\underline{A}$ and $\mathrm{Block}(A) = \{(i,(a_1,\dots,a_{\delta_i}));\ i \in I, (a_1,\dots,a_{\delta_i}) \in R_i(A)\}$. The edges are all pairs $[a,(i,(a_1,\dots,a_{\delta_i}))]$ such that $a \in (a_1,\dots,a_{\delta_i})$ (i.e. there exists an index k such that $a=a_k$). A is called a Δ-*tree* when $\mathrm{Inc}(A)$ is a tree.

As we mentioned in Section 1, A admits a dual iff it is a Δ-tree. A simple construction of duals for relational structures similar to the one in Definition 1 appeared in [7]:

Definition 3. Let A be a Δ-tree. Let D_A be relational structure with the base set $\underline{D_A} = \{f : \underline{A} \to \mathrm{Block}(A); [a,f(a)] \in E(\mathrm{Inc}(A)) \text{ for all } a \in \underline{A}\}$. The δ_i-tuple $(f_1,\dots,f_{\delta_i})$ belongs to $R_i(D_A)$ if and only if for every $(x_1,\dots,x_{\delta_i}) \in R_i(A)$ there exists $j \in \{1,\dots,\delta_i\}$ such that $f_j(x_j) \neq (i,(x_1,\dots,x_{\delta_i}))$.

Theorem 10. *[7] Let A be a Δ-tree. The structure D_A defined above is a dual of A.*

Analogously as in Theorem 5 we can prove easily that the core of D_A is a connected Δ-structure. But again we shall prove that even the structure D_A is connected after deleting all isolated vertices.

Throughout this section, we will use the following notation. Let $\widetilde{A}$ be Δ-tree, $\widetilde{D_A}$ its dual and let $\underline{\widetilde{D_A}} = \{\tilde{f}_1,\dots,\tilde{f}_j\}$. A will denote the Δ-tree obtained by inserting a new edge $b \in R_{i'}(A)$ which has a common vertex b_k with some $a \in R_i(\widetilde{A})$ (if b_k belongs to more edges, select one of them and call it a). Let D_A be the dual of A. It is not hard to see that D_A has vertices

$\{f_1, \dots, f_j, f'_1, \dots, f'_j\}$, such that the mappings f_t and f'_t for $t = 1, \dots, j$ are both derived from $\tilde{f}_t$ and they differ only on the vertex b_k. More precisely, f_t and f'_t coincide with $\tilde{f}_t$ on vertices of $\widetilde{A}$, they are defined in the only possible way on vertices of b different from b_k (that is, $f_t(u) = f'_t(u) = (i', b)$ for $u \in b$, $u \neq b_k$), and $f_t(b_k) = (i, a)$, while $f'_t(b_k) = (i', b)$. Notice that the elements $f'_1, \dots, f'_j$ are not necessarily distinct.

Also, c_l will denote the l-th vertex of the edge c and $R(A)$ will be the set of all edges of the Δ-system A, i.e. $R(A) = \bigcup_{j \in I} R_j(A)$.

The next lemma reveals a close relationship between D_A and $\widetilde{D_A}$: if we delete the vertices $f'_1, \dots, f'_j$ in D_A and all edges containing them, we get exactly a copy of $\widetilde{D_A}$. Proof follows immediately from Definition 3.

Lemma 11. *Let $i \in I$. $(\tilde{f}_1, \dots, \tilde{f}_{\delta_i})$ is an edge of $\widetilde{D_A}$ if and only if $(f_1, \dots, f_{\delta_i})$ is an edge of D_A.*

Sinks and sources together with the sets of their neighbors played a crucial role in characterizing the isolated vertices of duals of graphs. The classes of the equivalence defined below play similar role in characterizing the isolated vertices of duals of relational structures.

Definition 4. For every $i \in I$ and $k \in \{1, \dots, \delta_i\}$ we will define equivalence $\approx_{(i,k)}$ on $R_i(A)$: $a \approx_{(i,k)} b$ if there exists an integer $m \geq 1$ and a sequence of edges $a = c^1, c^2m, \dots, c^m = b$ with $c^1, \dots, c^m \in R_i(A)$ which satisfies the following: for every $j = 1, \dots, m-1$ there is an index $l_j \neq k$ such that the edges c^j and c^{j+1} share a vertex v, and v occupies the l_j-th position in both edges (that is, $c^j_{l_j} = c^{j+1}_{l_j}$).

