Charles University in Prague
Faculty of Mathematics and Physics

Doctoral Thesis

Homomorphisms and Structural Properties of Relational Systems

Jan Foniok

Department of Applied Mathematics
Supervisor: Prof. RNDr. Jaroslav Nešetřil, DrSc.
Branch 14: Discrete Models and Algorithms

Preface

My interest in graph homomorphisms dates back to the Spring School of Combinatorics in 2000. The School is traditionally organised by the Department of Applied Mathematics of the Charles University in a hilly part of the Czech Republic; in 2000 it was one of the few times when the Spring School was international not only in terms of its participants, but also in terms of its venue. Participants will never forget carrying a blackboard across the border between Germany and the Czech Republic and the exciting boat trip on the Vltava.

The study text on homomorphisms [28], specially prepared for the Spring School, aroused my curiosity that has eventually resulted in both my master's thesis and this doctoral dissertation.

The study of graph homomorphisms was pioneered by G. Sabidussi, Z. Hedrlín and A. Pultr in the 1960's. It was part of an attempt to develop a theory of general mathematical structures in the framework of algebra and category theory. Many nice and important results have emerged from their work and the work of their followers. Even so, until recently many graph theorists would not include homomorphisms among the topics of central interest in graph theory.

Nevertheless, graph homomorphisms and structural properties of graphs have recently attracted much attention of the mathematical community. The reason may be in part that the homomorphism point of view has proved useful in various areas ranging from colouring and graph reconstruction to applications in artificial intelligence, telecommunication, and even statistical physics. A book [18] now exists that introduces the topic and brings together the most important parts of the theory and its applications. This thesis surveys my small contribution to the ongoing research in this area.

I typeset the thesis using XeLaTeX. This new typesetting system based on Donald Knuth's TeX and the format LaTeX extends TeX's capabilities of handling Unicode input, multi-lingual and multi-script documents and especially modern OpenType fonts.

When I first saw it, I liked the typeface of the font Andulka very much and have soon decided to use it for my thesis. For typesetting mathematics, I

used AMS Euler [21], an upright cursive font that tries to emulate a mathematician's style of handwriting mathematical formulas on a blackboard, which is upright rather than italic.

Some results contained in this thesis have been published or accepted for publication.

[11] J. Foniok, J. Nešetřil, and C. Tardif. Generalised dualities and finite maximal antichains. In F. V. Fomin, editor, *Graph-Theoretic Concepts in Computer Science (Proceedings of WG 2006)*, volume 4271 of *Lecture Notes in Comput. Sci.*, pages 27–36. Springer-Verlag, 2006.

[12] J. Foniok, J. Nešetřil, and C. Tardif. On finite maximal antichains in the homomorphism order. *Electron. Notes Discrete Math.*, 29:389–396, 2007.

[13] J. Foniok, J. Nešetřil, and C. Tardif. Generalised dualities and maximal finite antichains in the homomorphism order of relational structures. *European J. Combin.*, to appear.

Acknowledgements

I thank my supervisor, Jaroslav Nešetřil, for introducing me to the world of homomorphisms and for persistent support of my research. Discussions with him have always been inspiring, but at the same time compelling and enjoyable.

Many thanks to Claude Tardif for wonderful cooperation and for his surrealistic attitudes. Claude's contribution to my being able to write this thesis is indeed valuable.

For numerous suggestions I am grateful to Manuel Bodirsky, Julia Böttcher and Zdeněk Hedrlín. I also acknowledge the generous support of the Institute for Theoretical Computer Science (ITI) in Prague and the EU Research Training Network COMBSTRU. Last but not least I thank Ida Švejdarová for lending me a pen when I needed one.

Contents

Notation

$\underline{A}$	...	base set of a Δ-structure A [see 1.2.4]
$A \to B$	...	A is homomorphic to B
$A \sim B$	...	A is homomorphically equivalent to B; $A \to B$ and $B \to A$
$A \parallel B$	...	A and B are incomparable; $A \not\to B$ and $B \not\to A$
$A/\approx$	...	factor structure [see 1.2.7]
$\mathrm{Block}(A)$	...	see 1.6.5
C^B	...	exponential structure [see 1.5.9]
$\mathcal{C}(\Delta)$	...	category of Δ-structures; the homomorphism order of Δ-structures
$D(F)$	...	dual of the Δ-tree F
$D(\mathcal{F})$	...	dual of a finite set $\mathcal{F}$ of Δ-trees
$\mathcal{D}(\mathcal{F})$	...	finite dual set of a finite set $\mathcal{F}$ of Δ-forests
$\mathrm{DSh}(A)$	...	directed shadow of A [see 1.6.1]
$\to D$	...	$\{A : A \to D\}$
$\to \mathcal{D}$	...	$\{A : A \to D \text{ for some } D \in \mathcal{D}\}$
$\not\to \mathcal{D}$	...	$\{A : A \not\to D \text{ for all } D \in \mathcal{D}\}$
Δ	...	a type of relational structures; $\Delta = (\delta_i : i \in I)$
$f : A \to B$	...	f is a homomorphism from A to B
$f[S]$	...	image of the set S under the mapping f, if S is a subset of the domain of f; $f[S] = \{f(s) : s \in S\}$
$f(e)$	...	if $e = (u_1, u_2, \ldots, u_k)$, then $f(e) = \big(f(u_1), f(u_2), \ldots, f(u_k)\big)$
$f \restriction T$	...	restriction of the function f to a subset T of the domain of f
$F \not\to$	...	$\{A : F \not\to A\}$
$\mathcal{F} \to$	...	$\{A : F \to A \text{ for some } F \in \mathcal{F}\}$
$\mathcal{F} \not\to$	...	$\{A : F \not\to A \text{ for all } F \in \mathcal{F}\}$
$\rfloor(x)$	...	height label of the vertex x [see 1.8.2]
I	...	set of indices; $\Delta = (\delta_i : i \in I)$
$\mathrm{Inc}(A)$	...	incidence graph of A [see 1.6.5]
K_k	...	complete graph on k vertices
$p \vee q$	...	supremum of p and q; join in a lattice
$p \wedge q$	...	infimum of p and q; meet in a lattice

$p \Rightarrow q$	...	Heyting operation [see 3.3.1]
$\vec{P}_k$	...	directed path with k edges
$R_i(A)$	...	the ith edge set of a Δ-structure A [see 1.2.4]
$S^{\downarrow}$	...	downset generated by S
$S^{\uparrow}$	...	upset generated by S
$\mathrm{Sh}(A)$	...	shadow of A [see 1.6.2]
$\vec{T}_k$	...	transitive tournament with k vertices
$[a]_{\approx}$	...	class of the equivalence $\approx$ that contains a
$\top$	...	Δ-structure with one vertex and all loops; $\underline{\top} = \{1\}$; $R_i(\top) = \underline{\top}^{\delta_i}$ for all $i \in I$
$\bot$	...	Δ-structure with one vertex and no edges; $\underline{\bot} = \{1\}$; $R_i(\bot) = \emptyset$ for all $i \in I$
$\coprod_{j\in J} A_j$; $A + B$	...	sum of relational structures [see 1.5.2]
$\prod_{j\in J} A_j$; $A \times B$	...	product of relational structures [see 1.5.5]

1 Introduction

Whatever you can do or
dream you can, begin it.
Boldness has genius, magic
and power in it. Begin it now.

(Johann Wolfgang von Goethe)

1.1 Motivation and overview

In this thesis, we study homomorphisms of finite relational structures. Finite relational structures can be viewed in several ways. The view we adopt consists in seeing them as a generalisation of graphs. Relational structures may actually be described as oriented uniform hypergraphs with coloured edges. There are three main differences from ordinary graphs: edges are ordered, they are tuples of possibly more than two vertices, and there are various kinds of edges.

Homomorphisms are mappings between vertex sets of relational structures. Homomorphisms preserve edges; so the image of an edge is an edge. Moreover, it is an edge of the same kind.

Thus homomorphisms endow graphs and relational structures with an algebraic structure that will be familiar to an algebraist or category theorist.

The unifying concept in the thesis is the question of existence of homomorphisms. It interconnects the two main topics presented here.

The first topic is homomorphism dualities. There the existence of a homomorphism between structures is equivalent to the non-existence of a homomorphism between other structures. In particular, we study situations where a class of relational structures is characterised both by the non-existence of a homomorphism from some *finite* set of structures, and by the existence of a homomorphism to some other *finite* set of structures. Such situations are called finite homomorphism dualities. We provide a full characterisation of finite homomorphism dualities.

The other topic is the homomorphism order, where the existence of a homomorphism defines a relation that turns out to induce a partial order on the class of relational structures. We examine especially finite maximal antichains in the homomorphism order. We find a surprising correspondence

between maximal antichains and finite dualities. Many finite maximal antichains have the splitting property; we derive a structural condition on those antichains that do not have this property.

The main results of the thesis are the characterisation of all finite homomorphism dualities (Theorem 2.4.26) and the splitting property of finite maximal antichains in the homomorphism order with described exceptions (Theorem 3.4.12 and Theorem 3.4.13).

Other results include a new construction of dual structures, which generalises two previous constructions (Section 2.2). Furthermore we extend our results on homomorphism dualities for relational structures into the context of lattices (Section 3.3). And finally we state several consequences of these results in the area of computational complexity (Chapter 4).

1.2 Relational structures

First things first. We study homomorphisms of relational structures, so let us first define relational structures.

1.2.1 Definition. A **type** Δ is a sequence $(\delta_i : i \in I)$ of positive integers; I is a finite set of indices. A (finite) **relational structure** A of type Δ is a pair $\big(X, (R_i : i \in I)\big)$, where X is a finite nonempty set and $R_i \subseteq X^{\delta_i}$; that is, R_i is a δ_i-ary relation on X. Relational structures of type Δ are denoted by capital letters A, B, C, ...

1.2.2. There are many natural examples of relational structures. Perhaps the simplest are digraphs (with loops allowed), which are simply Δ-structures of type $\Delta = (2)$. This example is also the motivation for our terminology. The class of all partially ordered sets is a subclass of the class of all Δ-structures for $\Delta = (2)$, requiring that the relation be reflexive, transitive and antisymmetric.

The class of all (2)-structures whose relation is symmetric is the class of all undirected graphs. This is an important example and we shall keep it in mind as we define properties and operations on Δ-structures; all of them carry over immediately to undirected graphs. For an undirected graph G, the base set is usually called the *vertex set* and denoted by $V(G)$.

1.2.3. In a logician's words, relational structures are models of theories with no function symbols; and finite relational structures are such models in the theory of finite sets.

1.2.4 Definition. If $A = \big(X, (R_i : i \in I)\big)$, the **base set** X is denoted by $\underline{A}$ and the relation R_i by $R_i(A)$. We often refer to a relational structure of

type Δ as **Δ-structure**. The type Δ is almost always fixed in the following text. The elements of the base set are called **vertices** and the elements of the relations R_i are called **edges**. For the set of all edges of a Δ-structure A we use the notation $R(A)$, that is

$$R(A) := \bigcup_{i \in I} R_i(A).$$

To distinguish between various relations of a Δ-structure we speak about **kinds of edges** (so the elements of $R_i(A)$ are referred to as the *edges of the ith kind*).

1.2.5 Unary relations. Some of the relations of Δ-structures may be unary (this is the same as saying that some of the numbers δ_i may be equal to one). In this thesis, however, we consider only relational structures with *no* unary relations. This is a rather technical assumption. All the results and proofs remain valid even for structures with unary relations, but some adjustments would have to be made.

For example, in Theorem 3.4.13 we suppose that there are at most two relations. In fact, we should suppose that there are at most two relations of arity greater than one and an arbitrary number of unary relations.

Elsewhere the statements would have to be slightly altered, like in Theorem 2.5.3. If we allow unary relations, the upper bound has to be replaced by $n^{n+|I|}$.

Such adjustments would, in our opinion, make the text less clear and more difficult to read, so we find it useful to assume that structures have no unary relations.

Substructures of relational structures are what we are familiar with from graph theory as *induced subgraphs*. It is possible to define "non-induced" substructures as well, but in our context it is more convenient not to do so.

1.2.6 Definition. A **substructure** of a Δ-structure A is any structure $\big(S, (R'_i : i \in I)\big)$ such that S is a subset of $\underline{A}$ and $R'_i = R_i \cap S^{\delta_i}$. It is also called the **substructure of** A **induced by** S; and it is a **proper substructure** if $S \subsetneq \underline{A}$.

Next we introduce a factorisation construction. This construction can be viewed as gluing equivalent vertices together. An example of the construction (for digraphs) is in Figure 1.1.

1.2.7 Definition. Let A be a Δ-structure and let $\approx$ be an equivalence relation on $\underline{A}$. We define the **factor structure** $A/\approx$ to be the Δ-structure whose base set is the set of all equivalence classes of the relation $\approx$, so $\underline{A/\approx} = \underline{A}/\approx$, and

a δ_i-tuple of equivalence classes is in the relation R_i of $A/\approx$ if we can find an element in each of the classes such that these elements form an edge of A,

$$R_i(A/\approx) = \Big\{ \big([a_1]_\approx, [a_2]_\approx, \ldots, [a_{\delta_i}]_\approx\big) : (a_1, a_2, \ldots, a_{\delta_i}) \in R_i(A) \Big\},\ i \in I.$$

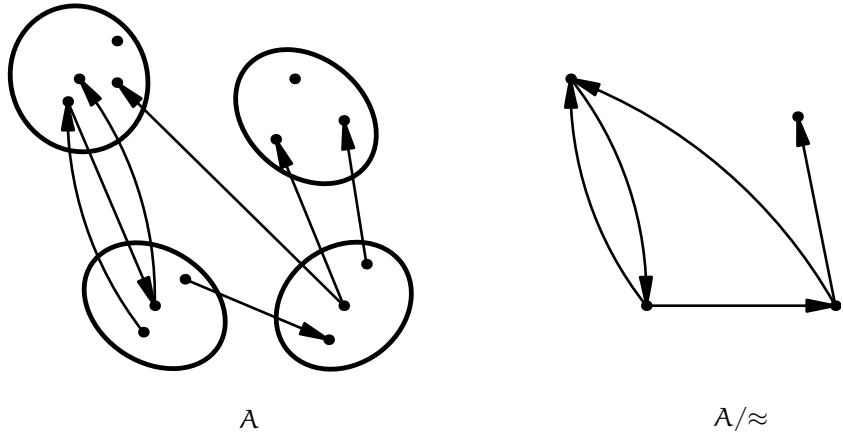

Figure 1.1: Factor structure

Finally, et us once more stress the two restrictions posed on relational structures in this thesis: All structures we consider are *finite* and have *no unary relations*.

1.3 Homomorphisms

Familiar with the notion of relational structures, we continue by defining homomorphisms. As one might expect, they are mappings of base sets that preserve all the relations.

1.3.1 Definition. Let A and A' be two relational structures of the same type Δ. A mapping $f : \underline{A} \to \underline{A'}$ is a **homomorphism** from A to A' if for every $i \in I$ and for every $u_1, u_2, \ldots, u_{\delta_i} \in \underline{A}$ the following implication holds:

$$(u_1, u_2, \ldots, u_{\delta_i}) \in R_i(A) \quad \Rightarrow \quad \big(f(u_1), f(u_2), \ldots, f(u_{\delta_i})\big) \in R_i(A').$$

1.3.2 Definition. The fact that f is a homomorphism from A to A' is denoted by

$$f : A \to A'.$$

If there exists a homomorphism from A to A', we say that A is **homomorphic** to A' and write $A \to A'$; otherwise we write $A \not\to A'$. If A is homomorphic to A' and at the same time A' is homomorphic to A, we say that A and A' are **homomorphically equivalent** and write $A \sim A'$. If on the other hand there exists no homomorphism from A to A' and no homomorphism from A' to A, we say that A and A' are **incomparable** and write $A \parallel A'$.

1.3.3 Definition. A homomorphism from A to itself is called an **endomorphism** of A.

Next we give two simple examples of homomorphisms.

1.3.4 Example. This is a trivial example: Let $\Delta = (2)$, so we consider digraphs. Let $\vec{P}_k$ be the directed path on vertices $\{0, 1, \ldots, k-1\}$ with edge set $\big\{(j, j+1) : j = 0, 1, \ldots, k-1\big\}$. Then $f : V(\vec{P}_k) \to V(G)$ is a homomorphism if and only if $f(0), f(1), \ldots, f(k)$ is a directed walk in G.

1.3.5 Example. For an undirected graph G, a homomorphism to the complete graph K_k is essentially a k-colouring of G: imagine the vertices of K_k as colours, and since edges of G have to be preserved, distinct vertices of G have to be mapped to distinct vertices of K_k, in other words they have to be assigned distinct colours.

1.3.6 Composition of homomorphisms. It is a very important aspect of homomorphisms that they compose – the composition of two homomorphisms is a homomorphism as well. This composition operation endows a set of Δ-structures with a structure of an algebraic flavour. Homomorphisms of relational structures share this property with morphisms of other structures, like topological spaces, semigroups, monoids, partial orders and many others. This flavour is discussed in more detail in Section 1.5 but it is omnipresent throughout the thesis.

1.3.7 Definition. As usual, a homomorphism from A to A' is an **isomorphism** if it is a bijection and f^{-1} is a homomorphism of A' to A. If there exists an isomorphism between A and A', we say that A and A' are **isomorphic** and write $A \cong A'$. An isomorphism of A with itself is called an **automorphism** of A.

1.4 Retracts and cores

In this section, we introduce the notion of *cores*. Cores are Δ-structures that are minimal in the following sense: A core is not homomorphically equivalent to any smaller structure.

For the formal definition, we use retractions. A retraction is an endomorphism that does not move any vertex in its image.

1.4.1 Definition. Let A be a Δ-structure. An endomorphism $f : A \to A$ is a **retraction** if it leaves its image fixed, in other words if $f(x) = x$ for all $x \in f[\underline{A}]$. A substructure B of A is called a **retract** of A if there exists a retraction of A onto B; a retract is **proper** if it is a proper substructure.

Later we will need the fact that a retract is homomorphically equivalent to the original structure.

1.4.2 Lemma. *If* B *is a retract of* A*, then* A *and* B *are homomorphically equivalent.*

Proof. If B is a retract of A, then B is a substructure of A and so the identity mapping is a homomorphism from B to A. On the other hand, the retraction is a homomorphism from A to B. □

1.4.3 Definition. A Δ-structure C is called a **core** if it has no proper retracts. A retract C of A is called a **core of** A if it is a core.

Several other conditions are equivalent to the one we chose for the definition of a core.

1.4.4 Lemma (Characterisation of cores)**.** *For a* Δ*-structure* C *the following conditions are equivalent.*

(1) C *is a core (that is,* C *has no proper retracts).*

(2) C *is not homomorphic to any proper substructure of* C*.*

(3) Every endomorphism of C *is an automorphism.*

Proof. (1) $\Rightarrow$ (2): Suppose that f is a homomorphism of C to a proper substructure of C. Then f is a permutation of its image $f[\underline{C}]$ and so there exists a positive integer k such that f^k restricted to $f[\underline{C}]$ is the identity mapping. Then f^k is a retraction of C onto a proper retract.

(2) $\Rightarrow$ (3): Let $f : C \to C$ be an endomorphism. The mapping f is surjective, so it is a bijection ($\underline{C}$ is finite). There exists a positive integer k such

that f^k is the identity mapping. Then $f^{-1} = f^{k-1}$ is a homomorphism, so f is an automorphism.

(3) $\Rightarrow$ (1): Every retraction is an endomorphism, so every retraction of C is an automorphism. Thus the image of every retraction of C is C, hence C has no proper retracts. □

Next we prove that every relational structure has exactly one core (up to isomorphism). For the proof we use two lemmas.

1.4.5 Lemma. *Let A and B be two Δ-structures. If there exist surjective homomorphisms $f : A \to B$ and $g : B \to A$, then A and B are isomorphic.*

Proof. The existence of f shows that $|\underline{A}| \geq |\underline{B}|$ and the existence of g shows $|\underline{A}| \leq |\underline{B}|$, so $|\underline{A}| = |\underline{B}|$. This means that f is a bijection and that $g \circ f$ is a bijection, so there exists a positive integer k such that $(g \circ f)^k$ is the identity mapping. Then $(g \circ f)^{k-1} \circ g = f^{-1}$, therefore f^{-1} is a homomorphism and f an isomorphism. □

1.4.6 Lemma. *Let C and C' be two cores. If C and C' are homomorphically equivalent, they are isomorphic.*

Proof. Let $f : C \to C'$ and $g : C' \to C$. The mapping $f \circ g$ is an endomorphism of C'. Since C' is a core, it is an automorphism, therefore both f and g are surjective. The existence of an isomorphism follows from Lemma 1.4.5. □

This lemma implies in particular that every Δ-structure is homomorphically equivalent to at most one core. The next proposition asserts that it is actually *exactly* one.

1.4.7 Proposition. *Every Δ-structure A has a unique core C (up to isomorphism). Moreover, C is the unique core to which A is homomorphically equivalent.*

Proof. Proof of existence: select the retract C of A of the smallest size. Any potential proper retract C' of C would be a smaller retract of A, and so C is a core.

Let C and C' be two distinct cores of G. By the definition of a core, there exist homomorphisms $f : G \to C$ and $f' : G \to C'$. The restrictions $f \restriction C'$ and $f' \restriction C$ show that C and C' are homomorphically equivalent; by Lemma 1.4.6 they are isomorphic. □

1.4.8 Corollary. *A Δ-structure C is a core if and only if it is not homomorphically equivalent to a Δ-structure with fewer vertices.*

Proof. Suppose that C and D are homomorphically equivalent and D has fewer vertices than C. Then the core of D has at most $|\underline{D}|$ vertices, but $|\underline{D}| < |\underline{C}|$, so C is not the unique core to which D is homomorphically equivalent. Hence C is not a core.

Conversely, if C is not a core, then it is homomorphically equivalent to its core C'. The core C' is a proper retract of C, hence it has fewer vertices. □

Thus we can summarise four equivalent definitions of a core.

1.4.9 Corollary (Characterisation of cores revisited)**.** *For a Δ-structure* C *the following conditions are equivalent.*

(1) C *is a core (that is,* C *has no proper retracts).*

(2) C *is not homomorphic to any proper substructure of* C*.*

(3) Every endomorphism of C *is an automorphism.*

(4) C *is not homomorphically equivalent to a Δ-structure with fewer vertices.* □

1.5 The category of relational structures and homomorphisms

It is little surprising that structures of an algebraic nature, like Δ-structures, with suitably selected mappings among them, should form a category. Here, we observe some basic properties of the category of Δ-structures and homomorphisms. The reader may consult [1] or [3] for an introduction to category theory.

1.5.1. For a fixed type $\Delta = (\delta_i : i \in I)$, let $\mathcal{C}(\Delta)$ be the category of all Δ-structures (objects) and their homomorphisms (morphisms).

1.5.2 Definition. Let J be a nonempty finite index set and let A_j, $j \in J$ be Δ-structures. We define the **sum** $\coprod_{j\in J} A_j$ to be the *disjoint union* of the structures A_j; formally the base set of the sum is defined by

$$\underline{\coprod_{j\in J} A_j} = \bigcup_{j\in J} \left(\{j\} \times \underline{A_j}\right),$$

and the relations by

$$R_i\left(\coprod_{j\in J} A_j\right) = \bigcup_{j\in J}\Big\{\big((j,x_1),(j,x_2),\dots,(j,x_{\delta_i})\big) : (x_1,x_2,\dots,x_{\delta_i}) \in R_i(A_j)\Big\}, \quad i \in I.$$

1.5.3 Proposition. *The Δ-structure $A = \coprod_{j\in J} A_j$ with embeddings $\iota_j : A_j \to A$ such that $\iota_j : x \mapsto (j,x)$ for $j \in J$ is the coproduct of the Δ-structures A_j in the category $\mathcal{C}(\Delta)$.*

Proof. Clearly the embeddings ι_j are homomorphisms. Suppose we have another structure A' and homomorphisms $g_j : A_j \to A'$. Then $f : (j,x) \mapsto g_j(x)$ is the unique homomorphism from A to A' such that $f\iota_j = g_j$ for all j in J. □

1.5.4 Corollary. *The sum $\coprod_{j\in J} A_j$ is homomorphic to a Δ-structure B if and only if $A_j \to B$ for every $j \in J$.*

For the sum of two Δ-structures, we use the notation $A + B$, or more generally, we can write $A_1 + A_2 + \cdots + A_n$ for the sum of n structures.

We remark here that the sum of graphs, as defined above, is the usual operation of disjoint union of graphs.

We go on to describe product in the category of Δ-structures.

1.5.5 Definition. Let $J = \{1,2,\dots,n\}$ be an index set and let A_j, $j \in J$ be Δ-structures. We define the **product** $\prod_{j\in J} A_j = A$ to be the Δ-structure whose base set is the Cartesian product of the vertex sets of the factors, and there is an edge of a kind if and only if there is an edge of the same kind in each projection. Formally,

$$\underline{A} = \prod_{j\in J} \underline{A_j},$$

$$R_i(A) = \Big\{\big((x_{1,1},x_{1,2},\dots,x_{1,n}),(x_{2,1},x_{2,2},\dots,x_{2,n}),\dots, (x_{\delta_i,1},x_{\delta_i,2},\dots,x_{\delta_i,n})\big) : (x_{1,j},x_{2,j},\dots,x_{\delta_i,j}) \in R_i(A_j) \text{ for all } j \in J\Big\}, \qquad i \in I.$$

1.5.6 Example. Let $\Delta = (2,2)$. Let A and B be the Δ-structures depicted in Figure 1.2. Figure 1.3 shows the product $A \times B$.

