

Generalised dualities and maximal finite antichains in the homomorphism order of relational structures^{*}

Jan Foniok¹, Jaroslav Nešetřil¹, and Claude Tardif^{2**}

¹ Department of Applied Mathematics
and

Institute of Theoretical Computer Science (ITI)^{***}
Charles University

Malostranské nám. 25, 118 00 Praha 1, Czech Republic

{foniok,nesetril}@kam.mff.cuni.cz

² Department of Mathematics and Computer Science

Royal Military College of Canada

PO Box 17000, Station 'Forces'

Kingston, Ontario K7K 7B4 Canada

Claude.Tardif@rmc.ca

Abstract. The motivation for this paper is three-fold. First, we study the connectivity properties of the homomorphism order of directed graphs, and more generally for relational structures. As opposed to the homomorphism order of undirected graphs (which has no non-trivial finite maximal antichains), the order of directed graphs has finite maximal antichains of any size. In this paper, we characterise explicitly all maximal antichains in the homomorphism order of directed graphs.

Quite surprisingly, these maximal antichains correspond to generalised dualities. The notion of generalised duality is defined here in the full generality as an extension of the notion of finitary duality, investigated in [17]. Building upon the results of the cited paper, we

^{*} This research was partially supported by the EU Research Training Network COMBSTRU.

^{**} The third author's research is supported by grants from NSERC and ARP.

^{***} The Institute for Theoretical Computer Science is supported as project 1M0021620808 by the Ministry of Education of the Czech Republic.

fully characterise the generalised dualities. It appears that these dualities are determined by forbidding homomorphisms from a finite set of forests (rather than trees).

Finally, in the spirit of [1], [12], [4] we shall characterise “generalised” Constraint Satisfaction Problems (defined also here) problems that are first order definable. These are again just generalised dualities corresponding to finite maximal antichains in the homomorphism order.

1 Introduction

Several classical colouring problems (such as bounding the chromatic number of graphs with given properties) can be treated more generally and sometimes more efficiently in the context of graphs and homomorphisms between them. Recall that, given graphs $G = (V, E)$, $G' = (V', E')$, a *homomorphism* is any mapping $f : V \rightarrow V'$ that preserves edges:

$$xy \in E \Rightarrow f(x)f(y) \in E'.$$

This is denoted by $f : G \rightarrow G'$. For a recent introduction to the topic of graphs and their homomorphisms, we refer the reader to the book [8].

Let H be a fixed graph (sometimes called a template). For an input graph G , the *H-colouring problem* asks whether there exists a homomorphism $G \rightarrow H$. Such a homomorphism is also called an *H-colouring*; the K_k -colouring problem is simply the question whether $\chi(G) \leq k$. Of course, the complexity of the *H-colouring problem* depends on H . This complexity was determined for undirected graphs in [7]. However, already for directed graphs the problem is unsolved.

The *H-colouring problem* is also (and perhaps more often) called the *constraint satisfaction problem (CSP(H))*. This is particularly fitting when the problem is generalised to relational structures and their homomorphisms, as these structures can encode arbitrary constraints. This setting, originally motivated by problems from Artificial Intelligence, leads to the important problem of dichotomy, general heuristic algorithms (consistency check) and, more recently, to an interesting and fruitful algebraic setting (pioneered by Bulatov, Jeavons and Krokhin, cf. [10], [3]).

Further work in the area of CSP complexity led to the following dichotomy conjecture.

Conjecture 1 ([6]). Let H be a finite relational structure. Then $\text{CSP}(H)$ is either solvable in polynomial time or NP-complete.

Some particular instances of CSP were studied intensively. This includes the case when the graphs for which an H -colouring exists are determined by well-described forbidden subgraphs (see [9], [14]) and as a special case, when they are determined by a finite family of forbidden subgraphs. Of course, in these cases we get polynomial instances of CSP.

A pair (F, D) of directed graphs is called a *duality pair* if for every directed graph G , we have $F \rightarrow G$ if and only if $G \not\rightarrow D$. Here, and from now on, $A \rightarrow B$ denotes the fact that there exists a homomorphism from A to B . The duality relationship is denoted by the equation

$$F \rightarrow = \not\rightarrow D$$

where $F \rightarrow$ denotes the class of graphs admitting a homomorphism from F and $\not\rightarrow D$ the class of graphs not admitting a homomorphism to D . The dualities in the category of directed graphs are characterised in [11], [17]:

Theorem 2 ([11], [17]). *Given a directed graph F , there exists a directed graph D_F such that (F, D_F) is a duality pair if and only if F is homomorphically equivalent to an orientation of a tree. For a Δ -tree F , such a Δ -structure D_F is unique up to homomorphism equivalence.*

We say that A and B are homomorphically equivalent if both $A \rightarrow B$ and $B \rightarrow A$.

Here we generalise the notion of a duality pair: for two finite sets of graphs \mathcal{F}, \mathcal{D} , we say that $(\mathcal{F}, \mathcal{D})$ is a *generalised duality* if for any graph G , there exists $F \in \mathcal{F}$ such that $F \rightarrow G$ if and only if $G \rightarrow D$ for no $D \in \mathcal{D}$; briefly

$$\bigcup_{F \in \mathcal{F}} F \rightarrow = \bigcap_{D \in \mathcal{D}} \not\rightarrow D.$$

Building upon the results of [17], we fully characterise the generalised dualities (Section 3). It appears that these dualities are determined by forbidding homomorphisms from a finite set of forests. In particular, we prove that (up to homomorphic equivalence) if $(\mathcal{F}, \mathcal{D})$ is a generalised duality, then \mathcal{F} is a set of forests and \mathcal{D} is uniquely determined by \mathcal{F} . We provide the construction of \mathcal{D} from any finite set of forests \mathcal{F} . In Section 5.2 we show that furthermore \mathcal{F} is also uniquely determined by a possible right-hand side \mathcal{D} .

As a consequence of this characterisation and using a recent result of [12], we are able to show that the decision problem whether for a finite set \mathcal{H}

of graphs there exists \mathcal{F} such that $(\mathcal{F}, \mathcal{H})$ is a generalised duality is NP-complete (Section 5.2).

The relation \rightarrow induces a partial order on the classes of homomorphic equivalence of graphs. This order is called the *homomorphism order*. The homomorphism order is actually a distributive lattice, with the disjoint union of graphs being the supremum and the categorical product being the infimum. (The standard order-theoretic terminology is applied here.)

Particular studied properties of the homomorphism order were density (solved for undirected graphs by Welzl [21] and for directed graphs by Nešetřil and Tardif [17]) and the description of finite maximal antichains (characterised for size 2 by [18]).

The description of generalised dualities shows a surprising link to maximal antichains. In this paper, we show that all finite maximal antichains in the homomorphism order of digraphs are in a 1-1 correspondence with generalised dualities (Section 4): up to finitely many described exceptions, finite maximal antichains are exactly the sets $\mathcal{F} \cup \mathcal{D}$, where $(\mathcal{F}, \mathcal{D})$ is a generalised duality.

Let us note that the problem is hard and captivating for infinite graphs. It has been proved in [16] that for every countable infinite graph G , G not equivalent to K_1 , K_2 , K_ω , there exists a graph H incomparable with G . In this case, infinitely many maximal antichains exist as well, but, as conjectured in [16], all maximal antichains seem to contain a finite graph.

The explicit description of finite maximal antichains allows us to show that it is decidable whether a finite set of directed graphs is a maximal antichain (Section 5.1).

Finally, we extend a recent result of Atserias [1]. We note that the problem whether an input graph is homomorphic to at least one of a finite set \mathcal{H} of graphs is definable by a first-order formula (in the language with equality and adjacency as relational symbols) if and only if the set \mathcal{H} is the right-hand side of a generalised duality (Section 5.4).

We believe that the interplay of order theoretic notions (such as maximal antichain) and descriptive complexity notions (such as generalised duality and first order definability) leads to further insight into the structure of CSP.

2 Preliminaries

2.1 Partial orders

Let $\mathcal{P} = (P, \leq)$ be a poset. We say that a subset Q of P is an *antichain* in \mathcal{P} , if neither $a \leq b$ nor $b \leq a$ for any two distinct elements a, b of Q (such elements are called *incomparable* and the fact is denoted by $a \parallel b$). An antichain Q is *maximal*, if any set S such that $Q \subsetneq S \subseteq P$ is not an antichain. A maximal antichain is also called a *MAC*; a *k-MAC* is a maximal antichain of size k . In this paper we deal only with finite antichains.