The relation $\approx_{(i,k)}$ is clearly an equivalence. $[x]_{\approx_{(i,k)}}$ will denote the class of the equivalence $\approx_{(i,k)}$ containing the edge x.

Thus if two edges of a tree A share a vertex and they belong to the same equivalence class, then the shared vertex occupies the same position in both edges and moreover this position is different from k.

Theorem 12 (Characterization of isolated vertices of duals)**.** *Let f be a vertex of D_A, $i \in I$ and $k \in \{1, \dots, \delta_i\}$. $f \neq f_k$ holds for every $(f_1, \dots, f_{\delta_i}) \in R_i(D_A)$ if and only if there exists an edge $x \in R_i(A)$ satisfying the following two conditions:*

(1) If two edges y and z share a vertex v and y is an element of $[x]_{\approx_{(i,k)}}$ while z is not, then v is the k-th coordinate of y.)

(2) $f(a_k) = (i, (a_1, \ldots, a_{\delta_i}))$ *for every* $(a_1, \ldots, a_{\delta_i}) \in [x]_{\approx_{(i,k)}}$.

We will call the class $[x]_{\approx_{(i,k)}}$ containing such edge x *problematic class* for f and k.

Proof. First suppose that there exists an edge $x \in R_i(A)$ satisfying both conditions. Suppose for contradiction that $f = f_k$ for some $(f_1, \ldots, f_{\delta_i}) \in R_i(D_A)$. By definition of edges of D_A there exists some vertex x_j in the edge x such that $f_j(x_j) = (i', y) \neq (i, x)$ for some $i' \in I$ and $y \in R_{i'}(A)$. y shares a vertex with x and if $y \notin [x]_{\approx_{(i,k)}}$, then by (1) this vertex occupies k-th position in x. Thus $f_k(x_k) = (i', y) \neq (i, x)$, which is contradicts the condition (2). So we have $y \in [x]_{\approx_{(i,k)}}$. We will denote $x^1 = x$, $x^2 = y$ and continue constructing the sequence $\{x^m; m \in \mathbb{N}\}$. Suppose $x^m \in [x]_{\approx_{(i,k)}}$ is already known for some m. Find $j \in \{1, \ldots, \delta_i\}$ such that $f_j(x^m_j) = (i', z) \neq (i, x^m)$ (such j exists for all $x^m \in R_i(A)$ because $(f_1, \ldots, f_{\delta_i})$ is an element of $R_i(D_A)$) and denote $x^{m+1} = z$. Since $x^m \in [x]_{\approx_{(i,k)}}$ and the two conditions hold, we again have $i = i'$ and $z \in [x]_{\approx_{(i,k)}}$. This way we obtain an infinite sequence of edges from $[x]_{\approx_{(i,k)}}$ such that every two consecutive edges are different and have a vertex in common. Can some elements repeat in this sequence? Suppose that $x^m = x^{m+l}$ for some $m, l \in \mathbb{N}$. If this is true for more than one pair m,l, choose the pair with the smallest value of l. Two subsequent edges are different, so $l \geq 2$. For every t, $x^t_{j_t}$ will denote the common vertex of x^t and x^{t+1}. If $l \geq 3$, then $(i, x^m), x^m_{j_m}, (i, x^{m+1}), x^{m+1}_{j_{m+1}}, \ldots, (i, x^{m+l-1}), x^{m+l-1}_{j_{m+l-1}}, (i, x^{m+l}) = (i, x^m)$ is a sequence of vertices of Inc(A) joined by edges such that no two of them are the same except for the first and the last. In other words, it is a cycle, which contradicts the definition of Δ-tree. The only remaining case is $l = 2$. Since x^{m+1} follows after x^m in our sequence, we have $f_{j_m}(x^m_{j_m}) = (i, x^{m+1})$. Also $x^{m+2} = x^m$ follows after x^{m+1}, so $f_{j_{m+1}}(x^{m+1}_{j_{m+1}}) = (i, x^m)$. But since both x^m and x^{m+1} are elements of $[x]_{\approx_{(i,k)}}$ and thus $j_m = j_{m+1}$, we get $(i, x^{m+1}) = f_{j_m}(x^m_{j_m}) = f_{j_{m+1}}(x^{m+1}_{j_{m+1}}) = (i, x^m)$, which contradicts our assumption. So there is no such l and the elements of our sequence never repeat. We obtained an infinite branch in a finite Δ-system, which is a contradiction.