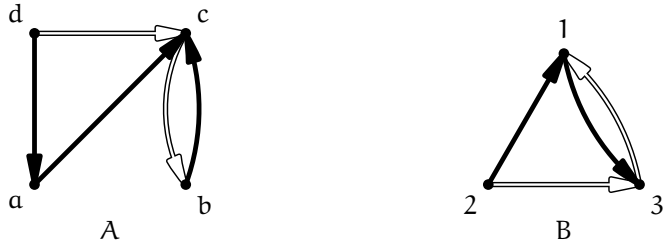

Figure 1.2: The Δ-structures A and B

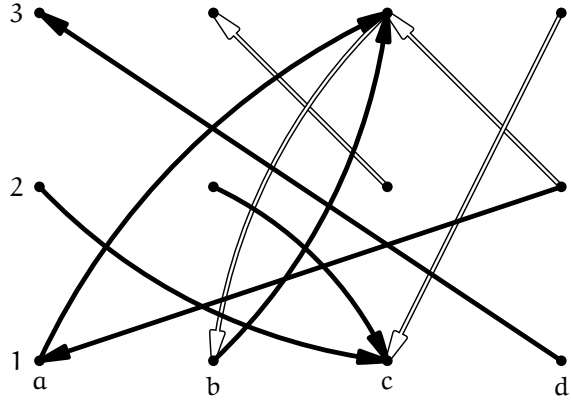

Figure 1.3: The product $A \times B$

1.5.7 Proposition. *The Δ-structure $A = \prod_{j\in J} A_j$ with projections $\pi_j : A \to A_j$ such that $\pi_j : (x_1, x_2, \ldots, x_n) \mapsto x_j$ for $j \in J$ is the product of the Δ-structures A_j in the category $\mathcal{C}(\Delta)$.*

Proof. The projections are indeed homomorphisms; and whenever A' is a Δ-structure such that $g_j : A' \to A_j$ are homomorphisms, then $f : x \mapsto \big(g_1(x), g_2(x), \ldots, g_n(x)\big)$ is the unique homomorphism from A' to A such that $\pi_j f = g_j$ for all j in J. □

1.5.8 Corollary. *A Δ-structure B is homomorphic to the product $\prod_{j\in J} A_j$ if and only if $B \to A_j$ for every $j \in J$.*

Analogously to sums, we use the convenient notation $A \times B$ and $A_1 \times A_2 \times \cdots \times A_n$ for products.

Finally, we introduce exponentiation in $\mathcal{C}(\Delta)$. The definition is somewhat technical, but exponentiation is important in the context of homomorphism dualities.

1.5.9 Definition. Let $\mathbf{B}$ and $\mathbf{C}$ be two Δ-structures. We define the **exponential structure** $\mathbf{C}^{\mathbf{B}}$ to be the Δ-structure whose base set is

$$\underline{\mathbf{C}^{\mathbf{B}}} = \{f : f \text{ is a mapping from } \underline{\mathbf{B}} \text{ to } \underline{\mathbf{C}}\},$$

and the ith relation is the set of all δ_i-tuples $(f_1, f_2, \ldots, f_{\delta_i})$ such that whenever the δ_i-tuple $(b_1, b_2, \ldots, b_{\delta_i})$ is an element of $R_i(\mathbf{B})$, then

$$\big(f_1(b_1), f_2(b_2), \ldots, f_{\delta_i}(b_{\delta_i})\big) \in R_i(\mathbf{C}).$$

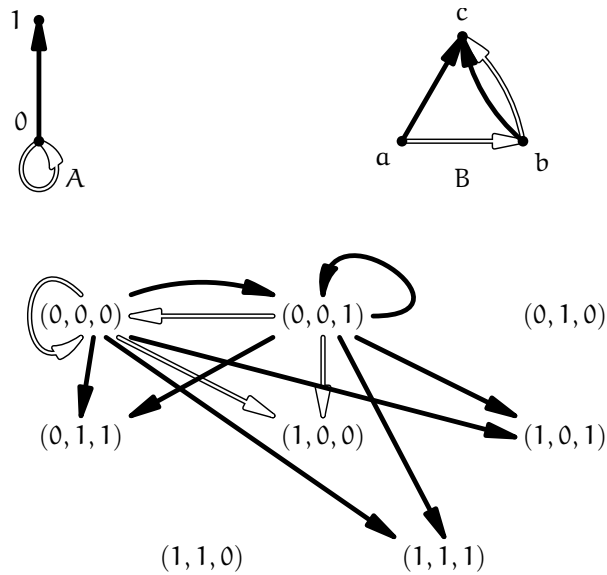

Figure 1.4: An example of an exponential structure

1.5.10 Example. An example of two Δ-structures $\mathbf{A}$ and $\mathbf{B}$ and the exponential structure $\mathbf{A}^{\mathbf{B}}$ for $\Delta = (2, 2)$ is shown in Fig. 1.4. A mapping $f : \underline{\mathbf{B}} \to \underline{\mathbf{A}}$ is represented by the triple $\big(f(a), f(b), f(c)\big)$. The existence of a unique black edge in $\mathbf{A}$ means that for two functions $f_1, f_2 : \underline{\mathbf{B}} \to \underline{\mathbf{A}}$ to be connected with an edge in $\mathbf{A}^{\mathbf{B}}$, it must hold that $f_1(a) = f_1(b) = 0$ (the initial vertices of

all black edges in B must be mapped to the initial vertex of the only black edge in A) and $f_2(c) = 1$ (a similar condition for the terminal vertices of the black edges). Similarly, for outlined edges, (f_1, f_2) is an outlined edge in A^B if and only if $f_1(a) = f_1(b) = f_2(b) = f_2(c) = 0$.

1.5.11 Proposition. *For Δ-structures B and C, the exponential structure C^B together with the homomorphism* $\mathrm{eval} : C^B \times B \to C$ *defined by* $\mathrm{eval}(f, b) := f(b)$ *is an exponential object in the category $\mathcal{C}(\Delta)$.*

Proof. If $\big((f_1, b_1), (f_2, b_2), \ldots, (f_{\delta_i}, b_{\delta_i})\big) \in R_i(C^B \times B)$, then

$$(f_1, f_2, \ldots, f_{\delta_i}) \in R_i(C^B) \quad \text{and} \quad (b_1, b_2, \ldots, b_{\delta_i}) \in R_i(B)$$

by the definition of product. It follows from the definition of edges in the exponential structure that

$$\big(f_1(b_1), f_2(b_2), \ldots, f_{\delta_i}(b_{\delta_i})\big) \in R_i(C).$$

So the mapping eval is indeed a homomorphism.

Let A be a structure and $g : A \times B \to C$ a homomorphism. Define $\lambda g : \underline{A} \to \underline{C^B}$ by setting $\lambda g(a)(b) := g(a, b)$. Suppose $(a_1, a_2, \ldots, a_{\delta_i}) \in R_i(A)$. Now, if $(b_1, b_2, \ldots, b_{\delta_i}) \in R_i(B)$, then

$$\big(\lambda g(a_1)(b_1), \lambda g(a_2)(b_2), \ldots, \lambda g(a_{\delta_i})(b_{\delta_i})\big) = \\ \big(g(a_1, b_1), g(a_2, b_2), \ldots, g(a_{\delta_i}, b_{\delta_i})\big) \in R_i(C)$$

because g is a homomorphism. Therefore $\big(\lambda g(a_1), \lambda g(a_2), \ldots, \lambda g(a_{\delta_i})\big) \in R_i(C^B)$ and λg is a homomorphism from A to C^B; it is easy to check that it is the only such homomorphism that satisfies $\mathrm{eval} \circ (\lambda g \times \mathrm{id}_B) = g$, that is, that the following diagram commutes.

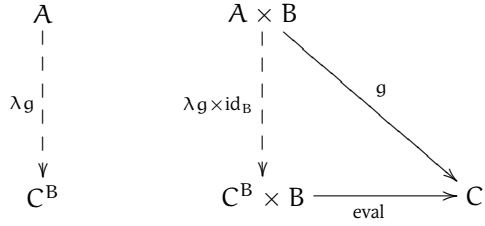

□

1.5.12 Corollary. *For any Δ-structures A, B and C,*

$$A \to C^B \quad \textit{if and only if} \quad A \times B \to C.$$

Proof. If $g : A \times B \to C$, then λg defined in the proof of Proposition 1.5.11 is a homomorphism from A to C^B. Conversely, if $h : A \to C^B$, then $\mathrm{eval} \circ (h \times \mathrm{id}_B)$ is a homomorphism from $A \times B$ to C. □

1.5.13. The base set of C^B consists of all mappings from $\underline{B}$ to $\underline{C}$, but not all of them are homomorphisms. It follows immediately from the definition of relations of the exponential structure, that homomorphisms from B to C are exactly those elements f of $\underline{C^B}$, for which the δ_i-tuple $(f, f, \ldots, f)$ is in $R_i(C^B)$ for all $i \in I$.

1.5.14 Proposition. *The category $\mathcal{C}(\Delta)$ of Δ-structures and homomorphisms is Cartesian closed.*

Proof. Let $V = \{1\}$ and let $\top$ be the Δ-structure such that

$$\underline{\top} = V \quad \text{and}$$
$$R_i(\top) = V^{\delta_i} \quad \text{for } i \in I.$$

There exists exactly one homomorphism from any Δ-structure to $\top$, namely the constant mapping to 1. Hence $\top$ is the terminal object of $\mathcal{C}(\Delta)$. Using Proposition 1.5.7, we see that there are finite products in $\mathcal{C}(\Delta)$; and by Proposition 1.5.11, there are exponential objects. □

1.6 Connectedness and irreducibility

In this section, we define connected relational structures. Our notion of connectedness generalises weak connectedness of digraphs. We define it in two principally different ways – using auxiliary undirected multigraphs (*shadows* and *incidence graphs*), and by specifying structural conditions. We show that these two ways are equivalent.

Irreducibility is a dual notion to connectedness in the category $\mathcal{C}(\Delta)$. We mention a famous problem connected to irreducibility: Hedetniemi's product conjecture.

1.6.1 Definition. The **directed shadow** of a Δ-structure A is the directed multigraph $\mathrm{DSh}(A)$ whose vertices are the elements of $\underline{A}$ and there is one edge from a to b for each occurrence of the vertices a, b in an edge in some $R_i(A)$ of arity $\delta_i \geq 2$ such that $(a_1, \ldots, a_{\delta_i}) \in R_i(A)$ with $a_j = a$, $a_{j+1} = b$ for some $1 \leq j < \delta_i$.

1.6.2 Definition. The **shadow** of a Δ-structure A is the undirected multigraph $\mathrm{Sh}(A)$ that is created from $\mathrm{DSh}(A)$ by replacing every directed edge with an undirected edge (the symmetrisation of DSh).

1.6.3 Example. Let $\Delta = (2,3)$, let A be a Δ-structure,

$$A = \Big(\{1,2,\dots,6\}, \big(\{(3,2),(6,3),(6,5)\}, \{(1,5,6),(4,4,1),(4,5,2)\}\big)\Big).$$

The the directed shadow $\mathrm{DSh}(A)$ and the shadow $\mathrm{Sh}(A)$ of the Δ-structure A are shown in Fig. 1.5. The loop at the vertex 4 is caused by the triple $(4,4,1) \in R_3(A)$.

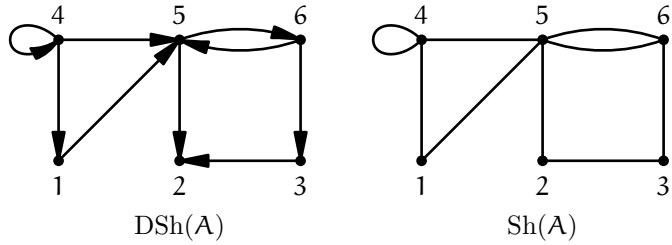

Figure 1.5: The directed shadow and the shadow of a Δ-structure

Shadows "preserve homomorphisms" – a homomorphism of Δ-structures is also a homomorphism of their shadows. The converse, however, is not true in general.

1.6.4 Lemma. *If* $f : A \to B$ *is a homomorphism of* Δ*-structures, then* f *is a graph homomorphism from* $\mathrm{Sh}(A)$ *to* $\mathrm{Sh}(B)$.

Proof. If $\{u,v\}$ is an edge of $\mathrm{Sh}(A)$, then by definition there is an edge $e \in R_i(A)$ for some $i \in I$ such that u and v appear as consecutive vertices in e. Therefore $f(u)$ and $f(v)$ appear as consecutive vertices in the edge $f(e)$ of B, hence $\{u,v\}$ is an edge of $\mathrm{Sh}(B)$. □

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

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

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

1.6.6 Example. The incidence graph $\mathrm{Inc}(A)$ of the Δ-structure A from Example 1.6.3 is shown in Figure 1.6.

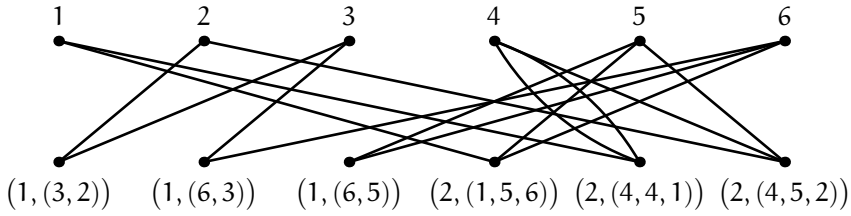

Figure 1.6: The incidence graph

Next we formulate three structural conditions and show that they are equivalent with each other as well as with the connectedness of the shadow and the incidence graph.

1.6.7 Lemma. *For a core* G*, the following are equivalent:*
(1) If $G \to A + B$ *for some structures* A, B*, then* $G \to A$ *or* $G \to B$*.*
(2) If $G \sim A + B$ *for some structures* A, B*, then* $B \to A$ *or* $A \to B$*.*
(3) If $G \sim A + B$ *for some structures* A, B*, then* $G \sim A$ *or* $G \sim B$*.*
(4) The shadow $\mathrm{Sh}(G)$ *is connected.*
(5) The incidence graph $\mathrm{Inc}(G)$ *is connected.*

Proof.
(1) $\Rightarrow$ (2): If $G \sim A + B$, then $G \to A + B$, and using (1) we have $G \to A$ or $G \to B$. In the first case $B \to A + B \sim G \to A$, hence $B \to A$. In the latter case $A \to A + B \sim G \to B$, and so $A \to B$.
(2) $\Rightarrow$ (3): Suppose $G \sim A + B$. By (2) we have $B \to A$, and therefore $A \sim A + B \sim G$; or we have $A \to B$, and then $B \sim A + B \sim G$.
(3) $\Rightarrow$ (1): Let $G \to A+B$. Using distributivity, we have $(G \times A)+(G \times B) \sim G \times (G + A) \times (G + B) \times (A + B) \sim G$, since G is homomorphic to all other factors.
(3) $\Rightarrow$ (5): Suppose that (3) holds but $\mathrm{Inc}(G)$ is disconnected. Let A' be a component of $\mathrm{Inc}(G)$ and $B' = \mathrm{Inc}(G) - A'$; let A be the substructure of G induced by $V(A') \cap \underline{G}$ and let B be the substructure of G induced by $V(B') \cap \underline{G}$. Then $G = A + B$ but both A and B are proper substructures of G, so if $G \sim A$ or $G \sim B$, then G is not a core, a contradiction.
(5) $\Rightarrow$ (4): In $\mathrm{Inc}(G)$, if any two vertices $u, v \in \underline{G}$ have a common neighbour e in $\mathrm{Block}(G)$, they belong to the same edge e of G, and so there is a path from u to v in $\mathrm{Sh}(G)$. Therefore the existence of a path from u to v in $\mathrm{Inc}(G)$ implies the existence of a path from u to v in $\mathrm{Sh}(G)$.
(4) $\Rightarrow$ (1): First observe that since every edge of $A + B$ is either an edge of A or an edge of B, we have that $\mathrm{Sh}(A + B) = \mathrm{Sh}(A) + \mathrm{Sh}(B)$. Let f :

$G \to A + B$. Then $f : \mathrm{Sh}(G) \to \mathrm{Sh}(A + B) = \mathrm{Sh}(A) + \mathrm{Sh}(B)$ by Lemma 1.6.4. Because $\mathrm{Sh}(G)$ is connected, $f[\underline{G}] \subseteq \underline{A}$ or $f[\underline{G}] \subseteq \underline{B}$: otherwise there is an edge $\{u, v\}$ of $\mathrm{Sh}(G)$ such that $f(u)$ is a vertex of $\mathrm{Sh}(A)$ and $f(v)$ is a vertex of $\mathrm{Sh}(B)$, but then $\{f(u), f(v)\}$ is not an edge of $\mathrm{Sh}(A) + \mathrm{Sh}(B)$, a contradiction with f being a homomorphism. Therefore $f : G \to A$ or $f : G \to B$. □

1.6.8. The conditions (1)–(3) above are equivalent even for structures that are not cores; if a Δ-structure G satisfies (1)–(3), its core satisfies all the conditions (1)–(5). Similarly, the conditions (4)–(5) are equivalent for all structures G.

1.6.9 Definition. A Δ-structure is called **connected** if it satisfies the equivalent conditions (4)–(5) of Lemma 1.6.7. Maximal connected substructures of a Δ-structure A are called the **components** of A.

1.6.10. It is easy to see that every Δ-structure is the sum of its components; and that the decomposition into components is unique.

A part of the previous lemma holds in the dual category too.

1.6.11 Lemma. *For a structure* G*, the following are equivalent:*
(1) If $A \times B \to G$ *for some structures* A, B*, then* $A \to G$ *or* $B \to G$*.*
(2) If $G \sim A \times B$ *for some structures* A, B*, then* $B \to A$ *or* $A \to B$*.*
(3) If $G \sim A \times B$ *for some structures* A, B*, then* $G \sim A$ *or* $G \sim B$*.*

Proof. Repeat the proof of Lemma 1.6.7; reverse all arrows and replace $+$ with $\times$. □

1.6.12 Definition. A Δ-structure is called **irreducible** if it satisfies the equivalent conditions of Lemma 1.6.11.

1.6.13. One might expect to find conditions involving $\mathrm{Inc}(G)$ and $\mathrm{Sh}(G)$, similar to (4) and (5) of Lemma 1.6.7 to be equivalent with irreducibility of a Δ-structure. But in spite of being dual to connectedness, irreducibility is much more tricky. For example, it is an easy exercise to devise an efficient algorithm for testing connectedness of a Δ-structure. On the other hand, irreducibility is not even known to be decidable.

1.6.14. Another difference from connectedness is that it is not true that every Δ-structure can be decomposed as a finite product of irreducible structures (dually to the decomposition into connected components).

Here it is worthwhile to mention that our terminology is inspired by algebra, and lattice theory in particular (see Section 1.9). In the past, other names for this property have been used; namely *productive* in [30] and *multiplicative* in [14] and by many other authors. We believe that the word *irreducible* fits the meaning of the property better.

The term *multiplicative* is motivated by a property of irreducible structures described in the next paragraphs.

Hedetniemi's product conjecture

1.6.15 Definition. Let $\mathcal{A}$ be a class of Δ-structures that is closed under homomorphic equivalence. We say that $\mathcal{A}$ is **multiplicative** if $\mathbf{A}_1 \in \mathcal{A}$ and $\mathbf{A}_2 \in \mathcal{A}$ implies that $\mathbf{A}_1 \times \mathbf{A}_2 \in \mathcal{A}$.

A famous conjecture of Hedetniemi [15] states that the chromatic number of the product of two (undirected) graphs is equal to the minimum of the chromatic numbers of the two graphs. This conjecture is equivalent to the following.

1.6.16 Conjecture. *The class $\bar{\mathcal{K}}_k$ of all graphs that are not k-colourable is multiplicative for every positive integer k.*

For some graph classes, such as the class of all graphs that are not homomorphic to a fixed graph, multiplicativity has an equivalent description. This is in fact not restricted to graphs, but holds for relational structures as well.

1.6.17 Proposition. *Let $\mathbf{H}$ be a Δ-structure. Then the class*

$$\mathcal{H} := \{\mathbf{X} : \mathbf{X} \not\to \mathbf{H}\}$$

of all Δ-structures that are not homomorphic to $\mathbf{H}$ is multiplicative if and only if $\mathbf{H}$ is irreducible.

Proof. Let $\mathbf{H}$ be irreducible. By condition (1) of Lemma 1.6.11, if the product $\mathbf{A}_1 \times \mathbf{A}_2$ is homomorphic to $\mathbf{H}$, then $\mathbf{A}_1 \to \mathbf{H}$ or $\mathbf{A}_2 \to \mathbf{H}$. Thus if $\mathbf{A}_1 \times \mathbf{A}_2 \notin \mathcal{H}$, then $\mathbf{A}_1 \notin \mathcal{H}$ or $\mathbf{A}_2 \notin \mathcal{H}$.

Conversely, if $\mathcal{H}$ is multiplicative, then condition (1) of Lemma 1.6.11 is satisfied. Hence the Δ-structure $\mathbf{H}$ is irreducible. □

Recall from Example 1.3.5 that a graph is k-colourable if and only if it is homomorphic to the complete graph $\mathbf{K}_k$. Thus Hedetniemi's conjecture has another equivalent formulation.

1.6.18 Conjecture. *The complete graph* K_k *is irreducible for every positive integer* k.

The conjecture is evidently true for $k = 1$. It merely says that if the product $A_1 \times A_2$ has no edges, then A_1 or A_2 has no edges either.

It is not difficult to show that if both A_1 and A_2 are non-bipartite, then $A_1 \times A_2$ is non-bipartite too; this is the case $k = 2$. El-Zahar and Sauer [6] have proved that K_3 is irreducible. Almost no other examples of irreducible (undirected) graphs are known, though. Some more irreducible graphs have been found by Tardif [41]. Nevertheless, Hedetniemi's product conjecture remains wide open for general k.

1.7 Paths, trees and forests

Trees are very important in the context of homomorphism dualities (see especially Theorem 2.1.12). We define them as structures whose shadow is a tree. Forests are then structures consisting of tree components. In addition, we define paths in Δ-structures as "linear trees": every edge has at most two neighbouring edges.

1.7.1 Definition. A Δ-structure A is called a **Δ-tree** or simply a **tree** if $\mathrm{Sh}(A)$ is a tree; it is called a **Δ-forest** or just a **forest** if $\mathrm{Sh}(A)$ is a forest.

1.7.2. Notice that A is a Δ-tree if and only if $\mathrm{Inc}(A)$ is a tree; and that A is a Δ-forest if and only if each component of A is a Δ-tree.

1.7.3 Definition. A Δ-tree P is called a **Δ-path** if every edge of P intersects at most two other edges and every vertex of P belongs to at most two edges.

1.7.4. In every Δ-path with at least two edges there are two edges (**end edges**) such that each of them shares a vertex with exactly one other edge. Any other edge (**middle edge**) shares two of its vertices, each with one other edge.

As a generalisation of acyclic directed graphs, that is directed graphs without directed cycles, we introduce acyclic relational structures.

1.7.5 Definition. A Δ-structure A is called **acyclic** if there is no directed cycle in its directed shadow $\mathrm{DSh}(A)$.

Evidently, every Δ-tree is acyclic. We need this fact especially in Section 2.2.

1.8 Height labelling and balanced structures

Balanced digraphs are digraphs that are homomorphic to a directed path. The name originates from the fact that the number of forward edges and the number of backward edges is the same along every cycle in a balanced digraph (and such cycles are called balanced). For digraphs being homomorphic to a directed path is the same as being homomorphic to an oriented forest. For relational structures it is not, and the forest definition is the suitable one.

1.8.1 Definition. We say that a Δ-structure A is **balanced** if A is homomorphic to a Δ-forest.

For a balanced digraph, the height of a vertex is defined as the length of the longest directed path ending in that vertex. As a generalisation, we define height labelling of relational structures.

1.8.2 Definition. Let A be a Δ-structure and let $\beth$ be a labelling of its vertices with $(\sum_{i\in I}\delta_i - |I|)$-tuples of integers, indexed by $(i,1),(i,2),\dots,(i,\delta_i-1)$, $i \in I$.

We say that $\beth$ is a **height labelling** of A if

- there exists a vertex x of A such that $\beth(x) = (0,0,\dots,0)$, and
- whenever $(x_1, x_2, \dots, x_{\delta_i}) \in R_i(A)$ and $1 \le j < \delta_i$, then

$$\begin{aligned} \big(\beth(x_{j+1})\big)_{(i,j)} &= \big(\beth(x_j)\big)_{(i,j)} + 1, \quad \text{and} \\ \big(\beth(x_{j+1})\big)_{(i',j')} &= \big(\beth(x_j)\big)_{(i',j')} \quad \text{for } (i',j') \neq (i,j). \end{aligned} \tag{1.1}$$

The first condition in the above definition is purely technical; it facilitates the exposition of arguments. It is possible to omit it without altering the essence of the definition.

1.8.3 Proposition. *If A is a balanced Δ-structure, then A has a height labelling. If a height labelling of a connected structure exists, it is unique up to an additive constant vector.*

Proof. First we prove that a height labelling exists for any Δ-tree T: pick an arbitrary vertex x and set $\beth(x) = (0,0,\dots,0)$. On all other vertices the labelling is defined recursively. If a neighbour y of x in the shadow $\mathrm{Sh}(T)$ has already been assigned a label, the label $\beth(x)$ is determined by the conditions (1.1). In this way, a label is assigned to every vertex, because $\mathrm{Sh}(T)$ is a tree.

By separately labelling each component, we get a height labelling also for any Δ-forest.

Now if $f : A \to T$ is a surjective homomorphism from a balanced Δ-structure A to a Δ-forest T, let us fix a height labelling of T and for a vertex x of A define

$$\beth(x) := \beth\big(f(x)\big).$$

The labelling defined in this way is a height labelling; it is not difficult to check the conditions (1.1).

Uniqueness follows from the fact that the difference between the labels of two vertices of A depends only on what edges of A generated the path between the two vertices in the shadow $\mathrm{Sh}(A)$. □

1.8.4. We may picture height labelling as if we replace each edge of the ith kind by $\delta_i - 1$ binary edges of various kinds and then count forward edges and backward edges. A height labelling for a structure exists if and only if each path (in the shadow) between any fixed pair of vertices counts the same difference of the numbers of forward edges and backward edges for all kinds. This is related to balanced structures and the height of a path, as defined for digraphs in [38].

1.9 Partial orders

Here we state the definitions and elementary facts about partial orders and lattices that we need later (chiefly in Chapter 3). Several introductory books on this topic are available, such as the recommended one by Davey and Priestley [4].

1.9.1 Definition. A **partial order** is a binary relation $\preccurlyeq$ over a set P which is reflexive, antisymmetric, and transitive, that is for all a, b, and c in P, we have that:

1. $a \preccurlyeq a$ (reflexivity);
2. if $a \preccurlyeq b$ and $b \preccurlyeq a$ then $a = b$ (antisymmetry); and
3. if $a \preccurlyeq b$ and $b \preccurlyeq c$ then $a \preccurlyeq c$ (transitivity).

