Duality Theorems for Finite Structures (characterising gaps and good characterisations)

Jaroslav Nešetřil
Charles University

Claude Tardif
University of Regina

October 26, 1998

Abstract

We provide a correspondence between the subjects of duality and density in classes of finite relational structures. Duality purposes to characterise the structures C that do not admit a homomorphism into a given target B by the existence of a homomorphism *from* a structure A into C. Density is the order-theoretic property of containing no covers (or 'gaps'). We show that the covers in the skeleton of a category of finite relational models correspond naturally to certain instances of duality statements, and we characterise these covers.

1 Introduction

The object of the theory of homomorphism duality is to characterise a family $\mathcal{C}$ of obstructions to the existence of a homomorphism into a given structure B. In a large sense, such a class $\mathcal{C}$ always exists; for instance, the class of all the structures not admitting a homomorphism to B has this property. However, it is desirable to seek a more tractable family of obstructions to make this characterisation meaningful. The classical examples of graph theory makes this point clear. A graph is bipartite if and only if it does not contain an odd cycle; hence, the odd cycles are a family of obstructions to the existence of a homomorphism into the complete graph K_2. However, the class of directed graphs provides a much more fertile ground for the theory, and numerous examples of tree dualities and of bounded treewidth dualities are known (see [10], [5]).

When the family $\mathcal{C}$ of obstructions is finite (or algorithmically "well behaving"), then such theorems clearly provide an example of *good characterisations* (in the sense of Edmonds [2]). Any instance of such good characterisation is called a *homomorphism duality.* This concept was introduced by Nešetřil and Pultr [10] and applied to various graph theoretical good characterisations in [9]. The simplest homomorphism dualities are those where the family of obstructions consists from just one structure. In the other words such homomorphism dualities are described by a pair A, B of structures (graphs) as follows:

Definition 1.1 *(Singleton) Homomorphism Duality Scheme*
G admits a homomorphism into B if and only if A does not admit a homomorphism into G

Despite of the fact that singleton homomorphism dualities are scarce for both undirected and directed graphs, for more general structures (such as oriented matroids with suitable version of strong maps) the (singleton) homomorphism duality may capture general theorems such as Farkas Lemma (see [6]). J. Nešetřil and A. Pultr described in [10] all singleton homomorphism dualities for undirected graphs and Komárek [7] and J. Nešetřil, Tardif [11] described all homomorphism dualities for directed graphs. In this paper we solve the problem in a surprising generality: we describe all singleton homomorphism dualities for finite relational structures in general. In view of the scarcity of examples that arise in the category of undirected graphs, it seems unlikely that the framework for such a generalisation would be found in this context rather than that of directed graphs. Yet paradoxically, this is precisely what happens. We are able to explain the absence of good characterisations for undirected graphs by an apparently unrelated result, that is, the density theorem of Welzl, which states that the class of undirected non-bipartite graphs is dense with respect to the homomorphism order. The argument is purely categorial and extends to all relational structures, as shown in Section 2. In this context, our main result is the correspondence between 'duality pairs' and 'gap pairs' described in Theorem 2.8. Using this correspondence we achieve simultaneously a characterisation of both singleton homomorphism dualities (for finite relational structures) and of gaps in the partial order of relational structures ordered by the existence of a homomorphism. In this way non-gap pairs are treated in Section 3 and duality pairs in Section 2. Together they give the full characterisation. As a consequence we also describe not only all singleton dualities but also all homomorphism dualities which are induced by finitely many obstructions (we call them *finitary hom-dualities*, see Theorem 2.9. We conclude the paper with some examples and open problems.

2 Duality: A correspondence

Relational Structures

A relational structure of a given type generalizes the notion of a relation and of a graph to more relations and to higher (non-binary) arities. The concept was isolated in thirties by logicians (e.g. Löwenheim, Skolem) who developed logical "static" theory. As we shall see this influenced terminology even today as we find useful to speak about models (of our chosen relational language). In the sixties new impulses (e.g. Isbell, Hedrlín, Pultr, Lovász) came from the study of algebraic categories and the resulting "dynamic" studies called for a more explicit approach, see [4], [14], [8]. We shall adopt here a later notation (with a stance of logical vocabulary).

A type Δ is a sequence $(\delta_i; i \in I)$ of positive integers. A *relational system* A of *type* Δ is a pair $(X, (R_i; i \in I))$ where X is a set and $R_i \in X^{\delta_i}$; that is R_i is a δ_i-nary relation on X. In this paper we shall always assume that X is a finite set (thus we consider finite relational systems only).

Relational systems (of type Δ) will be denoted by capital letters A,B,C,.... A relational system of type Δ is also called a Δ-system (or a model). If $A = (X, (R_i; i \in I))$ we also denote the base set X as $\underline{A}$ and the relation R_i by $R_i(A)$. Let $A = (X, (R_i; i \in I))$ and $B = (Y, (S_i; i \in I))$ be Δ-systems. A mapping $f : X \to Y$ is called a *homomorphism* if for each $i \in I$ holds: $(x_1, ..., x_{\delta_i}) \in R_i$ implies $(f(x_1), ..., f(x_{\delta_i})) \in S_i$.

In other words a homomorphism f is any mapping $F : \underline{A} \to \underline{B}$ which satisfies $f(R_i(A)) \subset R_i(B)$ for each $i \in I$. (Here we extended the definition of f by putting $f(x_1, ..., x_t) = (f(x_1), ..., f(x_t))$.)

For Δ-systems A and B we write $A \to B$ if there exists a homomorphism from A to B. Hence the symbol $\to$ denotes a relation that is defined on the class of all Δ-systems. This relation is clearly reflexive and transitive, thus induces a quasi-ordering of all Δ-systems. As is usual with quasi-orderings, it is convenient to reduce it to a partial order on classes of equivalent objects: Two Δ-systems A and B are called *homomorphically equivalent* if we have both $A \to B$ and $B \to A$; we then write $A \sim B$.

The relation $\to$ induces an order on the classes of homomorphically equivalent Δ-system, which we call the *homomorphism order*. The operations of sum, product and exponentiation reveal the rich categorical structure of the homomorphism order:

- The *sum* $A + B$ of A and B has the property that for any Δ-system C, we have $A + B \to C$ if and only if $A \to C$ and $B \to C$.

- The *product* $A \times B$ of A and B has the property that for any Δ-system C, we have $C \to A \times B$ if and only if $C \to A$ and $C \to B$.

- The *B-th exponent* A^B of A has the property that for any Δ-system C, we have $B \times C \to A$ if and only if $C \to A^B$.