2.2 Relational structures

Let $\Delta = (\delta_i; i \in I)$ be a finite sequence of positive integers. A *relational structure of type Δ* (or a Δ -structure) is a pair $A = (V, (R_i; i \in I))$, where V is a nonempty finite set and R_i are relations such that $R_i \subseteq V^{\delta_i}$ for all $i \in I$.

In this way, directed graphs are relational structures for $\Delta = (2)$. Therefore, the set V is usually called the *vertex set* of the Δ -structure, and its elements are called *vertices*; the sets R_i are called *edge sets* and their elements *edges*. When distinguishing edges of distinct edge sets, we usually speak about *colours of edges*.

The notation $V(A)$ and $R_i(A)$ is often used to denote the vertex set and the i -th edge set of a relational structure A , respectively.

The Δ -structure $B = (W, (S_i; i \in I))$ is a *substructure* of $A = (V, (R_i; i \in I))$, if $W \subseteq V$ and $S_i \subseteq R_i \cap W^{\delta_i}$ for all $i \in I$; B is the substructure of A *induced* by W if $S_i = R_i \cap W^{\delta_i}$ for all $i \in I$.

The *incidence graph* $\text{Inc}(A)$ of a Δ -structure A is the bipartite multigraph $(V_1 \cup V_2, E)$ with parts $V_1 = V(A)$ and

$$V_2 = \text{Block}(A) := \{(i, (a_1, \dots, a_{\delta_i})) : i \in I, (a_1, \dots, a_{\delta_i}) \in R_i(A)\},$$

and one edge between a and $(i, (a_1, \dots, a_{\delta_i}))$ for each occurrence of a as some a_k in an edge $(a_1, \dots, a_{\delta_i}) \in R_i(A)$.

A Δ -structure A is *connected* if $\text{Inc}(A)$ is connected; a *connected component* of A is each substructure induced by all the vertices of A in a connected component of $\text{Inc}(A)$.

A Δ -structure A is called a Δ -*tree* if $\text{Inc}(A)$ is a tree. Note that A is not a Δ -tree if multiple edges appear in $\text{Inc}(A)$, i.e. if a vertex appears in an edge of A more than once. The structure A is called a Δ -*forest* if all its connected components are Δ -trees.

2.3 Homomorphisms

Let $A = (V, (R_i; i \in I))$ and $B = (W, (S_i; i \in I))$ be two relational structures of the same type Δ . A function $f : V \rightarrow W$ is a *homomorphism* from A to B , if for any $i \in I$ and any edge $(v_1, v_2, \dots, v_{\delta_i}) \in R_i$ the δ_i -tuple $(f(v_1), f(v_2), \dots, f(v_{\delta_i}))$ is in S_i . We write $f : A \rightarrow B$.

We say that A is *homomorphic* to B and write $A \rightarrow B$, if there exists a homomorphism $f : A \rightarrow B$. The fact that A is not homomorphic to B is denoted by $A \not\rightarrow B$. If $A \rightarrow B$ and $B \rightarrow A$, we say that A and B are *hom-equivalent* and write $A \sim B$. Notice that this is by far not the same as being isomorphic; e.g. all directed graphs with a loop are pairwise hom-equivalent.

It is easy to see that the binary relation \rightarrow on the class of all relational structures of a fixed type Δ is reflexive (because the identity function is a homomorphism) and transitive (because the composition of homomorphisms is a homomorphism).

A relational structure A is called a *core* if it is not homomorphic to any of its proper substructures.

The following is a well-known fact (see, e.g. [8]).

Lemma 3. *Any relational structure G is hom-equivalent to a unique core C (up to isomorphism).*

Thus we can usually restrict our attention to cores without loss of generality.

As a consequence, we get that the set of all (non-isomorphic) cores with the relation \rightarrow is a partially ordered set, denoted by $\mathcal{C}(\Delta)$; we speak of the *homomorphism order* of relational structures.

We keep the slightly unusual notation $A \rightarrow B$ instead of the more common $A \leq B$ for the homomorphism partial order. Where convenient, however, we use $A < B$ to denote that $A \rightarrow B$ and at the same time $B \not\rightarrow A$.

Let A be a Δ -structure. If there exists a Δ -tree T such that $A \rightarrow T$, we say that A is *balanced*. It is evident that A is balanced if and only if it is homomorphic to a Δ -forest.

2.4 Sums and products

For a finite nonempty set $\mathcal{Q} = \{Q_1, Q_2, \dots, Q_t\}$ of Δ -structures (of the same type Δ), we define the *sum*

$$S = \sum_{j=1}^t Q_j = Q_1 + Q_2 + \dots + Q_t$$

to be the disjoint union of the structures in \mathcal{Q} . We define the *product*

$$P = \prod_{j=1}^t Q_j = Q_1 \times Q_2 \times \dots \times Q_t$$

to be the structure with

$$\begin{aligned} V(P) &= V(Q_1) \times V(Q_2) \times \dots \times V(Q_t), \\ R_i(P) &= \{((v_{1,1}, \dots, v_{1,t}), \dots, (v_{\delta_i,1}, \dots, v_{\delta_i,t})) : \\ &\quad (v_{1,j}, \dots, v_{\delta_i,j}) \in R_i(Q_j), j = 1, \dots, t\}, i \in I. \end{aligned}$$

The sum and product defined in this way are the sum and product in terms of category theory (see e.g. [8], [2]). In particular, the homomorphism order of Δ -structures is a distributive lattice, with the product of two structures being the infimum and the sum being the supremum.

2.5 Homomorphism dualities

A pair of Δ -structures (F, D) is a *duality pair* if for every Δ -structure X , there exists a homomorphism $X \rightarrow D$ if and only if there exists no homomorphism $F \rightarrow X$.

The following theorem, which provides a characterisation of homomorphism dualities, is one of our starting motivations.

Theorem 4 ([17]). *For a Δ -structure F there exists a Δ -structure D such that (F, D) is a duality pair if and only if F is a Δ -tree. For a Δ -tree F , such a Δ -structure D is unique up to homomorphism equivalence.*

The unique core D such that (F, D) is a duality pair is called *the dual* of the Δ -tree F . We use the notation $D = D(F)$.

In the following, we will need that every dual is *irreducible*, i.e. whenever $X \times Y \rightarrow D$, then $X \rightarrow D$ or $Y \rightarrow D$. (Such a structure is sometimes called multiplicative or productive ([14]), but the word “irreducible” seems to be more appropriate here.)

Lemma 5. *If (F, D) is a duality pair and $X \times Y \rightarrow D$, then $X \rightarrow D$ or $Y \rightarrow D$.*

Proof. If $X \not\rightarrow D$ and $Y \not\rightarrow D$, then $F \rightarrow X$ and $F \rightarrow Y$, hence $F \rightarrow X \times Y$ and so $X \times Y \not\rightarrow D$. \square

Remark 6. Irreducibility is the dual property to connectedness: a Δ -structure A is connected if and only if whenever $A \rightarrow X + Y$, we have $A \rightarrow X$ or $A \rightarrow Y$. Note that in a duality pair (F, D) , the structure F is connected and D is irreducible.

For Δ -structures F and D , let $(F \rightarrow)$ denote the set $\{X : F \rightarrow X\}$ and let $(\rightarrow D)$ denote the set $\{X : X \rightarrow D\}$. The sets $(F \leftrightarrow)$ and $(\leftrightarrow D)$ are defined analogously.

From now on, let $\mathcal{F} = \{F_1, F_2, \dots, F_m\}$, $\mathcal{D} = \{D_1, D_2, \dots, D_p\}$ be two sets of Δ -structures which are all cores, and let $\mathcal{F}_c = \{C_1, C_2, \dots, C_n\}$ be the set of all distinct connected components of the Δ -structures in \mathcal{F} .

We say that the pair $(\mathcal{F}, \mathcal{D})$ is a *generalised duality* if $F_i \parallel F_{i'}$ for $i \neq i'$, $D_k \parallel D_{k'}$ for $k \neq k'$ and

$$\bigcap_{i=1}^m (F_i \leftrightarrow) = \bigcup_{k=1}^p (\rightarrow D_k).$$

The special case $p = 1$ is characterised by the following theorem proved in [17].