A set with a partial order is called a **partially ordered set** or a **poset**. So formally, a partially ordered set is an ordered pair $(P, \preccurlyeq)$, where P is called the **base set** and $\preccurlyeq$ is a partial order over P.

1.9.2 Definition. A subset $Q \subseteq P$ of the poset $(P, \preccurlyeq)$ is a **downset** if, for all elements p and q, if p is less than or equal to q and q is an element of Q, then p is also in Q. For an arbitrary subset $S \subseteq P$ the **downset generated by** S is the smallest downset that contains S; it is denoted by $S^{\downarrow}$.

1.9.3 Definition. Similarly, a subset $Q \subseteq P$ of the poset $(P, \preccurlyeq)$ is an **upset** if, for all elements p and q, if $q \preccurlyeq p$ and q is an element of Q, then p is also in Q. For an arbitrary subset $S \subseteq P$ the **upset generated by** S is the smallest upset that contains S; it is denoted by $S^{\uparrow}$.

1.9.4 Definition. Suppose $\preccurlyeq$ is a partial order on a nonempty set P. Then the elements $p, q \in P$ are said to be **comparable** provided $p \preccurlyeq q$ or $q \preccurlyeq p$. Otherwise they are called **incomparable**. A subset Q of P is an **antichain** if all elements of Q are pairwise incomparable. An antichain is **maximal** if there is no other antichain strictly containing it.

1.9.5 Definition. Given a subset Q of a poset P, the **supremum** of Q is the least element of P that is greater than or equal to each element of Q; so the supremum of Q is an element u in P such that

1. $x \preccurlyeq u$ for all x in Q, and
2. for any v in P such that $x \preccurlyeq v$ for all x in Q it holds that $u \preccurlyeq v$.

Dually, the **infimum** of Q is the greatest element of P that is less than or equal to each element of Q.

1.9.6 Definition. A **lattice** is a poset whose subsets of size two all have a supremum (called **join**; the join of p and q is denoted by $p \vee q$) and an infimum (called **meet**; the meet of p and q is denoted by $p \wedge q$).

There are many natural examples of lattices.

1.9.7 Example. The natural numbers in their usual order form a lattice, under the operations of minimum and maximum.

The positive integers also form a lattice under the operations of taking the greatest common divisor and least common multiple, with divisibility as the order relation: $a \leq b$ if a divides b.

1.9.8 Example. For any set A, the collection of all subsets of A can be ordered via subset inclusion to obtain a lattice bounded by A itself and the empty set. Set intersection and union interpret meet and join, respectively.

1.9.9 Example. All subspaces of a vector space V form a lattice. Here the meet is the intersection of the subspaces and the join of two subspaces W and W' is the minimal subspace that contains the union $W \cup W'$.

An important property of lattices is distributivity. We meet it again in Lemma 3.3.2 and the following paragraphs.

1.9.10 Definition. A lattice L is **distributive** if the following identity holds for all x, y, and z in L:

$$x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z).$$

This says that the meet operation preserves non-empty finite joins. It is a basic fact of lattice theory that the above condition is equivalent to its dual:

$$x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z).$$

1.9.11 Example. Not every lattice is distributive. From our examples above, the three lattices in 1.9.7 and 1.9.8 are all distributive. However, the lattice of subspaces of a vector space from 1.9.9 is not distributive: Consider the space $\mathbb{R}^2$ and let A, B and C be three distinct subspaces of dimension 1 (straight lines). Then $A \vee B = \mathbb{R}^2$ and $A \cap C = B \cap C = 0$, and thus

$$C = (A \vee B) \cap C \neq (A \cap C) \vee (B \cap C) = 0.$$

Next we introduce the notion of a gap, which is an island of non-density in a partially ordered set.

1.9.12 Definition. A pair (p, q) of elements of a poset P is a **gap** if $p \prec q$, and for every $r \in P$, if $p \preccurlyeq r \preccurlyeq q$ then $r = p$ or $r = q$.

In a dense poset there exists r such that $p \prec r \prec q$ for any pair (p, q) with $p \prec q$. Thus a poset is dense if and only if it has no gaps. Gaps are further discussed in Sections 3.2 and 3.3.

2 Homomorphism dualities

> It was one more argument
> to support his theory that
> nice things are nicer
> than nasty ones.
>
> *(Kingsley Amis, Lucky Jim)*

Homomorphism dualities are situations where a class of relational structures is characterised in two ways: by the non-existence of a homomorphism *from* some fixed set of structures, and by the existence of a homomorphism *to* some other fixed set of structures.

For example, an (undirected) graph is bipartite if and only if it is homomorphic to the complete graph K_2; so the existence of a homomorphism to K_2 is determined by the non-existence of a homomorphism from odd cycles.

An important aspect of dualities is that in some cases they make the respective class more accessible. Duality guarantees the existence of a certificate for positive as well as negative answers to the membership problem. In both cases it is a homomorphism, either a homomorphism from the "forbidden" set or a homomorphism to the other ("dual") set of structures.

In special cases this provides an example of a *good characterisation* in the sense of Edmonds [5]. This may mean that an effective algorithm for testing membership is available.

Homomorphism dualities have been studied for many years now. Their roots go back to the early 1970's, appearing already in Nešetřil's textbook on graph theory [26]. The pioneering work was done by Nešetřil and Pultr [30]. They found the first instances of homomorphism dualities – the duality pairs of directed paths and transitive tournaments (see Example 2.1.9). More duality pairs were later discovered by Komárek [22], and his work led to a characterisation theorem for digraphs [23]. Nešetřil and Tardif [34] found a connection to the homomorphism order (Chapter 3) and generalised the notion for relational structures and together with the author of this thesis they have recently found a full characterisation [13] (Theo-

rem 2.4.26 here). Meanwhile, dualities have again been generalised in two different contexts [29, 31].

In this chapter, we first investigate dualities with homomorphisms to a single structure characterised by the non-existence of a homomorphism from a single structure (duality pairs) and show some properties of such structures. Then we fully characterise finite dualities. Finally, we address some extremal problems related to the size of the involved structures.

2.1 Duality pairs

Duality pairs are the simplest cases of homomorphism dualities. Here a class of relational structures is characterised by the existence of a homomorphism to a single structure $\mathbb{D}$ and at the same time by the non-existence of a homomorphism from a single structure $\mathbb{F}$.

2.1.1 Definition. Let $\mathbb{F}$, $\mathbb{D}$ be relational structures. We say that the pair $(\mathbb{F}, \mathbb{D})$ is a **duality pair** if for every structure $\mathbb{A}$ we have $\mathbb{F} \to \mathbb{A}$ if and only if $\mathbb{A} \not\to \mathbb{D}$.

One should view the Δ-structure $\mathbb{F}$ as a characteristic obstacle, which prevents a Δ-structure from being homomorphic to $\mathbb{D}$.

We introduce the following short notation for a symbolic description of duality pairs.

2.1.2. Let

$$\mathbb{F}\not\to := \{\mathbb{A} : \mathbb{F} \not\to \mathbb{A}\},$$
$$\to\mathbb{D} := \{\mathbb{A} : \mathbb{A} \to \mathbb{D}\}.$$

Then $(\mathbb{F}, \mathbb{D})$ is a duality pair if and only if

$$\mathbb{F}\not\to = \to\mathbb{D}.$$

For proving that a certain pair of structures is indeed a duality pair, the following characterisation is frequently convenient.

2.1.3 Lemma. *Let $\mathbb{F}$ and $\mathbb{D}$ be Δ-structures. Then $(\mathbb{F}, \mathbb{D})$ is a duality pair if and only if $\mathbb{F} \not\to \mathbb{D}$, and whenever $\mathbb{F} \not\to \mathbb{A}$, then $\mathbb{A} \to \mathbb{D}$.*

Proof. If $(\mathbb{F}, \mathbb{D})$ is a duality pair, then $\mathbb{F} \to \mathbb{F}$ and hence $\mathbb{F} \not\to \mathbb{D}$ by the definition of a duality pair. This proves one implication.

Next is the proof of the opposite implication. According to the definition, it has to be shown that if $\mathbb{F} \to \mathbb{A}$, then $\mathbb{A} \not\to \mathbb{D}$. Suppose that $\mathbb{F} \to \mathbb{A}$

and $A \to D$. Then by composition of homomorphisms $F \to D$. That is a contradiction. □

2.1.4. In our notation for duality pairs, the letter F stands for *forbidden* and the letter D for *dual*.

As a warm-up, we state a few results, which are not difficult to prove but they provide the first experience with homomorphism dualities.

The first proposition states the obvious: that replacing F or D in a duality pair with a homomorphically equivalent structure does not change the duality.

2.1.5 Proposition. *Let* (F, D) *be a duality pair. If* F' *is homomorphically equivalent to* F*, and* D' *is homomorphically equivalent to* D*, then* (F', D') *is also a duality pair.*

Proof. By Lemma 2.1.3 we have $F \not\to D$. By the same lemma it suffices to prove that $F' \not\to D'$, and that whenever $F' \not\to A$, then $A \to D'$.

First we prove that $F' \not\to D'$. Suppose on the contrary that $F' \to D'$. Then $F \to D$ because $F \to F' \to D' \to D$ since homomorphisms compose. That is a contradiction.

Now we show that whenever $F' \not\to A$, then $A \to D'$. Suppose that $F' \not\to A$, then also $F \not\to A$, so $A \to D$ because (F, D) is a duality pair. Since $D \to D'$, we have $A \to D'$ by composition of homomorphisms. □

Next we show that in a duality pair (F, D) the structure D is uniquely determined by F up to homomorphic equivalence. Vice versa, up to homomorphic equivalence the structure F is determined by D.

2.1.6 Proposition. *Let* F*,* F'*,* D*,* D' *be* Δ*-structures. If* (F, D) *and* (F, D') *are duality pairs, then* D *and* D' *are homomorphically equivalent. If* (F, D) *and* (F', D) *are duality pairs, then* F *and* F' *are homomorphically equivalent.*

Proof. Suppose that both (F, D) and (F, D') are duality pairs. Lemma 2.1.3 implies that $F \not\to D$ since (F, D) is a duality pair, so $D \to D'$ because (F, D') is a duality pair. Moreover $F \not\to D'$ because (F, D') is a duality pair, so $D \to D'$ because (F, D) is a duality pair. Hence $D \sim D'$.

The proof of the second part is analogous. □

2.1.7 Corollary. *If* (F, D) *is a duality pair, then there exists a unique core* D' *such that* (F, D') *is a duality pair.* □

This corollary motivates the following definition.

2.1.8 Definition. Let D be a core and let (F, D) be a duality pair. Then the Δ-structure D is called **the dual** of F; it is denoted by $D(F)$.

Now we present the first example of duality pairs.

2.1.9 Example ([30])**.** Let us first consider digraphs, that is Δ-structures with $\Delta = (2)$. Let $\vec{P}_k$ denote the directed path with k edges and $\vec{T}_k$ the transitive tournament on k vertices. Then for $k \geq 1$ the pair $(\vec{P}_k, \vec{T}_k)$ is a duality pair.

Proof. We proceed by induction on k. For $k = 1$, clearly $\vec{P}_1 \nrightarrow \vec{T}_1$, and if $\vec{P}_1 \nrightarrow A$, then A has no edges, whence $A \to \vec{T}_1$. So $(\vec{P}_1, \vec{T}_1)$ is a duality pair by Lemma 2.1.3.

Now we prove the induction step for $k \geq 2$. Again, obviously $\vec{P}_k \nrightarrow \vec{T}_k$. Moreover, suppose that $\vec{P}_k \nrightarrow A$. Then the digraph A is acyclic (it contains no directed cycles). Therefore there exists a vertex of A with out-degree zero, that is with no outward edges going from it. Let A' be the digraph created from A by deleting all vertices with out-degree zero. It is easy to see that $\vec{P}_{k-1} \nrightarrow A'$, so by induction we have that $f' : A' \to \vec{T}_{k-1}$. We can extend the homomorphism f' to a homomorphism from A to $\vec{T}_k$: the vertices of A' are mapped in the same way as by f', to the subgraph of $\vec{T}_k$ consisting of non-zero out-degree vertices, which is isomorphic to $\vec{T}_{k-1}$; and the remaining vertices of A are mapped to the terminal vertex (sink) of $\vec{T}_k$. □

The next example is due to Komárek [22]. It is presented here without proof. However, let us point out that the example is the essence of the mosquito construction (see Section 2.2) and historically it was an important milestone on the way to the description of all duality pairs.

2.1.10 Example ([22])**.** For two positive integers m, n, define the digraph $P_{m,n} = (V, E)$ to be the oriented path with vertices

$$V = \{a_0, a_1, \ldots, a_m, b_0, b_1, \ldots, b_n\}$$

and edges

$$E = \{(a_j, a_{j+1}) : j = 0, 1, \ldots, m-1\} \cup \{(b_j, b_{j+1}) : j = 0, 1, \ldots, n-1\} \cup \{(b_0, a_m)\}$$

(see Figure 2.1).

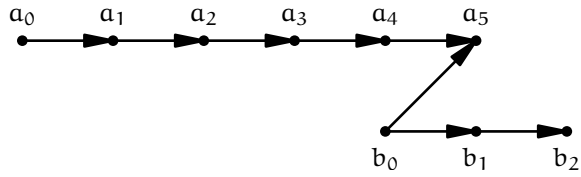

Figure 2.1: The path $P_{5,2}$

Furthermore, let $D_{m,n} = (W, F)$ be the digraph with the vertex set defined by

$$W = \{(i,j) : i \geq 0,\ j \geq 0,\ 0 \leq i+j \leq m+n-2\}$$

and edges by

$$F = \Big\{\big((i,j),(i',j')\big) : i < i',\ j > j',\ \text{and } i < m \text{ or } j < n\Big\}.$$

An example is in Figure 2.2.

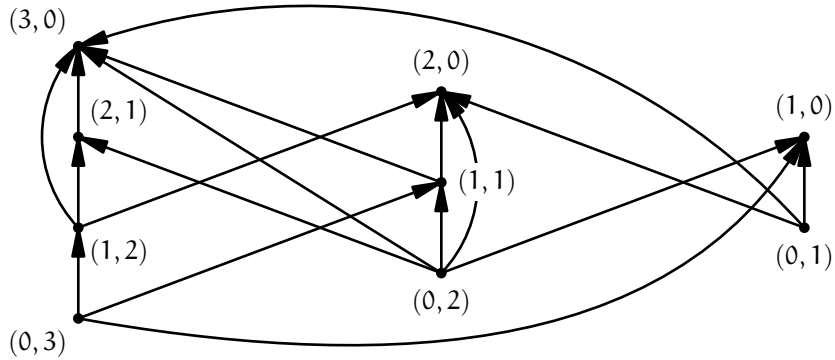

Figure 2.2: The digraph $D_{2,3}$

Then $(P_{m,n}, D_{m,n})$ is a duality pair.

In the next example, we consider Δ-structures with more than one relation. It shows that if some relations of a Δ-structure F are empty, then in the dual structure the corresponding relations contain all possible tuples of vertices.

2.1.11 Example. Let $\Delta = (\delta_i : i \in I)$, let $\mathbf{F} = \big(\underline{F}, (R_i : i \in I)\big)$ be a Δ-structure and let $I' \subseteq I$ be a set of indices such that $R_i = \emptyset$ for all $i \in I \setminus I'$. Define $\Delta' = (\delta_i : i \in I')$ and let $\mathbf{F}'$ be the Δ'-structure $\big(\underline{F}, (R_i : i \in I')\big)$. Suppose that there exists a Δ'-structure $\mathbf{D}'$ such that $(\mathbf{F}', \mathbf{D}')$ is a duality pair. By adding complete relations to $\mathbf{D}'$ we get the Δ-structure $\mathbf{D}$ with

$$\underline{D} = \underline{D'},$$
$$R_i(\mathbf{D}) = \begin{cases} R_i(\mathbf{D}') & \text{if } i \in I', \\ (\underline{D})^{\delta_i} & \text{if } i \in I \setminus I'. \end{cases}$$

Then the pair $(\mathbf{F}, \mathbf{D})$ is a duality pair.

Proof. We use Lemma 2.1.3. First, $\mathbf{F} \nrightarrow \mathbf{D}$ because $\mathbf{F}' \nrightarrow \mathbf{D}'$. Now let $\mathbf{A}$ be a Δ-structure such that $\mathbf{F} \nrightarrow \mathbf{A}$. This means that $\mathbf{F}' \nrightarrow \mathbf{A}' = \big(\underline{A}, (R_i(\mathbf{A}) : i \in I')\big)$ and therefore there exists a homomorphism $f' : \mathbf{A}' \to \mathbf{D}'$. Obviously $f : a \mapsto f'(a)$ is a homomorphism from $\mathbf{A}$ to $\mathbf{D}$. □

The following theorem characterises all duality pairs.

2.1.12 Theorem ([34]). *If* $(\mathbf{F}, \mathbf{D})$ *is a duality pair, then* $\mathbf{F}$ *is homomorphically equivalent to a* Δ*-tree. Conversely, if* $\mathbf{F}$ *is a* Δ*-tree with more than one vertex, then there exists a unique (up to homomorphic equivalence) structure* $\mathbf{D}$ *such that* $(\mathbf{F}, \mathbf{D})$ *is a duality pair.*

Proof. The proof is split into two parts. Here we prove that if $(\mathbf{F}, \mathbf{D})$ is a duality pair and $\mathbf{F}$ is a core, then $\mathbf{F}$ is a tree. The second part of the proof can be found in Section 2.2, which contains a construction of the dual structure for any tree $\mathbf{F}$. Uniqueness follows from Proposition 2.1.6.

Our proof uses an idea of Komárek [23], who proved the characterisation of duality pairs for digraphs. We assume that in a duality pair $(\mathbf{F}, \mathbf{D})$ the Δ-structure $\mathbf{F}$ is a core that is not a Δ-tree. The idea of the proof is to construct an infinite sequence of Δ-structures $\mathbf{F}_k$ such that $\mathbf{F}$ is not homomorphic to any of them. By duality, all the structures $\mathbf{F}_k$ are homomorphic to $\mathbf{D}$, and we will show that this implies that $\mathbf{F} \to \mathbf{D}$. That is a contradiction with the definition of a duality pair (see Lemma 2.1.3).

Hence if $(\mathbf{F}, \mathbf{D})$ is a duality pair and $\mathbf{F}$ is a core, then $\mathbf{F}$ is a Δ-tree. If $(\mathbf{F}', \mathbf{D})$ is a duality pair and $\mathbf{F}'$ is arbitrary, then by Proposition 2.1.6 also the pair $(\mathbf{F}, \mathbf{D})$ is a duality pair, where $\mathbf{F}$ is the core of $\mathbf{F}'$. So $\mathbf{F}$ is a Δ-tree, and therefore $\mathbf{F}'$ is homomorphically equivalent to a Δ-tree.

This part of the proof is split into eight steps (2.1.13–2.1.20).

2.1.13. F *is connected.* Because: If $F \to A + B$, then $A + B \nrightarrow D$, and so either A or B is not homomorphic to D, and so $F \to A$ or $F \to B$. This is one of the equivalent descriptions of connectedness in Lemma 1.6.7.

We proceed to define the structures F_k for all positive integers k.

2.1.14. We suppose that F is not a Δ-tree; thus $\mathrm{Sh}(F)$ contains a cycle. Let $\{u, v\}$ be an edge of $\mathrm{Sh}(F)$ that lies in a cycle. This edge was added to the shadow $\mathrm{Sh}(F)$ because there is an edge $e \in R_j(F)$ that contains u and v as consecutive vertices. Now, F_1 will be the Δ-structure constructed from F by removing the edge e. Furthermore, F_2 is constructed by taking two copies of F_1 and by joining them by two edges of the jth kind in the way depicted in Figure 2.3.

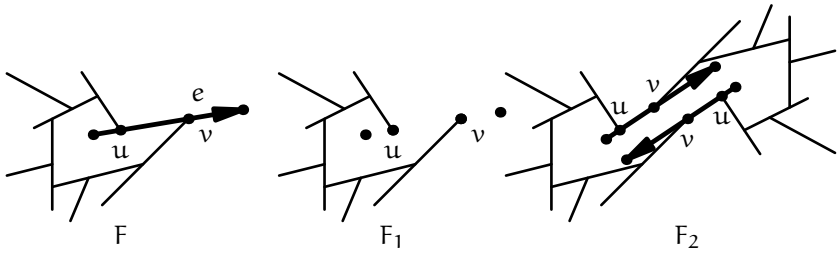

Figure 2.3: The construction of F_1 and F_2

In general, for $k \geq 2$, take k copies of F_1 and join each with all other copies by edges of the jth kind. Two edges are used to connect each pair of copies of F_1, so altogether $2\binom{k}{2}$ new edges are introduced. The resulting structure is F_k.

Formally, if $e = (e_1, \ldots, u, v, \ldots, e_{\delta_j})$, define F_k in the following way:

$$\begin{aligned}
\underline{F_k} &:= \{1, 2, \ldots, k\} \times \underline{F}, \\
R_j(F_k) &:= \Big\{ \big((q, x_1), (q, x_2), \ldots, (q, x_{\delta_j})\big) : \\
&\qquad e \neq (x_1, x_2, \ldots, x_{\delta_j}) \in R_j(F),\ 1 \leq q \leq k \Big\} \\
&\quad \cup \Big\{ \big((q, e_1), (q, e_2), \ldots, (q, u), (q', v), \ldots, (q', e_{\delta_j})\big) : \\
&\qquad 1 \leq q, q' \leq k,\ q \neq q' \Big\},
\end{aligned}$$

$$R_i(F_k) := \Big\{ \big((q,x_1),(q,x_2),\ldots,(q,x_{\delta_i})\big) : (x_1,x_2,\ldots,x_{\delta_i}) \in R_i(F),\ 1 \le q \le k \Big\} \qquad \text{for } i \ne j.$$

Next we prove several properties of the structures F_k.

2.1.15. $F_k \to F$. Because: The identity mapping is indeed a homomorphism from F_1 to F. For $k \ge 2$, mapping each vertex (q,x) to the corresponding vertex x of F provides a homomorphism (let us call this homomorphism h; so $h : (q,x) \mapsto x$).

2.1.16. $F \not\to F_1$. Because: The structures F and F_1 have the same number of vertices but F_1 has fewer edges. So a potential homomorphism from F to F_1 cannot be injective and so it would actually map F to a proper substructure of F. This is impossible because F is a core.

In the following four paragraphs, let $n := |\underline{F}|$ and let the vertices be enumerated in such a way that $\underline{F} = \{x_1, x_2, \ldots, x_n\}$ with $u = x_1$ and $v = x_2$.

2.1.17. *If $f : F \to F_k$ is a homomorphism, then the image $f[F]$ consists of exactly one copy of each vertex of F_1 in F_k. In particular, any homomorphism $f : F \to F_k$ is injective.* Because: If $f : F \to F_k$ is a homomorphism, then the image $f[F]$ contains vertices of at least two distinct copies of F_1 in F_k because of 2.1.16. If some vertex appears in more than one copy, say $f(x_i) = (p, x_l)$ and $f(x_{i'}) = (q, x_l)$ for some $p \ne q$, then the composed homomorphism $hf : F \to F$ is not an automorphism. This is a contradiction, since F is a core. Here h is the homomorphism $(q,x) \mapsto x$ defined in 2.1.15.

2.1.18. *If $F \to F_k$ for some k, then $F \to F_2$.* Because: Let $f : F \to F_k$. By 2.1.17, the image $f[F]$ contains exactly one vertex (p, x_1) and exactly one vertex (q, x_2). As F is connected, the image contains vertices of only at most two copies of F_1 in F_k, namely vertices in the form (p, x_i) and (q, x_i). So the image $f[F]$ lies entirely in a substructure of F_k that is isomorphic to F_2. Thus $F \to F_2$.

2.1.19. $F \not\to F_k$. Because: Suppose $F \to F_k$. Then there is a homomorphism $f : F \to F_2$ because of 2.1.18. We may assume that $(1, x_1)$ and $(2, x_2)$ are images under f of some vertices of F.

Now, $f[\underline{F}] = V_1 \cup V_2$, where V_1 contains vertices in the first copy and V_2 vertices in the second copy of F_1 in F_2, that is $V_l := \{x : x \in f[\underline{F}] \text{ and } x = (l, x_i) \text{ for some } i\}$ for $l = 1, 2$. Then $(1, x_1) \in V_1$ and $(2, x_2) \in V_2$.

For a Δ-structure A, let $r(A)$ denote the number of all edges of A, so

$$r(A) := \Big|\bigcup_{i\in I} R_i(A)\Big|.$$

Let G be the substructure induced in F_2 by $V_1 \cup V_2$. The homomorphism f is injective by 2.1.17, and as a result it maps every edge of F to an edge of G. Thus $r(G) \geq r(F)$.

The edges of G are split into three groups: the edges induced by V_1 (let their number be m_1), the edges induced by V_2 (let their number be m_2) and the edge

$$\big((1,e_1),(1,e_2),\ldots,(1,u),(2,v),\ldots,(2,e_{\delta_j})\big).$$

Hence $r(G) = m_1 + m_2 + 1$.

Let the base set $\underline{F}$ be partitioned into W_1 and W_2 by setting $W_l := \big\{x : (l,x) \in V_l\big\}$. So W_1 consists of those vertices x of F whose copy $(1,x)$ is in the image $f[\underline{F}]$; similarly for W_2.

Because of 2.1.17 we indeed have that $\underline{F} = W_1 \cup W_2$ and $W_1 \cap W_2 = \emptyset$. Observe that the set W_1 induces m_1 edges in F and the set W_2 induces m_2 edges. Recall that to get F_1 we deleted the edge $e = (e_1,\ldots,u,v,\ldots,e_{\delta_j})$ of F. The edge e contains vertices $e_1,\ldots,u \in W_1$ and $v,\ldots,e_{\delta_j} \in W_2$, so it is not induced by either W_1 or W_2. Moreover, u and v are vertices of a cycle in $Sh(F)$, so there has to be another edge in the cycle with one end in W_1 and the other end in W_2. This edge appears in yet another edge of F. Therefore $r(F) \geq m_1 + m_2 + 2$.