To prove the other implication, suppose that the right side does not hold; for every equivalence class $[x]_{\approx_{(i,k)}}$ either

- exists $a \in [x]_{\approx_{(i,k)}}$ such that the common vertex of a and the rest of A is j-th in a for $j \neq k$

  or

- exists $a \in [x]_{\approx_{(i,k)}}$ such that $f(a_k) \neq (i, a)$.

We will define mappings $f_1, \ldots, f_{k-1}, f_{k+1}, \ldots, f_{\delta_i}$ such that $(f_1, \ldots, f_{\delta_i})$ is an edge of D_A for $f_k = f$. First we define $f_l(c_l)$ with $l = 1, \ldots, \delta_i$ for all edges $c \in R_i(A)$: for every class $[x]_{\approx_{(i,k)}}$ of the equivalence $\approx_{(i,k)}$ we construct set S of edges c such that $f_l(c_l)$ is already known for every l. We use the following algorithm:

(1) If the first conditions holds (there exists $a \in [x]_{\approx_{(i,k)}}$ such that for some $j \neq k$ a_j belongs to some edge $y \in R_{i'}(A)$ and $y \notin [x]_{\approx_{(i,k)}}$), we put $f_j(a_j) = (i', y)$. We can always do this. It would be impossible only if $f_j(a_j)$ was already defined while using the algorithm for some edge z in a different class of equivalence $\approx_{(i,k)}$. But this algorithm defines $f_t(a_s)$ only for $t = s$. Therefore if f_j is defined for some vertex $z_l = a_j$, then $l = j$ and $z \in R_i(A)$ and thus z is in the same equivalence class as a. Next put $f_l(a_l) = (i, a)$ for all $l = 1, \ldots, \delta_i$ such that $f_l(a_l)$ is not defined yet and let $S = \{a\}$.

If the second condition holds, the situation is even simpler: we put $f_l(a_l) = (i, a)$ for $l \neq k$ and again $S = \{a\}$.

(2) If there is some edge $b \in [x]_{\approx_{(i,k)}} \setminus S$ which has a common vertex with some edge in S, then for all l such that $f_l(b_l)$ is undefined put $f_l(b_l) = (i, b)$ and add b to S. Repeat this step if possible.

We claim that after this algorithm has finished, we have $[x]_{\approx_{(i,k)}} = S$. Suppose contrary: there is some $y \in [x]_{\approx_{(i,k)}}$ which does not belong to S. Take the shortest sequence $a = c^0, c^1, \ldots, c^m = y$ of edges from $[x]_{\approx_{(i,k)}}$ such that each of them shares a vertex with its successor. Let $p \in \{1, \ldots, m\}$ be the smallest index such that $c^p \notin S$. c^p shares a vertex $c^p_{j_{p-1}}$ with the edge $c^{p-1} \in S$, so we can define $f_l(c^p_l) = (i, c^p)$ for $l \neq j_{p-1}$ and add c^p to S. Thus indeed $[x]_{\approx_{(i,k)}} = S$.

Also for every $y \in [x]_{\approx_{(i,k)}}$ there exists j such that $f_j(y_j) \neq (i, y)$. Again, consider y for which this is not true and the shortest sequence $a = c^0$, $c^1, \ldots,$ $c^m = y$ such that c^t and c^{t+1} share vertex $c^t_{j_t}$. For $t = 0, \ldots, m-1$ we have $j_t \neq j_{t+1}$ (the sequence is shortest possible) and thus $f_t(c^t_{j_t}) = f_t(c^{t+1}_{j_t}) = (i, c^t)$ holds for every t, particularly for $t + 1 = m$, which is a contradiction.

After we use the algorithm for all classes of equivalence $\approx_{(i,k)}$, we define the mappings arbitrarily on the rest of the vertices. The edges that are not in $R_i(A)$ will not prevent the existence of the edge $(f_1, \ldots, f_{\delta_i})$ and all edges in $R_i(A)$ satisfy the condition for the existence of such edge. Thus we have $(f_1, \ldots, f_{\delta_i}) \in R_i(D_A)$. $\square$