- The two distributive laws hold between the sum and the product:

$$\begin{aligned} A \times (B + C) &\sim (A \times B) + (A \times C), \\ A + (B \times C) &\sim (A + B) \times (A + C). \end{aligned}$$

Thus, the homomorphism order is a distributive lattice with exponentiation. This categorical description will be more relevant to us that the actual (i.e. inner) description of sums, products and exponents, which is bit technical though standard. The sum $A + B$ of two Δ-systems A and B is just their disjoint union. Their product $A \times B$ has base set $\underline{A} \times \underline{B}$, and for $i \in I$, we have $((a_1, b_1), \ldots, (a_{\delta_i}, b_{\delta_i})) \in R_i(A \times B)$ if and only if $(a_1, \ldots, a_{\delta_i}) \in R_i(A)$ and $(b_1, \ldots, b_{\delta_i}) \in R_i(B)$. The B-th exponent A^B of A has the set of all functions from $\underline{B}$ to $\underline{A}$ as base set, and for $i \in I$, we have $(f_1, \ldots, f_{\delta_i}) \in R(A^B)$ if and only if we have $(f_1(b_1), \ldots, f_{\delta_i}(b_{\delta_i})) \in R(A)$ whenever $(b_1, \ldots, b_{\delta_i}) \in R(B)$.

These definitions of products and exponents will not be needed again until Section 4, where they are used in the construction of examples. The sum has an additional descriptive function, as it embodies the standard notion of connectedness: a Δ-system is *connected* if it cannot be represented as a sum of two nonempty Δ-systems. It is easy to see from this that if A, B, C

are Δ-systems such that A is connected and $A \to B + C$, then $A \to B$ or $A \to C$, but note that this is actually a consequence of the distributive lattice structure of the order homomorphism.

Finally, in the context of finite structures, the concepts of retracts and cores are quite useful. Let A, B be Δ-systems with $\underline{B} \subseteq \underline{A}$. Then B is called a *retract* of A if there exists a homomorphism $f : A \mapsto B$ whose restriction to $\underline{B}$ is the identity. In particular, if A is finite and $f : A \mapsto A$ is a homomorphism, then for a sufficiently large n, $f^n(A)$ is a retract of A. A finite Δ-system A is called a *core* if it has no proper retracts, or equivalently, if every homomorphism $f : A \mapsto A$ is an automorphism of A. Any finite Δ-system A has a retract A' which is a core, as is easily seen by selecting A' as a retract of A with the smallest cardinality. The question of uniqueness is easily settled by the following observation.

Lemma 2.1 *Let A' be a core which is homomorphically equivalent to A. Then for any homomorphism $\phi : A' \mapsto A$, there exists a homomorphism $\phi' : A \mapsto A'$ such that $\phi' \circ \phi$ is the identity on A'. Conversely, for any homomorphism $\psi : A \mapsto A'$, there exists a homomorphism $\psi' : A' \mapsto A$ such that $\psi \circ \psi'$ is the identity on A'*

Proof. Let $\phi : A' \mapsto A$ and $\psi : A \mapsto A'$ be arbitrary homomorphisms. Then, $\gamma = \psi \circ \phi$ is an automorphism of A'. Thus, the maps $\phi' = \gamma^{-1} \circ \psi$ and $\psi' = \phi \circ \gamma^{-1}$ satisfy $\phi' \circ \phi = \psi \circ \psi' = \mathrm{id}_{A'}$. ■

As a consequence of this result, all the retracts of A which are cores must be isomorphic, and it makes sense to think of A' as *the* core of A. Furthermore, all the Δ-systems which are homomorphically equivalent to A must have isomorphic cores. Thus, in our investigations of homomorphisms between finite Δ-systems, we can usually restrict our attention to cores without loss of generality.

Duality pairs and gap pairs

Singleton good characterisations are those where the family of obstructions consists of just one structure. This leads to the following: Given two Δ-systems A and B, we call B the *dual* of A if the following holds.

> For every Δ-system C, there exists a homomorphism from A to C if and only if there does not exist a homomorphism from C to B.

This statement admits a natural interpretation in terms of ideals and filters in the homomorphism order: Let $\to A$ denote the class of Δ-systems which admit a homomorphism into A, and similarly for $\not\to A$, $A \to$ and $A \not\to$. Then, $\to A$ is just the principal ideal generated by A in the order homomorphism, $\not\to A$ is its complement, $A \to$ is the principal filter generated by A and $A \not\to$ is its complement. The statement above is just the equality $A \to = \not\to B$.

Definition 2.2 Let A, B be Δ-systems. We say that the couple (A, B) is a *duality pair* if we have the equality

$$A \to = \not\to B.$$

In this section, we present an alternative characterisation of duality pairs based on the following observation.

Lemma 2.3 *Let (A, B) be a duality pair, where A and B are cores. Then A is connected, $A \times B \to A$ and for every Δ-system C such that $A \times B \to C \to A$, we have either $C \sim A \times B$ or $C \sim A$.*

Proof. We first show that A must be connected. Suppose that $A = A_1 + \ldots + A_n$. Then, $A \not\to A_i$ implies $A_i \to B$ for $i = 1, \ldots, n$. Therefore, $A = A_1 + \ldots + A_n \to B$, and this implies $B \not\to B$, which is absurd.

Thus, A is connected. We clearly have $A \times B \to A$, and for any Δ-system C such that $A \times B \to C \to A$, we either have $A \sim C$, or $A \not\to C$. In the latter case, we have $C \to B$, whence $C \to A \times B$. ■

This result motivates the following definition.

Definition 2.4 Let A, B be Δ-systems. We say that the couple (A, B) is a *gap pair* if $A \to B$, $B \not\to A$ and every Δ-system C such that $A \to C \to B$ satisfies $C \sim A$ or $C \sim B$.

Hence, a gap pair is just a cover in the homomorphism order. Lemma 2.3 shows how gap pairs are derived from duality pairs. The converse is the following.

Lemma 2.5 *Let (A, B) be a gap pair, where B is connected. Then (B, A^B) is a duality pair.*

Proof. For every model C of Δ-system, we have $A \to A + (B \times C) \to B$. Since (A, B) is a gap pair, this implies that we have either $A + (B \times C) \sim A$, or $A + (B \times C) \sim B$. However, we have $A + (B \times C) \sim A$ if and only if $B \times C \to A$, that is $C \to A^B$. Also, since B is connected and $B \not\to A$, we have $A + (B \times C) \sim B$ if and only if $B \to B \times C$, that is, $B \to C$. This shows that the classes $B \to$ and $\to A^B$ are complementary. However, we know that the complement of the class $\to A^B$ is the class $\not\to A^B$. Thus,

$$B \to = \not\to A^B.$$

■

Hence there is a natural correspondence between the duality pairs (A, B) and the gap pairs (C, D) where D is connected. Starting from a duality pair (A, B), we find the gap pair $(A \times B, A)$ by Lemma 2.3, whence $(A, (A \times B)^A)$ is a duality pair by Lemma 2.5. We then have $\to B = \to (A \times B)^A$, and thus $B \sim (A \times B)^A$. Conversely, if (A, B) if a gap pair and B is connected, then (B, A^B) is a duality pair by Lemma 2.5, whence $(B \times A^B, B)$ is a gap pair by Lemma 2.3. We clearly have $A \sim B \times A^B$. This shows that up to homomorphic equivalence, the correspondence described in Lemmas 2.3 and 2.5 is one-to-one and onto.