We conclude that $m_1 + m_2 + 1 = r(G) \geq r(F) \geq m_1 + m_2 + 2$, a contradiction.

Now we can derive a contradiction, thus disproving the assumption that F is *not* a Δ-tree.

2.1.20. As a consequence of duality, $F_k \to D$ for all k, because $F \twoheadrightarrow F_k$ by 2.1.19. If $e = (e_1,\ldots,e_t = u, v = e_{t+1},\ldots,e_{\delta_j})$ is the deleted edge, let $k > |\underline{D}|^t$ and let $f : F_k \to D$ be a homomorphism. Then for some p, q such that $p \neq q$ we have $f(p,e_l) = f(q,e_l)$ for all l satisfying that $1 \leq l \leq t$.

Define $g : \underline{F} \to \underline{D}$ by setting $g(x) := f(p,x)$. If $(x_1,x_2,\ldots,x_{\delta_i}) \in R_i(F)$ is an edge distinct from e, then $\big((p,x_1),(p,x_2),\ldots,(p,x_{\delta_i})\big) \in R_i(F_k)$ and therefore

$$\big(g(x_1),g(x_2),\ldots,g(x_k)\big) \in R_i(D).$$

Besides,

$$\big(g(e_1),g(e_2),\ldots,g(e_t)\big) \in R_j(D)$$

because f is a homomorphism,

$$\begin{aligned}\big(g(e_1), g(e_2), \dots, g(e_t)\big) &= \big(f(p, e_1), f(p, e_2), \dots, f(p, e_t)\big)\\ &= \big(f(q, e_1), \dots, f(q, e_t), f(p, e_{t+1}), \dots, f(p, e_{\delta_j})\big)\end{aligned}$$

and

$$\big((q, e_1), \dots, (q, e_t), (p, e_{t+1}), \dots, (p, e_{\delta_j})\big) \in R_j(F_k).$$

Hence g is a homomorphism from F to D, a contradiction with duality.

This proves that F is a Δ-tree. □

2.2 Three constructions

In this section, we first show two ways to construct the dual of a Δ-tree F. The first construction is by Komárek [23] for digraphs (we will call it the *mosquito construction*), the second by Nešetřil and Tardif [37, 36] for general Δ-structures (called here the *bear construction*).

This section's main purpose is to present a more general construction, of which both cited constructions are special cases.

Mosquito construction

For the mosquito construction, we consider digraphs, so let $\Delta = (2)$. Let F be a core tree. The duals of directed paths are described in Example 2.1.9, so here we suppose that F is a core digraph other than a directed path.

First we introduce some notation.

2.2.1. We define the function $\mu : \underline{F} \to \mathbb{N}^2$ by $\mu(x) := (d, u)$, where d is the length of the longest directed path ending in the vertex x and u is the length of the longest directed path starting from x.

Further, let $r(F)$ be the length of the longest directed path in F, that is $r(f) := \max\{r : \vec{P}_r \to F\}$.

Note that $r(F) = \max\big\{d + u : (d, u) = \mu(x),\ x \in \underline{F}\big\}$.

2.2.2. Let $p(F)$ be the **height** of F defined as the length of the shortest directed path that F is homomorphic to:

$$p(F) := \min\{r : F \to \vec{P}_r\}.$$

For any $p, q \geq 0$ define

$$\Phi(p, q) := \big\{a \in \underline{F} : \mu(a) = (d, u),\ d \leq p,\ u \leq q\big\}.$$

Now follows the definition of the mosquito dual of the oriented tree F.

2.2.3 Definition. Let F be an oriented tree. The vertices of the **mosquito dual** $D_m(F)$ are triples (p, q, ϕ), where $0 \leq p, q < p(F)$ and $p+q < p(F)$ and $\phi : \Phi(p, q) \to \underline{F}$ is a function such that $\phi(a)$ is a neighbour (in-neighbour or out-neighbour) of a for all $a \in \Phi(p, q)$.

The edges of the dual $D_m(F)$ are pairs $\big((p, q, \phi), (p', q', \phi')\big)$ such that

(i) $p < p'$ and $q > q'$, and

(ii) there is no edge (a, b) of F such that $\phi(a) = b$ and $\phi'(b) = a$.

The correctness of the construction is asserted by the following theorem.

2.2.4 Theorem ([23]). *For an oriented tree F that is not a directed path, the pair $\big(F, D_m(F)\big)$ is a duality pair.* □

Bear construction

Next we present a construction of the dual by Nešetřil and Tardif, which works for Δ-structures of an arbitrary type Δ.

2.2.5 Definition ([36]). Let F be a core Δ-tree. Remember the definition of $\mathrm{Block}(F)$ and $\mathrm{Inc}(F)$ from 1.6.5. Define the **bear dual** $D_b(F)$ as the Δ-structure on the base set

$$\underline{D_b(F)} := \Big\{ f : \underline{F} \to \mathrm{Block}(F) : \big\{x, f(x)\big\} \in E\big(\mathrm{Inc}(F)\big) \text{ for all } x \in \underline{F} \Big\}$$

with relations

$$R_i\big(D_b(F)\big) := \big\{(f_1, f_2, \ldots, f_{\delta_i}) : \text{for all } e = (x_1, x_2, \ldots, x_{\delta_i}) \in R_i(F) \\ \text{there exists } j \in \{1, \ldots, \delta_i\} \text{ such that } f_j(x_j) \neq (i, e)\big\}.$$

2.2.6 Theorem ([36]). *For any Δ-tree F, the pair $\big(F, D_b(F)\big)$ is a duality pair.*

A generalisation: The animal construction

In this subsection we provide a common framework for both constructions. We present a construction with a parameter called a *positional-function family*. Depending on what family we take, we get distinct dual structures. Later we will show what positional-function families have to be considered to give the mosquito construction and the bear construction.

We start with the definition of a positional-function family.

Recall the definition of an acyclic Δ-structure from 1.7.5.

2.2.7 Definition. Let $(Q, \preccurlyeq)$ be a partially ordered set. A **positional-function family** is a family

$$\{\mu_A : A \text{ is an acyclic structure}\}$$

of functions such that

(i) $\mu_A : \underline{A} \to Q$ for all acyclic structures A,

(ii) whenever A and B are acyclic structures and there exists a homomorphism $f : A \to B$ such that $f(x) = y$, then $\mu_A(x) \preccurlyeq \mu_B(y)$,

(iii) for any non-empty finite downset S in Q (see 1.9.2) there exists an acyclic **representing structure** $\Theta = \Theta(S)$ with

1. a mapping $\Omega : \underline{\Theta} \to S$ such that for any homomorphism f from an acyclic structure A to Θ and for any vertex $a \in \underline{A}$, we have $\mu_A(a) \preccurlyeq \Omega\big(f(a)\big)$, and
2. a mapping $\Psi : S \to \underline{\Theta}$ such that for every acyclic structure A with $\mu_A[A] \subseteq S$ the mapping given by $a \mapsto \Psi\big(\mu_A(a)\big)$ is a homomorphism from A to Θ.

In order to construct the dual of a tree F, we need a positional-function family satisfying a certain finiteness condition, as we will see shortly. This condition is trivially satisfied if the poset Q is finite, as is the case of the bear construction. The condition is as follows.

2.2.8 Definition. Let $(Q, \preccurlyeq)$ be a partially ordered set. Let us have a positional-function family $M = \{\mu_A : A \text{ acyclic}\}$ with $\mu_A : \underline{A} \to Q$ and with representing structures $\Theta(S)$ for all non-empty finite downsets S.

Let F be a Δ-tree. Let

$$T := \bigcup_{\substack{F \nrightarrow A \\ A \text{ acyclic}}} \mu_A[A]$$

and let $S(F) := T^{\downarrow}$. So $S(F)$ is the smallest downset that contains T as a subset. We say that the positional-function family M is **suitable** for the Δ-tree F if the following condition is satisfied:

$$\text{The downset } S(F) \text{ is finite.} \tag{2.1}$$

2.2.9. We define the mapping $\Phi : \underline{\Theta} \to 2^{\underline{F}}$ by

$$\Phi(\theta) := \big\{y \in \underline{F} : \mu_F(y) \preccurlyeq \Omega(\theta)\big\}.$$

It is obvious from the definition that if $y \in \underline{F}$ is mapped to θ by some homomorphism from F to Θ, then $y \in \Phi(\theta)$.

Now we proceed to define the animal dual of a Δ-tree F with respect to a positional-function family $M = \{\mu_A : A \text{ acyclic}\}$ that is suitable for F.

2.2.10 Definition. The base set $\underline{D_a}$ of the **animal dual** D_a consists of pairs (θ, ϕ) such that θ is a vertex of $\underline{\Theta}$ and $\phi : \Phi(\theta) \to \text{Block}(F)$ is a mapping such that y and $\phi(y)$ are adjacent in $\text{Inc}(F)$, that is y appears in the edge given by $\phi(y)$ in F.

The δ_i-tuple $\big((\theta_1, \phi_1), (\theta_2, \phi_2), \ldots, (\theta_{\delta_i}, \phi_{\delta_i})\big)$ is an element of the relation $R_i(D_a)$ if and only if $(\theta_1, \theta_2, \ldots, \theta_{\delta_i}) \in R_i(\Theta)$ and there is no edge $e = (y_1, y_2, \ldots, y_{\delta_i}) \in R_i(F)$ such that $\phi_1(y_1) = \phi_2(y_2) = \cdots = \phi_{\delta_i}(y_{\delta_i}) = (i, e)$. (Some $\phi_j(y_j)$ may be undefined but that does not matter; in such a case the edge e does not satisfy the equality.)

To prove that the animal construction is correct and the structure D_a forms indeed a duality pair with F, we will use Lemma 2.1.3. First we prove that F is not homomorphic to D_a; this is shown by the following two lemmas. The other part of the proof follows the statement of Theorem 2.2.13.

2.2.11 Lemma. *Let A be an arbitrary Δ-structure and let $f : A \to D_a$ be a homomorphism. Let the mapping $g : \underline{A} \to \underline{\Theta}$ be defined by $g(a) := \theta$ such that $f(a) = (\theta, \phi)$. Then g is a homomorphism from A to Θ.*

Proof. If $(a_1, a_2, \ldots, a_{\delta_i}) \in R_i(A)$, then $\big(f(a_1), f(a_2), \ldots, f(a_{\delta_i})\big) \in R_i(D_a)$ because f is a homomorphism. Hence $\big(g(a_1), g(a_2), \ldots, g(a_{\delta_i})\big) \in R_i(\Theta)$ by the definitions of g and the dual structure. □

2.2.12 Lemma. *The tree F is not homomorphic to D_a.*

Proof. Reductio ad absurdum: Suppose that there is a homomorphism $f : F \to D_a$. Let y be an arbitrary element of $\underline{F}$ and let $f(y) = (\theta, \phi)$. By Lemma 2.2.11, there is a homomorphism $g : F \to \Theta$ that maps y to θ, whence $y \in \Phi(\theta)$ by 2.2.9. So $\phi(y)$ is defined; let

$$\phi(y) =: \big(i, (y_1, y_2, \ldots, y_{\delta_i})\big)$$

and let

$$e := (y_1, y_2, \ldots, y_{\delta_i}) \in R_i(F).$$

Since f is a homomorphism, $\big(f(y_1), f(y_2), \ldots, f(y_{\delta_i})\big) \in R_i(D_a)$. By the definition of edges of D_a, there is an index j such that $\phi_j(y_j) \neq (i, e)$ if $f(y_j) = (\theta_j, \phi_j)$. Let $\phi_j(y_j) = (i', e')$. So we have a walk of length four in $\text{Inc}(F)$, namely y, e, y_j, e', with the property that $y \neq y_j$ and $e \neq e'$.

Repeating the procedure we can get an arbitrarily long walk

$$z_1, e_1, z_2, e_2, \ldots, z_n, e_n$$

in Inc(F) such that $z_j \neq z_{j+1}$ and $e_j \neq e_{j+1}$. This is a contradiction, because Inc(F) is a tree. □

Now we prove the correctness of the animal construction.

2.2.13 Theorem. *Let* $(Q, \preccurlyeq)$ *be a partially ordered set. Let* $\{\mu_A : A \text{ acyclic}\}$ *be a positional-function family with representing structure* $\Theta(S)$ *for every non-empty finite downset* S *in* Q*. If* F *is a* Δ*-tree such that the condition* (2.1) *holds, and* D_a *is the* Δ*-structure defined in 2.2.10, then the pair* (F, D_a) *is a duality pair.*

Proof. We have just seen that $F \nrightarrow D_a$, and so by Lemma 2.1.3 it remains to show that whenever $F \nrightarrow X$, then $X \to D_a$. The proof uses an idea of the proof for the bear dual [36]. This idea is reworked so as to fit the animal construction.

Fix a labelling ℓ of the vertices of Inc(F) by positive integers such that different vertices get different labels and the subgraph of Inc(F) induced by $\{u : \ell(u) \geq n\}$ is a connected subtree for all positive integers n. Such a labelling can be defined by repeatedly labelling and deleting the leaves of Inc(F).

For a vertex $y \in \underline{F}$ and for its neighbour $b = (i, e) \in \text{Block}(F)$ in Inc(F), let $T_{y,b}$ be the maximal subtree of Inc(F) that contains y and b but no other neighbour of y. Let $F_{y,b}$ be the Δ-tree such that $\text{Inc}(F_{y,b}) = T_{y,b}$ (so $F_{a,b}$ is a substructure of F). For a vertex y and for $b \neq b'$, the subtrees $F_{y,b}$ and $F_{y,b'}$ intersect in exactly one vertex: the vertex y.

Let X be a Δ-structure such that $F \nrightarrow X$; we will define a mapping $f : \underline{X} \to \underline{D_a}$ and prove that it is a homomorphism.

For every $x \in \underline{X}$ and $y \in \underline{F}$ we define

$$K(x, y) := \big\{b \in \text{Block}(F) : \{b, y\} \in E(\text{Inc}(F)),\\ \text{and there is no homomorphism } g : F_{y,b} \to X \text{ s.t. } g(y) = x\big\}.$$

For every $x \in \underline{X}$ and $y \in \underline{F}$, the set $K(x, y)$ is non-empty. Otherwise there would be a homomorphism $g_b : F_{y,b} \to X$ for all edges b incident with y, satisfying $g_b(y) = x$, and so their union would define a homomorphism from F to X.

Let $x \in \underline{X}$ be an arbitrary vertex. Define $f(x) := (\theta_x, \phi_x)$ by setting

$$\theta_x := \Psi\big(\mu_X(x)\big)$$

and let $\phi_x : \Phi(\theta_x) \to \text{Block}(F)$ be defined by letting $\phi_x(y)$ be the element b of $K(x, y)$ with the smallest label $\ell(b)$. The element θ_x is well-defined because $\mu_X[X] \subseteq S$ by the definition of S.

Suppose that $(x_1, x_2, \ldots, x_{\delta_i}) \in R_i(X)$. We want to show that

$$\big(f(x_1), f(x_2), \ldots, f(x_{\delta_i})\big) \in R_i(D_\alpha).$$

Let $f(x_j) = (\theta_j, \phi_j)$. Then

$$(\theta_1, \theta_2, \ldots, \theta_{\delta_i}) = \Big(\Psi\big(\mu_X(x_1)\big), \Psi\big(\mu_X(x_2)\big), \ldots, \Psi\big(\mu_X(x_{\delta_i})\big)\Big) \in R_i(\Theta)$$

by the definition of a representing structure.

Thus it remains to prove that there is no edge $e = (y_1, y_2, \ldots, y_{\delta_i}) \in R_i(F)$ such that $\phi_1(y_1) = \phi_2(y_2) = \cdots = \phi_{\delta_i}(y_{\delta_i}) = (i, e)$. For the sake of contradiction, suppose that such an edge exists. Let $N(y_j)$ denote the set of all neighbours of y_j in $\mathrm{Inc}(F)$ different from (i, e) and let $N := \bigcup_{1 \le j \le \delta_i} N(y_j)$. Among the elements of N, there may be at most one with its ℓ-label bigger than the label of (i, e) because of the way ℓ was defined.

If there is no such element, then for any j, no element of $N(y_j)$ belongs to $K(x_j, y_j)$, because otherwise we would not have selected (i, e) as the value of $\phi_j(y_j)$. Therefore for all j and all $b \in N(y_j)$ there is a homomorphism $g_{j,b} : F_{y_j,b} \to X$ such that $g_{j,b}(y_j) = x_j$ and the union of all these homomorphisms defines a homomorphism from F to X, a contradiction.

Thus there is a unique element $b' \in N$ such that $\ell(b') > \ell(i, e)$. Hence $b' \in N(y_{j'})$ and for all $b \in N$ different from b' we can find a homomorphism $g_{j,b}$ as above.

Then the mapping g such that

$$g(y) := \begin{cases} x_{j'} & \text{if } y = y_{j'}, \\ g_{j,b}(y) & \text{if } y \in \underline{F_{y_j,b}} \text{ and } y \neq y_{j'}, \end{cases}$$

is a homomorphism from $F_{y_j,(i,e)}$ to X, proving that $(i, e) \notin K(x_{j'}, y_{j'})$ and so contradicting the value of $\phi_{j'}$. □

Bear and mosquito are animals

In the beginning of this section we promised a generalisation of both the bear and the mosquito constructions. We have seen a *metaconstruction*: it produces different results depending on what positional-function family we plug in. Here we show what to plug in to get the two previous constructions. That proves that indeed the bear construction and the mosquito construction are special cases of the animal construction.

2.2.14 Example (bear)**.** Let $Q = \{\diamondsuit\}$ be a one-element poset. For an acyclic structure A define μ_A to be the constant mapping that maps all vertices of A to $\diamondsuit$. The conditions (i) and (ii) of Definition 2.2.7 are trivially satisfied. Let Θ be the Δ-structure defined by $\underline{\Theta} = Q$ and $R_i(\Theta) = \underline{\Theta}^{\delta_i}$ for all $i \in I$, and let Ω and Ψ be the identity mapping on $Q = \underline{\Theta}$. The structure Θ is the representing structure $\Theta(Q)$ for the only downset Q, since the constant mapping from any Δ-structure to Θ is a homomorphism. Also the condition (2.1) is satisfied trivially.

It is easy to see that for any Δ-tree F the dual structures $D_a(F)$ and $D_b(F)$ are isomorphic; the mapping defined by $(\theta, \phi) \mapsto \phi$ is an isomorphism.

2.2.15 Example (mosquito)**.** Now, let $(Q, \preccurlyeq)$ be the product $(\mathbb{N}, \leq) \times (\mathbb{N}, \leq)$, that is $Q = \mathbb{N} \times \mathbb{N}$ and $(d, u) \preccurlyeq (d', u')$ if and only if $d \leq d'$ and $u \leq u'$. For an acyclic structure A, the positional function μ_A is defined as in 2.2.1: $\mu_A(a) = (d, u)$, where d is the length of the longest directed path ending in the vertex a and u is the length of the longest directed path starting in a. As a homomorphism to an acyclic structure maps a directed path bijectively, the condition (ii) of Definition 2.2.7 is satisfied.

For a downset S, the representing digraph is $\Theta = (S, E)$, where

$$\big((p, q), (p', q')\big) \in E \text{ if and only if } p < p' \text{ and } q > q'.$$

Furthermore, we define $\Omega = \Psi = \mathrm{id}_S$. It can be checked easily that this correctly defines a representing structure.

Because (in the context of digraphs) every tree is homomorphic to a directed path, for any tree F the set $S(F)$ contains pairs (p, q) with $p + q$ bounded from above by the height of F. Therefore the condition (2.1) is satisfied for any oriented tree F.

Evidently, $D_a(F)$ and $D_m(F)$ are isomorphic.

2.2.16 Problem. Find other suitable positional-function families and representing structures to get essentially new constructions of the dual structure for a general type Δ.

2.3 Properties of the dual

Two particular properties of dual structures are worthwhile to mention: connectedness and irreducibility.

Dual is irreducible

Recall that a Δ-structure D is called irreducible if $A \times B \to D$ implies that $A \to D$ or $B \to D$ for every two structures A, B (see Lemma 1.6.11).

We show that the dual of a Δ-tree is irreducible. This statement is dual to the connectedness of the left-hand side of a duality pair (see 2.1.13).

2.3.1 Proposition. *If* (F, D) *is a duality pair, then the* Δ*-structure* D *is irreducible.*

Proof. Let A, B be structures such that $A \times B \to D$. By duality, $F \nrightarrow A \times B$. Therefore $F \nrightarrow A$ or $F \nrightarrow B$; and using duality once more, it follows that $A \to D$ or $B \to D$. □

Dual is connected

We have seen that in a duality pair (F, D) the core of F is connected (2.1.13). Now we prove that the dual D is connected too.

2.3.2 Proposition. *If* (F, D) *is a duality pair and* D *is a core, then the* Δ*-structure* D *is connected.*

Proof. Suppose that F and D are core Δ-structures, the pair (F, D) is a duality pair and D is not connected. Hence F has edges of all kinds; otherwise the dual is connected, see 2.1.11.

Let $\Delta = (\delta_i : i \in I)$. Let K be the Δ-structure that consists of isolated edges, one edge of each kind; formally

$$\underline{K} = \big\{(i, u) : i \in I \text{ and } 1 \le u \le \delta_i\big\},$$
$$R_i(K) = \Big\{\big((i,1),(i,2),\dots,(i,\delta_i)\big)\Big\}.$$

Now, the structure J is obtained from K by gluing all edges at the first vertex, and J' is obtained by gluing them at the last vertex. In other words, we have two equivalence relations on $\underline{K}$, namely $\approx$ and $\approx'$, defined by

$$(i, u) \approx (j, v) \text{ if } u = v = 1 \text{ or } (i, u) = (j, v),$$
$$(i, u) \approx' (j, v) \text{ if } u = \delta_i \text{ and } v = \delta_j \text{ or } (i, u) = (j, v).$$

Define J and J' as factor structures (see 1.2.7): $J := K/\approx$ and $J' := K/\approx'$. The construction is illustrated in Figure 2.4.

Suppose F is homomorphic to both J and J'. The structure F is connected, so it has at least two incident edges e, e' of distinct kinds. As $F \to J$, the vertex these two edges share is the starting vertex of each of them, because all

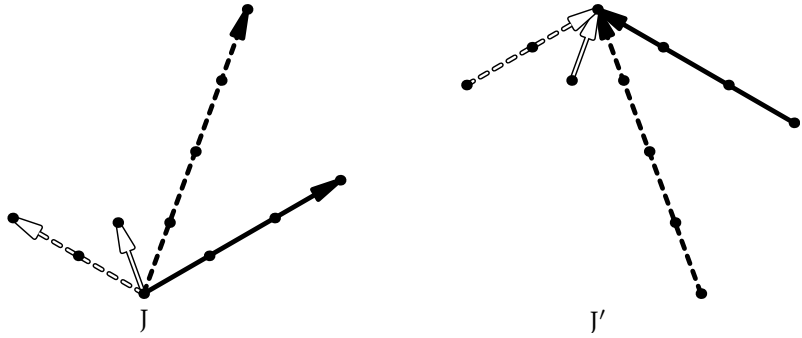

Figure 2.4: The construction of J and J′

other vertices are distinct as they are homomorphically mapped to distinct vertices of J. But since $F \to J'$, the starting vertices of e and e' are distinct; a contradiction.

Therefore either J or J′ is homomorphic to D if at least two relations have arity greater than two.

In any case, D has a component that contains edges of all kinds. This is satisfied trivially if there is only one kind of edges, that is if $|I| = 1$.

So we can connect all components of D with long zigzags, see Figure 2.5. We get a connected Δ-structure D′; if the zigzags are long enough, any substructure E of D′ induced by at most $|\underline{F}|$ vertices contains vertices of only one of the components of D; hence E is homomorphic to D. Because F is *not* homomorphic to D, it is not homomorphic to D′ either.

But then, by duality, the connected structure D′ is homomorphic to D, so it is homomorphic to a component of D. Therefore D is homomorphic to a proper substructure of D, a contradiction with D being a core. □

Connectedness of duals was originally proved for digraphs by Nešetřil and Švejdarová [33] in a different way, by examining the bear construction.

Dually, we would expect all trees to be irreducible. But this shows another difference between the dual notions of connectedness and irreducibility; not all trees, not even all paths are irreducible, as the following example demonstrates.

2.3.3 Example. Let $\Delta = (2)$. Figure 2.6 shows two oriented paths P_1 and P_2 and their product $P_1 \times P_2$. We can see that the core P of $P_1 \times P_2$ is an oriented path. Thus P is homomorphically equivalent to the product of two

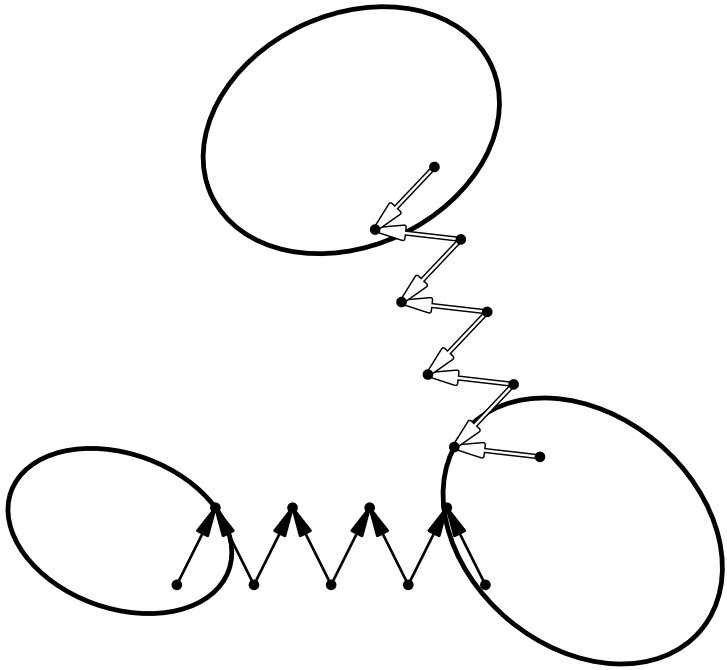

Figure 2.5: Connecting the components of a dual with zigzags

incomparable structures P_1 and P_2. It follows from Lemma 1.6.11 that P is not irreducible.