Corollary 13. *Let g be a vertex of D_A. If g is the k-th coordinate of some edge $(g_1, \ldots, g_{\delta_i}) \in R_i(D_A)$ but $\tilde{g}$ is isolated in $\widetilde{D_A}$, then there exists single problematic class $[x]_{\approx_{(i,k)}}$ for $\tilde{g}$ and k in $\widetilde{D_A}$ (i.e. a class satisfying the conditions (1) and (2) from Theorem 12), and this class is not problematic any more (for g and k) after inserting the edge b.*

D_A can contain vertices that do not belong to any edge (isolated vertices) and also vertices that are only in unary relations. To simplify notation, we will extend the definition of isolated vertex so that it includes also the latter kind of vertices: from now on, a vertex u will be *isolated* if and only if for every $i \in I$ with $\delta_i > 1$ and for every $a \in R_i(A)$ we have $u \notin a$. We will prove that after removing such isolated vertices, D_A is a connected Δ-system.

The proof of connectedness of duals for graphs (Theorem 8) can be generalized for relational structures. However, we chose a different approach, one that shows how the dual changes when we modify the original tree, and thus provides additional insight into its structure. This proof gives an alternative proof of Theorem 8 (without the bound on diameter).

We will need the following lemma. Recall that if we insert a new edge b into a Δ-tree $\widetilde{A}$ with $\widetilde{D_A} = \{\tilde{f}_1, \ldots, \tilde{f}_j\}$, we obtain a new Δ-tree A with $D_A = \{f_1, \ldots, f_j, f'_1, \ldots, f'_j\}$.

Lemma 14. *If B is a component of D_A which contains more than one vertex, then some of the vertices $f_1, \ldots, f_j$ belong to B (that is, B does not contain only $f'_1, \ldots, f'_j$).*

Proof. Suppose that we inserted a new edge $b \in R_{i'}(A)$ and it has a common vertex $a_{j_1} = b_{j_2}$ with some $a \in R_i(A)$. Let $\delta_l > 1$ and let $(f'_1, \ldots, f'_{\delta_l})$ be an edge that does not contain any of the vertices $f_1, \ldots, f_j$. One can easily see that this can happen only if $l \neq i'$. If $a \notin R_l(A)$, then we can define $g(b_{j_2}) = (i, a)$, $g(u) = f'_1(u)$ for $u \neq b_{j_2}$. $(g, f'_2, \ldots, f'_{\delta_l})$ is an edge: if for every $c \in R_l(A)$ there exists index j such that $f'_j(c_j) \neq (l, c)$, then such j exists for every $c \in R_l(A)$ for this new set of mappings as well. If $a \in R_l(A)$, then we will choose $m \neq j_1$ (we can do this because $\delta_l > 1$) and define $g(b_{j_2}) = (i, a)$, $g(u) = f'_m(u)$ for $u \neq b_{j_2}$. Then $(f'_1, \ldots, f'_{m-1}, g, f'_{m+1}, \ldots, f'_{\delta_l})$ is also an edge in $R_l(D_A)$. Moreover, in both cases $g \in \{f_1, \ldots, f_j\}$. $\square$

So called "zigzag paths", i.e. paths with alternating directions of edges, play an important role in the proof of Theorem 8. The equivalence classes defined below are analogues of zigzag paths for relational structures.

Definition 5. For every i we will define an equivalence $\sim_i$ on $R_i(A)$: $a \sim_i b$ if there exists a sequence $a = c^1, c^2, \ldots, c^m = b$ of edges from $R_i(A)$ such that for every j there exists index $l_j \in \{1, \ldots, \delta_i\}$ such that $c^j_{l_j} = c^{j+1}_{l_j}$.

Contrary to the definition of $\approx_{(i,k)}$, now the index l_j can be arbitrary. For every x and k we have $[x]_{\approx_{(i,k)}} \subseteq [x]_{\sim_i}$.

Theorem 15. *If the Δ-system obtained from $\widetilde{D_A}$ by removing isolated vertices is connected, then D_A is also connected after removing isolated vertices.*

Proof. By deleting the vertices $f'_1, \ldots, f'_j$ from D_A (and all edges incident with them) we get a copy of $\widetilde{D_A}$ (Lemma 11). By assumption, this copy has at most one nontrivial connected component (i.e. connected component with more than one vertex), say C. For contradiction suppose that there exists some nontrivial component C' in D_A such that $C' \cap C = \emptyset$. Necessarily, some of the vertices $f'_1, \ldots, f'_j$ belong to C'. But C' also contains some $g \in \{f_1, \ldots, f_j\}$ (Lemma 14) and since $g \notin C$, $\tilde{g}$ was isolated in $\widetilde{D_A}$. Thus, to prove the theorem, it suffices to find a path in $\mathrm{Inc}(D_A)$ beginning with g and ending in C for every $g \in \{f_1, \ldots, f_j\}$ such that $\tilde{g}$ is isolated in $\widetilde{D_A}$ but $g = g_k$ for some k, some $i \in I$ with $\delta_i > 1$ and some $(g_1, \ldots, g_{\delta_i}) \in R_i(D_A)$. This will contradict the existence of C'.