Theorem 7 ([17]). *Let $\mathcal{F} = \{F_1, F_2, \dots, F_m\}$ be a finite nonempty set of Δ -structures. The pair $(\mathcal{F}, \{D\})$ is a generalised duality if and only if D is homomorphically equivalent to $\prod_{i=1}^m D_i$ and (F_i, D_i) is a duality pair for $i = 1, 2, \dots, m$.*

When $p = |\mathcal{D}| = 1$, the generalised duality $(\mathcal{F}, \mathcal{D})$ is also called a *finitary homomorphism duality* in [17]. The theorem states that the finitary dual is the product of the duals of the Δ -trees F_1, \dots, F_m ; this product will be denoted by $D(F_1, \dots, F_m)$ or $D(M)$ if $M = \{F_1, \dots, F_m\}$.

We can also consider the case $\mathcal{F} = \emptyset$. Then $\mathcal{D} = \{\mathbf{1}\}$, where $\mathbf{1}$ is a single vertex with all loops of all arities, i.e. $\mathbf{1} = (V, (V^{\delta_i} : i \in I))$ and V is the one-element set $\{1\}$.

3 Generalised dualities

In this section, we characterise all generalised dualities. We restrict ourselves to the case $|\mathcal{F}| \geq 2$, as the other cases are described in the previous section. First, we present a construction of generalised dualities from a family of forests (rather than just trees).

3.1 The construction

Let $\mathcal{F} = \{F_1, F_2, \dots, F_m\}$ be an arbitrary fixed nonempty finite set of core Δ -forests that are pairwise incomparable (in $\mathcal{C}(\Delta)$).

Consistently with the above notation, let $\mathcal{F}_c = \{C_1, \dots, C_n\}$ be the set of all distinct connected components of the structures in \mathcal{F} ; each of these components is a core Δ -tree.

Remark 8. All incomparable sets of core forests can be constructed (some-what more explicitly) in the following way. (This also shows interesting relationships among trees, forests and the homomorphism order.)

Let $(\mathcal{T}, \rightarrow)$ be the suborder of the homomorphism order induced on the class \mathcal{T} of all core Δ -trees and let \mathcal{A} be the set of all nonempty finite antichains in $(\mathcal{T}, \rightarrow)$. The set \mathcal{A} is in a 1-1 correspondence with core Δ -forests.

We define a binary relation \trianglelefteq on \mathcal{A} : for $A, A' \in \mathcal{A}$, we have $A \trianglelefteq A'$ if and only if for each Δ -structure $T \in A$ there exists a Δ -structure $T' \in A'$ such that $A \rightarrow A'$. Obviously, \trianglelefteq is a partial order on \mathcal{A} . It can be seen that \trianglelefteq is isomorphic to the homomorphism order $\mathcal{C}(\Delta)$ restricted to the set of all Δ -forests.

Let \mathcal{N} be an arbitrary nonempty finite subset of \mathcal{A} that is an antichain with respect to \trianglelefteq , i.e. $N \trianglelefteq N'$ for no distinct elements $N, N' \in \mathcal{N}$. This condition expresses the fact that all the forests in the set \mathcal{F} are pairwise incomparable.

Suppose $\mathcal{N} = \{N_1, N_2, \dots, N_m\}$ and set

$$F_i = \sum_{T \in N_i} T,$$

$\mathcal{F} = \{F_1, F_2, \dots, F_m\}$. Then

$$\mathcal{F}_c = \bigcup_{i=1}^m N_i = \{C_1, \dots, C_n\}.$$

Thus we can see that core forests are in a 1-1 correspondence with antichains in the set \mathcal{T} of all core Δ -trees.

A subset $M \subseteq \mathcal{F}_c$ is a *quasitransversal* if it satisfies

(T1) M is an antichain, i.e. for every $C \neq C' \in M$ we have $C \parallel C'$, and

(T2) M supports \mathcal{F} , i.e. for every $F \in \mathcal{F}$ there exists $C \in M$ such that $C \rightarrow F$.

For two quasitransversals M, M' we define $M \preceq M'$ if and only if for every $C' \in M'$ there exists $C \in M$ such that $C \rightarrow C'$. Note that this order is different from the homomorphism order of forests corresponding to the quasitransversals. On the other hand, we have:

Lemma 9. *Let M, M' be two quasitransversals. Then the dual structures $D(M)$ and $D(M')$ exist and $D(M) \rightarrow D(M')$ if and only if $M \preceq M'$.*

Proof. By Theorem 7, the dual structures $D(M)$ and $D(M')$ exist and

$$D(M) = \prod_{C \in M} D(C), \quad D(M') = \prod_{C' \in M'} D(C').$$

By the infimum property of the product, it suffices to show that $D(M) \rightarrow D(C')$ for any $C' \in M'$. So, let $C' \in M'$. Because $M \preceq M'$, there exists $C \in M$ such that $C \rightarrow C'$. By the definition of a duality pair, $C \rightarrow C'$ implies that $C' \nrightarrow D(C)$ and this implies that $D(C) \rightarrow D(C')$. We conclude that $D(M) \rightarrow D(C) \rightarrow D(C')$. \square

Lemma 10. *The relation \preceq is a partial order on the set of all quasitransversals.*

Proof. Obviously, \preceq is both reflexive and transitive. Suppose now that $M \preceq M'$ and $M' \preceq M$, and let $B \in M$. Then there exists $B' \in M'$ such that $B' \rightarrow B$ and there exists $B'' \in M$ such that $B'' \rightarrow B'$. Consequently $B'' \rightarrow B$, hence by (T1) we have $B = B' = B''$, so $M \subseteq M'$. Similarly we get that $M' \subseteq M$. \square

A quasitransversal M is a *transversal* if

(T3) M is a maximal quasitransversal in \preceq .

Set $\mathcal{D} = \mathcal{D}(\mathcal{F}) = \{D(M) : M \text{ is a transversal}\}$.

We have:

Theorem 11. *The pair $(\mathcal{F}, \mathcal{D})$ is a generalised duality.*

Before presenting the proof, we illustrate the construction by three examples.

Example. First, let $\mathcal{F} = \{T_1, T_2, \dots, T_n\}$ be a set of pairwise incomparable trees and D_1, D_2, \dots, D_n their respective duals. By (T2), every transversal contains all these trees. Therefore there exists only one transversal $M = \{T_1, T_2, \dots, T_n\}$ and $\mathcal{D} = \{D(M)\} = \{D_1 \times D_2 \times \dots \times D_n\}$. This case also shows that the finitary duality is a special case of the generalised duality.

Now, let T_1, T_2, T_3 and T_4 be pairwise incomparable trees with duals D_1, D_2, D_3, D_4 . Let $\mathcal{F} = \{T_1 + T_2, T_1 + T_3, T_4\}$. Then we have two transversals $\{T_1, T_4\}$ and $\{T_2, T_3, T_4\}$; and $\mathcal{D} = \{D_1 \times D_4, D_2 \times D_3 \times D_4\}$.

Finally, let $T_1 \rightarrow T_3$ and $\mathcal{F} = \{T_1 + T_2, T_3 + T_4\}$. The transversals are $\{T_1\}$, $\{T_2, T_3\}$ and $\{T_2, T_4\}$. Hence $\mathcal{D} = \{D_1, D_2 \times D_3, D_2 \times D_4\}$.

Proof of Theorem 11. By the definition of \mathcal{F} , the Δ -forests F_i and $F_{i'}$ are incomparable in $\mathcal{C}(\Delta)$ for any $i \neq i'$. Any two distinct elements of \mathcal{D} are incomparable, because any two transversals are incomparable with respect to \preceq (they are all maximal in this order) and because of Lemma 9.

Let X be a Δ -structure such that $X \rightarrow D$ for some $D \in \mathcal{D}$. We want to prove that $F_i \rightarrow X$ for $i = 1, \dots, m$. For contradiction, assume that $F_i \rightarrow X$ for some i . Let M be the transversal for which $D(M) = D$. By (T2), there exists $C \in M$ such that $C \rightarrow F_i \rightarrow X$, therefore $X \rightarrow D(C)$. This is a contradiction with the assumption that $X \rightarrow D \rightarrow D(C)$.