2.4 Finite dualities

Introduction

In this section, we generalise duality pairs in a natural way: instead of forbidding homomorphisms from a single structure, we forbid homomorphisms from a finite set of structures; and on the other side, we allow structures to map to some of a finite number of structures.

2.4.1 Definition. Let $\mathcal{F}$ and $\mathcal{D}$ be two finite sets of core Δ-structures such that no homomorphisms exist among the structures in $\mathcal{F}$ and among the structures in $\mathcal{D}$. We say that $(\mathcal{F}, \mathcal{D})$ is a **finite homomorphism duality**

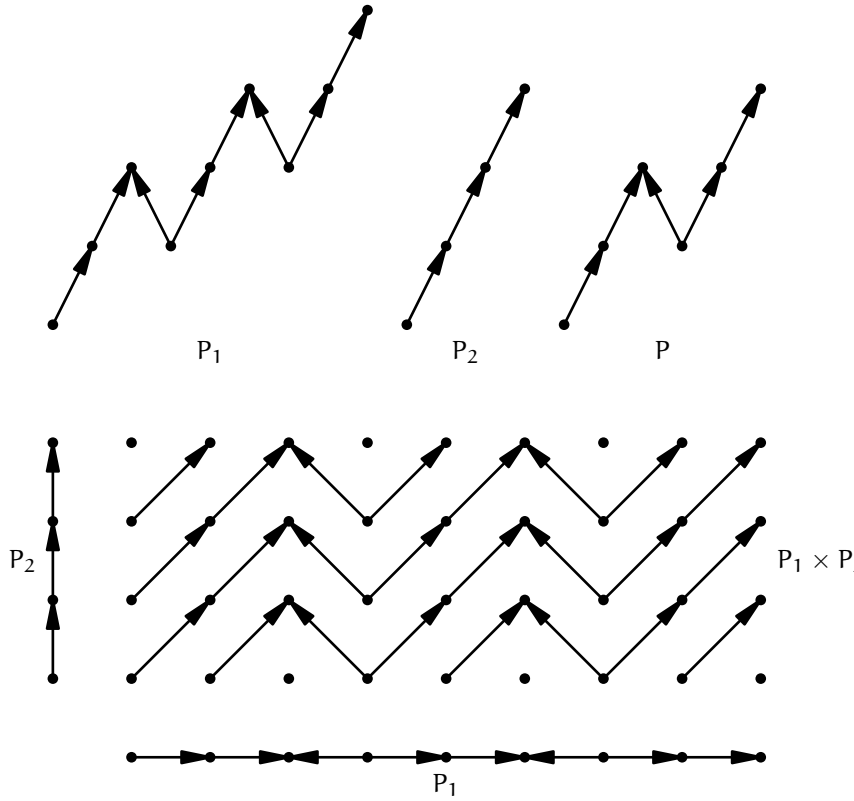

Figure 2.6: A non-irreducible path

(often just a **finite duality**) if for every Δ-structure A there exists $F \in \mathcal{F}$ such that $F \to A$ if and only if for all $D \in \mathcal{D}$ we have $A \not\to D$.

2.4.2. For a symbol like $\to\mathcal{D}$, two definitions are possible. It may denote either the set of all structures that admit a homomorphism to *some* structure in $\mathcal{D}$, or the set of structures that map to *each* structure in $\mathcal{D}$.

It is convenient for us to use the positive symbols $\to\mathcal{D}$ and $\mathcal{F}\to$ in the former sense, while the negative symbols $\mathcal{F}\not\to$ and $\not\to\mathcal{D}$ will be used in the latter sense. So we define

$$\mathcal{F}\to := \{A : F \to A \text{ for some } F \in \mathcal{F}\},$$
$$\mathcal{F}\not\to := \{A : F \not\to A \text{ for all } F \in \mathcal{F}\},$$

$$\to\mathcal{D} := \{A : A \to D \text{ for some } D \in \mathcal{D}\},$$
$$\not\to\mathcal{D} := \{A : A \not\to D \text{ for all } D \in \mathcal{D}\}.$$

In this notation, $(\mathcal{F}, \mathcal{D})$ is a finite homomorphism duality if and only if

$$\mathcal{F}\not\to = \to\mathcal{D}.$$

Since the classes $\mathcal{F}\to$ and $\mathcal{F}\not\to$ are complementary and so are the classes $\to\mathcal{D}$ and $\not\to\mathcal{D}$, the pair $(\mathcal{F}, \mathcal{D})$ is a duality pair if and only if

$$\mathcal{F}\to = \not\to\mathcal{D}.$$

We begin exploring the world of finite dualities by considering situations where the set $\mathcal{D}$ has only one element. Here we remark that forbidding homomorphisms from a finite set of structures is equivalent to forbidding a finite number of substructures. This is again characteristic of finite dualities, since all the classes $\to D$ are characterised by forbidden subgraphs (although in many cases by infinitely many of them).

2.4.3 Proposition. *Let* D *be a core* Δ*-structure. Then the following are equivalent:*

(1) There exists a finite set $\mathcal{F}$ *of* Δ*-structures such that the pair* $(\mathcal{F}, \{D\})$ *is a finite homomorphism duality.*

(2) There exists a finite set $\mathcal{F}'$ *of* Δ*-structures such that any* Δ*-structure* A *is homomorphic to* D *if and only if it contains no element of* $\mathcal{F}'$ *as its substructure.*

Proof. Suppose (1) holds and set $\mathcal{F}'$ to be the set of all homomorphic images of structures in $\mathcal{F}$. Then (2) follows from the definition of duality.

Conversely, if (2) holds, let $\mathcal{F}$ be the set of all cores of the structures in $\mathcal{F}'$. If $A \not\to D$, then A contains some element $F \in \mathcal{F}'$ as a substructure, so the core of F is homomorphic to A. And if $A \to D$ and $F \in \mathcal{F}$, then $F \not\to A$, for otherwise $F \to D$, a contradiction with (2). Therefore $(\mathcal{F}, \{D\})$ is a finite duality. □

Next we characterise all finite dualities $(\mathcal{F}, \mathcal{D})$ with a singleton right-hand side, that is dualities such that $|\mathcal{D}| = 1$.

2.4.4 Theorem ([34])**.** *If* $(\mathcal{F}, \{D\})$ *is a finite homomorphism duality, then all elements of* $\mathcal{F}$ *are* Δ*-trees and*

$$D \sim \prod_{F \in \mathcal{F}} D(F). \tag{2.2}$$

Conversely, for any finite collection $\mathcal{F}$ of Δ-trees, if (2.2) *holds, then the pair* $(\mathcal{F}, \{D\})$ *is a finite homomorphism duality.*

Proof. For the proof we need some terminology and results of Section 3.2.

Let $(\mathcal{F}, \{D\})$ be a finite duality. Suppose some $F \in \mathcal{F}$ is disconnected, $F = F_1 + F_2$. Then no element of $\mathcal{F}$ is homomorphic to any of F_1 and F_2, so $F_1 \to D$ and $F_2 \to D$, and also $F = F_1 + F_2 \to D$, a contradiction with the fact that $(\mathcal{F}, \{D\})$ is a duality. Therefore all elements of $\mathcal{F}$ are connected.

Let $F \in \mathcal{F}$. Suppose that $D \to A \to F + D$, but that $A \not\to D$. By duality, there exists $F' \in \mathcal{F}$ such that $F' \to A$. Since F' is connected, $F' \not\to D$, and $A \to F + D$, we have that $F' \to F$. Hence $F' = F$ because distinct elements of $\mathcal{F}$ are incomparable. Therefore $A \not\to D$ and $D \to A \to F + D$ implies that $F + D \to A$. This proves that $(D, F + D)$ is a gap, as $D < F + D$ because $F \not\to D$ by duality.

By Proposition 3.2.2, there exists a duality pair (T', D') such that $T' \to F + D \to T' + D'$ and $D \sim (F + D) \times D'$ (see the following diagram, in which gaps are marked by double arrows).

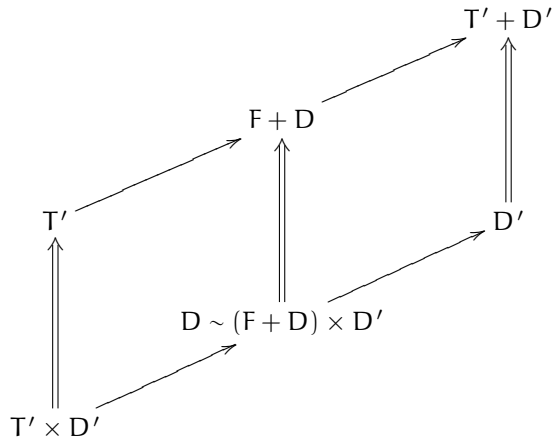

If $F \to D'$, then by duality $T' \not\to F$, and because T' is connected (it is in fact a Δ-tree), $T' \to D \to D'$, a contradiction with (T', D') being a duality. So $F \not\to D'$.

Since F is connected and $F \not\to D'$, we get that $F \to T'$; moreover, from duality we know that $T' \to F$. We conclude that $T' \sim F$ and $D' \sim D(F)$ is the dual of F. In this way, we have proved that $D \to D(F)$ for all $F \in \mathcal{F}$, and hence also $D \to \prod_{F \in \mathcal{F}} D(F)$.

On the other hand, $F \not\to \prod_{F \in \mathcal{F}} D(F)$ for any $F \in \mathcal{F}$, and so the product

$\prod_{F\in\mathcal{F}} D(F)$ is homomorphic to D because $(\mathcal{F},\{D\})$ is a finite duality. Therefore $D \sim \prod_{F\in\mathcal{F}} D(F) \to D$.

Conversely, if (2.2) holds then $F \nrightarrow \prod_{F\in\mathcal{F}} D(F)$ for any $F \in \mathcal{F}$, and in addition if $F \nrightarrow A$ for any $F \in \mathcal{F}$ then $A \to \prod_{F\in\mathcal{F}} D(F)$. Hence $\left(\mathcal{F}, \prod_{F\in\mathcal{F}} D(F)\right)$ is a finite duality. □

2.4.5 Corollary. *If $\mathcal{F}$ is a finite set of Δ-trees, then there exists a unique core* D *such that $(\mathcal{F},\{D\})$ is a finite homomorphism duality.* □

2.4.6. This uniquely determined dual core D is denoted by $D(\mathcal{F})$.

Transversal construction

Having characterised all finite dualities with a singleton right-hand side, we carry on by providing a construction of finite dualities. Later we will see that all finite dualities result from this construction.

The construction, which we will call the *transversal construction*, starts with a finite set of Δ-forests. The forests are decomposed into components and we consider sets consisting of the components. Some of these sets satisfy certain properties and are called *transversals*. Each transversal is a set of Δ-trees. The dual side of the finite duality is then constructed by taking the dual structures for each transversal (structures from Theorem 2.4.4).

2.4.7. Let $\mathcal{F} = \{F_1, F_2, \ldots, F_m\}$ be an arbitrary fixed non-empty finite set of core Δ-forests that are pairwise incomparable ($F_j \nrightarrow F_k$ for $j \neq k$). Let $\mathcal{F}_c = \{C_1, \ldots, C_n\}$ be the set of all distinct connected components of the structures in $\mathcal{F}$; each of these components is a core Δ-tree.

First we define *quasitransversals* to be certain sets of components appearing in the structures in $\mathcal{F}$.

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

(T1) any two distinct elements of $\mathcal{M}$ are incomparable, and

(T2) $\mathcal{M}$ **supports** $\mathcal{F}$, that is for every $F \in \mathcal{F}$ there exists $C \in \mathcal{M}$ such that $C \to F$.

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

2.4.10 Lemma. *Let $\mathcal{M}$, $\mathcal{M}'$ be two quasitransversals. Then the dual structures $D(\mathcal{M})$ and $D(\mathcal{M}')$ exist, and $D(\mathcal{M}) \to D(\mathcal{M}')$ if and only if $\mathcal{M} \preccurlyeq \mathcal{M}'$.*

Proof. By Theorem 2.4.4, the dual structures $D(\mathcal{M})$ and $D(\mathcal{M}')$ exist and

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

Let $\mathcal{M} \preccurlyeq \mathcal{M}'$; we want to show that $D(\mathcal{M}) \to D(\mathcal{M}')$. By the properties of product, it suffices to show that $D(\mathcal{M}) \to D(C')$ for any $C' \in \mathcal{M}'$. So, let $C' \in \mathcal{M}'$. Because $\mathcal{M} \preccurlyeq \mathcal{M}'$, there exists $C \in \mathcal{M}$ such that $C \to C'$. By the definition of a duality pair, $C \to C'$ implies that $C' \nrightarrow D(C)$ and this implies that $D(C) \to D(C')$. We conclude that $D(\mathcal{M}) \to D(C) \to D(C')$.

For the converse implication, let $D(\mathcal{M}) \to D(\mathcal{M}')$. We want to show that for any C' in $\mathcal{M}'$ there is C in $\mathcal{M}$ with $C \to C'$. Indeed, for $C' \in \mathcal{M}'$ we have $D(\mathcal{M}) \to D(\mathcal{M}') \to D(C')$; using duality, $C' \nrightarrow D(\mathcal{M})$, and therefore $C' \nrightarrow D(C)$ for some $C \in \mathcal{M}$. By duality $C \to C'$. □

2.4.11 Lemma. *The relation $\preccurlyeq$ is a partial order on the set of all quasitransversals.*

Proof. Obviously, $\preccurlyeq$ is both reflexive and transitive (a preorder).

Suppose now that $\mathcal{M} \preccurlyeq \mathcal{M}'$ and $\mathcal{M}' \preccurlyeq \mathcal{M}$, and let $C \in \mathcal{M}$. Then there exists $C' \in \mathcal{M}'$ such that $C' \to C$ and there exists $C'' \in \mathcal{M}$ such that $C'' \to C'$. Consequently $C'' \to C$, hence by (T1) we have $C = C' = C''$, so $\mathcal{M} \subseteq \mathcal{M}'$. Similarly we get that $\mathcal{M}' \subseteq \mathcal{M}$. Therefore $\mathcal{M} = \mathcal{M}'$ whenever $\mathcal{M} \preccurlyeq \mathcal{M}'$ and $\mathcal{M}' \preccurlyeq \mathcal{M}$. So $\preccurlyeq$ is antisymmetric; it is a partial order. □

2.4.12 Definition. A quasitransversal $\mathcal{M}$ is a **transversal** if

(T3) $\mathcal{M}$ is maximal with respect to the order $\preccurlyeq$.

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

2.4.14 Lemma (Transversal construction works)**.** *The pair $(\mathcal{F}, \mathcal{D})$ is a finite homomorphism duality.*

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

2.4.15 Example. First, suppose that $\mathcal{F} = \{T_1, T_2, \ldots, T_n\}$ is a set of pairwise incomparable trees and let $D_1, D_2, \ldots, D_n$ be their respective duals. By (T2), every quasitransversal contains all these trees. Therefore there exists only one quasitransversal $\mathcal{M} = \{T_1, T_2, \ldots, T_n\}$ and it is a transversal. So $\mathcal{D} = \{D(\mathcal{M})\} = \{D_1 \times D_2 \times \cdots \times D_n\}$. This corresponds to the situation in Theorem 2.4.4.

2.4.16 Example. 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 $\mathcal{F} = \{T_1 + T_2, T_1 + T_3, T_4\}$. The partial order $\preccurlyeq$ of quasitransversals is depicted in the following diagram:

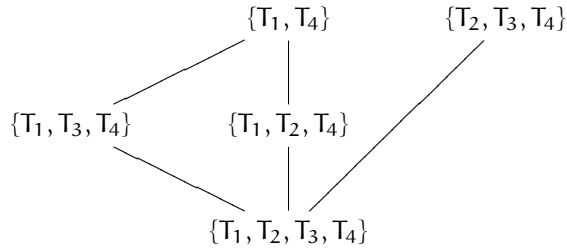

We have two transversals $\{T_1, T_4\}$ and $\{T_2, T_3, T_4\}$; and $\mathcal{D} = \{D_1 \times D_4, D_2 \times D_3 \times D_4\}$.

2.4.17 Example. Finally, let $T_1 \to T_3$ and $\mathcal{F} = \{T_1 + T_2, T_3 + T_4\}$. This time, we get the following order of quasitransversals:

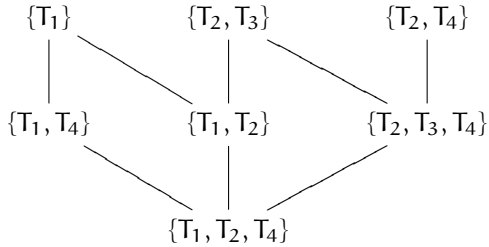

The transversals are $\{T_1\}$, $\{T_2, T_3\}$ and $\{T_2, T_4\}$. Hence $\mathcal{D} = \{D_1, D_2 \times D_3, D_2 \times D_4\}$.

Proof of Lemma 2.4.14. By the definition of $\mathcal{F}$, any two distinct elements of $\mathcal{F}$ are incomparable. Any two distinct elements of $\mathcal{D}$ are incomparable too, because any two transversals are incomparable with respect to $\preccurlyeq$ (they are all maximal in this order) and because of Lemma 2.4.10.

Let X be a Δ-structure such that $X \to D$ for some $D \in \mathcal{D}$. We want to prove that $F_i \not\to X$ for $i = 1, \ldots, m$. To obtain contradiction, assume that $F_i \to X$ for some i. Let $\mathcal{M}$ be the transversal for which $D(\mathcal{M}) = D$. By (T2), there exists $C \in \mathcal{M}$ such that $C \to F_i \to X$, therefore $X \not\to D(C)$. This is a contradiction with the assumption that $X \to D \to D(C)$ (here $D \to D(C)$ because D is the product of the duals of the structures in $\mathcal{M}$, the component C is an element of $\mathcal{M}$, and the projection is a homomorphism).

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

Characterisation

We will now prove that actually all finite homomorphism dualities are obtained from the transversal construction.

2.4.18. Let $(\mathcal{F}, \mathcal{D})$ be a finite homomorphism duality. Suppose that $\mathcal{F} = \{F_1, F_2, \ldots, F_m\}$ and $\mathcal{D} = \{D_1, D_2, \ldots, D_p\}$. By definition, 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 $\mathcal{F}_c = \{C_1, C_2, \ldots, C_n\}$ 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 their definition nor for proving Lemma 2.4.11 we needed the fact that the elements of $\mathcal{F}_c$ are trees.

2.4.19. For a quasitransversal $\mathcal{M}$, let $\overline{\mathcal{M}} = \{C' \in \mathcal{F}_c : C \in \mathcal{M} \Rightarrow C \nrightarrow C'\}$ be the set of all components "not supported" by $\mathcal{M}$.

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

(1) $C \nrightarrow D$ *for every* $C \in \mathcal{M}$,

(2) $C' \to D$ *for every* $C' \in \overline{\mathcal{M}}$.

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

Obviously, such D satisfies (2).

Further, we will prove that D satisfies (1) as well. For the sake of contradiction, suppose that there is $C \in \mathcal{M}$ such that $C \to D$.

Consider $\mathcal{M}' = \mathcal{M} \setminus \{C\}$. The set $\mathcal{M}'$ is not a quasitransversal, because otherwise we would have $\mathcal{M} \prec \mathcal{M}'$ and $\mathcal{M}$ would not satisfy (T3). Hence $\mathcal{M}'$

fails to satisfy (T2), and we can find $F \in \mathcal{F}$ which is not supported by $\mathcal{M}'$. It follows that $C \to F$.

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

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

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

It remains to prove uniqueness. If $D, D' \in \mathcal{D}$ both satisfy (1) and (2) and $D \neq D'$, that is $D \parallel D'$, then $D + D'$ violates the definition of finite homomorphism duality: $D + D'$ is homomorphic to no $\check{D}$ in $\mathcal{D}$, otherwise the elements of $\mathcal{D}$ would not be incomparable, contradicting the definition of finite duality; at the same time no F in $\mathcal{F}$ is homomorphic to $D + D'$, because (by the definition of a transversal) for every $F \in \mathcal{F}$ there is $C \in M$ such that $C \to F$, but $C \not\to D + D'$, because C is connected and by (1) it is homomorphic to neither D nor D'. □

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

2.4.22 Lemma. $\mathcal{D} = \{\mathrm{d}(\mathcal{M}) : \mathcal{M} \textit{ is a transversal}\}$.

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

Let $\mathcal{M}$ be a transversal such that $\mathcal{M}' \preccurlyeq \mathcal{M}$. To prove that $D = \mathrm{d}(\mathcal{M})$, it suffices to check conditions (1) and (2) of Lemma 2.4.20.

If $C \in \mathcal{M}$, then there exists $C' \in \mathcal{M}'$ such that $C' \to C$. Therefore $C \not\to D$, so condition (1) is satisfied.

Now condition (2): Suppose on the contrary that there exists $\check{C} \in \overline{\mathcal{M}}$ such that $\check{C} \not\to D$. Consider the Δ-structure $X = \check{C} + D$. If $F \to X$ for some $F \in \mathcal{F}$, then by the property (T2) of $\mathcal{M}$ there exists $C \in \mathcal{M}$ that is homomorphic to F. But since $\check{C} \in \overline{\mathcal{M}}$, we have that $C \not\to \check{C}$, hence $C \to D$. This is a contradiction with the condition (1). It follows that $X \to \check{D}$ for

some $\check{D} \in \mathcal{D}$, hence $D \to \check{D}$, so $D = \check{D}$. That is a contradiction with $\check{C} \not\to D$ and $\check{C} \to \check{D}$. □

2.4.23 Lemma. *For two distinct transversals $\mathcal{M}_1$, $\mathcal{M}_2$, we have*

(a) $\overline{\mathcal{M}_1} \cap \mathcal{M}_2 \neq \emptyset$,

(b) $\mathrm{d}(\mathcal{M}_1) \not\to \mathrm{d}(\mathcal{M}_2)$.

Proof.

(a) By (T3), $\mathcal{M}_1 \not\subseteq \mathcal{M}_2$, and therefore there exists $C_2 \in \mathcal{M}_2$ such that $C_1 \not\to C_2$ for any $C_1 \in \mathcal{M}_1$. Obviously $C_2 \in \overline{\mathcal{M}_1} \setminus \overline{\mathcal{M}_2} \subseteq \overline{\mathcal{M}_1}$. Since we selected $C_2 \in \mathcal{M}_2$, we have that $C_2 \in \overline{\mathcal{M}_1} \cap \mathcal{M}_2$.

(b) Let $C_2 \in \overline{\mathcal{M}_1} \cap \mathcal{M}_2$, as above. Then $C_2 \to \mathrm{d}(\mathcal{M}_1)$ and $C_2 \not\to \mathrm{d}(\mathcal{M}_2)$. Consequently $\mathrm{d}(\mathcal{M}_1) \not\to \mathrm{d}(\mathcal{M}_2)$. □

2.4.24 Lemma. *If $\mathcal{M}$ is a transversal, then the pair $\big(\mathcal{M}, \{\mathrm{d}(\mathcal{M})\}\big)$ is a finite homomorphism duality, and consequently $\mathrm{d}(\mathcal{M}) = \mathrm{D}(\mathcal{M})$.*

Proof. We want to prove that

$$\mathcal{M}{\not\to} = {\to}\mathrm{d}(\mathcal{M}).$$

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

(1) $A \in \mathcal{M}{\not\to} = \bigcap_{C \in \mathcal{M}} (C \not\to)$
(2) $C \not\to A$ for any $C \in \mathcal{M}$
(3) $C \not\to A + \coprod_{\check{C} \in \overline{\mathcal{M}}} \check{C}$ for any $C \in \mathcal{M}$
(4) $A + \coprod_{\check{C} \in \overline{\mathcal{M}}} \check{C} \to \mathrm{d}(\mathcal{M})$
(5) $A \to \mathrm{d}(\mathcal{M})$
(6) $A \in {\to}\mathrm{d}(\mathcal{M})$

Because: (1) $\Leftrightarrow$ (2) and (5) $\Leftrightarrow$ (6) by definition. (4) $\Rightarrow$ (5) immediately. (5) $\Rightarrow$ (2) by Lemma 2.4.20(1). (2) $\Rightarrow$ (3) follows from the definition of $\overline{\mathcal{M}}$ and the fact that C is connected.

It remains to prove that (3) $\Rightarrow$ (4): Let $X = A + \coprod_{\check{C} \in \overline{\mathcal{M}}} \check{C}$. If $F \to X$ for some $F \in \mathcal{F}$, then by (T2) there exists $C \in \mathcal{M}$ such that $C \to F \to X$, a contradiction. Thus no element of $\mathcal{F}$ is homomorphic to X, hence $X \to D$ for some $D \in \mathcal{D}$. By Lemma 2.4.22, $D = \mathrm{d}(\mathcal{M}')$ for a transversal $\mathcal{M}'$; by Lemma 2.4.20 and Lemma 2.4.23(a), $\mathcal{M}' = \mathcal{M}$.

The equivalence (1) $\Leftrightarrow$ (6) is precisely the definition of finite duality.

By Theorem 2.4.4, the dual is uniquely determined if it is a core, so $\mathrm{d}(\mathcal{M}) = \mathrm{D}(\mathcal{M})$. □

Lemma 2.4.24 and Theorem 2.1.12 imply that any element of a transversal is a Δ-tree, but we have not proved that every structure in $\mathcal{F}_c$ is an element of some transversal. However, we have the following lemma, for whose proof we will once again use the characterisation of gaps in Section 3.2.

2.4.25 Lemma. *Each component $C \in \mathcal{F}_c$ is a Δ-tree.*

Proof. Suppose that $C \in \mathcal{F}_c$ is not a tree. By Lemma 2.4.24, Theorem 2.1.12, and Theorem 2.4.4, C is an element of no transversal. Set

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

Clearly, $A < C$ because all the summands are less than C and C is connected. As C is not a tree, it has no dual; because it is connected, the pair (A, C) is not a gap by Lemma 3.2.1. Let X be a structure satisfying that $A < X < C$.

Then for any $C' \in \mathcal{F}_c$ such that $C \neq C'$, we have $C' \to X$ if and only if $C' \to C$ and $X \to C'$ if and only if $C \to C'$. Indeed: if $C' \to C$, then $C' \to A \to X$; if $C \to C'$, then $X \to C \to C'$. On the other hand, if $C \parallel C'$, then $X \to C'$ implies $X \to C \times C' \to A$ (because $C \times C'$ is one of the summands in the above definition of A), a contradiction with $A < X$. Moreover $C' \to X$ implies $C' \to C$.