Let g be such vertex. By Corollary 13 there was only one problematic class $[x]_{\approx_{(i,k)}}$ for $\tilde{g}$ in $\tilde{A}$ and after adding the edge b this class is not problematic any more. This could happen only if $a \in [x]_{\approx_{(i,k)}}$ and thus $a \in R_i(A)$. In this situation we distinguish two cases.

(1) $[a]_{\sim_i} \neq R(\widetilde{A})$. Then there is an edge $y \in R_{i'}(\widetilde{A})$ (for i' not necessarily distinct from i) such that $y \notin [a]_{\sim_i}$ and y has a common vertex with some edge in $[a]_{\sim_i}$. Let $y = c^1, c^2, \ldots, c^m = a$ be the shortest sequence of edges such that c^t and c^{t+1} have a common vertex for $t = 1, \ldots, m-1$. Let G^1 denote the edge $(g_1, \ldots, g_{\delta_i})$ and let c^2 share its s-th vertex with $y = c^1$. We will define g^*_s: $g^*_s(c^2_s) = (i', c^1)$, $g^*_s(u) = g_s(u)$ for $u \in \underline{A} \setminus \{c^2_s\}$. Then $G^2 = (g_1, \ldots, g_{s-1}, g^*_s, g_{s+1}, \ldots, g_{\delta_i})$ is also an edge in $R_i(D_A)$ because:

- if $y \notin R_i(A)$, then the newly defined mappings coincide with the original ones on $R_i(A)$ (except for c^2) and nothing changed about the fact that $(\forall x \in R_i(A))(\exists l \in \{1, \dots, \delta_i\})(G_l^2(x_l) \neq (i, x))$ and
- if $y \in R_i(A)$, then the common vertex of y and c^2 is s-th in c^2 and s'-th in y, $s \neq s'$ and again, the condition still holds.

Next we will define edges $G^3, \dots, G^m$: suppose G^{t-1} is already known and the edges c^{t-1} and c^t share their s-th vertex, then $G_l^t = G_l^{t-1}$ for $l \neq s$, $G_s^t(c_s^t) = (i, c^{t-1})$ and $G_s^t(v) = G_s^{t-1}(v)$ for $v \in \underline{A} \setminus \{c_s^t\}$. There can only be one edge in $R_i(A)$ that violates the condition for G^t being an edge of D_A, and that is the edge c^{t-1}. But we have $G_{s'}^{t-1}(c_{s'}^{t-1}) = c^{t-2}$ for the vertex $c_{s'}^{t-1}$ shared by c^{t-1} and c^{t-2}. And because $s \neq s'$ (we selected the shortest path from a to y), we have also $G_{s'}^t(c_{s'}^{t-1}) = (i, c^{t-2})$. Thus this edge is also harmless, and $G^t \in R_i(D_A)$. Last, let us define $G_s^{m+1}(a_s) = (i, a)$ (where a_s is the common vertex of a and the newly inserted edge b), $G_s^{m+1} = G_s^m$ for the rest of the vertices and $G_{s'}^{m+1} = G_{s'}^m$ for $s \neq s'$. G^{m+1} is an edge (for reasons similar to those above), so all δ_i-tuples $G^1, \dots, G^{m+1}$ are elements of $R_i(D_A)$. Every two subsequent edges in this sequence differ only in one coordinate and since $\delta_i > 1$, they share at least one vertex. Let $G^{m+1} = (h_1, \dots, h_{\delta_i})$ and let us define $G^{m+2} = (h_1^*, \dots, h_{\delta_i}^*)$ by putting $h_t^*(v) = h_t(v)$ for $v \in \underline{A} \setminus \{a_s\}$ (again, a_s is the common vertex of a and b) and $h_t^*(a_s) = (i, a)$ for $t = 1, \dots, \delta_i$. Obviously $G^{m+2} \in R_i(D_A)$ and G^{m+2} has a common vertex $h_s^* = h_s$ with G^{m+1}. Mappings h_t^* coincide with $\tilde{h}_t$ on all vertices of A except for the new vertices $b \setminus \{a_s\}$ and thus, by Lemma 11, $(\tilde{h}_1, \dots, \tilde{h}_{\delta_i})$ is an edge of $\widetilde{D_A}$. So we found a sequence that begins with the original edge G^1 and ends with some edge G^{m+2} belonging to the component C.