It remains to characterise the other gap pairs in the homomorphism order, namely those where the second member is not connected. We use the following observation.

Lemma 2.6 *Let (A, B) be a gap pair, where B is connected. Then for every Δ-system C such that $C \to B$ and $B \not\to C$, we have $C \to A$.*

Proof. We have $A \to A + C \to B$, but since B is connected, we have $B \not\to A + C$, whence $A + C \sim A$, that is, $C \to A$. ■

Lemma 2.7 *Let (A, B) be a gap pair, where B is connected. Then for any C such that $A \to C \to A^B$, $(C, C + B)$ is a gap pair. Moreover, for each gap pair (C, D), there exists a gap pair (A, B) such that B is connected, $A \to C \to A^B$ and $D \sim C + B$.*

Proof. By Lemma 2.5, if B is connected and (A, B) is a gap pair, then (B, A^B) is a duality pair. Hence if $A \to C \to A^B$, then $B \not\to C$, thus $C + B \not\to C$. Suppose that we have $C \to D \to C + B$ for some Δ-system D. Then either $B \to D$, in which case $D \sim C + B$, or every connected component of D that admits a homomorphism to B also admits a homomorphism to A by Lemma 2.6. Since $A \to C$, this implies $D \sim C$.

It remains to show that every gap pair has this structure. Let (C, D) be an arbitrary gap pair. Then for every connected component B of D, we have $C \to C + B \to D$, which implies that either $C + B \sim C$ or $C + B \sim D$. Since $D \not\to C$, the second alternative must be true of at least one connected component B of D. We then have $D \sim C + B$. No Δ-system E can satisfy $C \times B \to E \to B$ and $B \not\to E \not\to C \times B$, for then we would have $C \to C + E \to C + B$ and $C + B \not\to C + E \not\to C$, a contradiction to the fact that $(C, C + B)$ is a gap pair. Hence, putting $A = C \times B$, we have that (A, B) is a gap pair. By Lemma 2.5, (B, A^B) is then a duality pair. Since $B \not\to C$, we then have $A \to C \to A^B$. ■

Hence, we can provide a complete description of the correspondence between duality pairs and gap pairs.

Theorem 2.8 *Let Δ be a fixed type. Then the gap pairs in the class of all finite Δ-systems are the pairs (C, D) such that there exists a duality pair (A, B) with $A \times B \to C \to B$ and $D \sim C + A$. Conversely, the duality pairs in the class of all finite Δ-systems are the pairs (B, A^B) where (A, B) is a gap pair and B is connected.* ■

Thus we have the following characterisation of finitary hom-dualities:

Theorem 2.9 *Let Δ be a fixed type. Then there exists a finite family $\mathcal{C} = \{A_1, \ldots, A_n\}$ of Δ-systems such that*

$$\bigcup_{i=1}^{n} (A_i \to) = \not\to B \tag{1}$$

if and only if $B = \times_{i=1}^{n} B_i$, where (A_i, B_i) is a duality pair for $i = 1, \ldots, n$.

Proof. Let $(A_1, B_1), \ldots, (A_n, B_n)$ be duality pairs, and $B = \times_{i=1}^{n} B_i$. Then for any Δ-system C, we have $C \not\to B$ if and only $C \not\to B_i$ for some i, that is, if and only if $A_i \to C$. This shows that (1) holds.

Conversely, let B be a Δ-system such that (1) holds for some family $\mathcal{C} = \{A_1, \ldots, A_n\}$ of Δ-systems. We can assume that $A_i \not\to A_j$ for every $i \neq j$, and from this follows that each A_i is connected just as in the proof of Lemma 2.3. Moreover, $(B, B + A_i)$ is easily seen to be a gap pair. Thus by Lemma 2.7, there exists a gap pair (C_i, A_i) such that $C_i \to B \to C_i^{A_i}$. Putting $B_i = C_i^{A_i}, i = 1, \ldots, n$, we then have $\bigcup_{i=1}^n (A_i \to) = \not\to \times_{i=1}^n B_i$, whence $B \sim \times_{i=1}^n B_i$. ■

Density

A partially ordered set P is called *dense* if it has the property that for any $x, y \in P$ such that $x < y$, there exists $z \in P$ such that $x < z < y$. Therefore, a homomorphism order is dense if and only if it does not contain any gaps. Theorem 2.8 shows that duality and density are just two aspects of the same question. Both have been investigated in the case of directed and undirected graphs, but these subjects have been treated independently up to now.

Duality is essentially a void concept in the category of undirected graphs, since (as shown in [10]) $(\emptyset, K_1), (K_1, K_2)$ are the only duality pairs. On the positive side, this implies that the class of undirected graphs is dense except for these 'trivial' gaps. This property was eventually acknowledged, and Welzl [16] was the first to give a proof of what became known as the 'density theorem' for undirected graphs. The original argument was a long and involved ad hoc construction. It seems natural that such a result would be difficult to prove, since the question of the existence of homomorphisms between non-bipartite graphs is NP-complete. However, a short and elegant proof of the density theorem, based on exponentiation, was later found independently by Perles and by Nešetřil (see e.g. [9]).

This unexpected proof opened the way for new investigations on the subject of density. In [17], Welzl had attracted the attention on the density problem for vertex-transitive graphs, which has recently been solved by Tardif [15] and independently by Perles. In another direction, Nešetřil and Zhu [13] investigated the class of oriented paths, and proved a density result similar to that of Welzl. In this context, the structure of the gaps is more intricate, and their complete characterisation was a feat. It turns out that the gaps in the class of oriented paths are also gaps in the class of all directed graphs. On the other hand, the Nešetřil-Perles proof of the density theorem adapts to some classes of directed graphs (such as unbalanced graphs, se [9]), but not all. Thus, the problem of characterising the gaps in the category of directed graphs remained open for a long time, with no simple solution in view. All the while, the duality pairs in the class of directed graphs had already been characterised by Komárek: For a directed (core) graph G, there exists a directed graph D_G such that (G, D_G) is a duality pair if and only if G is an orientation of a tree. Thus, modulo the correspondence presented here, the problem of density for directed graphs was solved even before it was formulated, as we mentioned in [11].