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.

Suppose $F \to G$. Then $C \to G$. Because C is connected, it is homomorphic to a component of G. Since $C \not\to X$, it is homomorphic to some other component of F, contradicting that F is a core. Therefore $F \not\to G$.

In addition, $F' \not\to G$ for any $F \neq F' \in \mathcal{F}$, because $F' \to G$ implies $F' \to F$. Therefore $G \to D$ for some $D \in \mathcal{D}$. Let M be the transversal such that $D = D(M)$. Recall that C is an element of no transversal, so $C \notin M$. The structure D is a product of duals and hence $C' \not\to G$ for any $C' \in M$; therefore $C' \not\to X$ and $C' \not\to C$ for any $C' \in M$. Consequently $C \to D$. 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 $F \to D$, a contradiction. □

We finish this section by a theorem that characterises all finite dualities.

2.4.26 Theorem (Characterisation of finite dualities)**.** *If $(\mathcal{F}, \mathcal{D})$ is a finite homomorphism duality, then all elements of $\mathcal{F}$ are Δ-forests and $\mathcal{D} = \mathcal{D}(\mathcal{F})$ results from the transversal construction. In particular, $\mathcal{D}$ is determined by $\mathcal{F}$ uniquely up to homomorphic equivalence.*

Conversely, for any finite collection $\mathcal{F}$ of core Δ-forests, $\big(\mathcal{F}, \mathcal{D}(\mathcal{F})\big)$ is a finite homomorphism duality.

Proof. All elements of $\mathcal{F}$ are forests because of of Lemma 2.4.25. The set $\mathcal{D}$ is uniquely determined as a consequence of Lemma 2.4.22 and because of Lemma 2.4.24 and Theorem 2.4.4 it is determined by the transversal construction.

The second part is Lemma 2.4.14. □

Now we can view the notation $(\mathcal{F}, \mathcal{D})$ from a different perspective: the letter $\mathcal{F}$ stands for *forbidden*, as we mentioned in 2.1.4, but it may also be understood to stand for *forests*.

2.4.27 Back from duals to forests. Our construction of duals from forests relied heavily on the fact that every finite Δ-structure is a finite sum of components, structures that are connected. Although we mentioned in 1.6.14 that some structures are not a finite product of irreducible structures. However, it can be shown that the set $\mathcal{F}$ in a duality pair $(\mathcal{F}, \mathcal{D})$ is determined uniquely by $\mathcal{D}$ too.

The characterisation of finite homomorphism dualities implies that dual structures can be factored into a product of irreducible structures. A construction dual to the transversal construction produces the forests from the dual set. This is covered in more detail when we discuss a complexity issue in Section 4.2.

2.5 Extremal aspects of duality

Extremal theories are concerned with questions how large can an object be if it satisfies certain conditions, or has certain properties.

2.5.1. In the context of homomorphism dualities, we are interested in the following four questions.

(1) Given a Δ-tree, how large can its dual be?

(2) Given a right-hand side of a duality pair, how large can the corresponding Δ-tree be?

(3) Given a finite set $\mathcal{F}$ of Δ-forests, how large can the set $\mathcal{D}(\mathcal{F})$ be?

(4) Given a right-hand side of a finite duality, how large can the corresponding left-hand side be?

Naturally, one has to define a suitable notion of size for this purpose.

Results have been published about questions (1) and (2), partially also about (4). Finite dualities have been studied in full generality only recently, so questions (3) and (4) have not yet been thoroughly investigated.

2.5.2. When examining extremal problems about homomorphism dualities, we consider only cores. That is, we ask how large the *core* of the dual of a *core* Δ-tree is, and analogously for the other questions.

The reason for this is clear: in every class of homomorphic equivalence there are arbitrarily large structures. The smallest structure in such a class is the (unique) core in it. By taking sums of an arbitrary number of disjoint copies of the core, we can produce arbitrarily large homomorphically equivalent structures.

Concerning question (1), an upper bound on the size of a Δ-tree's dual follows from the bear construction. The bound on the size of the base set of the dual is exponential in terms of the size of the base set of the Δ-tree.

2.5.3 Theorem ([36])**.** *Let* $\Delta = (\delta_i : i \in I)$*. Let* (F, D) *be a duality pair such that* D *is a core and let* $n := |\underline{F}|$*. Then*

$$|\underline{D}| \leq n^n.$$

Proof. Let F be a Δ-tree. Consider the bear construction from 2.2.5. A vertex of the dual $D_b(F)$ is a function that assigns each vertex x of F an edge e of F such that x appears in e. Since F is a tree, no two distinct edges contain more than one vertex in common. Thus the number of edges containing a fixed vertex x is at most n. Hence the number of vertices of $D_b(F)$ is at most n^n, as we were supposed to prove. □

Nešetřil and Tardif [36] also provide a construction of paths whose duals indeed have exponential size.

2.5.4 Theorem ([36, Theorem 8])**.** *For any sufficiently large positive integer* N *there exists a core* Δ*-tree* F *such that* $|\underline{F}| \geq N$ *and if* D *is the core for which* (F, D) *is a duality pair, then*

$$|\underline{D}| \geq 2^{n/7 \log_2 n},$$

where $n := |\underline{F}|$*.*

For question (2), it is easier to measure the size of the forbidden Δ-tree in terms of its diameter.

We assume that the reader knows that the *distance* of two vertices in an undirected graph is the number of edges of the shortest path connecting them, and that a graph's *diameter* is the maximum distance of a pair of its vertices. The notion of diameter of an undirected graph is used to define the diameter for Δ-structures.

2.5.5 Definition. The **diameter** of a Δ-structure A is half the diameter of its incidence graph $\mathrm{Inc}(A)$.

Larose, Loten and Tardif [24] proved an upper bound on the diameter of forbidden trees in a finite duality with a singleton right hand side.

2.5.6 Theorem. *If* $\big(\mathcal{F}, \{D\}\big)$ *is a finite duality,* $F \in \mathcal{F}$ *is a core, and* $n = |\underline{D}|$*, then the diameter of* F *is at most* n^{n^2}*.* □

This theorem implies a bound on the number of edges of such forbidden trees, see Lemma 4.2.2. However, this bound is very rough, even though no examples are known that have an exponential number of edges in terms of the number of vertices of the dual.

3 Homomorphism order

> There is nothing more difficult to take in hand, more perilous to conduct or more uncertain in its success than to take the lead in the introduction of a new order of things.
>
> *(Niccolo Machiavelli)*

The relation of existence of a homomorphism on the class of all Δ-structures induces a partial order, called the *homomorphism order*.

Properties of this partial order have been widely studied in algebraic, category theory, random and combinatorial context. The homomorphism order motivates several questions linked to problems of existence of homomorphisms, see Section 4.1.

Density- and universality-related issues have attracted special attention. Universality of the homomorphism order of undirected graph was proved already in 1969 by Hedrlín [16]; in particular, it was shown that any countable partial order is an induced suborder of the homomorphism order. Universality was also studied for special classes of digraphs, and it has recently been proved that even the relatively small class of all directed paths induces a universal countable partial order [19, 20].

The examination of density has a long history too. A complete description of all non-dense parts, called *gaps*, for the homomorphism order of undirected graphs, was given in 1982 by Welzl [43].

For directed graphs and general relational structures, density has a non-obvious link to duality. It was shown by Nešetřil and Tardif [34] that all gaps correspond to duality pairs; we survey their results in Section 3.2. This connection can be extended from the homomorphism order to lattices satisfying some extra axioms (Heyting algebras with finite connected decompositions). Such an extension is presented in Section 3.3. Some of the ideas are contained in a paper of Nešetřil, Pultr and Tardif [31]. We add the description of finite dualities.

Our further interest concentrates on another issue. In Section 3.4 we study finite maximal antichains in the homomorphism order. In particular, we show that with a few characterised exceptions finite maximal antichains have the splitting property. This by itself provides a connection to finite homomorphism dualities but in the case of relational structures with at most two relations we can prove that even the exceptional antichains are formed from dualities. For structures with more than two relations this question remains open.

3.1 Homomorphism order

3.1.1. The relation $\to$ of being homomorphic is reflexive, as the identity mapping is a homomorphism from a Δ-structure to itself, and it is transitive, since the composition of two homomorphisms, if possible, is a homomorphism too. Thus $\to$ is a preorder.

There are standard ways to transform a preorder into a partial order. It may be done by identifying equivalent objects, or by choosing a particular representative for each equivalence class. The resulting partial order is identical in both cases.

For $\to$, a suitable representative for each equivalence class is a core Δ-structure. We have already observed in 1.4.7 that there is a unique core in each class of homomorphic equivalence; unique up to isomorphism.

3.1.2 Proposition. *Let Δ be a fixed type. Then the relation $\to$ of being homomorphic is a partial order on the set of all core Δ-structures (taken up to isomorphism).*

Proof. Follows from the discussion above. □

3.1.3 Definition. The partial order $\to$ from Proposition 3.1.2 is called the **homomorphism order** and denoted by $\mathcal{C}(\Delta)$.

Any treatise on the homomorphism order is substantially simplified by talking about the order of Δ-structures rather than cores or equivalence classes. For instance, when we say that $\mathbf{A}$ is less than $\mathbf{B}$ in the homomorphism order, we mean that *the core of* $\mathbf{A}$ is less than *the core of* $\mathbf{B}$ in the homomorphism order. Similarly, when we (soon) say that $\mathbf{A} \times \mathbf{B}$ is the infimum of $\mathbf{A}$ and $\mathbf{B}$ in the homomorphism order, we mean that *the core of* $\mathbf{A} \times \mathbf{B}$ is actually the infimum. This approach is fairly standard in algebra.

With all this in mind, we observe that the homomorphism order is a nice partial order: it is a lattice, and moreover a Heyting algebra.

3.1.4 Proposition. *The homomorphism order $\mathcal{C}(\Delta)$ is a Heyting algebra. In particular, for $\mathbb{A}, \mathbb{B} \in \mathcal{C}(\Delta)$*

(1) the product $\mathbb{A} \times \mathbb{B}$ is the infimum (meet) of $\mathbb{A}$ and $\mathbb{B}$,

(2) the sum $\mathbb{A} + \mathbb{B}$ is the supremum (join) of $\mathbb{A}$ and $\mathbb{B}$,

(3) one vertex with no edges $\big(\{1\}, (\emptyset, \emptyset, \dots \emptyset)\big) =: \bot$ is the least element, and one vertex with all loops $\big(\{1\}, (\{1\}^{\delta_i} : i \in I)\big) =: \top$ is the greatest element in $\mathcal{C}(\Delta)$,

(4) the exponential structure $\mathbb{B}^{\mathbb{A}}$ is the Heyting operation $\mathbb{A} \Rightarrow \mathbb{B}$.

Proof. (1) By 1.5.8, $\mathbb{C} \to \mathbb{A} \times \mathbb{B}$ if and only if $\mathbb{C} \to \mathbb{A}$ and $\mathbb{C} \to \mathbb{B}$. So $\mathbb{A} \times \mathbb{B}$ is the infimum of $\mathbb{A}$ and $\mathbb{B}$.

(2) By 1.5.4, $\mathbb{A} + \mathbb{B} \to \mathbb{C}$ if and only if $\mathbb{A} \to \mathbb{C}$ and $\mathbb{B} \to \mathbb{C}$. So $\mathbb{A} + \mathbb{B}$ is the supremum of $\mathbb{A}$ and $\mathbb{B}$.

(3) Let $\mathbb{A}$ be an arbitrary Δ-structure. By definition, $\underline{\mathbb{A}}$ is non-empty and clearly any function mapping v to an arbitrary element of $\underline{\mathbb{A}}$ is a homomorphism from $\bot$ to $\mathbb{A}$. So $\bot$ is the least element.

On the other hand, let f be the constant function from $\underline{\mathbb{A}}$ to $\{v\}$ such that $f(a) = v$ for all $a \in \underline{\mathbb{A}}$. Since $\top$ has all loops, all edges of $\mathbb{A}$ are preserved by f and hence it is a homomorphism from $\mathbb{A}$ to $\top$. Therefore $\top$ is the greatest element.

(4) By 1.5.12, $\mathbb{C} \to \mathbb{B}^{\mathbb{A}}$ if and only if $\mathbb{A} \times \mathbb{C} \to \mathbb{B}$. Since products are infima, this is exactly the Heyting axiom (see Definition 3.3.1). □

To illustrate the homomorphism order's power, we give (without proof) one more example of its properties. The homomorphism order is a universal countable partial order. Several proofs of this can be found in the literature [16, 19, 20, 27, 39].

3.1.5 Theorem. *Let Δ be a type with at least one relation of arity at least two. Every countable partial order is an induced suborder of the homomorphism order of Δ-structures.* □

3.2 Gaps and dualities

In this section we briefly survey the results of [34] about a connection between duality pairs (see Section 2.1) and gaps in the homomorphism order. The explicit description of gaps, besides being of interest by itself, provides a different proof of the characterisation of duality pairs (Theorem 2.1.12). We have also used it for proving Theorem 2.4.26.

The first fact we state is that the top of a gap, if connected, is the left-hand side of a duality pair.

3.2.1 Lemma ([34]). *Let* (A, B) *be a gap pair and let* B *be connected. Then* (B, A^B) *is a duality pair.* □

Hence a connected top of a gap is (homomorphically equivalent) to a Δ-tree.

The following proposition characterises all gaps.

3.2.2 Proposition ([34]). *Gaps are exactly all the pairs* (A, B) *such that there exists a duality pair* (F, D) *with* $F \to B \to F + D$ *and* $A \sim B \times D$. *Moreover,* $B \sim A + F$. □

The correspondence is depicted in the following two diagrams, in which double arrows denote gaps. Here (F, D) is a duality pair.

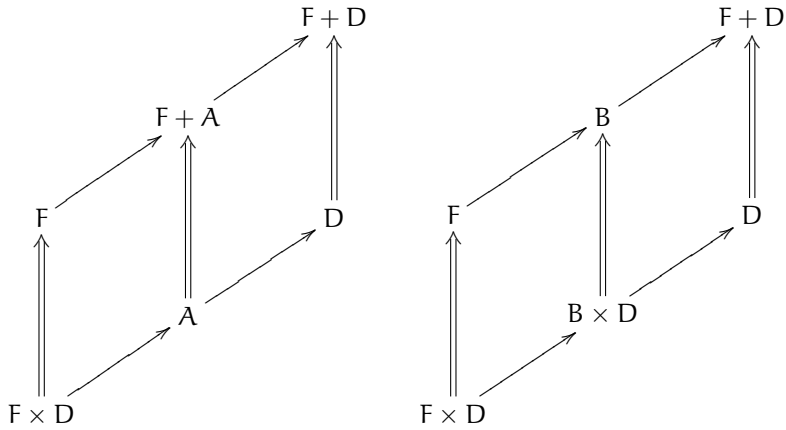

3.3 Dualities and gaps in Heyting algebras

The previous section presents a connection between finite dualities and gaps in the homomorphism order. However, few properties typical of the homomorphism order were used to prove them. Here we look at a more general case. We provide conditions under which a theory of gaps and dualities can be developed for partially ordered sets.

Gaps, duality pairs and combined dualities (which correspond to finite dualities with a singleton right-hand side) in Heyting algebras have been studied by Nešetřil, Pultr and Tardif [31]. We extend their results to a complete description of dualities.

At the same time, Proposition 3.1.4 implies that dualities for relational structures are a special case of this general theory.

3.3.1 Definition. A lattice P with an additional binary operation $\Rightarrow$ is a **Heyting algebra** if a least element and a greatest element exist in P and for all $p, q, r \in P$,

$$p \preccurlyeq q \Rightarrow r \quad \text{if and only if} \quad p \wedge q \preccurlyeq r.$$

Of course, not every lattice is distributive. However, it is well known that every Heyting algebra is distributive.

3.3.2 Lemma. *Every Heyting algebra is a distributive lattice.*

Proof. First we show that $(a \wedge b) \vee (a \wedge c) \preccurlyeq a \wedge (b \vee c)$. This is true in every lattice. Clearly $a \wedge b \preccurlyeq a$ and $a \wedge b \preccurlyeq b \vee c$, so $a \wedge b \preccurlyeq a \wedge (b \vee c)$. Similarly $a \wedge c \preccurlyeq a \wedge (b \vee c)$. Hence the inequality holds.

Next, it suffices to prove that whenever $a \wedge b \preccurlyeq y$ and $a \wedge c \preccurlyeq y$, then $a \wedge (b \vee c) \preccurlyeq y$. In connection with the previous paragraph, it implies that $a \wedge (b \vee c)$ is the supremum of $a \wedge b$ and $a \wedge c$.

So suppose that $a \wedge b \preccurlyeq y$ and $a \wedge c \preccurlyeq y$. Then $b \preccurlyeq a \Rightarrow y$ and $c \preccurlyeq a \Rightarrow y$. Thus $b \vee c \preccurlyeq a \Rightarrow y$. Hence $a \wedge (b \vee c) \preccurlyeq y$. □

An important property for the development of duality theory for relational structures was the existence of a decomposition of every relational structure into connected components. We generalise connectedness in the context of Heyting algebras.

3.3.3 Definition. Let L be a lattice. An element a of L is **connected** if the equality $a = b \vee c$ implies that $a = b$ or $a = c$.

Next we observe that connectedness in distributive lattices (and thus in Heyting algebras) has the same equivalent descriptions (1)–(3) as in Lemma 1.6.7.

3.3.4 Lemma. *Let a be an element of a distributive lattice L. Then the following conditions are equivalent.*

(1) If $a \preccurlyeq b \vee c$ for some elements b, c of L, then $a \preccurlyeq b$ or $a \preccurlyeq c$.

(2) If $a = b \vee c$ for some elements b, c of L, then $b \preccurlyeq c$ or $c \preccurlyeq b$.

(3) The element a is connected.

Proof.

(1) ⇒ (2): If $a = b \vee c$, then $a \preccurlyeq b \vee c$, and using (1) we have $a \preccurlyeq b$ or $a \preccurlyeq c$. In the first case $c \preccurlyeq b \vee c = a \preccurlyeq b$, hence $c \preccurlyeq b$. In the latter case $b \preccurlyeq b \vee c = a \preccurlyeq c$, and so $b \preccurlyeq c$.

(2) $\Rightarrow$ (3): Suppose $a = b \vee c$. By (2) we have $b \preccurlyeq c$, and therefore $c = b \vee c = a$; or we have $c \preccurlyeq b$, and then $b = b \vee c = a$.

(3) $\Rightarrow$ (1): Let $a \preccurlyeq b \vee c$. Here we need distributivity: $(a \wedge b) \vee (a \wedge c) = a \wedge (a \vee b) \wedge (a \vee c) = a$ since $a \preccurlyeq a \vee b$ and $a \preccurlyeq a \vee c$. □

The existence of connected components is then generalised by the following notion.

3.3.5 Definition. We say that a lattice L **has finite connected decompositions** if each element x of L is a supremum of a finite set of connected elements.

Analogously as for relational structures (Definition 2.1.1) we define duality pairs for lattices.

3.3.6 Definition. A pair (f, d) of elements of a lattice L is a **duality pair** if for any element $x \in L$,

$$f \preccurlyeq x \quad \text{if and only if} \quad x \not\preccurlyeq d.$$

3.3.7 Definition. An element f of a lattice L is called a **primal** if there exists $d \in L$ such that (f, d) is a duality pair. An element d of a lattice L is called a **dual** if there exists $f \in L$ such that (f, d) is a duality pair.

The next proposition is an analogue of 2.1.13.

3.3.8 Proposition ([31])**.** *In a distributive lattice, every primal is connected.*

Proof. We prove that if (f, d) is a duality pair and $f \preccurlyeq b \vee c$ for some elements $b, c \in L$, then $f \preccurlyeq b$ or $f \preccurlyeq c$. By Lemma 3.3.4 it follows that f is connected.

So suppose that $f \preccurlyeq b \vee c$. By duality, $b \vee c \not\preccurlyeq d$, thus (by a property of join) $b \not\preccurlyeq d$ or $c \not\preccurlyeq d$. Using duality once again we get that $f \preccurlyeq b$ or $f \preccurlyeq c$. □

Recall that a gap in a poset L is a pair (p, q) of elements of L such that $p \prec q$ and no element r satisfies that $p \prec r \prec q$ (Definition 1.9.12). The connection between gaps and duality pairs (Proposition 3.2.2 for relational structures) is as follows.

3.3.9 Theorem ([31])**.** *The gaps in a Heyting algebra* L *with finite connected decompositions are exactly the pairs* (a, b) *such that for some duality pair* (f, d)

$$f \wedge d \preccurlyeq a \preccurlyeq d \quad \textit{and} \quad b = a \vee f.$$

□

3.3.10 Definition. A pair (F, D) of finite subsets of a lattice L is a **finite duality** if

1. $f \not\preccurlyeq f'$ if $f, f' \in F$ and $f \neq f'$,
2. $d \not\preccurlyeq d'$ if $d, d' \in D$ and $d \neq d'$, and
3. for any $x \in L$ there exists $f \in F$ such that $f \preccurlyeq x$ if and only if $x \not\preccurlyeq d$ for any $d \in D$.

The following theorem is an analogue of Theorem 2.4.4. It describes finite dualities with a singleton right-hand side. The proof is just a translation of the relational-structure proof of Theorem 2.4.4 into the language of Heyting algebras, therefore we do not repeat it here.

3.3.11 Theorem. *Let L be a Heyting algebra with finite connected decompositions. For a finite subset $F = \{f_1, f_2, \dots, f_n\}$ of L and for an element $d \in L$, the pair $(F, \{d\})$ is a finite duality if and only if there exist elements $d_1, d_2, \dots, d_n$ such that (f_i, d_i) is a duality pair for $i = 1, 2, \dots, n$ and $d = d_1 \wedge d_2 \wedge \cdots \wedge d_n$.* □

The **transversal construction** of dualities in a lattice L with finite connected decompositions is defined analogously to the definition in the context of Δ-structures, contained in paragraphs 2.4.7–2.4.13.

3.3.12 Definition. Let L be a lattice with finite connected decompositions.

Let $F = \{f_1, f_2, \dots, f_m\}$ be an arbitrary fixed non-empty finite set of pairwise incomparable elements of L. For each element f_i of F fix a finite connected decomposition

$$f_i = \bigvee_{j=1}^{k_i} c_{i,j}.$$

Now let F_c be the set of all connected elements appearing in the decompositions, that is

$$F_c := \bigcup_{i=1}^{m} \{c_{i,j} : 1 \leq j \leq k_i\}.$$

Then F_c is called a **set of components** for F.

Quasitransversals are defined analogously to quasitransversals for relational structures (see 2.4.8).

3.3.13 Definition. A subset $M \subseteq F_c$ is a **quasitransversal** if it satisfies

(T1) any two distinct elements of M are incomparable, and

(T2) M **supports** F, that is for every $f \in F$ there exists $c \in M$ such that $c \preccurlyeq f$.

3.3.14. For two quasitransversals M, M' we define that $M \trianglelefteq M'$ if and only if for every $c' \in M'$ there exists $c \in M$ such that $c \preccurlyeq c'$.

3.3.15 Lemma. *The relation $\trianglelefteq$ is a partial order on the set of all quasitransversals.*

Proof. Obviously, $\trianglelefteq$ is both reflexive and transitive (a preorder).

Suppose now that $M \trianglelefteq M'$ and $M' \trianglelefteq M$, and let $c \in M$. Then there exists $c' \in M'$ such that $c' \preccurlyeq c$ and there exists $c'' \in M$ such that $c'' \preccurlyeq c'$. As a result $c'' \preccurlyeq c$, thus by (T1) we have $c = c' = c''$, hence $M \subseteq M'$. Similarly we get that $M' \subseteq M$. Hence $M = M'$ whenever $M \trianglelefteq M'$ and $M' \trianglelefteq M$. Therefore $\trianglelefteq$ is antisymmetric; it is a partial order. □

3.3.16 Definition. A quasitransversal M is a **transversal** if

(T3) M is maximal with respect to the order $\trianglelefteq$.

A characterisation similar to Theorem 2.4.26 follows.

3.3.17 Theorem. *Let L be a Heyting algebra with finite connected decompositions.*

Let $F \subseteq L$ be finite. If each element of F is a finite join of primals, and D is the result of the transversal construction, then (F, D) is a finite duality.

Conversely, if (F, D) is a finite duality, then each element of F decomposes into a finite join of primals and D is the result of the transversal construction. □

In contrast to the homomorphism order, this does not in general mean that the right-hand (dual) side of a finite duality is uniquely determined by the left-hand (primal) side. That is so because the decomposition into connected components may not be unique, so the transversal construction produces a different result for different decompositions.

We do not give a detailed proof because the proof is a translation of the proof of Theorem 2.4.26. It suffices to check that for proving the lemmas in Section 2.4 we did not use any other properties of relational structures than the homomorphism order's being a Heyting algebra with finite connected decompositions. That is no longer true in the next section.

3.4 Finite maximal antichains

In this section we study finite maximal antichains in the homomorphism order. In particular, we are interested in the *splitting property* of these antichains.

If Q is an arbitrary maximal antichain in a poset P, then every element of P is comparable with some element q of Q. In other words, the poset P is the union of the downset generated by Q and the upset generated by Q, that is $P = Q^{\uparrow} \cup Q^{\downarrow}$. The splitting property of the antichain Q means that Q can be split into two subsets Q_1, Q_2 such that the poset P is the union of the upset generated by Q_1 and the downset generated by Q_2. So any element of P is either above some element in Q_1 or below some element in Q_2 (see Figure 3.1).

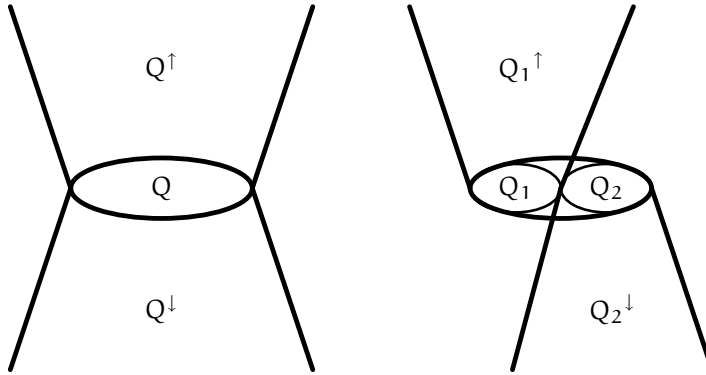

Figure 3.1: The splitting of an antichain

A formal definition follows.