(2) $[a]_{\sim_i} = R(\widetilde{A})$. In this case if $b \in R_i(A)$ and $a_s = b_s$ holds for some s (that is, the common vertex of a and b occupies the same position in both edges), then $R_i(D_A) = \emptyset$. This is because A is homomorphically equivalent to the Δ-system B with $\underline{B} = \{b_1, \dots, b_{\delta_i}\}$ and $R_i(B) = R(B) = \{b\}$ and clearly $R_i(D_B) = \emptyset$. If there exists some $i' \neq i$ such that $\delta_{i'} > 1$, then for any $\delta_{i'}$-tuple $f_1, \dots, f_{\delta_i'} \in \underline{D_A}$ we have $(f_1, \dots, f_{\delta_i'}) \in R_{i'}(D_A)$ (since $R_{i'}(A) = \emptyset$, nothing can prevent the existence of such edge), otherwise all vertices of D_A are isolated.

If $b \in R_i(A)$, but $b_s = a_{s'}$ for $s \neq s'$ (the common vertex occupies different positions in the two edges) or $b \notin R_i(A)$, then describing $R_i(D_A)$ is also relatively easy. First, let us label the edges of A recursively according to their distance from b and let $c(x)$ denote the label given to the edge x. This way we get $c(b) = 0$, $c(a) = 1$, $c(y) = 2$ for edges that have a common vertex with a etc. Since A is a Δ-tree, such labeling exists and is unique. Now define sets $H_k \subseteq \underline{D_A}$ for $k = 1, \ldots, \delta_i$: f will belong to H_k if and only if it maps every x_k that is k-th in some edge $x \in R_i(A) \setminus \{b\}$ to the edge with the smallest label of all edges containing x_k. Formally $H_k = \{f; (\forall x \in R_i(A) \setminus \{b\})(\forall d \in R(A))(x_k \in d \Rightarrow c(f(x_k)) \leq c(d))\}$. Plus, we have an additional requirement for f in H_s: $f(b_s) = (i, a)$ (b_s is the vertex shared by a and b). We will prove that $R_i(D_A) = H_1 \times H_2 \times \cdots \times H_{\delta_i}$. The inclusion $\supseteq$ is obvious. Let's prove the other inclusion. For contradiction suppose that $(g_1, \ldots, g_{\delta_i}) \in R_i(D_A)$ but $g_k \notin H_k$ for some k. This means that there is some $x^0 \in R_i(A)$ such that $g_k(x^0_k) = (i, x^0)$ but there exists some $y \in R_i(A)$, $x^0_k \in y$, which is closer to b than x^0. Since $(g_1, \ldots, g_{\delta_i})$ is an edge, there is some $l_0 \neq k$ for which $g_{l_0}(x^0_{l_0}) = (i, x^1)$ for $x^1 \neq x^0$. Since A is a Δ-tree, y is the only edge incident to x^0 such that $c(y) < c(x^0)$, and therefore $c(x^1) > c(x^0)$. Analogously there exists $l_1 \neq l_0$ for which $g_{l_1}(x^1_{l_1}) = (i, x^2)$ for $x^2 \neq x^1$. Again $c(x^2) > c(x^1)$. The sequence $x^1, x^2, \ldots$ can finish only if it reaches b at some point. However, considering that $c(x^1) < c(x^2) < \ldots$, this will never happen. The system A is finite, thus we obtain contradiction. If there exists an $i' \neq i$ with $\delta_{i'} > 1$, then again all vertices of D_A belong to a single nontrivial connected component (if $b \in R_{i'}(A)$, then $(f_1, \ldots, f_{\delta_{i'}})$ is an edge whenever $f_s(b_s) = (i, a)$, and if $b \notin R_{i'}(A)$, then all $\delta_{i'}$-tuples are edges). If there is no such i', then the nontrivial connected component contains exactly the elements of $R_i(A)$, and $H_1 \times H_2 \times \cdots \times H_{\delta_i}$ has only one connected component.