Now, let X be a Δ -structure such that $F_i \rightarrow X$ for $i = 1, \dots, m$. We want to prove that there exists $D \in \mathcal{D}$ such that $X \rightarrow D$. Let C_{j_i} be a component of F_i such that $C_{j_i} \rightarrow X$ for $i = 1, \dots, m$. Let $M' = \min_{\rightarrow} \{C_{j_i} : i = 1, \dots, m\}$, where by $\min_{\rightarrow} S$ we mean the set of all elements of S that are minimal with respect to the homomorphism order \rightarrow . Because M' is a quasitransversal, there exists a transversal M such that $M' \preceq M$. We have that $C \rightarrow X$ for each $C \in M$, and thus $X \rightarrow D(M) \in \mathcal{D}$. \square

3.2 The characterisation

We will now prove that actually all generalised dualities are of the form presented in Section 3.1.

Theorem 12. *If $(\mathcal{F}, \mathcal{D})$ is a generalised duality, then all elements of \mathcal{F} are forests and $\mathcal{D} = \mathcal{D}(\mathcal{F})$; in particular, \mathcal{D} is uniquely determined by \mathcal{F} .*

Proof. We split the proof into five steps. Suppose that $\mathcal{F} = \{A_1, A_2, \dots, A_m\}$ and $\mathcal{D} = \{D_1, D_2, \dots, D_p\}$. We assume that all the structures in \mathcal{F} and also all the structures in \mathcal{D} are pairwise incomparable cores. Consistently

with the above notation, let $\mathcal{F}_c = \{C_1, C_2, \dots, C_n\}$ be the set of all distinct connected components of the structures in \mathcal{F} . Quasitransversals and transversals are defined in the same way as above; notice that neither for the definition nor for proving Lemma 10 we needed the fact that the elements of \mathcal{F}_c are trees.

For a quasitransversal M , let $\overline{M} = \{C' \in \mathcal{F}_c : C \in M \Rightarrow C \not\rightarrow C'\}$.

Fact 1. *If $M \subseteq \mathcal{F}_c$ is a transversal, then there exists a unique Δ -structure $D \in \mathcal{D}$ that satisfies*

- (1) $C \rightarrow D$ for every $C \in M$,
- (2) $C' \rightarrow D$ for every $C' \in \overline{M}$.

Proof. If $\overline{M} = \emptyset$, let $D \in \mathcal{D}$ be arbitrary. Otherwise set $S = \sum_{C' \in \overline{M}} C'$. Because $(\mathcal{F}, \mathcal{D})$ is a generalised duality, either there exists $F \in \mathcal{F}$ such that $F \rightarrow S$ or there exists $D \in \mathcal{D}$ such that $S \rightarrow D$. If $F \rightarrow S$, by (T2) some $C \in M$ satisfies $C \rightarrow F \rightarrow S$, and since C is connected, $C \rightarrow C'$ for some $C' \in \overline{M}$, which is a contradiction with the definition of \overline{M} . Therefore there exists $D \in \mathcal{D}$ that satisfies $S \rightarrow D$.

Obviously, D satisfies (2).

Let $C \in M$ such that $C \rightarrow D$. Consider $M' = M \setminus \{C\}$. M' is not a quasitransversal, because otherwise we would have $M \prec M'$ and M would not satisfy (T3). Hence M' fails to satisfy (T2), and we can find $A \in \mathcal{F}$ which is not supported by M' . It follows that $C \rightarrow A$.

Consider Q' , the set of all elements of \mathcal{F} that are not supported by M' . We know that Q' is nonempty because $A \in Q'$.

There exists $A' \in Q'$ such that C is a connected component of A' : otherwise let M^* be the set of all components C^* of Δ -structures in Q' such that $C \rightarrow C^*$, and let $M'' = \min_{\rightarrow}(M' \cup M^*)$ be the set of all structures in the union of M' and M^* that are minimal with respect to the homomorphism order $\mathcal{C}(\Delta)$. The set M'' is a quasitransversal but $M \prec M''$, contradicting the fact that M is a transversal.

All the components of A' are elements of $\overline{M} \cup \{C\}$. The assumption that $C \rightarrow D$ leads, using (2), to the conclusion that $A' \rightarrow D$. That is a contradiction with the definition of generalised duality.

It remains to prove uniqueness: If $D, D' \in \mathcal{D}$ both satisfy (1) and (2) and $D \neq D'$, i.e. $D \parallel D'$, then $D + D'$ violates the definition of generalised duality. \square

For a transversal M , the unique $D \in \mathcal{D}$ satisfying the conditions (1) and (2) above is denoted by $d(M)$.

Fact 2. $\mathcal{D} = \{d(M) : M \text{ is a transversal}\}$.

Proof. Let $D \in \mathcal{D}$. We want to show that $D = d(M)$ for a transversal M . Let $M' = \min_{\rightarrow} \{C' \in \mathcal{F}_c : C' \rightarrow D\}$ be the set of all $\mathcal{C}(\Delta)$ -minimal components that are not homomorphic to D . The set M' is a quasitransversal: if some $F \in \mathcal{F}$ is not supported by M' , then all its components are homomorphic to D , and also $F \rightarrow D$, a contradiction.

Let M be a transversal such that $M' \preceq M$. To prove that $D = d(M)$, it suffices (by the uniqueness part of Fact 1) to check conditions (1) and (2).

If $C \in M$, then there exists $C' \in M'$ such that $C' \rightarrow C$. Therefore $C \rightarrow D$, so the condition (1) is satisfied.

Now suppose there exists $\check{C} \in \overline{M}$ such that $\check{C} \rightarrow D$. Consider the Δ -structure $X = \check{C} + D$. If $A \rightarrow X$ for some $A \in \mathcal{F}$, then by the property (T2) of M there exists $C \in M$ that is homomorphic to A . But since $\check{C} \in \overline{M}$, we have that $C \rightarrow \check{C}$, hence $C \rightarrow D$. This is a contradiction with condition (1). It follows that $X \rightarrow \check{D}$ for some $\check{D} \in \mathcal{D}$, hence $D \rightarrow \check{D}$, so $D = \check{D}$. That is a contradiction with $\check{C} \rightarrow D$ and $\check{C} \rightarrow \check{D}$. \square

Fact 3. For two distinct transversals M_1, M_2 , we have (a) $\overline{M_1} \cap M_2 \neq \emptyset$, (b) $d(M_1) \neq d(M_2)$.

Proof.

(a) By (T3), $M_1 \not\preceq M_2$, and therefore there exists $C_2 \in M_2$ such that $C_1 \rightarrow C_2$ for any $C_1 \in M_1$. Obviously $C_2 \in \overline{M_1} \setminus \overline{M_2} \subseteq \overline{M_1}$. Since we selected $C \in M_2$, we have that $C \in \overline{M_1} \cap M_2$.

(b) Let $C_2 \in \overline{M_1} \cap M_2$, as above. Then $C_2 \rightarrow d(M_1)$ and $C_2 \rightarrow d(M_2)$. \square

Fact 4. If M is a transversal, then the pair $(M, \{d(M)\})$ is a finitary homomorphism duality, and consequently $d(M) = D(M)$.

Proof. We want to prove that

$$\bigcap_{C \in M} (C \rightarrow) = (\rightarrow d(M)).$$

We claim that for a Δ -structure G , the following statements are equivalent:

(1) $G \in \bigcap_{C \in M} (C \rightarrow)$

- (2) $C \dashv G$ for any $C \in M$
- (3) $C \dashv G + \sum_{\check{C} \in \overline{M}} \check{C}$ for any $C \in M$
- (4) $G + \sum_{\check{C} \in \overline{M}} \check{C} \rightarrow d(M)$
- (5) $G \rightarrow d(M)$
- (6) $G \in (\rightarrow d(M))$

Because: (1) \Leftrightarrow (2) and (5) \Leftrightarrow (6) by definition. (4) \Rightarrow (5) immediately. (5) \Rightarrow (2) by Fact 1(1). (2) \Rightarrow (3) follows from the definition of \overline{M} and the fact that C is connected.

It remains to prove that (3) \Rightarrow (4): Let $X = G + \sum_{\check{C} \in \overline{M}} \check{C}$. If $A \rightarrow X$ for some $A \in \mathcal{F}$, then by (T2) there exists $C \in M$ such that $C \rightarrow A \rightarrow X$, a contradiction. Thus no element of \mathcal{F} is homomorphic to X , hence $X \rightarrow D$ for some $D \in \mathcal{D}$. By Fact 2, $D = d(M')$ for a transversal M' ; by Fact 1 and Fact 3(a), $M' = M$.