We will show that the case of directed graphs is a faithful reflection of the general situation in relational structures: The structures that are first members of duality pairs are 'trees' in a certain sense. According to Theorem 2.8, we may choose to confront the problem from the point of view of density instead of that of duality. This is indeed the approach adopted in the next section.

3 Density: A characterisation

Shadows of relational structures

Let A be a Δ-system. The *shadow* of A is the unoriented multigraph $G(A)$ whose vertices are the elements of $\underline{A}$, and containing an edge e joining a and b whenever there exists a relation R of arity $n \geq 2$ in A such that $(a_1, \ldots, a_n) \in R(A)$ with $a_i = a, a_{i+1} = b$ for some i.

The full structure of relational models determines which maps are homomorphisms, but their shadows are sufficient for a description of those which admit a dual. The 1-ary relations do not play any part in the definition of shadows, while the relations of higher arities may contribute many edges and loops. A cycle of $G(A)$ can be a 1-cycle (i.e., a loop), a 2-cycle (i.e., two parallel edges), or an ordinary n-cycle with $n \geq 3$. We will call A a *tree* if $G(A)$ is a tree (and thus has neither multiple edges nor loops). The purpose of this section is to prove the following.

Theorem 3.1 *Let A be a connected core. Then there exists a Δ-system B such that (B, A) is a gap pair if and only if A is a tree.*

The proof consists of two constructions. The first is the arrow construction discussed in [12], which is used to find Δ-systems that are homomorphically 'in between' two given Δ-systems. We next characterise the 'gap below a tree' with a new construction based on local inversions of homomorphisms.

The arrow construction

Definition 3.2 Let A be a Δ-system, and $\mathcal{P}$ a partition of $\underline{A}$.(Here we wiew a partition $\mathcal{P}$ as a set of disjoint subsets $P_1, ..., P_n$.) The *quotient* $A|\mathcal{P}$ is the Δ-system defined on the set $\mathcal{P}$ by putting, for every relation R_i of A, $(P_1, \ldots, P_{\delta_i}) \in R_{\delta_i}(A|\mathcal{P})$ if and only if there exists $a_i \in A_i, i = 1, \ldots, \delta_i$ such that $(a_1, \ldots, a_{\delta_i}) \in R_i(A)$. We denote $\phi_{\mathcal{P}}$ the quotient map from A to $A|\mathcal{P}$.

The quotient $A|\mathcal{P}$ has the least structure which makes the quotient map $\phi_{\mathcal{P}}$ a homomorphism. This explains why quotients arise naturally in connection with homomorphisms. In particular, when $\mathcal{P}$ is a partition of $\underline{A}$, then a homomorphism $\phi : A \mapsto B$ can be factored through $A|\mathcal{P}$, that is, expressed as $\phi = \psi \circ \phi_{\mathcal{P}}$, where $\psi : A|\mathcal{P} \mapsto B$ is a homomorphism, if and only if $\mathcal{P}$ refines the partition $\{\phi^{-1}(b) : b \in \underline{B}\}$ of $\underline{A}$.

The operation presented next can be described informally as replacing the arcs of a directed graph by copies of a Δ-system. Formally, this procedure is best presented in terms of quotients.

Definition 3.3 Let K be a directed graph and D a Δ-system. For $a, b \in \underline{D}$, the arrow product $K * D(a, b)$ is the Δ-system $(E(K) \cdot D)|\mathcal{P}$, where

- $E(K) \cdot D$ is the disjoint union of $|E(K)|$ copies of D, defined on the base set $E(K) \times \underline{D}$ by putting $((e_1, d_1), \ldots, (e_{\delta_i}, d_{\delta_i})) \in R_i(E(K) \cdot D)$ if and only if $e_1 = \ldots = e_{\delta_i}$ and $(d_1, \ldots, d_{\delta_i}) \in R_i(D)$ for every $i \in I$.

- $\mathcal{P} = \mathcal{P}_0 \cup \mathcal{P}_1$ is a partition of $\underline{E(K) \cdot D}$, where $\mathcal{P}_0$ contains one set V_u for each vertex u of K, defined by

$$V_u = \{((u, v), a) : (u, v) \in E(K)\} \cup \{((v, u), b) : (v, u) \in E(K)\},$$

  and $\mathcal{P}_1$ consists of singletons containing the remaining elements:

$$\mathcal{P}_1 = \{\{(e, c)\} : e \in E(K) \text{ and } c \in \underline{D} \setminus \{a, b\}\}.$$

Thus, $K * D(a, b)$ is the Δ-system obtained by replacing every directed edge (u, v) of K by a copy of D, with a taking the role of u and b taking the role of v. Independently of the structure of K, we then have $K * D(a, b) \to D|\mathcal{P}_{ab}$, where the partition $\mathcal{P}_{ab}$ only identifies a and b (in all our applications elements a and b will be distinct). This observation will be the basis of our construction.

Proposition 3.4 *Let A be a connected core such that $G(A)$ contains a cycle. Then for any Δ-system B such that $A \not\to B$, there exists a Δ-system C such that $C \to A$, $A \not\to C$ and $C \not\to B$.*

Proof. We first 'split up' an element of $\underline{A}$, that is, express A as a quotient $D|\mathcal{P}_{ab}$, where the partition $\mathcal{P}_{ab}$ only identifies two elements a and b. We need to choose this element carefully. Since $G(A)$ contains a cycle, there exists $i \in I$ and $(a_1, \ldots, a_{\delta_i}) \in R_i(A)$ such that at least one of the corresponding edges $[a_j, a_{j+1}]$ is contained in a cycle of $G(A)$ (to simplify the notation, we use brackets to denote edges even though $G(A)$ is a multigraph). Let j be the minimum index such that the edge $[a_j, a_{j+1}]$ is contained in a cycle. Put $a = a_j$ and $a' = a_{j+1}$. Let D be the Δ-system defined on the base set $\underline{D} = \underline{A} \cup \{b\}$ (we assume $b \notin A$) by putting,

- $R_i(D) = R_i(A) \setminus (a_1, \ldots, a_{\delta_i}) \cup (a_1, \ldots, a_{i-1}, b, a_{i+1}, \ldots, a_{\delta_i})$,
- $R_{i'}(D) = R_{i'}(A)$ if $i' \neq i$.

We then have $A = D|\mathcal{P}_{ab}$. The edges of $G(A)$ correspond naturally to the edges of $G(D)$. The vertex b of $G(D)$ is incident to at most two edges, and by minimality, the edge $[a, a']$, which is in a cycle of $G(A)$, correspond to $[b, a']$ which is not contained in any cycle of $G(D)$. (Note that if $a = a'$, $[a, a]$ is a loop of $G(A)$ while $[b, a]$ is not a loop of $G(D)$.)

