Embedding dualities for set partitions and for relational structures*

Vít Jelínek† Martin Klazar‡

September 3, 2010

Abstract

We show that for a set F of forbidden set partitions and an integer k there is a finite collection D of partitions of ordinals, such that any finite partition with at most k blocks avoids all the elements of F if and only if it is contained in at least one element of D. Using this result, we reprove rationality of the generating function enumerating a hereditary class of set partitions with a bounded number of blocks. We show that this result does not extend to partitions with an unbounded number of blocks.

We also consider hereditary classes of relational structures. We give a characterization of those classes that can be expressed as classes of finite substructures of a finite collection of (possibly infinite) relational structures.

1 Introduction

In this paper, we investigate the concept of *dualities* among relational structures. For a family F of finite 'forbidden' structures, we say that a family D of (possibly infinite) structures is *dual* to F, if for every finite structure s it holds that s does not contain any structure $f \in F$ as substructure if and only if s is a substructure of some $d \in D$. The pair (F, D) is then known as a *duality pair*. This terminology is motivated by an analogous concept of homomorphism dualities. For an overview of the topic of homomorphism dualities see, e.g., Hell and Nešetřil [13].

With a duality pair (F, D) we may naturally associate a hereditary class C of structures, where C is the set of all the finite structures that do not contain any forbidden substructure $f \in F$, or equivalently, C is the set of all finite substructures of elements of D. Conversely, it is easy to see any hereditary class

*Supported by project MSM0021620838 and by grant 1M0545 of the Czech Ministry of Education.

†Department of Applied Mathematics, Charles University in Prague, and The Mathematics Institute, School of Computer Science, Reykjavík University, jelinek@kam.mff.cuni.cz

‡Department of Applied Mathematics, Charles University in Prague and The Institute for Theoretical Computer Science, Charles University in Prague, klazar@kam.mff.cuni.cz

C arises in this way from a (not necessarily unique) duality pair (F, D). It is natural to ask whether a given class C has a 'nice' description in terms of F or D, e.g., whether there is a description in which F or D is a finite set.

In this paper, our main goal is to study hereditary classes C of relational structures that can be described by a duality pair (F, D) in which D is a finite set of (possibly infinite) relational structures. In such situation, we call the set D a *finite dual* of C.

This paper has two main parts. In the first part, we consider set partitions whose ground set is an ordinal number. We show that in this setting, every hereditary class of set partitions with a bounded number of blocks can be characterized by a finite dual. The proof of this result is based on the Myhill–Nerode characterization of regular languages. Our technique also allows us to re-derive a known enumerative result on the growth of hereditary classes of set partitions. We also prove a negative result, by presenting a class of partitions defined by a single forbidden pattern that has no finite dual.

In the second part of this paper, we turn to general relational structures. The main result of this part is a general criterion characterizing hereditary classes of relational structures admitting finite duals.

Let us now introduce necessary definitions and notation in order to state our results precisely. The proofs are relegated to the remaining two sections.

By a *partition* we mean a set partition, i.e., a set π of nonempty and disjoint sets, called *blocks*, whose union is the *ground set* $X = \bigcup \pi$ of the partition π; we say that π is a *partition of* X. A *normalized* partition has ground set $[n] = \{1, 2, \ldots, n\}$ for some $n \geq 0$ (we define $[0] = \emptyset$). The *size* of a partition is the size of its ground set. In particular, a *finite* partition is a partition of a finite set. We will consider both finite and infinite partitions.

We say that a partition π of $[m]$ is *contained* in a partition σ of $[n]$ if there is an increasing injection $f\colon [m] \to [n]$ with the property that i and j are in the same block of π if and only if $f(i)$ and $f(j)$ are in the same block of σ. This relation will be denoted by $\pi \subseteq \sigma$; it is a partial ordering on the set of all normalized partitions. If $\pi \nsubseteq \sigma$, we say that σ *avoids* π. The notion of containment can be naturally extended to partitions of an arbitrary linearly ordered ground set, including infinite partitions.