The equivalence (1) \Leftrightarrow (6) is precisely the definition of finitary duality. \square

Fact 4 and Theorem 4 imply that any element of a transversal is a Δ -tree. In fact, we have:

Fact 5. *Each component $C \in \mathcal{F}_c$ is a tree.*

For the proof, we will need the following density result:

Theorem 13 ([17]). *Let A and C be relational structures such that $A < C$, and C is a connected structure that is not a tree. Then there exists a structure X such that $A < X < C$.*

Proof of Fact 5. Suppose that $C \in \mathcal{F}_c$ is not a tree. By Fact 4 and Theorem 4, C is an element of no transversal. Set

$$A = \sum_{\substack{C' \in \mathcal{F}_c \\ C' < C}} C' + \sum_{\substack{C' \in \mathcal{F}_c \\ C' \parallel C}} (C \times C').$$

Clearly, $A < C$. Let X be a structure such that $A < X < C$, as in Theorem 13. Then for any $C' \neq C \in \mathcal{F}_c$, we have $C' \rightarrow X$ if and only if $C' \rightarrow C$ and $X \rightarrow C'$ if and only if $C \rightarrow C'$. Indeed: if $C' \rightarrow C$, then $C' \rightarrow A \rightarrow X$; if $C \rightarrow C'$, then $X \rightarrow C \rightarrow C'$. On the other hand, if $C \parallel C'$, then $X \rightarrow C'$ implies $X \rightarrow A$, a contradiction with $A < X$. Moreover $C' \rightarrow X$ implies $C' \rightarrow C$.

Let $F \in \mathcal{F}$ be such that C is a component of F and let G be the structure obtained from F by replacing C with X .

We have that $F \dashv G$ since $C \dashv G$, because otherwise F would not be a core. In addition, $F' \dashv G$ for any $F \neq F' \in \mathcal{F}$, because $F' \rightarrow G$ implies $F' \rightarrow F$. Therefore $G \rightarrow D$ for some $D \in \mathcal{D}$. Let M be the transversal such that $D = D(M)$. Note that we assume $C \notin M$. The structure D is a finitary dual and hence $C' \dashv G$ for any $C' \in M$; therefore $C' \dashv X$ and $C' \dashv C$ for any $C' \in M$. Consequently $C \rightarrow D$. We know that all components of G are homomorphic to D , so all components of F are homomorphic to D as well. We conclude that $F \rightarrow D$, a contradiction. \square

Now we can finish the proof of Theorem 12: All elements of \mathcal{F} are forests by virtue of Fact 4, Fact 5 and Theorem 7. The set \mathcal{D} is uniquely determined as a consequence of Fact 2 and owing to Fact 4 and Theorem 7 it is determined by the transversal construction. \square

Remark 14. We used the fact that every Δ -structure can be uniquely expressed as the sum of its connected components. It is not true in general that every Δ -structure can be expressed as the product of a finite number of irreducible (multiplicative) Δ -structure, i.e. atoms in the homomorphism order lattice.

However, if \mathcal{D} is a possible right-hand side of a generalised duality, we know that each of its elements is a product of atoms. This allows to construct \mathcal{F} from \mathcal{D} . We provide the construction in Section 5.2.

4 Finite maximal antichains

4.1 Maximal antichains of size 1

An earlier result characterises all 1-MACs in the homomorphism order of directed graphs.

Theorem 15 ([20]). *The only maximal antichains of size 1 in the homomorphism order of directed graphs are directed paths of length 0, 1, and 2 and a single vertex with a loop.*

Here, we present the characterisation of all 1-MACs for relational structures of other types. We prove that only trivial 1-MACs exist for relational structures with $|\Delta| > 1$. Note that a *loop* is an edge in the form (x, x, \dots, x) for a vertex x .

Theorem 16. *Let $t \geq 2$, $\Delta = (\delta_1, \delta_2, \dots, \delta_t)$ and $\delta_1 \geq 2$. Then the only two cores that form maximal antichains of size one in $\mathcal{C}(\Delta)$ are a vertex with no edges $\mathbf{0}$ and a vertex with all loops $\mathbf{1}$.*

Proof. Clearly, $\mathbf{0} = (\{v\}, (\emptyset, \emptyset, \dots, \emptyset))$ is the least and $\mathbf{1} = (\{v\}, (V^{\delta_1}, V^{\delta_2}, \dots, V^{\delta_t}))$ the greatest element of $\mathcal{C}(\Delta)$: for any $X \in \mathcal{C}(\Delta)$, we have $\mathbf{0} \rightarrow X$ and $X \rightarrow \mathbf{1}$.

Let $A = (V, (R_1, R_2, \dots, R_t))$ be a Δ -structure. We need to show that unless $A \sim \mathbf{0}$ or $A \sim \mathbf{1}$, there is a Δ -structure B such that $A \parallel B$. Thus suppose that $A \not\sim \mathbf{0}$ and $A \not\sim \mathbf{1}$.

First, if there exists $1 \leq i \leq t$ such that $R_i = \emptyset$, then the Δ -structure B with $V(B) = \{u\}$, $R_i(B) = \{u\}^{\delta_i}$ and $R_j(B) = \emptyset$ for all $j \neq i$, is incomparable with A . The same Δ -structure is incomparable with A if A has edges of all colours but there is no loop $(a, a, \dots, a) \in R_i$.

Now suppose that A has loops of all colours. As $A \not\sim \mathbf{1}$, no vertex in $V(A)$ has all loops. Let M be the set of vertices that have loops in all colours $2, 3, \dots, t$, i.e.

$$M = \{u \in V : (u, u, \dots, u) \in R_i \text{ for all } i \in \{2, 3, \dots, t\}\},$$

and let $m = |M|$. Further let

$$\begin{aligned} V(B) &= \{0, 1, 2, \dots, m\}, \\ R_1(B) &= \{(a, b, b, \dots, b) : 0 \leq a < b \leq m\}, \\ R_i(B) &= V(B)^{\delta_i} \text{ for } 2 \leq i \leq t. \end{aligned}$$

We know that $A \not\rightarrow B$ because no loop appears in $R_1(B)$. If a homomorphism $f : B \rightarrow A$ existed, it would have to map all vertices of B to the subset $M \subseteq V(A)$. The mapping f cannot be injective, because $|V(B)| > |M|$. Therefore there exist $u, v \in V(B)$ such that $u < v$ and $f(u) = f(v) = x \in M$. As $(u, v, v, \dots, v) \in R_1(B)$ and f is a homomorphism, $(f(u), f(v), f(v), \dots, f(v)) = (x, x, \dots, x) \in R_1(A)$ and the vertex x has all loops, contradicting the fact that A is not hom-equivalent to $\mathbf{1}$. \square

For relational structures with one relation, all 1-MACs are characterised by the following theorem. For $\Delta = (k)$, we define the Δ -structure P_1 by $V(P_1) = \{1, 2, \dots, k\}$ and $E(P_1) = \{(1, 2 \dots, k)\}$ (a single edge). Similarly to the previous case, let $\mathbf{0}$ have a single vertex and no edges and $\mathbf{1}$ have one vertex v and a loop (v, v, \dots, v) .

Theorem 17. *Let $\Delta = (k)$ and $k \geq 3$. Let A be a core structure of type Δ that forms a 1-MAC in $\mathcal{C}(\Delta)$. Then A is one of $\mathbf{0}$, P_1 and $\mathbf{1}$.*

Proof. It is obvious that each of $\mathbf{0}$, P_1 and $\mathbf{1}$ form a 1-MAC. We shall show that if A is not hom-equivalent to one of these three structures, there exists a structure incomparable with A .

For $i = 1, 2, \dots, k-1$ let L_i be the Δ -structure with $V(L_i) = \{1, 2, \dots, k-1\}$ and $E(L_i) = \{(1, 2, \dots, i-1, i, i, i+1, \dots, k-1)\}$. Notice that $L_i \not\rightarrow L_j$ for any $i \neq j$.