It is possible that we already have $D \not\to B$, and we are done (as $D \to A$ and $A \not\to D$ as one can check easily - if not see the argument below). In any case, no homomorphism from D to B can identify a and b, since $D|\mathcal{P}_{ab} = A \not\to B$. Hence, such homomorphisms from D to B can be thought of as colourings of a and b with different elements of B. Let K be an arbitrary oriented graph with chromatic number greater than the cardinality of $\underline{B}$, and put $C = K * D(a, b)$. We shall prove that C satisfies the required conditions.

- $C \to A$ as noted above. The natural homomorphism $\phi : C \mapsto A$ coincides with the quotient map $\phi_{\mathcal{P}_{ab}}$ on every canonical copy of D in C.
- $C \not\to B$ since any homomorphism from C to B should map V_u and V_v to different elements of $\underline{B}$ whenever (u, v) is an edge of K, which is impossible because $\chi(K) > |\underline{B}|$.

It only remains to show that $A \not\to C$. Note that $A \to C$ holds if and only if A is the core of C. Supposing that this is the case, there exists a homomorphism $\phi' : A \mapsto C$ such that $\phi \circ \phi' = \mathrm{id}_A$ by Lemma 2.1, where $\phi : C \mapsto A$ is the canonical homomorphism defined above. This means that ϕ' maps a to the set V_u for some vertex u of K, and maps every other element c of $\underline{A}$ to a singleton $\{(e, c)\}$. Note that ϕ' induces an embedding of $G(A)$ into $G(C)$. In particular, the edge $[a, a']$ of $G(A)$ is mapped to some edge $[V_u, \{(e, a')\}]$ in $G(C)$, and the cycle containing $[a, a']$ must be mapped to some cycle in $G(C)$. However, since $[b, a']$ is not contained in any cycle of $G(D)$, any cycle of $G(C)$ containing $[V_u, \{(e, a')\}]$ must pass through other sets V_v corresponding to vertices of K. This is a contradiction, since these sets do not belong to the image of ϕ'. Note that if $a = a'$, the contradiction is immediate since the edges of C corresponding to $[a, a]$ are not loops. ■

In particular, if $B \to A$ and $A \not\to B$, then with C as in Proposition 3.4, we have $B \to B \cup C \to A$ and $A \not\to B \cup C \not\to B$. Thus, we have the following.

Corollary 3.5 *Let A be a connected core such that $G(A)$ contains a cycle. Then for any Δ-system B such that $B \to A$ and $A \not\to B$, there exists a Δ-system C such that $B \to C \to A$ and $A \not\to C \not\to B$.* ■

This completes the first part of the proof of Theorem 3.1

The gap below a tree

The concept of 'tree' is quite descriptive, even in the case of relational Δ-systems. A Δ-system which is a tree will be called shortly Δ-tree. However, it is not entirely clear what should be meant by a 'subtree' of a Δ-system A which is a tree. Indeed, the Δ-system A may contain 1-ary relations, and these are not represented in the structure of $G(A)$. Our construction makes an extensive use of subtrees, and it is necessary to give a precise definition.

Definition 3.6 Let A be a Δ-tree. A *subtree* of A is a Δ-system B which is a tree such that $\underline{B} \subseteq \underline{A}$ and the inclusion is a homomorphism from B to A. B is a *proper subtree* of A if the inclusion is not an isomorphism from B to A.

It can happen that different subtrees of A share the same shadow. In particular, a proper subtree B of A can have $G(A)$ as its shadow. In this case, there exists a 1-ary relation R_i such that $a \in R_i(A)$ and $a \notin R_i(B)$ for some $a \in \underline{A} = \underline{B}$. This shows that the standard set-theoretic notation is not convenient to represent subtree inclusion. We will reserve the set-theoretic notation for the base sets, and use lattice-theoretic notations to order subtrees, as detailed in the following.

Definition 3.7 Let A be a Δ-tree.

- The family of all Δ-subtrees of A is denoted $\mathcal{T}_A$. We will use the symbol $\leq$ to denote the order relation "is a subtree of" on $\mathcal{T}_A$.
- $(\mathcal{T}_A, \leq)$ is a lattice, with $B_1 \wedge B_2$ defined by $\underline{B_1 \wedge B_2} = \underline{B_1} \cap \underline{B_2}$ and $R_i(B_1 \wedge B_2) = R_i(B_1) \cap R_i(B_2)$ for every $i \in I$.
- Conversely, we write $B_1 \vee B_2$ to denote the supremum of B_1 and B_2 in $(\mathcal{T}_A, \leq)$.

This lattice ordering of the subtrees of a tree parallels the situation in graphs, but note that $G(B_1 \vee B_2)$ can be strictly greater than the smallest subtree of $G(A)$ containing both $G(B_1)$ and $G(B_2)$. This follows from the fact that not all the subtrees of $G(A)$ induce subtrees of A. Given a subtree T of $G(A)$, there exists a subtree B of A with $G(B) = T$ if and only if for every $i \in I$ $\Delta_i \geq 2$, $(a_1, \ldots, a_{\delta_i}) \in R_i(A)$ implies that the path $P = a_1, \ldots, a_{\delta_i}$ of $G(A)$ is either contained in T or intersects T in at most one vertex.

Following these preliminary definitions, we can present the concept that will be the basis of our construction.

Definition 3.8 Let A be a tree. For $a \in \underline{A}$, a set $\mathcal{I}$ of proper subtrees of A containing a is called a a-ideal if

- $B' \leq B \in \mathcal{I}$ implies $B' \in \mathcal{I}$ whenever $a \in \underline{B'}$,

- $B_1 \vee B_2 \in \mathcal{I}$ whenever $B_1, B_2 \in \mathcal{I}$ and $\underline{B_1} \cap \underline{B_2} = \{a\}$.

We can now present our construction of the predecessor of a tree.

Definition 3.9 Let A be a Δ-tree. The Δ-system $A^{\downarrow}$ is defined on the base set

$$\underline{A^{\downarrow}} = \{(a, \mathcal{I}) : a \in \underline{A}, \mathcal{I} \text{ is an } a\text{-ideal of } \mathcal{T}_A\}$$

as follows:

- If R_i is a 1-ary relation, i.e. if $\delta_i = 1$, put $(a, \mathcal{I}) \in R_i(A^{\downarrow})$ if $a \in R_i(B)$ for some $B \in \mathcal{I}$.
- If R_i is a δ-ary relation with $\delta > 1$, put $((a_1, \mathcal{I}_1), \ldots, (a_\delta, \mathcal{I}_\delta)) \in R_i(A^{\downarrow})$ if and only if $(a_1, \ldots, a_\delta) \in R_i(A)$ and $\bigvee_{j=1}^{\delta} B_j \in \bigcap_{j=1}^{\delta} \mathcal{I}_j$ for every family $B_j \in \mathcal{I}_j, j = 1, \ldots, \delta$ such that $\underline{B_{j'}} \cap \underline{B_j} = \emptyset$ whenever $j' \neq j$.