□

Proof of Theorem 4. Any Δ-tree can be built in a finite number of steps from the empty Δ-tree (i.e. a Δ-tree B with $\underline{B} = \emptyset$) by inserting leaves in such a way that the Δ-systems obtained in each step are Δ-trees. Thus we can proceed by induction, with the inductive step being the essence of the previous theorem. □

Distance $d(u, v)$ of vertices u and v in a relational structure is defined as the smallest k for which there exists a sequence $u = u_0, \dots, u_k = v$ such that u_i and u_{i+1} belong to the same edge. A closer look at the proof of Theorem 15 gives a polynomial upper bound on diameter of A.

Lemma 16. *If A is a Δ-tree with n vertices, then the diameter of D_A, after removing isolated vertices, is at most $O(n^2)$.*

Proof. In the proof of Theorem 15 we constructed a sequence $G^1, \dots, G^{m+1}$ such that G^1 contains g and G^{m+1} contains a vertex in C. The distance of g from C is therefore at most $m+1$. This holds for every non-isolated g that is not in C. If f and h are non-isolated vertices of D_A, d denotes distance and $\mathrm{diam}(C)$ is the diameter of C, then

$$d(f,h) \leq d(f,C) + \mathrm{diam}(C) + d(h,C) \leq \mathrm{diam}(C) + 2(m+1) \tag{2}$$

We can strengthen the definition of m and let it be the *shortest* distance of a from some $y \notin [a]_{\sim_i}$. Then $m-1$ is bounded by the number of edges in $[a]_{\sim_i}$, which cannot exceed the number of all non-unary edges of A (i.e. those with $\delta_i > 1$). Hence $m + 1 \leq |\bigcup_{\delta_i > 1} R_i(A)| + 2$. A Δ-tree with n vertices can have at most $n-1$ non-unary edges. Now, let A be such Δ-tree. As in the proof of Theorem 4, we can build A by inserting non-unary edges one by one in at most $n-1$ steps so that in each step the resulting Δ-structure is a Δ-tree, and concluding with adding all unary edges. The presence of unary edges cannot increase the diameter, so we only need to estimate the diameter before performing the last step (of equivalently we may without loss of generality assume that A has no unary edges). Using (2) repeatedly, we get

$$\mathrm{diam}(D_A) \leq 1 + \sum_{k=1}^{n-2} 2(k+2) \leq 1 + 4(n-2) + (n-1)(n-2) = O(n^2). \tag{3}$$

□

4 Concluding remarks

The linearity of diameter suggests the existence of fast algorithms for D_T. The above proof of the connectivity of D_T yields an algorithm which finds a path from f to g in at most $2\Delta(T)n^2$ steps. Is there a linear algorithm?

Knowing that D_T is connected, one might also try to determine its connectivity. Are there always vertices of small degree in D_T? How does the minimal degree in D_T depend on the height of T?

References

[1] P. Erdős, *Graph theory and probability*, Canad. J. Math., 11 (1959), pp. 34–38.

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

[3] J. Nešetřil, *On maximal finite antichains in the homomorphism order of directed graphs*, Discussiones Mathematicae Graph Theory, 23 (2003), pp. 325–332.

[4] J. Nešetřil and P. Ossona de Mendez, *Colouring and homomorphisms of minor-closed classes*, Discrete and Computational Geometry, 25 (2003), pp. 651–664.

[5] J. Nešetřil and C. Tardif, *A dualistic approach to bounding the chromatic number of a graph*, ITI Series, 2001-036.

[6] ——, *Duality theorems for finite structures (characterizing gaps and good characterizations)*, J. Combin. Theory B, 80 (2000), pp. 80–97.

[7] ——, *Short answers to exponentially long questions: Extremal aspects of homomorphism duality*, KAM-DIMATIA Series, (2004).

[8] J. Nešetřil and X. Zhu, *On sparse graphs with given colorings and homomorphisms*, J. Combin. Theory B, 80 (2004), pp. 161–172.

[9] D. Slater, R. Tarjan, and W. Thurston, *Rotation distance, triangulations and hyperbolic geometry*, J. Amer. Math. Soc., 1 (1988), pp. 647–681.

[10] I. Švejdarová, *Coloring of graphs and dual structures*, B.Sc. thesis, Charles University, Prague, 2003.