First suppose that there exists $i \in \{1, 2, \dots, k-1\}$ such that $L_i \not\rightarrow A$. If $A \not\rightarrow L_i$, then A and L_i are incomparable. Otherwise, $L_j \not\rightarrow A$ for any j , as $L_j \rightarrow A \rightarrow L_i$ is a contradiction. If $A \not\rightarrow L_1$ or $A \not\rightarrow L_{k-1}$, we are done. If not, let $f : A \rightarrow L_1$ and $g : A \rightarrow L_{k-1}$. Define C to be the structure with

$$\begin{aligned} V(C) &= \{1, 2, \dots, k-1\}^2, \\ E(C) &= \{(\langle 1, 1 \rangle, \langle 1, 2 \rangle), \langle 2, 3 \rangle, \dots, \langle k-2, k-1 \rangle, \langle k-1, k-1 \rangle\}; \end{aligned}$$

C consists of an edge and several isolated vertices, and thus it is hom-equivalent to P_1 . Set $h : V(A) \rightarrow V(C)$, $h(u) = \langle f(u), g(u) \rangle$. Clearly h is a homomorphism of A to C , so $A \rightarrow P_1$, and therefore $A \sim \mathbf{0}$ or $A \sim P_1$.

Now assume that all $L_i \rightarrow A$. We use an argument similar to the proof of Theorem 16. For $i = 1, 2, \dots, k-1$ define $V_i \subseteq V(A)$ to be the set of homomorphic images of the i -th vertex in L_i , i.e.

$$V_i = \{u \in V(A) : f(i) = u \text{ for a homomorphism } f : L_i \rightarrow A\}.$$

Let

$$M = \bigcap_{i=2}^{k-1} V_i.$$

Let $m = |M|$, and let B be the structure with

$$\begin{aligned} V(B) &= \{0, 1, 2, \dots, m+1\}, \\ E(B) &= \{(a, b, b, \dots, b) : 0 \leq a < b \leq m+1\}. \end{aligned}$$

Clearly $L_1 \not\rightarrow B$, hence $A \not\rightarrow B$. If $B \not\rightarrow A$, we have found a structure incomparable with A .

Let $f : B \rightarrow A$. By the definition of M and B , the homomorphism f maps the vertices $1, 2, \dots, m+1$ to M , so it must identify some two of them; say $f(a) = f(b) = u$, $1 \leq a < b \leq m+1$. The edge (a, b, b, \dots, b) of B is mapped by f to the loop (u, u, \dots, u) of A , and hence $A \sim \mathbf{1}$. \square

4.2 Generalised dualities as maximal antichains

In this section, let $\Delta = (k)$ be a fixed type, $k \geq 2$, i.e. we only consider relational structures with just one relation that is not unary. (Note that for structures with one unary relation, only two cores exist and the homomorphism order is isomorphic to the total order on a 2-element set.)

This section is motivated by the following:

Theorem 18 ([18]). *The 2-MACs in the homomorphism order of directed graphs are precisely the pairs $\{T, D_T\}$, where T is a core tree different from P_0 , P_1 and P_2 , and D_T is its dual.*

First, we discuss when a generalised duality forms a maximal antichain; precisely, for what families \mathcal{F} of incomparable forests is $\mathcal{Q} = \mathcal{F} \cup \mathcal{D}(\mathcal{F})$ a maximal antichain in the homomorphism order of Δ -structures.

Obviously, if a generalised duality forms an antichain, then it is maximal. It is also evident that $F \not\rightarrow D$ for any $F \in \mathcal{F}$, $D \in \mathcal{D}$. So, a generalised duality does not form an antichain if and only if there exist $D \in \mathcal{D}$ and $F \in \mathcal{F}$ such that $D \rightarrow F$.

Let $P_1 = (\{1, 2, \dots, k\}, \{(1, 2, \dots, k)\})$ be the Δ -structure consisting of a single edge. If $P_1 \in \mathcal{F}_c$, then obviously $\mathcal{F} = \{P_1\}$ and $\mathcal{D} = \{\mathbf{0}\}$. So for the rest, we may assume that $P_1 \notin \mathcal{F}_c$.

Let $\mathcal{S} = \{T_1, T_2, \dots, T_q\}$ be the set of all core Δ -trees with two edges.

Lemma 19. *Let \mathcal{F} be a set of pairwise incomparable core Δ -forests, $\mathcal{F} \neq \{\mathbf{0}\}$, $\mathcal{F} \neq \{P_1\}$. Then $\mathcal{F} \cup \mathcal{D}(\mathcal{F})$ is not an antichain if and only if \mathcal{F} is the set \mathcal{S} of all core Δ -trees with two edges.*

Proof. We have just observed that if $\mathcal{F} \cup \mathcal{D}(\mathcal{F})$ is not an antichain, then there exist $D \in \mathcal{D}$ and $F \in \mathcal{F}$ such that $D \rightarrow F$. Fix such F and D .

Since F is a Δ -forest, we have that D is balanced. Moreover, by Theorem 12, $D = D(M)$ for a transversal $M \subseteq \mathcal{F}_c$.

Let $T \in \mathcal{S}$; T has two edges, $e_1 = (u_1, \dots, u_k)$ and $e_2 = (v_1, \dots, v_k)$. Without loss of generality, we may assume that u_1 is not a vertex of e_2 and v_k is not a vertex of e_1 . For a positive integer s , we define the *thunderbolt* $Z_{T,s}$ to be the structure constructed from T by adding a path with $2s - 1$ edges (a zigzag) and by identifying its first vertex with u_1 and its last vertex with v_k , see Fig. 1.

Notice in particular, that any proper substructure of $Z_{T,s}$ is homomorphic to T .




Fig. 1. An example of thunderbolts for structures with a ternary relation

The thunderbolts are not balanced, hence for any s and T , $Z_{T,s} \not\rightarrow D$.

Let $T \in \mathcal{S}$ be fixed now. Because of the finitary duality of M and D and the fact that $Z_{T,s} \not\rightarrow D$, for every positive integer s there exists $C \in M$ such that $C \rightarrow Z_{T,s}$. Therefore some $C \in M$ is homomorphic to $Z_{T,s}$ for infinitely many values of s . Consequently C is homomorphic to some proper substructure of $Z_{T,s}$ and thus it is homomorphic to T . Since T has only two edges and $C \not\rightarrow P_1$, we have that $C = T$. Applying the argument to all $T \in \mathcal{S}$, we get that $M = \mathcal{S}$. Accordingly, $D = D(\mathcal{S}) = P_1$.

Seeing that $P_1 \in \mathcal{D}$ and that by definition \mathcal{D} is an antichain, we get $\mathcal{D} = \{P_1\}$. Hence there exists only one transversal, and that is \mathcal{S} . That is only possible if $\mathcal{F} = \mathcal{S}$.

For the other implication, if $\mathcal{F} = \mathcal{S}$, then $\mathcal{D}(\mathcal{F}) = \{P_1\}$, so $\mathcal{F} \cup \mathcal{D}(\mathcal{F})$ is not an antichain. \square

We have now observed that only two generalised dualities that are not antichains exist: $(\{P_1\}, \{\mathbf{0}\})$ and $(\mathcal{S}, \{P_1\})$. Let us now examine the question when a maximal antichain is not a generalised duality.

Observe that a finite maximal antichain \mathcal{Q} is a generalised duality if and only if there exist disjoint sets \mathcal{F}, \mathcal{D} such that $\mathcal{Q} = \mathcal{F} \cup \mathcal{D}$ and for an arbitrary Δ -structure X there exists $F \in \mathcal{F}$ such that $F \rightarrow X$ or there exists $D \in \mathcal{D}$ such that $X \rightarrow D$.

Lemma 20. *Let \mathcal{Q} be a finite maximal antichain different from $\{\mathbf{0}\}$, $\{P_1\}$ and $\{\mathbf{1}\}$. Then the following are equivalent:*

1. \mathcal{Q} is not formed from a generalised duality, i.e. whenever $\mathcal{Q} = \mathcal{F} \cup \mathcal{D}$, the pair $(\mathcal{F}, \mathcal{D})$ is not a generalised duality,
2. \mathcal{Q} is the set \mathcal{S} of all core Δ -trees with two edges.

Proof. The set \mathcal{S} is obviously an antichain that is not a generalised duality.