3.4.1 Definition. We say that a maximal antichain $Q \subset P$ **splits** if there exists a partition of S into disjoint subsets Q_1 and Q_2 such that $P = Q_1{}^{\uparrow} \cup Q_2{}^{\downarrow}$. In such a case we say that (Q_1, Q_2) is a **splitting** of the maximal antichain Q.

The problem of splitting maximal antichains in the homomorphism order took on significance when it was observed by Nešetřil and Tardif [35] that except for two small exceptions finite maximal antichains of size two split in the order of digraphs. That result is extended here to Δ-structures and to finite maximal antichains of any size; however, the description of exceptions is more involved.

Our approach is direct. In 3.4.2 we define a partition of any finite maximal antichain and prove that – apart from exceptional cases described later – this partition is a splitting of the antichain.

At the end of this section, we suggest an alternative approach that may lead to a different proof of the splitting property of finite maximal antichains. It is based on a general condition for the splitting of antichains in arbitrary posets by Ahlswede, Erdős, Graham and Soukup [2, 9].

There is also a link to homomorphism dualities because a splitting of a finite maximal antichain is trivially a finite homomorphism duality. In the case of relational structures with at most two relations, we show that the link is stronger: even those finite maximal antichains that do not split correspond to homomorphism dualities.

For more than two relations this is unknown, but there is a significant increase in the complexity of the homomorphism order. This suggests that the property may not hold in this case.

Splitting finite antichains

We would like to partition a finite maximal antichain $\mathcal{Q}$ in the homomorphism order $\mathcal{C}(\Delta)$ into disjoint sets $\mathcal{F}$ and $\mathcal{D}$ in such a way that $\mathcal{F}^{\uparrow} \cup \mathcal{D}^{\downarrow} = \mathcal{C}(\Delta)$. A partition is defined in the next paragraph. In the following we show that in many cases it satisfies the equality.

3.4.2 Splitting a finite antichain. Let $\mathcal{Q} = \{Q_1, Q_2, \ldots, Q_n\}$ be a finite maximal antichain in $\mathcal{C}(\Delta)$. Recursively, define the sets $\mathcal{F}_0$, $\mathcal{F}_1$, ..., $\mathcal{F}_n$ in this way:

1. Let $\mathcal{F}_0 = \emptyset$.
2. For $i = 1, 2, \ldots, n$: check whether there exists a Δ-structure X satisfying
   (i) $Q_i < X$,
   (ii) $F \nrightarrow X$ for any $F \in \mathcal{F}_{i-1}$, and
   (iii) $Q_j \nrightarrow X$ for any $j > i$.
   If such a structure X exists, let $\mathcal{F}_i = \mathcal{F}_{i-1} \cup \{Q_i\}$, otherwise let $\mathcal{F}_i = \mathcal{F}_{i-1}$.
3. Finally, let $\mathcal{F} = \mathcal{F}_n$ and $\mathcal{D} = \mathcal{Q} \setminus \mathcal{F}$.

Because $\mathcal{Q}$ is a maximal antichain, $\mathcal{Q}^{\uparrow} \cup \mathcal{Q}^{\downarrow} = \mathcal{C}(\Delta)$. In addition, $\mathcal{F}^{\uparrow} \subseteq \mathcal{Q}^{\uparrow}$ and $\mathcal{D}^{\downarrow} \subseteq \mathcal{Q}^{\downarrow}$ since $\mathcal{F} \subseteq \mathcal{Q}$ and $\mathcal{D} \subseteq \mathcal{Q}$. Therefore the equality

$$\mathcal{F}^{\uparrow} \cup \mathcal{D}^{\downarrow} = \mathcal{C}(\Delta),$$

which characterises the splitting of the antichain $\mathcal{Q}$, is equivalent to the pair of equalities

$$\mathcal{F}^{\uparrow} = \mathcal{Q}^{\uparrow},$$
$$\mathcal{D}^{\downarrow} = \mathcal{Q}^{\downarrow}.$$

The following lemma asserts that $\mathcal{F}^{\uparrow} = \mathcal{Q}^{\uparrow}$.

3.4.3 Lemma. *Let $\mathcal{Q}$ be a finite maximal antichain and $\mathcal{F}$, $\mathcal{D}$ be defined in 3.4.2. If $Q \in \mathcal{Q}$, X is a Δ-structure, and $Q < X$, then there exists $F \in \mathcal{F}$ such that $F < X$.*

Proof. Among the elements of $\mathcal{Q}$ that are homomorphic to X, let Q_i be the element of $\mathcal{Q}$ with the greatest index i. Then either $F \to X$ for some $F \in \mathcal{F}_{i-1}$, or all the conditions (i), (ii), (iii) are satisfied and $Q_i \in \mathcal{F}$. So we have found $F \in \mathcal{F}$ such that $F \to X$.

If $F = Q$, then $X \nrightarrow F$ by the assumption that $Q < X$. If on the other hand $F \neq Q$, then the existence of a homomorphism from X to F would imply that $Q \to F$. This is a contradiction because F and Q are distinct elements of an antichain. Hence $X \nrightarrow F$ and therefore $F < X$. □

To prove that $(\mathcal{F}, \mathcal{D})$ is a splitting of $\mathcal{Q}$, it remains to show that $\mathcal{D}^{\downarrow} = \mathcal{Q}^{\downarrow}$. However, this is not true for all finite maximal antichains. The following lemma provides a simple description of antichains for which $(\mathcal{F}, \mathcal{D})$ is not a splitting.

3.4.4 Lemma. *Let $\mathcal{Q}$ be a finite maximal antichain and $\mathcal{F}$, $\mathcal{D}$ be defined in 3.4.2. Then exactly one of the following conditions holds:*

(1) The pair $(\mathcal{F}, \mathcal{D})$ is a splitting of $\mathcal{Q}$.

(2) There exists a structure Y such that $Q \nrightarrow Y$ for any $Q \in \mathcal{Q}$ and $Y \nrightarrow D$ for any $D \in \mathcal{D}$.

Proof. If $(\mathcal{F}, \mathcal{D})$ is a splitting and Y is an arbitrary structure such that $Q \nrightarrow Y$ for any $Q \in \mathcal{Q}$, then $Y \to Q$ for some $Q \in \mathcal{Q}$ because $\mathcal{Q}$ is a maximal antichain. Moreover, $Q \in \mathcal{D}$ because of splitting.

Conversely, suppose $(\mathcal{F}, \mathcal{D})$ is not a splitting. So there exists a structure Y that violates the definition of a splitting: $\mathcal{F} \nrightarrow Y$ and $Y \nrightarrow \mathcal{D}$. Because of Lemma 3.4.3 we have $\mathcal{Q} \nrightarrow Y$. □

If (1) holds, then the antichain $\mathcal{Q}$ splits. Now we investigate those maximal antichains that satisfy (2). The structure Y has to be comparable with some element of the maximal antichain $\mathcal{Q}$, and because of the condition (2)

there exists $F \in \mathcal{F}$ such that $Y < F$. We show that all such Y's are bounded from above by a fairly simple structure.

However, first we need some preparation. The recursive definition 3.4.2 assures that for every element F of $\mathcal{F}$ there is a witness X that forces F to be added to $\mathcal{F}$. This is formally expressed in the following lemma. The witness for F is denoted by $\check{F}$.

3.4.5 Lemma. *Let $\mathcal{Q}$ be a finite maximal antichain and $\mathcal{F}$, $\mathcal{D}$ be defined in 3.4.2. For every F in $\mathcal{F}$ there exists a Δ-structure $\check{F}$ such that $F < \check{F}$ and moreover F is the only element of $\mathcal{F}$ that is homomorphic to $\check{F}$.*

Proof. The structure X satisfying the properties (i), (ii), (iii), which caused $F = Q_i$ to be an element of $\mathcal{F}_i$ has the required properties for $\check{F}$. □

We use a tool, which is a generalisation of a famous theorem of Erdős [7] (this was one of the first applications of the then emerging probabilistic method).

3.4.6 Theorem ([8, 25, 32])**.** *Let Δ be an arbitrary type, and let g and k be positive integers. Then there exists a Δ-structure $G = G(g, k)$ such that*

- *every substructure of G induced by at most g vertices is a Δ-forest, and*
- *whenever the vertices of G are coloured by fewer than k colours, there exists a colour κ that induces an edge of each kind; that is, for each kind $i \in I$ there is an edge $e \in R_i(G)$ such that all vertices of e have colour κ.* □

Recall from Definition 1.8.1 that a balanced structure is a structure that is homomorphic to a forest.

3.4.7 Lemma. *Let $\mathcal{Q}$ be a finite maximal antichain and $\mathcal{F}$, $\mathcal{D}$ be defined in 3.4.2. If $F \in \mathcal{F}$, then F is balanced.*

Proof. Let $F \in \mathcal{F}$ be arbitrary and let $\check{F}$ be the structure whose existence is guaranteed by Lemma 3.4.5. Furthermore, let $k := \max\{|\underline{Q}|^{|\underline{\check{F}}|} : Q \in \mathcal{Q}\} + 1$ and let G be a Δ-structure such that any substructure of G on at most $|\underline{F}|$ vertices is a Δ-forest and whenever the vertices of G are coloured by fewer than k colours, there exists a colour that induces an edge of each kind (Theorem 3.4.6).

Consider the structure $H = \check{F} \times G$. Suppose that $f : H \to Q$ for some $Q \in \mathcal{Q}$. For every vertex u of G, the mapping $f_u : \underline{\check{F}} \to \underline{Q}$ is defined by $f_u(x) = f(u, x)$. We consider this assignment of mappings to vertices of G as a colouring of the vertices. Since there are only $|\underline{Q}|^{|\underline{\check{F}}|} < k$ possible mappings

from $\underline{\check{F}}$ to $\underline{Q}$, there exists a colour that induces an edge of every kind; so there exists a mapping $g : \underline{\check{F}} \to \underline{Q}$ satisfying the condition that

$$\text{for every } i \in I \text{ there exists an edge } (u_1, u_2, \dots, u_{\delta_i}) \text{ of } G \text{ such that } f_{u_1} = f_{u_2} = \cdots = f_{u_{\delta_i}} = g. \quad (3.1)$$

Then g is a homomorphism from $\check{F}$ to Q: whenever $(x_1, x_2, \dots, x_{\delta_i}) \in R_i(\check{F})$ for some $i \in I$, we have that

$$\begin{aligned}\big(g(x_1), g(x_2), \dots, g(x_{\delta_i})\big) &= \big(f_{u_1}(x_1), f_{u_2}(x_2), \dots, f_{u_{\delta_i}}(x_{\delta_i})\big) \\ &= \big(f(u_1, x_1), f(u_2, x_2), \dots, f(u_{\delta_i}, x_{\delta_i})\big) \in R_i(Q),\end{aligned}$$

because $\big((u_1, x_1), (u_2, x_2), \dots, (u_{\delta_i}, x_{\delta_i})\big) \in R_i(H)$ and f is a homomorphism from H to Q; here $(u_1, u_2, \dots, u_{\delta_i})$ is the edge of G from (3.1). That is a contradiction, because $F < \check{F}$ and $\mathcal{Q}$ is an antichain containing F. We conclude that $H \nrightarrow Q$ for any $Q \in \mathcal{Q}$.

By Lemma 3.4.3 and because $\mathcal{Q}$ is a maximal antichain, there exists $F' \in \mathcal{F}$ that is homomorphic to $H = \check{F} \times G \to \check{F}$. But F is the only element of $\mathcal{F}$ that is homomorphic to $\check{F}$, so we have $F' = F$, and consequently $F \to H \to G$. The image of a homomorphism from F to G has no more than $|\underline{F}|$ vertices, whence it is a forest. This concludes the proof. □

The elements of $\mathcal{F}$ are balanced for all finite maximal antichains, even for those for which $(\mathcal{F}, \mathcal{D})$ is a splitting.

In the following, we closely investigate the "non-splitting" antichains. We derive properties of the structures Y that violate the splitting (as in condition (2) of Lemma 3.4.4). In particular, we show that some paths – called *forbidden paths* – are not homomorphic to Y.

3.4.8 Definition. Every Δ-path has a height labelling; we say that a core Δ-path P is a **forbidden path** if it has two edges of the same kind whose vertices are not labelled the same. (This property does not depend on what height labelling we choose, see Proposition 1.8.3.)

3.4.9 Lemma. *Let $\mathcal{Q}$ be a finite maximal antichain in $\mathcal{C}(\Delta)$ and let $\mathcal{F}$, $\mathcal{D}$ be defined in 3.4.2. If Y is a Δ-structure such that $Y \nrightarrow D$ for any $D \in \mathcal{D}$ and $Y < F$ for some $F \in \mathcal{F}$, and P is a forbidden path, then $P \nrightarrow Y$.*

Proof. We may suppose that the two edges of the same kind $i \in I$ that are not labelled the same (which prove that P is indeed a forbidden path) are the end edges of P. Otherwise we could take a subpath of P (the smallest

substructure of P that contains both these edges) and show that it is not homomorphic to Y; consequently P is not homomorphic to Y either.

Let the two end edges of P be $(x_1, x_2, \ldots, x_{\delta_i}), (y_1, y_2, \ldots, y_{\delta_i}) \in R_i(P)$. At most one of the vertices x_1, x_{δ_i} is contained in another edge of P, and so is at most one of y_1, y_{δ_i}; if a vertex is contained in only one edge, we call it **free**.

Let Z be a long zigzag: a path with $2m$ or $2m+1$ edges, depending on the end edges of P. If both x_1 and y_1 are free or if both x_{δ_i} and y_{δ_i} are free, we use an even number of edges; otherwise we use an odd number of edges. All edges of the zigzag are of the same kind as the end edges of P.

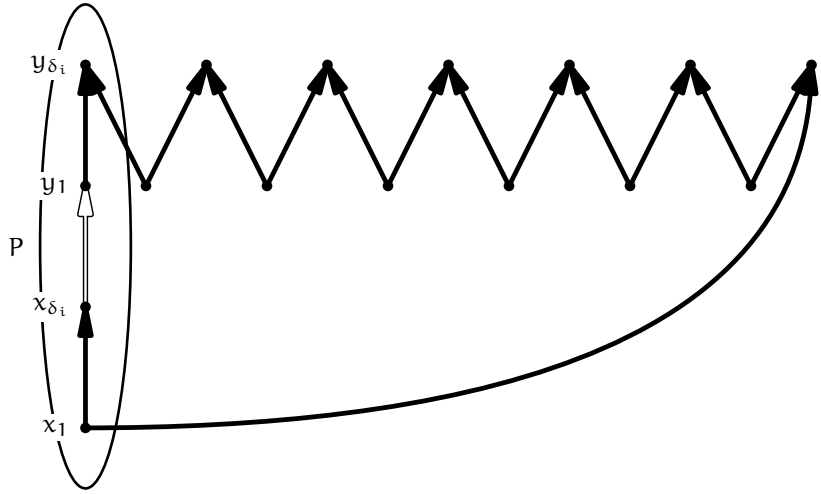

Figure 3.2: Constructing Z from a forbidden path P

Even though the definition of Z and W should be clear from Figure 3.2, we may also define them formally here. Suppose x_1 and y_1 are free vertices in P. Then

$$\begin{aligned} \underline{Z} &:= \{1, 2, \ldots, 2m(\delta_i - 1) + 1\}, \\ R_i(Z) &:= \{(k, k+1, \ldots, k+\delta_i - 1), (k+\delta_i - 1, k+\delta_i, \ldots, k+2\delta_i - 2) : \\ &\qquad k = 1, 1 + 2(\delta_i - 1), 1 + 4(\delta_i - 1), \ldots, 1 + 2(m-1)(\delta_i - 1)\}, \end{aligned}$$

$\approx$ is an equivalence relation on $\underline{P+Z}$ with $x_1 \approx 1$, $y_1 \approx 2m(\delta_i - 1) + 1$, and $a \approx a$, and finally

$$W := (P + Z)/\approx.$$

If other vertices in the end edges of P are free, the description is analogous.

Clearly, any proper substructure of W that does not contain all vertices of Z is homomorphic to P. We choose the length of the zigzag (by a suitable choice of m) in such a way that the number of vertices of Z is bigger than the number of vertices of any structure in $\mathcal{F}$.

Now observe that if there exists a height labelling of W, the vertices in the end edges of P have the same labels because they are joined by the zigzag; at the same time, they have distinct labels, because they are joined by the forbidden path P. Therefore no height labelling of W exists and by Proposition 1.8.3 the structure W is not balanced.

Now consider the sum $W + Y$. It is comparable with some some element of the maximal antichain $\mathcal{Q}$. However, $W + Y \not\to D$ for any $D \in \mathcal{D}$, because $Y \not\to D$ by our assumption on Y; also $W + Y \not\to F$ for any $F \in \mathcal{F}$, because W is not balanced and F is (so F is homomorphic to a forest, but W is not). Therefore $F \to W + Y$ for some $F \in \mathcal{F}$.

However, because the zigzag Z was very long, the image of a homomorphism from F to $W + Y$ does not contain all vertices of Z. As we have observed, therefore $F \to P + Y$. As $F \not\to Y$ (by the definition of Y), necessarily $P \not\to Y$. □

In the next lemma we prove that complex structures admit homomorphisms from forbidden paths, and correspondingly structures that admit no homomorphisms from forbidden paths are simple.

3.4.10 Lemma. *Let C be a connected Δ-structure. If no forbidden path is homomorphic to C, then C is homomorphic to a tree with at most one edge of each kind.*

Proof. Suppose that no forbidden path is homomorphic to C. If no height labelling of C exists, then there exist vertices u and v connected by two distinct paths in $\mathrm{Sh}(C)$ such that counting forward steps minus backward steps on these paths gives a different result (see 1.8.4). Let B be the minimal structure such that its shadow $\mathrm{Sh}(B)$ contains both of these paths ($\underline{B} \subseteq \underline{C}$, but it need not be an *induced* substructure; include only those edges whose shadow edges lie in the two paths). Let us "unfold" this structure B, which in a way resembles a cycle: choose an edge e of B, it intersects two other edges. Then construct a path P: its end edges are two copies of e, and the middle edges are the remaining edges of B.

Then $P \to C$ but P is a forbidden path, because the copies of e get different labels. Therefore C has a height labelling ↴.

Next observe that any two edges of the same kind are labelled in the same way. If there were two edges with differently labelled vertices, there would

be a path in C with these two edges as end edges (because C is connected) and that would be a forbidden path.

Let T be the structure with base set $\beth[\underline{C}]$, all labels used on the vertices of C. The edges of T are such that the identity mapping is a height labelling of T; in other terms

$$R_i(T) = \big\{(x_1, x_2, \dots, x_{\delta_i}) : (x_{j+1})_{(i,j)} = (x_j)_{(i,j)} + 1, \\ (x_{j+1})_{(i',j')} = (x_j)_{(i',j')} \text{ for } (i',j') \neq (i,j)\big\}.$$

Because all the edges of the same kind in C have the same labelling, T has at most one edge of each kind. Moreover, any cycle in $\mathrm{Sh}(T)$ would violate the height labelling of T, so T is a tree. And finally, the height labelling $\beth$ of C is a homomorphism from C onto T. □

3.4.11 Corollary. *Let* D^* *be the sum of all* Δ*-trees with at most one edge of each kind. If* Y *satisfies the condition (2) of 3.4.4, then* $Y \to D^*$.

Proof. The claim follows immediately from Lemmas 3.4.9 and 3.4.10. □

This shows that the cases when the antichain does not split are very specific (and one would like to say they are rather rare).

3.4.12 Theorem. *Let* $\mathcal{Q}$ *be a finite maximal antichain in* $\mathcal{C}(\Delta)$. *Let* D^* *be the sum of all* Δ*-trees with at most one edge of each kind. Suppose that every element* $Q \in \mathcal{Q}$ *has the property that whenever* $Y < Q$ *and* $Y \to D^*$ *then there exists a* Δ*-structure* X *such that* $Y < X < Q$ *and* $X \nrightarrow D^*$. *Then the antichain* $\mathcal{Q}$ *splits; the pair* $(\mathcal{F}, \mathcal{D})$ *defined in 3.4.2 is a splitting of* $\mathcal{Q}$.

Proof. For the sake of contradiction, suppose that the pair $(\mathcal{F}, \mathcal{D})$ is not a splitting of $\mathcal{Q}$. By Lemma 3.4.4, there exists a structure Y such that $Q \nrightarrow Y$ for any $Q \in \mathcal{Q}$ and $Y \nrightarrow D$ for any $D \in \mathcal{D}$. Since $\mathcal{Q}$ is a maximal antichain, Y is comparable with an element Q of $\mathcal{Q}$; thus there exists $Q \in \mathcal{Q}$ with $Y < Q$.

Since Y satisfies the condition (2) of 3.4.4, by Corollary 3.4.11 we have $Y \to D^*$. Thus by assumption there exists X such that $Y < X < Q$ and $X \nrightarrow D^*$. The structure X is not homomorphic to any $D \in \mathcal{D}$, because otherwise we would have $Y \to D$ by composition. Hence X satisfies the condition (2) of 3.4.4 as well, and by Corollary 3.4.11 it is homomorphic to D^*, a contradiction. □

The assumption on the elements of $\mathcal{Q}$ posed in the previous theorem means that no element of $\mathcal{Q}$ is "too small". In particular, it is neither homomorphic to D^* nor "immediately" above a structure homomorphic to D^*, that is, there is no gap (Y, Q) such that $Y \to D^*$.

In fact, the assumption can be weakened (and the theorem strengthened) by requiring that only elements of $\mathcal{F}$ constructed from $\mathcal{Q}$ by 3.4.2 have the property that whenever $Y < Q$ and $Y \to D^*$ then there exists a Δ-structure X such that $Y < X < Q$ and $X \nrightarrow D^*$. This is obvious from the proof, where we exploited the property only for elements of $\mathcal{F}$.

Connection to finite dualities

Further examination reveals that in the case of structures with at most two relations there are no infinite increasing chains below D^*. From that we can conclude that all elements of $\mathcal{F}$ are Δ-forests and thus we get the following theorem.

3.4.13 Theorem. *Let $\Delta = (\delta_i : i \in I)$ be a type such that $|I| \leq 2$. Then all finite maximal antichains in the homomorphism order $\mathcal{C}(\Delta)$ are exactly the sets*

$$\mathcal{Q} = \mathcal{F} \cup \{D \in \mathcal{D} : D \nrightarrow F \text{ for any } F \in \mathcal{F}\} \tag{3.2}$$

where $(\mathcal{F}, \mathcal{D})$ is a finite homomorphism duality.

Proof. If $(\mathcal{F}, \mathcal{D})$ is a finite duality, then for all $F \in \mathcal{F}$ and $D \in \mathcal{D}$ we have $F \nrightarrow D$, and so $\mathcal{Q} := \mathcal{F} \cup \{D \in \mathcal{D} : D \nrightarrow F \text{ for any } F \in \mathcal{F}\}$ is a finite antichain. Moreover, if $F \nrightarrow X$ for all $F \in \mathcal{F}$, then X is homomorphic to some element D of $\mathcal{D}$. Either $D \in \mathcal{Q}$ or $D \to F$ for some $F \in \mathcal{F}$; in any case, X is comparable with some element of $\mathcal{Q}$. Hence $\mathcal{Q}$ is a finite maximal antichain.

Conversely, suppose that $\mathcal{Q}$ is a finite maximal antichain in $\mathcal{C}(\Delta)$. Let $(\mathcal{F}, \mathcal{D})$ be the partition of $\mathcal{Q}$ defined in 3.4.2. One of the conditions of Lemma 3.4.4 is satisfied. If (1) is satisfied, then $(\mathcal{F}, \mathcal{D})$ is a splitting of $\mathcal{F}$, and so it is a finite duality in which no element of $\mathcal{D}$ is homomorphic to an element of $\mathcal{F}$ and $\mathcal{Q} = \mathcal{F} \cup \mathcal{D}$. It remains to examine the case that (2) is satisfied.

In the case (2), we first prove that all elements of $\mathcal{F}$ are forests. That implies that there exists a finite duality $(\mathcal{F}, \mathcal{D}')$; we then prove that this is the duality that satisfies (3.2).

Suppose that condition (2) of Lemma 3.4.4 is satisfied, $F \in \mathcal{F}$ and C is a component of F that is not a tree. By Lemma 3.2.1 there is no gap below C. Thus there exist infinitely many structures $X_1, X_2, \ldots$ such that $X_1 < X_2 < \cdots < C$. A simple case analysis reveals that the downset $D^{*\downarrow}$ generated by D^* contains no infinite increasing chain; this is only true for structures with at most two relations. Hence $C \nrightarrow D^*$.

Consequently if $Y \to D^*$ and $Y \to F$ for some structure Y, then there exists X such that $Y < X < F$ and $X \nrightarrow D^*$ (just like in Theorem 3.4.12). It follows that no structure Y exists such that $Y \to F$ but $Y \nrightarrow D$ for all $D \in \mathcal{D}$.

Now it is time to reuse the trick that served to prove Lemma 2.4.25. As there is no gap below C, we can find a structure B such that $B < C$, the structure B is homomorphic to exactly those components of structures in $\mathcal{Q}$ as C is, and exactly the same components of structures in $\mathcal{Q}$ are homomorphic to B as to C. Let Y be the structure constructed from F by replacing its component C with B. Clearly $Y \to F$, but Y is homomorphic to no $D \in \mathcal{D}$, because F is homomorphic to no $D \in \mathcal{D}$. This is a contradiction. Therefore all components of F are trees, and all elements of $\mathcal{F}$ are forests.

Invoke the transversal construction on $\mathcal{F}$ and get a finite duality $(\mathcal{F}, \mathcal{D}')$ (remember that $\mathcal{D}$ is defined by splitting the antichain $\mathcal{Q}$). We want to prove that $\mathcal{D}$ contains exactly the elements of $\mathcal{D}'$ that are not homomorphic to any element of $\mathcal{F}$.