We clearly have $A^{\downarrow} \to A$, since the projection $\pi : A^{\downarrow} \mapsto A$ defined by $\pi(a, \mathcal{I}) = a$ is a homomorphism. The next two lemmas will show that $(A^{\downarrow}, A)$ is a gap.

Lemma 3.10 *Let A be a Δ-tree. Then for any Δ-system X such that $X \to A$ and $A \not\to X$, we have $X \to A^{\downarrow}$.*

Proof. Let X be a Δ-system such that $X \to A$ and $A \not\to X$. Let $\phi : X \mapsto A$ be a homomorphism. For $x \in X$, let $\mathcal{I}_x$ be the family of subtrees B of A such that $\phi(x) \in B$ and ϕ admits a 'local inverse' around x with domain B, that is, a homomorphism $\psi_B : B \mapsto X$ such that $\psi_B(\phi(x)) = x$ and $\phi \circ \psi_B = \mathrm{id}_B$.

We first show that $\mathcal{I}_x$ is a $\phi(x)$-ideal. The fact that $A \not\to X$ implies that each element of $\mathcal{I}_x$ is a proper subtree of A. We clearly have $B' \in \mathcal{I}_x$ whenever $B' \leq B \in \mathcal{I}_x$ and $\phi(x) \in \underline{B'}$, since the restriction of a local inverse $\psi_B : B \mapsto X$ to B' is also a local inverse of ϕ. It remains to show that for $B_1, B_2 \in \mathcal{I}_x$ such that $\underline{B_1} \cap \underline{B_2} = \{\phi(x)\}$, we have $B_1 \vee B_2 \in \mathcal{I}_x$. Let $\psi_1 : B_1 \mapsto X, \psi_2 : B_2 \mapsto X$ be local inverses of ϕ around x. Note that $\underline{B_1 \vee B_2} = \underline{B_1} \cup \underline{B_2}$, and $\psi_1(\phi(x)) = \psi_2(\phi(x)) = x$. Hence we can define $\psi : B_1 \vee B_2 \mapsto X$ by $\psi(b) = \psi_i(b)$ if $b \in \underline{B_i}$. This is a homomorphism, since for any relation R_i $i \in I$, $(b_1, \ldots, b_{\delta_i}) \in R_i(B_1 \vee B_2)$ implies that $b_1, \ldots, b_{\delta_i}$ all belong to the same B_j, whence $(\psi_i(b_1), \ldots, \psi_i(b_{\delta_i})) \in R_i(X)$. Thus, $\mathcal{I}_x$ is a $\phi(x)$-ideal.

Therefore, we can define a map $\hat{\phi} : \underline{X} \mapsto \underline{A^{\downarrow}}$ by $\hat{\phi}(x) = (\phi(x), \mathcal{I}_x)$, and it only remains to show that $\hat{\phi}$ is a homomorphism from X to $A^{\downarrow}$. For $(x_1, \ldots, x_{\delta_i}) \in R_i(X)$, we have $(\phi(x_1), \ldots, \phi(x_{\delta_i})) \in R_i(A)$ since ϕ is a homomorphism. Let $\{B_1, \ldots, B_{\delta_i}\}$ be a family of pairwise disjoint subtrees of A such that $B_j \in \mathcal{I}_{x_j}, j = 1, \ldots, \delta_i$, and $\psi_j : B_j \mapsto X$ the associated local inverses. Then, $\underline{\bigvee_{j=1}^{\delta_i} B_j} = \bigcup_{j=1}^{\delta_i} \underline{B_j}$, and the map $\psi : \bigvee_{j=1}^{\delta_i} B_j \mapsto X$ defined by $\psi(b) = \psi_j(b)$ if $\overline{b \in \underline{b_j}}$ is a homomorphism by the same argument as the one used above. Thus, $\overline{\bigvee}_{j=1}^{\delta_i} B_j \in \bigcap_{j=1}^{\delta_i} \mathcal{I}_j$. Therefore, $((x_1, \mathcal{I}_{x_1}), \ldots, (x_{\delta_i}, \mathcal{I}_{x_{\delta_i}})) \in R_i(A^{\downarrow})$, and $\hat{\phi}$ is a homomorphism. ■

Lemma 3.11 *Let A be a tree and a core. Then $A \not\mapsto A^{\downarrow}$.*

Proof. Suppose that $\phi : A \mapsto A^{\downarrow}$ is a homomorphism. Since A is a core, we can assume by Lemma 2.1 that $\pi \circ \phi = \mathrm{id}_A$, where $\pi : A^{\downarrow} \mapsto A$ is the natural projection. Thus, for every $a \in \underline{A}$, there exists an a-ideal $\mathcal{I}_a$ such that $\phi(a) = (a, \mathcal{I}_a)$. We will show that $A \in \mathcal{I}_a$, for some $a \in \underline{A}$, which is a contradiction since $\mathcal{I}_a$ should only contain proper subtrees of A.

Note that the set of edges of $G(A)$ admits a canonical partition into paths $P_j = a_1^j, \ldots, a_{\delta_j}^j, j = 1, \ldots, m$ such that for every path P_j, there exists a relation R_j of arity $\delta_j \geq 2$ such that $({a^j}_1, \ldots {a^j}_{\delta_j}) \in R_j(A)$; these are called the *elementary paths* of A. If T is a subtree of $G(A)$ which is a union of some elementary paths, we denote A_T the subtree of A induced by T, that is, the maximal subtree of A having T as its shadow. We will prove that if a is a vertex of T, then $A_T \in \mathcal{I}_a$ by induction on the number of elementary paths of T. This will provide our contradiction since $A_{G(A)} = A$.

First step of our induction is the case where T contains no elementary path. Then, T consists of a single vertex a, and A_T is the Δ-system with base set $\{a\}$ such that for every 1-ary relation R_i, we have $a \in R_i(A_T)$ if and only if $a \in R_i(A)$. Let $R_1, \ldots, R_n$ be the 1-ary relations of A such that $a \in R_i(A)$. Since ϕ is a homomorphism, we have $(a, \mathcal{I}_a) \in R_i(A^{\downarrow})$ which means that $\mathcal{I}_a$ contains at least one tree B_i with base set $\{a\}$ such that $a \in R_i(B_i)$ for $i = 1, \ldots, n$. Therefore, $\vee_{i=1}^{n} B_i = A_T \in \mathcal{I}_a$ since $\mathcal{I}_a$ is an a-ideal.