Now suppose that $\mathcal{Q} = \{Q_1, \dots, Q_q\}$ is a finite maximal antichain that is not a generalised duality. We will split \mathcal{Q} into two disjoint sets \mathcal{F} , \mathcal{D} by the following procedure: First set $\mathcal{F} = \mathcal{D} = \emptyset$. In the l -th step, add Q_l to \mathcal{F} if and only if there exists a Δ -structure X such that $Q_l \rightarrow X$, $X \not\rightarrow Q_l$, $F \rightarrow X$ for any F that is already in \mathcal{F} and $Q_{l'} \not\rightarrow X$ for any $l' > l$; otherwise add Q_l to \mathcal{D} . Repeat until after q steps each element of \mathcal{Q} belongs either to \mathcal{F} or to \mathcal{D} .

For an element $F \in \mathcal{F}$, the Δ -structure X that caused F to be added to \mathcal{F} will be denoted by \check{F} .

Clearly, if $Q < X$ for a Δ -structure X and some $Q \in \mathcal{Q}$, then there exists $F \in \mathcal{F}$ such that $F \rightarrow X$.

Let G be a Δ -structure that is an orientation of a k -uniform hypergraph with no short cycles and with a high chromatic number. Precisely, we want that every substructure of G on at most N vertices is a Δ -forest, where $N = \max\{|V(F)| : F \in \mathcal{F}\}$; and under any colouring of vertices of G with K^L colours, G has a monochromatic edge; here $K = \max\{|V(Q)| : Q \in \mathcal{Q}\}$ and $L = \max\{|V(\check{F})| : F \in \mathcal{F}\}$. The existence of such a hypergraph was proved in [5] and in [13]; [15] provides a simple construction.

Fix arbitrary $F \in \mathcal{F}$. Let $H = G \times \check{F}$.

let $Q \in \mathcal{Q}$. If $f : H \rightarrow Q$ is a homomorphism and v is any vertex of G , let $f_v : V(\check{F}) \rightarrow V(Q)$ be the mapping defined by $f_v(x) = f(v, x)$. Because of the high chromatic number of G , there exists an edge (v_1, v_2, \dots, v_k) of G such that $f_{v_1} = f_{v_2} = \dots = f_{v_k} =: g$. If (x_1, x_2, \dots, x_k) is an edge of \check{F} , then $(g(x_1), g(x_2), \dots, g(x_k)) = (f(v_1, x_1), f(v_2, x_2), \dots, f(v_k, x_k))$ is an edge of Q since f is a homomorphism and $((v_1, x_1), (v_2, x_2), \dots, (v_k, x_k))$ is an edge of H . Therefore $g : \check{F} \rightarrow Q$ is a homomorphism, a contradiction. Consequently, $H \not\rightarrow Q$.

Thus there exists $F' \in \mathcal{F}$ such that $F' \rightarrow H$. As $H \rightarrow \check{F}$, we get that $F' = F$ because of the definition of \check{F} . Consequently, $F \rightarrow H \rightarrow G$, and so every element F of \mathcal{F} is balanced because of the high girth of G .

Since \mathcal{Q} is not formed from a generalised duality, there exists a Δ -structure Y such that $Y < F$ for some $F \in \mathcal{F}$ but $Y \not\rightarrow D$ for any $D \in \mathcal{D}$.

Recall the definition of thunderbolts $Z_{T,s}$ from the proof of Lemma 19, and recall that any Δ -structure containing $Z_{T,s}$ as its substructure is not balanced.

Let $T \in \mathcal{S}$. The Δ -structure $Y + Z_{T,s}$ is not balanced, therefore it is homomorphic to no element of \mathcal{F} ; it is not homomorphic to any $D \in \mathcal{D}$, because Y is not. It must be comparable, however, so $F' \rightarrow Y + Z_{T,s}$ for some $F' \in \mathcal{F}$. Therefore there exists $F \in \mathcal{F}$ such that F is homomorphic

to $Y + Z_{T,s}$ for infinitely many values of s , and thus F is homomorphic to $Y + T$. Since $F \dashv Y$, we have $Y + T \dashv Y$, hence $T \dashv Y$. We conclude that no $T \in \mathcal{S}$ is homomorphic to Y , consequently $Y \rightarrow P_1$.

As a consequence, $Y + T = T$ and so for every $T \in \mathcal{S}$ there exists $F \in \mathcal{F}$ such that $F \rightarrow T$. The assumption on \mathcal{Q} implies that $F = T$, therefore $\mathcal{Q} = \mathcal{F} = \mathcal{S}$. \square

Realising that 1-MACs characterised in the previous subsection are also formed from dualities, we come to the astonishing correspondence between generalised dualities and MACs.

Theorem 21. *Let $\Delta = (k)$. The correspondence*

$$(\mathcal{F}, \mathcal{D}) \mapsto \mathcal{Q} = \mathcal{F} \cup \{D \in \mathcal{D} : D \dashv F \text{ for any } F \in \mathcal{F}\}$$

is a one-to-one correspondence between generalised dualities and finite maximal antichains in the homomorphism order of Δ -structures.

Proof. Follows immediately from Theorems 15 and 17 and Lemmas 19 and 20. \square

Question 1. How do the results of this section generalise for Δ -structures with more than one relation?

5 Decidability, complexity and first order definability

5.1 MAC decidability

We are interested in the following decision problem, called the *MAC decision problem*: given a finite nonempty set \mathcal{Q} of Δ -structures, decide whether \mathcal{Q} is a maximal antichain. The results of the previous section allow us to state the following result.

Theorem 22. *Let $\Delta = (k)$, $k \geq 2$. Then the MAC decision problem is decidable. Moreover, it is NP-hard.*

Proof. The decision procedure is as follows: For each element of \mathcal{Q} , check whether its core is a forest; let $\mathcal{F} \subseteq \mathcal{Q}$ be the set of all such structures. Find all transversals over \mathcal{F} . For each transversal, construct its finitary dual (see, e.g. [19]). Check whether $\mathcal{Q} \setminus \mathcal{F}$ is formed exactly by the duals of transversals.

To prove NP-hardness, we will use the fact that for any type Δ there exists a Δ -tree T such that $\text{CSP}(T)$ is NP-complete. We have the following reduction of $\text{CSP}(T)$ to the MAC decision problem: for an input structure G of $\text{CSP}(T)$, let $\mathcal{Q}(G) = \{G + T, D(T)\}$. The set $\mathcal{Q}(G)$ can be constructed from G in polynomial time. By Theorem 21, $\mathcal{Q}(G)$ is a MAC if and only if $G \rightarrow T$. \square

Question 2. Is the MAC decision problem in NP? What is the complexity of the MAC decision problem if we restrict the input to sets of cores? (Compare [12].)

Another consequence of Theorem 21 is the following.

Theorem 23. *Let \mathcal{Q} be a finite maximal antichain in $\mathcal{C}(\Delta)$, $\Delta = (k)$, $k \geq 2$. An element of \mathcal{Q} that is comparable with an input structure A can be found in polynomial time.*

Proof. Due to Theorem 21, we know that $\mathcal{Q} = \mathcal{F} \cup \mathcal{D}$. Let \mathcal{F}_c be the set of all components of structures in \mathcal{F} . All members of \mathcal{F}_c are trees by Theorem 12. The existence of a homomorphism from a Δ -structure can be determined in polynomial time even by the greedy algorithm; to decide whether $A \rightarrow B$, it suffices to check whether any of the $|V(B)|^{|V(A)|}$ mappings of vertex sets is a homomorphism, and this number is a polynomial in the size of the input structure B .

For us, it suffices to check which trees in \mathcal{F}_c are homomorphic to A . Knowing the structure of \mathcal{F} and \mathcal{D} we either find some $F \in \mathcal{F}$ such that all its component are homomorphic to A , or there is a component in each of the forests in \mathcal{F} which is not homomorphic to A . These components form a quasitransversal M and there exists a transversal M' satisfying $M \preceq M'$. Then $A \rightarrow D(M')$ and at the same time $D(M') \in \mathcal{D}$. \square

5.2 Duality decidability

Using a recent result of [12], we can deduce that it is decidable whether for a set \mathcal{H} of Δ -structures there exists a set \mathcal{F} of Δ -structures such that $(\mathcal{F}, \mathcal{H})$ is a generalised duality.

Because of Theorem 12, \mathcal{H} is the right-hand side of a generalised duality if and only if each structure in \mathcal{H} is a finitary dual and they are pairwise

incomparable. The former is decidable (and even in NP) due to [12], the latter is obviously in NP. It also follows from [12] that in general, the problem is NP-complete.