First, let D' be an element of $\mathcal{D}'$ such that D' is homomorphic to no $F \in \mathcal{F}$. As $\mathcal{Q}$ is a maximal antichain and D' is incomparable to every element of $\mathcal{F}$, we have $D' \to D$ for some $D \in \mathcal{D}$. If $D \not\to D'$, then some element of $\mathcal{F}$ is homomorphic to D by duality, a contradiction with $\mathcal{Q}$ being an antichain. Hence $D \sim D'$, and so $D = D'$ since both D and D' are cores. Therefore $\mathcal{D}$ contains all elements of $\mathcal{D}'$ that are not homomorphic to anything in $\mathcal{F}$.

Finally, we show that $\mathcal{D}$ contains no other elements. Suppose that $D \in \mathcal{D}$. Because $\mathcal{Q}$ is an antichain, no $F \in \mathcal{F}$ is homomorphic to D. Thus by duality $D \to D'$ for some $D' \in \mathcal{D}'$. However, D' is homomorphic to no $F \in \mathcal{F}$ (otherwise D would also be), and so we know from the previous paragraph that $D' \in \mathcal{D}$ and consequently $D' \in \mathcal{Q}$. Once again using the fact that $\mathcal{Q}$ is an antichain we conclude that $D' = D$.

In this way, we have found a finite duality $(\mathcal{F}, \mathcal{D}')$ such that

$$\mathcal{Q} = \mathcal{F} \cup \{D \in \mathcal{D}' : D \not\to F \text{ for any } F \in \mathcal{F}\}. \qquad \square$$

The case of three or more relations ($|I| \geq 3$) is presently open. There may be a "quantum leap" here as indicated by the following result, which can be deduced from [19]. It implies that for more than two relations we cannot rely on the fact that the suborder induced by preimages of D^* is simple.

3.4.14 Proposition. *Let $\Delta = (2, 2, 2)$. Then the suborder of $\mathcal{C}(\Delta)$ induced by all structures homomorphic to D^* is a universal countable partial order; that is, any countable partial order is an induced suborder of this order.* $\square$

Cutting points

Finally, we show a connection of splitting antichains and cutting points.

3.4.15 Definition. Let P be a poset. An element $y \in P$ is called a **cutting point** if there are $x, z \in P$ such that $x \prec y \prec z$ and $[x, z] = [x, y] \cup [y, z]$. (The interval $[x, z] := \{y \in P : x \preccurlyeq y \preccurlyeq z\}$.)

The connection is that every finite maximal antichain without a cutting point splits. The following follows from [9, Theorem 2.10].

3.4.16 Theorem. *If S is a finite maximal antichain that does not contain a cutting point, then S splits.* □

Thus a characterisation of all cutting points may potentially provide another proof of the splitting property for finite maximal antichains. Some cutting points are actually connected to dualities.

3.4.17 Proposition. *Let T be a Δ-tree and let D be its dual. Then the Δ-structures $T + D$ and $T \times D$ are cutting points in the homomorphism order $\mathcal{C}(\Delta)$.*

Proof. Consider the interval $[\bot, T]$, which is equal to the downset generated by T. Suppose that X is a Δ-structure such that $X < T$. Then $X \to D$, because $T \not\to X$. Thus $X \to T \times D$. Hence the interval $[T \times D, T]$ contains only its endpoints, that is $[T \times D, T] = \{T \times D, T\}$. Moreover, $[\bot, T \times D] \cup [T \times D, T] = [\bot, T]$, so $T \times D$ is a cutting point.

Similarly, if $D < X$, then $T + D \to X$. Hence $[D, T + D] \cup [T + D, \top] = [D, \top]$ and so $T + D$ is a cutting point. □

However, at present the general problem remains open.

3.4.18 Problem. Characterise all cutting points in the homomorphism order.

4 Complexity

Out of intense complexities
intense simplicities emerge.

(Winston Churchill)

This chapter contains remarks and results on complexity issues; most of them are implied by the previous chapters. It leaves many questions unanswered, though.

First, we introduce a generalisation of the *Constraint Satisfaction Problem (CSP)*, which is a decision problem whether a homomorphism exists between two relational structures. We consider a parametrisation of CSP, where the target structure is fixed and the input is the domain. Our generalisation fixes a finite set $\mathcal{H}$ of structures as the parameter and we ask whether there exists a homomorphism from the input structure into some structure in the set $\mathcal{H}$. We observe that if the set $\mathcal{H}$ is the right-hand (dual) side of a finite homomorphism duality, the problem is solvable by a polynomial-time algorithm.

Next we examine the problem of deciding whether an input finite set of relational structures is the right-hand side of a finite homomorphism duality. The complexity of this problem with inputs restricted to sets containing a single structure has recently been determined by Larose, Loten and Tardif [24]. Using their result, we are able to prove that this problem is decidable.

Finally we consider the decision problem whether an input finite set of relational structures is a maximal antichain in the homomorphism order. Our characterisation of finite maximal antichains from Section 3.4 implies that this problem is decidable for structures with at most two relations. Moreover, we show that the problem is NP-hard. It is not known at present whether it belongs to the class NP.

In this chapter, by a **tractable problem** we mean a decision problem that can be solved by a deterministic Turing machine using a polynomial amount of computation time, that is a problem belonging to the class P.

4.1 Constraint satisfaction problem

First, we define the constraint satisfaction problem.

4.1.1 Definition. Let H be a fixed Δ-structure (called a **template**). The **constraint satisfaction problem** $\mathrm{CSP}(H)$ is the problem to decide for an input Δ-structure G whether there exists a homomorphism $G \to H$.

Several classical computational problems can be formulated as constraint satisfaction, as the following three examples show.

4.1.2 Example (k-colouring)**.** Recall from Example 1.3.5 that a homomorphism to the complete graph K_k is the same as a k-colouring of G. Therefore $\mathrm{CSP}(K_k)$ is nothing but k-colourability. This is well-known to be tractable if $k \leq 2$ and NP-complete if $k \geq 3$.

4.1.3 Example (3-SAT)**.** The widely known *3-SAT* or *3-satisfiability problem* takes as its input a propositional formula in conjunctive normal form such that each clause contains three literals; the question is whether the input formula is satisfiable, that is whether logical values can be assigned to its variables in a way that makes the formula true.

This problem is equivalent to $\mathrm{CSP}(H)$ for the following template H: let the type $\Delta = (3,3,3,3)$ and let $\underline{H} = \{0,1\}$. The Δ-structure H has four relations, namely R_0, R_1, R_2, and R_3. Let

$$\begin{aligned} R_0 &= \underline{H}^3 \setminus \{(0,0,0)\}, \\ R_1 &= \underline{H}^3 \setminus \{(1,0,0)\}, \\ R_2 &= \underline{H}^3 \setminus \{(1,1,0)\}, \text{ and} \\ R_3 &= \underline{H}^3 \setminus \{(1,1,1)\}. \end{aligned}$$

For an input formula ϕ, we construct the Δ-structure G_ϕ in such a way that $\underline{G_\phi}$ will be the set for all variables appearing in ϕ and for each clause of ϕ we add a triple to one of the relations: if there are exactly i negated literals in the clause, we add a triple to $R_i(G_\phi)$ consisting of the three variables appearing in the clause, with the negated variables first. For instance, for the clause $x_1 \wedge \neg x_2 \wedge x_3$ add the triple (x_2, x_1, x_3) to $R_1(G_\phi)$.

It is straightforward that an edge is preserved by a mapping f from $\underline{G_\phi}$ to $\underline{H}$ if and only if the corresponding clause is true in the assignment induced by f of logical values to variables. Therefore this assignment makes ϕ true if and only if the mapping f is a homomorphism from G_ϕ to H, and consequently ϕ is satisfiable if and only if G is homomorphic to H.

The next example is taken from Tsang's book [42]. It is often used for illustrating algorithms for solving CSP.

4.1.4 Example (N-queen problem)**.** Given any integer N, the problem is to position N queens on N distinct squares in an $N \times N$ chessboard, in such a way that no two queens should threaten each other. The rule is such that a queen can threaten any other pieces on the same row, column or diagonal.

This may be reformulated as the problem of assigning each of N variables (one for each row) a value from the set $\{1, 2, \ldots, N\}$, marking the column in which the queen is positioned. It is possible to define suitable template and relations that ensure that a homomorphism from the input to the template exists if and only if the input encodes a non-threatening position of queens on the chessboard.

Some of the examples are to an extent artificial. However, many common problems can be formulated as constraint satisfaction very naturally. These problems appear in numerous areas such as scheduling, planning, vehicle routing, networks, and bioinformatics. For more details see, for example, the book [40].

As the examples show, the complexity of $\mathrm{CSP}(H)$ depends on the template H. Considerable effort has recently gone into classifying the complexity of all templates. This complexity was determined for undirected graphs by Hell and Nešetřil [17]. However, already for directed graphs the problem is unsolved. Various results have led to the following conjecture.

4.1.5 Conjecture ([10])**.** *Let* H *be a finite relational structure. Then* $\mathrm{CSP}(H)$ *is either solvable in polynomial time or NP-complete.*

The following definition is motivated by finite dualities.

4.1.6 Definition. As an analogy to CSP, we define the **generalised constraint satisfaction problem** $\mathrm{GCSP}(\mathcal{H})$ to be the following decision problem: 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 $G \to H$.

The existence of a finite duality for a template set $\mathcal{H}$ ensures the existence of a polynomial-time algorithm for solving the particular generalised constraint satisfaction problem. This is a classical observation for the standard constraint satisfaction. We restate it here in view of the fact that the description of finite dualities is a principal result of the thesis.

4.1.7 Theorem. *If* $(\mathcal{F}, \mathcal{D})$ *is a finite homomorphism duality, then* $\mathrm{GCSP}(\mathcal{D})$ *is solvable by a polynomial-time algorithm.*

Proof. The key to the proof is to observe that the algorithm that checks for all possible mappings from $\underline{F}$ to $\underline{G}$ whether they are a homomorphism or not runs in time $O(|G|^{|F|})$, polynomial in the size of G.

Therefore for an input Δ-structure G, it is possible to check for all $F \in \mathcal{F}$ whether $F \to G$; if the response is negative for all F, then G is homomorphic to some H in $\mathcal{H}$, otherwise it is not. Clearly this testing can be done in time polynomial in the size of the input structure G. □

As in Conjecture 4.1.5, 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.

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

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

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

4.2 Deciding finite duality

We are interested in the following decision problem: For an input finite set $\mathcal{D}$ of Δ-structure, determine whether there exists a set $\mathcal{F}$ of Δ-structure such that $(\mathcal{F}, \mathcal{D})$ is a finite homomorphism duality.

The complexity of the problem was established by Larose, Loten and Tardif [24] in the special case where the input set $\mathcal{D}$ is a singleton, that is $|\mathcal{D}| = 1$.

4.2.1 Theorem ([24, Theorem 5.1])**.** *The problem of determining whether for a relational structure D there exists a finite set $\mathcal{F}$ of relational structures such that $(\mathcal{F}, \{D\})$ is a finite duality, is NP-complete.* □

This special case turns out to be essential. As a consequence we get that the general problem is decidable, as we show in the rest of this section.

Theorem 2.5.6 claims that in a duality pair $(\mathcal{F}, \{D\})$, the diameter of the elements of $\mathcal{F}$, which are cores by the definition of finite duality, is at most n^{n^2}, where $n = |\underline{D}|$. We would like to generate all core trees with small diameter. It has been established that their number is finite [24, Lemma 2.3]. By modifying the proof in the cited paper we get a rough recursive estimate for the number of such trees and the number of their edges.

4.2.2 Lemma. *Let $\Delta = (\delta_i : i \in I)$ be a type and let $s := |I|$ be the number of relations and $r := \max\{\delta_i : i \in I\}$ the maximum arity of a relation.*

Let $\mathcal{T}_d$ be the set of all core trees with a root such that the distance of any vertex from the root is at most d. Let $t_d := |\mathcal{T}_d|$ be the number of such trees and let m_d be the maximum number of edges of a tree in $\mathcal{T}_d$.

Then

$$t_0 = 1, \qquad m_0 = 0,$$
$$t_d \leq 2^{sr \cdot t_{d-1}^{r-1}}, \qquad m_d \leq sr \cdot t_{d-1}^{r-1} \cdot \big(1 + (r-1) \cdot t_{d-1} \cdot m_{d-1}\big).$$

Proof. There is exactly one rooted tree with all vertices in distance at most 0 from the root: the tree $\perp$ with one vertex and no edges. So $t_0 = 1$. The tree $\perp$ has no edges, hence $m_0 = 0$.

For a rooted tree with maximum distance at most d from the root v, we can encode every edge $(u_1, u_2, \ldots, u_{\delta_i})$ that contains the root by the name of the relation $i \in I$, the index k such that $v = u_k$ and the trees rooted at u_j for $u_j \neq u_k$. Such trees belong to $\mathcal{T}_{d-1}$, thus the number of possible labels for edges is at most $sr \cdot t_{d-1}^{r-1}$. In a core tree, all edges containing the root have pairwise distinct labels. Hence the number of such trees is $t_d \leq 2^{sr \cdot t_{d-1}^{r-1}}$.

In a tree with maximum distance at most d from the root v there are at most $sr \cdot t_{d-1}^{r-1}$ edges that contain the root. Hence there are at most $(r-1) \cdot sr \cdot t_{d-1}^{r-1}$ vertices other than the root v. In each of these vertices, no more than t_{d-1} trees from $\mathcal{T}_{d-1}$ are rooted. Therefore the number of edges in the tree is at most $sr \cdot t_{d-1}^{r-1} + (r-1) \cdot sr \cdot t_{d-1}^{r-1} \cdot t_{d-1} \cdot m_{d-1} = sr \cdot t_{d-1}^{r-1} \cdot \big(1 + (r-1) \cdot t_{d-1} \cdot m_{d-1}\big)$. □

As a consequence, in the duality pair $(\mathcal{F}, \{D\})$ the set $\mathcal{F}$ is computable from D.

4.2.3 Lemma. *There exists an algorithm that computes a set $\mathcal{F}$ of Δ-structures from an input Δ-structure D so that $(\mathcal{F}, \{D\})$ is a finite homomorphism duality, provided that such a set exists.*

Proof. According to Theorem 4.2.1 there is an algorithm that decides whether such a set $\mathcal{F}$ exists.

If it exists, by Theorem 2.5.6 the diameter of all elements of $\mathcal{F}$ is bounded by $d := n^{n^2}$, where $n = |\underline{D}|$. Thus there is a bound $m := m_d$ on the number of edges of the elements of $\mathcal{F}$ that is computable from Lemma 4.2.2. Let $\mathcal{F}'$ be the set of all core Δ-trees with at most m edges that are not homomorphic to D. Theorem 2.5.6 and Lemma 4.2.2 imply that for any Δ-structure X that is not homomorphic to D there exists a Δ-tree F with at most m edges such that $F \to X$ but $F \not\to D$. Hence

$$\mathcal{F}'^{\not\to} = \to D.$$

Therefore $\mathcal{F}$ is the set of all homomorphism-minimal elements of $\mathcal{F}'$.

It is fairly straightforward to design an algorithm for constructing all core trees with a bounded number of edges, as well as an algorithm for determining the homomorphism-minimal elements of a finite set of structures. □

We can conclude that the problem to determine whether an input set of structures is a right-hand side of a finite duality is decidable.

4.2.4 Theorem. *There exists an algorithm that determines whether for an input finite set $\mathcal{D}$ of Δ-structures there exists a finite set $\mathcal{F}$ of Δ-structures such that $(\mathcal{F}, \mathcal{D})$ is a finite homomorphism duality.*

Proof. The algorithm is as follows:

1. For each element D of $\mathcal{D}$, determine whether it is a core. If not, then no such duality can exist by definition.
2. For each element D of $\mathcal{D}$, determine whether it is a product of duals, that is whether there exists a finite set $\mathcal{F}$ such that $\big(\mathcal{F}, \{D\}\big)$ is a finite duality (see Theorem 4.2.1). If an element of $\mathcal{D}$ is not a product of duals, then $\mathcal{D}$ is not a right-hand side of a finite duality because of Theorem 2.4.26.
3. For each element D of $\mathcal{D}$, compute the set $\mathcal{F}(D)$ such that $\big(\mathcal{F}, \{D\}\big)$ is a finite duality (see Lemma 4.2.3).
4. Check whether $\{\mathcal{F}_D : D \in \mathcal{D}\}$ is the set of all transversals for some finite set $\mathcal{F}$ of core Δ-forests. This can be done greedily by considering all sets of core Δ-forests whose components appear in the sets $\mathcal{F}_D$, and by constructing the transversals (directly from Definition 2.4.12).

Theorem 2.4.26 implies that if the algorithm finds a set $\mathcal{F}$ in the last step, then $(\mathcal{F}, \mathcal{D})$ is a finite homomorphism duality; otherwise $\mathcal{D}$ is the right-hand side of no finite duality. □

4.3 Deciding maximal antichains

We consider the problem of deciding whether an input finite set of relational structures forms a finite maximal antichain in the homomorphism order. The problem is called the *MAC problem*; the letters MAC stand for "Maximal Antichain".

4.3.1 Definition. The **MAC problem** is to decide whether an input finite non-empty set $\mathcal{Q}$ of Δ-structures is a maximal antichain in the homomorphism order $\mathcal{C}(\Delta)$.

The characterisation of finite maximal antichains for types with at most two relations (Theorem 3.4.13) implies decidability of the MAC problem.

4.3.2 Theorem. *Let $\Delta = (\delta_i : i \in I)$ be a type such that $|I| \leq 2$. Then the MAC problem is decidable.*

Proof. The algorithm is as follows:

1. For each element of $\mathcal{Q}$, check whether its core is a forest. The core of a Δ-structure is computable (by checking the existence of a retraction to every substructure). Deciding whether a Δ-structure is a forest is possible even in polynomial time.
2. Let $\mathcal{F} \subseteq \mathcal{Q}$ be the set of all such structures. Find all transversals over $\mathcal{F}$. This can be done directly from Definition 2.4.12.
3. For each transversal $\mathcal{M}$, construct its dual $D(\mathcal{M})$. First use the bear construction (2.2.5) to construct the dual of each element of $\mathcal{M}$ and then take the product of all these duals.
4. Check whether $\mathcal{Q} \setminus \mathcal{F}$ is formed exactly by structures homomorphically equivalent to the duals of transversals constructed in the previous step.

Theorem 3.4.13 implies that the algorithm is correct. □

4.3.3 Theorem. *Let $\Delta = (\delta_i : i \in I)$ be a type such that $|I| \leq 2$. Then the MAC problem is NP-hard.*

Proof. We will use the fact that for any type Δ there exists a Δ-tree T such that $\mathrm{CSP}(T)$ is NP-complete. We construct the following reduction of $\mathrm{CSP}(T)$ to the MAC problem: For an input structure G of $\mathrm{CSP}(T)$, let $\mathcal{Q}(G) := \{G + T, D(T)\}$. The set $\mathcal{Q}(G)$ can be constructed from G in polynomial time. By Theorem 3.4.13, $\mathcal{Q}(G)$ is a finite maximal antichain if and only if $G \to T$. □

However, the algorithm given in the proof of Theorem 4.3.2 does not ensure that the MAC problem is in the class NP. This is not known at present.

4.3.4 Problem. Is the MAC problem in NP?

The hard part of the problem may actually consist in finding the cores of the involved structures. Also, in our proof of NP-hardness we actually reduce the decision whether the core of $G + T$ is T to the MAC problem. So it makes sense to ask whether the complexity of the MAC problem changes when inputs are restricted to cores.

4.3.5 Problem. What is the complexity of the MAC problem if input is restricted to sets of cores?

Bibliography

A classic is a book which
people praise and don't read.

(Mark Twain)

[1] J. Adámek, H. Herrlich, and G. E. Strecker. *Abstract and Concrete Categories. The Joy of Cats*. John Wiley and Sons, 1990. 16

[2] R. Ahlswede, P. L. Erdős, and N. Graham. A splitting property of maximal antichains. *Combinatorica*, 15(4):475–480, 1995. 72

[3] M. Barr and C. Wells. *Category Theory for Computing Science*. Les Publications CRM, Montréal, 3rd edition, 1999. 16

[4] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 2nd edition, 2002. 28

[5] J. Edmonds. Paths, trees, and flowers. *Canad. J. Math.*, 17:449–467, 1965. 31

[6] M. El-Zahar and N. Sauer. The chromatic number of the product of two 4-chromatic graphs is 4. *Combinatorica*, 5(2):121–126, 1985. 26

[7] P. Erdős. Graph theory and probability. *Canad. J. Math.*, 11:34–38, 1959. 74

[8] P. Erdős and A. Hajnal. On chromatic number of graphs and set-systems. *Acta Math. Hungar.*, 17(1–2):61–99, 1966. 74

[9] P. L. Erdős and L. Soukup. How to split antichains in infinite posets. *Combinatorica*, 27(2):147–161, 2007. 72, 81

[10] 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, 1998. 85

[11] J. Foniok, J. Nešetřil, and C. Tardif. Generalised dualities and finite maximal antichains. In F. V. Fomin, editor, *Graph-Theoretic Concepts in Computer Science (Proceedings of WG 2006)*, volume 4271 of *Lecture Notes in Comput. Sci.*, pages 27–36. Springer-Verlag, 2006. 4

[12] J. Foniok, J. Nešetřil, and C. Tardif. On finite maximal antichains in the homomorphism order. *Electron. Notes Discrete Math.*, 29:389–396, 2007. 4

[13] J. Foniok, J. Nešetřil, and C. Tardif. Generalised dualities and maximal finite antichains in the homomorphism order of relational structures. *European J. Combin.*, to appear. 4, 31

[14] R. Häggkvist, P. Hell, D. J. Miller, and V. Neumann Lara. On multiplicative graphs and the product conjecture. *Combinatorica*, 8(1):63–74, 1988. 25

[15] S. T. Hedetniemi. Homomorphisms of graphs and automata. University of Michigan Technical Report 03105-44-T, University of Michigan, 1966. 25

[16] Z. Hedrlín. On universal partly ordered sets and classes. *J. Algebra*, 11(4):503–509, 1969. 63, 65

[17] P. Hell and J. Nešetřil. On the complexity of H-coloring. *J. Combin. Theory Ser. B*, 48(1):92–119, 1992. 85

[18] P. Hell and J. Nešetřil. *Graphs and Homomorphisms*, volume 28 of *Oxford Lecture Series in Mathematics and Its Applications*. Oxford University Press, 2004. 3

[19] J. Hubička and J. Nešetřil. Finite paths are universal. *Order*, 22(1):21–40, 2005. 63, 65, 80

[20] J. Hubička and J. Nešetřil. Universal partial order represented by means of oriented trees and other simple graphs. *European J. Combin.*, 26(5):765–778, 2005. 63, 65

[21] D. E. Knuth and H. Zapf. Euler – A new typeface for mathematics. *Scholarly Publishing*, 20:131–157, 1989. 4

[22] P. Komárek. Some new good characterizations for directed graphs. *Časopis Pěst. Mat.*, 109(4):348–354, 1984. 31, 34

[23] P. Komárek. *Good characterisations in the class of oriented graphs*. PhD thesis, Czechoslovak Academy of Sciences, Prague, 1987. In Czech (Dobré charakteristiky ve třídě orientovaných grafů). 31, 36, 40, 41

[24] B. Larose, C. Loten, and C. Tardif. A characterisation of first-order constraint satisfaction problems. In *Proceedings of the 21st IEEE Symposium on Logic in Computer Science (LICS'06)*, pages 201–210. IEEE Computer Society, 2006. 62, 83, 86, 87

[25] L. Lovász. On chromatic number of finite set-systems. *Acta Math. Hungar.*, 19(1–2):59–67, 1968. 74

[26] J. Nešetřil. *Theory of Graphs*. SNTL, Prague, 1979. In Czech (Teorie grafů). 31

[27] J. Nešetřil. The coloring poset and its on-line universality. KAM-DIMATIA Series 2000-458, Charles University, Prague, 2000. 65

[28] J. Nešetřil. Combinatorics of mappings. KAM-DIMATIA Series 2000-472, Charles University, Prague, 2000. 3

[29] J. Nešetřil and P. Ossona de Mendez. Grad and classes with bounded expansion III. Restricted dualities. KAM-DIMATIA Series 2005-741, Charles University, Prague, 2005. 32

[30] J. Nešetřil and A. Pultr. On classes of relations and graphs determined by subobjects and factorobjects. *Discrete Math.*, 22(3):287–300, 1978. 25, 31, 34

[31] J. Nešetřil, A. Pultr, and C. Tardif. Gaps and dualities in Heyting categories. *Comment. Math. Univ. Carolin.*, 48(1):9–23, 2007. 32, 63, 66, 68

[32] J. Nešetřil and V. Rödl. A short proof of the existence of highly chromatic hypergraphs without short cycles. *J. Combin. Theory Ser. B*, 27(2):225–227, 1979. 74

[33] J. Nešetřil and I. Švejdarová. Diameters of duals are linear. KAM-DIMATIA Series 2005-729, Charles University, Prague, 2005. 48

[34] 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. 31, 36, 51, 63, 65, 66

[35] J. Nešetřil and C. Tardif. On maximal finite antichains in the homomorphism order of directed graphs. *Discuss. Math. Graph Theory*, 23(2):325–332, 2003. 71

[36] 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. 40, 41, 44, 61

[37] J. Nešetřil and C. Tardif. A dualistic approach to bounding the chromatic number of a graph. *European J. Combin.*, to appear. 40

[38] J. Nešetřil and X. Zhu. Path homomorphisms. *Math. Proc. Cambridge Philos. Soc.*, 120:207–220, 1996. 28

[39] A. Pultr and V. Trnková. *Combinatorial, Algebraic and Topological Representations of Groups, Semigroups and Categories*, volume 22 of *North-Holland Mathematical Library*. North-Holland, Amsterdam, 1980. 65

[40] F. Rossi, P. van Beek, and T. Walsh, editors. *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*. Elsevier, 2006. 85

[41] C. Tardif. Multiplicative graphs and semi-lattice endomorphisms in the category of graphs. *J. Combin. Theory Ser. B*, 95(2):338–345, 2005. 26

[42] E. Tsang. *Foundations of Constraint Satisfaction*. Academic Press, 1993. 85

[43] E. Welzl. Color families are dense. *Theoret. Comput. Sci.*, 17(1):29–41, 1982. 63