Now suppose that our assumption is true for every subtree of $G(A)$ with at most k elementary paths. Let $T \subseteq G(A)$ be the union of $k+1$ elementary paths. Then a belongs to some elementary path $P = a_1, \ldots, a_n$, of A. Let T_i be the connected component of $T - E(P)$ which contains a_i, $i = 1, \ldots, n$.

By our induction hypothesis, we have $A_{T_i} \in \mathcal{I}_{a_i}$, $i = 1, \ldots, n$. Since ϕ is a homomorphism and $(a_1, \ldots, a_n) \in R(A)$ for some $R \in \mathcal{L}$, we then have $((a_1, \mathcal{I}_{a_1}), \ldots, (a_n, \mathcal{I}_{a_n})) \in R(A^{\downarrow})$. By the definition of $A^{\downarrow}$, this implies

$$A_T = \bigvee_{i=1}^{n} A_{T_i} \in \bigcap_{i=1}^{n} \mathcal{I}_{a_i}.$$

In particular, $A_T \in \mathcal{I}_a$. ■

Combining the two previous lemmas, we get the following.

Proposition 3.12 *Let A be a Δ-system that is a tree and a core. Then $A^{\downarrow} \to A$, and for every $\mathcal{L}$-model B such that $A^{\downarrow} \to B \to A$, we have $B \sim A$ or $B \sim A^{\downarrow}$.* ■

This completes the second part of the proof of Theorem 3.1.

4 Examples

As a consequence of Theorem 3.1, the dualities that are characterised in Theorem 2.8 are of the type

$$A \to \; = \; \not\to (A^{\downarrow})^{A},$$

where A is a core tree. The 'good characterisations' of Theorem 2.9 are obtained by combining some of these dualities. In this section, we present a few examples to illustrate the use of the predecessor construction and exponentiation, and point out some questions raised by this characterisation.

Types Δ with binary relations allow for the construction of meaningful examples that are not too large. In our first example, the type Δ has two binary relations R_b and R_g. Let A be the Δ-system defined by $\underline{A} = \{1, 2, 3\}$, $R_b(A) = \{(1, 2)\}$ and $R_g(A) = \{(2, 3)\}$. If we interpret the elements of R_b as blue arcs and those of R_g as green arcs, then A is a path with the first arc blue and the second green. Thus, A is a core tree, and admits a predecessor $A^{\downarrow}$. The proper subtrees of A are the singletons and $B_1 = \{1, 2\}, B_2 = \{2, 3\}$ (we can identify subtrees with subsets of $\underline{A}$ since $\mathcal{L}$ does not contain 1-ary relations). Hence, the only nontrivial 1-ideal is $\mathcal{I}_1 = \{\{1\}, B_1\}$, the only nontrivial 3-ideal is $\mathcal{I}_3 = \{\{3\}, B_2\}$, and there are three nontrivial 2-ideals, namely $\mathcal{I}_2 = \{\{2\}, B_1\}$, $\mathcal{I}_2' = \{\{2\}, B_2\}$ and $\mathcal{I}_2'' = \{\{2\}, B_1, B_2\}$. The elements of $A^{\downarrow}$ corresponding to trivial ideals are necessarily isolated, hence can be omitted. Thus, the elements of $\underline{A^{\downarrow}}$ are $1_b = (1, \mathcal{I}_1)$, $2_b = (2, \mathcal{I}_2)$, $2_g = (2, \mathcal{I}_2')$, $2_\emptyset = (2, \mathcal{I}_2'')$, and $3_g = (3, \mathcal{I}_3)$. Since $\{x\} \vee B_i = A$ whenever $x \notin \underline{B_i}$, $R_b(A^{\downarrow})$ contains the only the arc $(1_b, 2_b)$, and $R_g(A^{\downarrow})$ contains the only the arc $(2_g, 3_g)$. It turns out that $2_\emptyset$ is also isolated; essentially, $A^{\downarrow}$ consists of the blue arc $(1_b, 2_b)$ and the green arc $(2_g, 3_g)$.

The base set of the dual $(A^{\downarrow})^A$ of A consists of functions from $\underline{A}$ to $\underline{A^{\downarrow}}$. There are 64 such functions, if we restrict the image to the core of $A^{\downarrow}$. However, most of them do not belong to the core of $(A^{\downarrow})^A$, and we can restrict our attention to the functions which carry the most structure.

- A function f is the beginning of a blue arc only if $f(1) = 1_b$.
- A function f is the end of a blue arc only if $f(2) = 2_b$.
- A function f is the beginning of a green arc only if $f(2) = 2_g$.
- A function f is the end of a green arc only if $f(3) = 3_g$.

Only two functions satisfy at least three of these conditions, namely f_b defined by $f_b(1) = 1_b, f_b(2) = 2_b, f_b(3) = 3_g$ and f_g defined by $f_g(1) = 1_b, f_g(2) = 2_g, f_g(3) = 3_g$. We have $(f_b, f_b), (f_g, f_b) \in R_b((A^{\downarrow})^A)$ and $(f_g, f_b), (f_g, f_g) \in R_g((A^{\downarrow})^A)$. We will not need to look any further. Given a $\mathcal{L}$-model C, define a map $\phi : \underline{C} \mapsto \underline{(A^{\downarrow})^A}$ by

$$\phi(c) = \begin{cases} f_b \text{ if } c \text{ is the end of a blue arc,} \\ f_g \text{ otherwise.} \end{cases}$$

Then, whenever $(c, d) \in R_b(C)$, we have $\phi(d) = f_b$, thus $(\phi(c), \phi(d)) \in R_b((A^{\downarrow})^A)$. If ϕ is not a homomorphism, then there exists $(c, d) \in R_g(C)$ such that $(\phi(c), \phi(d)) \notin R_g((A^{\downarrow})^A)$, that is, $\phi(c) = f_b$. The element c of C is then the beginning of a green arrow and the end of a blue arrow, whence $A \to C$. This shows that $\{f_b, f_g\}$ is the core of $(A^{\downarrow})^A$.

In this example, $(A^{\downarrow})^A$ turns out to have a relatively small core, even though its characterisation involves two exponential constructions, namely the predecessor construction and exponentiation. Is this always the case? Of course, the precise meaning of 'small' depends on the type Δ, since a tree with one element can have a dual with n elements whenever the type Δ has n 1-ary relations. Our question can therefore be formulated as follows.

Problem 1 *Given a fixed type Δ, does there exist a polynomial p_Δ such that for every type Δ-system A that is a core tree with $|\underline{A}| \leq n$, the core of $(A^{\downarrow})^A$ has at most p_Δ elements?*

We conclude with an example which shows that the size of the core of $(A^{\downarrow})^A$ can grow exponentially with respect to $|\underline{A}| + |\Delta|$. For a given integer $n \geq 2$, let Δ contains one n-ary relation R_0 and n 1-ary relations $R_1, \ldots, R_n$. We define the Δ-system A by $\underline{A} = \{1, \ldots, n\}$, $R_0(A) = \{(1, \ldots, n)\}$ and $R_i(A) = \{i\}, i = 1, \ldots, n$. Then A is a core tree and admits a dual $(A^{\downarrow})^A$. Instead of constructing $A^{\downarrow}$ and $(A^{\downarrow})^A$ directly from the definitions, it is possible to guess a plausible candidate for the dual of A and verify that it satisfies the required conditions. This turns out to be more practical in this case, and we will show that the dual of A has 2^n elements.