The next proposition follows from the results of [12], although it is not explicitly stated there.

Proposition 24 ([12]). *If $(\mathcal{F}, \{D\})$ is a finitary duality, $T \in \mathcal{F}$ is a core, and $n = |V(D)|$, then the number of edges of T is at most n^{n^2} .*

Given a right-hand side D of a generalised duality, we can compute the left-hand side \mathcal{F} with the following algorithm.

For every $D \in \mathcal{D}$ find the set

$$m(D) = \{T : T \text{ is a } \Delta\text{-tree, } T \rightarrow D, |V(T)| \leq n^{n^2}\}.$$

In each of these sets, determine the minimal elements in the homomorphism order, setting

$$M(D) = \min_{\rightarrow} m(D).$$

It is easy to see that

$$D = \prod_{T \in M(D)} D(T).$$

In this way, we have factored each $D \in \mathcal{D}$ into a product of irreducible Δ -structures, each of them being the dual of a tree in $M(D)$.

Let

$$\mathcal{D}_c = \bigcup_{D \in \mathcal{D}} \{D(T) : T \in M(D)\}$$

be the set of all factors appearing in these factorisations.

For determining \mathcal{F} , we will make use of a tool dual to transversals. The algorithm goes along the construction provided in section 3.1. We define a *quasicotransversal* to be any subset N of \mathcal{D}_c satisfying

- N is an antichain, i.e. for every $E \neq E' \in N$ we have $E \parallel E'$, and
- for every $D \in \mathcal{D}$ there exists $E \in N$ such that $D \rightarrow E$.

Dually to quasitransversals, for two quasicotransversals N, N' we define $N \preceq N'$ if and only if for every $E \in N$ there exists $E' \in N'$ such that $E \rightarrow E'$. Just like before, \preceq is a partial order on quasicotransversals and a *cotransversal* is a minimal quasicotransversal with respect to this order.

Let D be a Δ -structure. If there exists a Δ -tree T such that (T, D) is a duality pair, then T is determined by D uniquely up to homomorphism equivalence. Let the Δ -tree be denoted by $T(D)$.

Obviously, $T(D)$ is the (unique) maximal element of the set $\{T : T \not\rightarrow D\}$ of all trees that are not homomorphic to D . It can be computed because of Proposition 24.

For a cotransversal N , let

$$F(N) = \sum_{D \in N} T(D)$$

be the forest whose components form duality pairs with the elements of the cotransversal. Finally, let

$$\mathcal{F} = \{F(N) : N \text{ is a cotransversal}\}.$$

It can be proved, along the lines of the proof of Theorem 11, that in this way we have constructed the left-hand side \mathcal{F} of the duality pair $(\mathcal{F}, \mathcal{D})$.

5.3 GCSP dichotomy

As an analogy to CSP, we define *GCSP*, the *generalised constraint satisfaction problem*, as the following: given a finite set \mathcal{H} of Δ -structures, decide for an input Δ -structure G whether there exists $H \in \mathcal{H}$ such that $G \rightarrow H$.

Note that if $(\mathcal{F}, \mathcal{D})$ is a generalised duality, then $\text{GCSP}(\mathcal{D})$ is polynomially solvable.

As in Conjecture 1, one could ask whether there is a dichotomy for GCSP. However, this problem is not very captivating, as the positive answer to the dichotomy conjecture for CSP would imply a positive answer here as well:

Theorem 25. *Let \mathcal{H} be a finite nonempty set of pairwise incomparable Δ -structures.*

1. *If $\text{CSP}(H)$ is tractable for all $H \in \mathcal{H}$, then $\text{GCSP}(\mathcal{H})$ is tractable.*
2. *If $\text{CSP}(H)$ is NP-complete for some $H \in \mathcal{H}$, then $\text{GCSP}(\mathcal{H})$ is NP-complete.*

Proof. The first claim is evident. For the second claim, there exists a polynomial reduction of $\text{CSP}(H)$ to $\text{GCSP}(\mathcal{H})$. For an input G of $\text{CSP}(H)$, construct $G + H$ as an input for $\text{GCSP}(\mathcal{H})$. Using the pairwise incomparability of structures in \mathcal{H} , it is obvious that $G \rightarrow H$ if and only if there exists $H' \in \mathcal{H}$ such that $G + H \rightarrow H'$. \square

Thus from the complexity (and dichotomy) point of view, generalised CSP is equivalent to CSP. But their first-order definability is another matter: it is both interesting and more involved.

5.4 First-order definable GCSP

We remark that $\text{GCSP}(\mathcal{H})$ is first-order definable if and only if there exists a set \mathcal{F} such that $(\mathcal{F}, \mathcal{H})$ is a generalised duality. This result is an extension of a similar theorem for CSP contained in [1], and its proof follows the same way.

Thus we have the following:

Theorem 26. *Let \mathcal{H} be a finite set of core Δ -structures which are pairwise incomparable. Then the following are equivalent:*

1. *GCSP(\mathcal{H}) is first-order definable;*
2. *the existence of a homomorphism to some structure in \mathcal{H} is determined by a finite set of obstructions;*
3. *there exists a finite family \mathcal{F} of Δ -forests such that $\mathcal{H} = \mathcal{D}(\mathcal{F})$.*

References

1. A. Atserias. On digraph coloring problems and treewidth duality. In *Proceedings of the 20th IEEE Symposium on Logic in Computer Science (LICS'05)*, pages 106–115. IEEE Computer Society, 2005.
2. M. Barr and C. Wells. *Category Theory for Computing Science*. Les Publications CRM, Montréal, 3rd edition, 1999.
3. A. Bulatov, A. Krokhin, and P. G. Jeavons. Constraint satisfaction problems and finite algebras. In *Proceedings of the 27th International Colloquium on Automata, Languages and Programming (ICALP'00)*, volume 1853 of *Lecture Notes in Computer Science*, pages 272–282. Springer-Verlag, 2000.
4. V. Dalmau, A. Krokhin, and B. Larose. First-order definable retraction problems for posets and reflexive graphs. In *Proceedings of the 19th IEEE Symposium on Logic in Computer Science (LICS'04)*, pages 232–241. IEEE Computer Society, 2005.
5. P. Erdős and A. Hajnal. On chromatic number of graphs and set-systems. *Acta Math. Hungar.*, 17(1-2):61–99, 1966.
6. 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.

7. P. Hell and J. Nešetřil. On the complexity of H -coloring. *J. Combin. Theory Ser. B*, 48(1):92–119, 1992.
8. P. Hell and J. Nešetřil. *Graphs and Homomorphisms*, volume 28 of *Oxford Lecture Series in Mathematics and Its Applications*. Oxford University Press, 2004.
9. P. Hell, J. Nešetřil, and X. Zhu. Duality and polynomial testing of tree homomorphisms. *Trans. Amer. Math. Soc.*, 348(4):1281–1297, 1996.
10. P. G. Jeavons. On the algebraic structure of combinatorial problems. *Theoret. Comput. Sci.*, 200(1–2):185–204, 1998.
11. P. Komárek. *Good characterisations in the class of oriented graphs*. PhD thesis, Czechoslovak Academy of Sciences, Prague, 1987. In Czech.
12. B. Larose, C. Loten, and C. Tardif. A characterisation of first-order constraint satisfaction problems. *Submitted*, 2006.
13. L. Lovász. On chromatic number of finite set-systems. *Acta Math. Hungar.*, 19(1–2):59–67, 1968.
14. J. Nešetřil and A. Pultr. On classes of relations and graphs determined by subobjects and factorobjects. *Discrete Math.*, 22:287–300, 1978.
15. J. Nešetřil and V. Rödl. A short proof of the existence of highly chromatic hypergraphs without short cycles. *J. Combin. Theory Ser. B*, 27:225–227, 1979.
16. J. Nešetřil and S. Shelah. On the order of countable graphs. *European J. Combin.*, 24(6):649–663, 2003.
17. 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.
18. J. Nešetřil and C. Tardif. On maximal finite antichains in the homomorphism order of directed graphs. *Discuss. Math. Graph Theory*, 23:325–332, 2003.
19. 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.
20. J. Nešetřil and X. Zhu. Path homomorphisms. *Math. Proc. Cambridge Philos. Soc.*, 120:207–220, 1996.
21. E. Welzl. Color families are dense. *Theoret. Comput. Sci.*, 17:29–41, 1982.