Let B be the Δ-system whose base set consists of all the subsets of $\{1, \ldots, n\}$, defined by putting $S \in R_i(B)$ if $i \in S$ for $i = 1, \ldots, n$, and $(S_1, \ldots, S_n) \in R_0(S)$ if we have $i \notin S_i$ for some i. Then for any Δ-system C, we can define a map $\phi : \underline{C} \mapsto \underline{B}$ by $\phi(c) = \{i : c \in R_i(C)\}$. We then have $\phi_i(c) \in R_i(B)$ whenever $c \in R_i(C)$ for $i = 1, \ldots, n$. If $A \not\to C$, then for every $(c_1, \ldots, c_n) \in R_0(C)$, there exists an index i such that $c_i \notin R_i(C)$, whence $i \notin \phi(c_i)$ and $(\phi(c_1), \ldots, \phi(c_n)) \in R_0(B)$. Therefore, $A \not\to C$ implies $C \to B$. This shows that B is the dual of A, since $A \not\to B$.

We next show that B is a core. Let $\phi : B \mapsto B$ be a homomorphism. Then $S \subseteq \phi(S)$ for all $S \in \underline{B}$, since ϕ must preserve the 1-ary relations $R_1, \ldots, R_n$. Suppose that S is a proper subset of $\phi(S)$ for some $S \in \underline{B}$. Define $S_1, \ldots, S_n$ by $S_i = S$ if $i \in \phi(S)$ and $S_i = \{1, \ldots, n\}$ if $i \notin \phi(S)$. Then $(S_1, \ldots, S_n) \in R_0(B)$ while $(\phi(S_1), \ldots, \phi(S_n)) \notin R_0(B)$, a contradiction. This shows that ϕ must be the identity, whence B is a core.

The predecessor of A is homomorphically equivalent to $A \times B$. $R_0(A \times B)$ consists of the n-tuples $((1, S_1), \ldots, (n, S_n))$ such that $i \notin S_i$ for at least one index i, and $(i, S) \in R_j(A \times B)$ if and only if $j = i \in S$. The core C of $A \times B$ is obtained by collapsing, for each $i \in \{1, \ldots, n\}$, all the couples (i, S) such that $i \in S$ onto one element labelled i', and all the elements (i, S) such that $i \notin S$ onto one element labelled i''. If C is viewed as the core of $A^{\downarrow}$, then i' corresponds to the i-ideal which contains *all* proper subtrees of A, and i'' corresponds to the principal i-ideal generated by the subtree obtained by removing i from $R_i(A)$. Note that B can then be viewed as the set of functions $f \in \underline{C^A}$ such that $f(i) \in \{i', i''\}$ for $i = 1, \ldots, n$.

In both of our examples, the core of $A^{\downarrow}$ admits a unique homomorphism to A, and the core of $(A^{\downarrow})^A$ consists of the functions which map each element of $\underline{A}$ into its preimage by this homomorphism. This seems to suggest a general simplification for the construction of the core of $(A^{\downarrow})^A$. However, such a simplification could only make sense if the homomorphism from the core of $A^{\downarrow}$ to A was always unique, and this is not the case. Moreover, even when the homomorphism is unique, the core of $(A^{\downarrow})^A$ does not necessarily consist of functions mapping each element to its preimage. It would be interesting to know up to which extent can the characterisation of duals be simplified, and whether the indirect approach via density is optimal.

Acknowledgements This research was done while the second author was visiting Charles University. The support of DIMATIA is gratefully acknowledged. The first author was supported by GAČR 0194 and GAUK 194 grants.

References

[1] D. Duffus, N. Sauer, *Lattices arising in categorial investigations of Hedetniemi's conjecture*, Discrete Math. **152** (1996), 125–139.

[2] J. Edmonds, *Paths, Trees and Flowers*, Canad. J. Math. 17(1965), 449-467.

[3] R. Frucht, *Herstellung von Graphen mit vorgegebener abstrakter Gruppe*, Compos. Math. 6(1938), 239-250

[4] Z. Hedrlín, A. Pultr, *Symmetric Relations (Undirected Graphs) with Given Semigroups*, Monatshefte f. Mathematik 69,4(1965), 318-322

[5] P. Hell, J. Nešetřil, X. Zhu, *Duality and polynomial testing of tree homomorphisms*, Trans. Amer. Math. Soc. **348** (1996), 1281–1297.

[6] W. Hochstätter, J. Nešetřil, *Linear Programming Duality and Morphisms*, KAM Series 95-298

[7] P. Komárek, *Good charactersizations of graphs*, PhD thesis, Charles University, Prague 1988

[8] L. Lovász, *Operations with structures*, Acta Math. Acad. Sci. Hungar.(1967), 321-328

[9] J. Nešetřil, *Structure of graph homomorphisms*, to appear in J. of Combinatorics, Probability and Computing

[10] J. Nešetřil, A. Pultr, *On classes of relations and graphs determined by subobjects and factorobjects*, Discrete Math. **22** (1978), 287–300.

[11] J. Nešetřil, C. Tardif, *Density via duality*, preprint, 1998.

[12] J. Nešetřil, C. Tardif, *A new proof of the density of directed and undirected graphs*, preprint, 1998.

[13] J. Nešetřil, X. Zhu, *Path homomorphisms*, Math. Proc. Cambridge Philos. Soc. **120** (1996), 207–220.

[14] A. Pultr, V. Trnkova: *Combinatorial, Algebraical and Topological Structures*, North Holland, 1980

[15] C. Tardif, *Fractional multiples of graphs and the density of vertex-transitive graphs*, preprint, to appear in J. Algebraic Combin.

[16] E. Welzl, *Color-families are dense*, J. Theoret. Comput. Sci. **17** (1982), 29-41.

[17] E. Welzl, *Symmetric graphs and interpretations*, J. Combinatorial Theory Ser. B **37** (1984), 235-244.

Jaroslav Nešetřil
Department of Applied Mathematics, Charles University
Malostranské nám. 25, 11800 Praha 1
Czech Republic
nesetril@kam.ms.mff.cuni.cz

Claude Tardif
Department of Mathematics and Statistics

University of Regina
Regina, Saskatchewan
Canada, S4S 0A2.
tardif@math.uregina.ca