A *hereditary class* of partitions is a (typically infinite) set P of finite normalized partitions such that for any normalized partition σ of $[n]$ in P, every partition π of $[m]$ contained in σ lies in P as well. Hereditary classes of partitions have been previously studied, e.g., by Klazar [18] or by Balogh, Bollobás and Morris [4]. Several authors have also proposed different ways to define the containment relation of partitions, and obtained enumerative results in these settings, see for instance de Mier [8], Goyt [12], Jelínek and Mansour [16], or Sagan [22].

For a hereditary class P of partitions, we define the *basis* of P to be the set $F = F_P$ of minimal forbidden structures of P; formally,

$$F_P = \{\pi \mid \pi \notin P, \pi \text{ is normalized, and every proper subpartition of } \pi \text{ is in } P\}.$$

If F_P is finite, then P is called *finitely based*, if F_P is a singleton set, then P is called *principal*.

The *growth function* (or *speed*) of a hereditary class P is the function $n \mapsto |P_n|$ where P_n is the set of partitions of $[n]$ in P and $|X|$ denotes the cardinality of a set X. An important goal in the study of hereditary classes of partitions, as well as hereditary classes of other structures, is the characterization of possible growth functions of these classes. Recently, several authors have obtained strong results related to growth functions of classes of graphs [5, 6], ordered graphs [4], or permutations [1, 17, 24, 25]. For more details, we refer the reader to the survey of Bollobás [7].

To characterize the growth function, we need a handle on $\subseteq$. The description of a hereditary class P in terms of its basis F does not seem to offer enough insight into the growth rate of the class. A different approach is to characterize a hereditary class in terms of the 'dual' object, i.e., to describe a class P as the class of finite subpartitions of a small number of (typically infinite) partitions from a set D. This is similar to the approach of Atkinson, Murphy and Ruškuc [3] and Huczynska and Ruškuc [15] to permutation classes (which they represent by bijections between infinite linearly ordered sets) and is also motivated by the theory of homomorphism dualities (Hell and Nešetřil [13]). In the first part of our paper, we focus on the situation when the elements of the dual set D are partitions of ordinal numbers. We say that a class P of partitions has *finite ordinal dual* if there is a finite set D of partitions of ordinals such that the elements of P are exactly the finite partitions contained in at least one member of D. The set D is not necessarily unique.

As a simple example of a class with a finite ordinal dual, consider the class P consisting of the (normalized) partitions with at most one block. Then the basis contains a single partition $\pi = \{\{1\}, \{2\}\}$ and the dual also contains a single (infinite) partition $\sigma = \{\omega\}$, where $\omega = \{0, 1, 2, \dots\}$ denotes the smallest infinite ordinal, i.e., the set of nonnegative integers.

Another example is for $F_P = \{\pi\}$ where $\pi = \{\{1, 3\}, \{2\}\}$. Then P consists of (normalized) partitions whose blocks are intervals and a dual partition is $\sigma = \{\{0\}, \{1, 2\}, \{3, 4, 5\}, \dots\}$ with ground set ω (or any other interval partition of ω with unbounded lengths of intervals); we denote this partition as σ_I. In general, even for a principal class P, one needs more than one dual partition, as well as ordinals larger than ω for ground sets. As an example, consider the hereditary class P with $F_P = \{\pi\}$ where $\pi = \{\{1, 3\}, \{2\}, \{4\}\}$. We leave for the interested reader to verify that P admits as a dual a pair of partitions, the partition of 2ω

$$\sigma_1 = \sigma_I \cup \{\{\omega, \omega + 2, \omega + 4, \dots\}, \{\omega + 1, \omega + 3, \omega + 5, \dots\}\}$$

and the partition of 4ω

$$\sigma_2 = \sigma_I \cup \{\{\omega, \omega + 1, \omega + 2, \dots; 3\omega, 3\omega + 1, 3\omega + 2, \dots\}\} \cup (2\omega + \sigma_I).$$

Unfortunately, there are principal hereditary classes P which have no finite ordinal dual (see Theorem 1.3). Nevertheless, finite ordinal duals do exist for many hereditary classes, as shown by our first main result:

Theorem 1.1. *Let P be a hereditary class of partitions such that the number of blocks in $\pi \in P$ is bounded. Then the basis F_P is finite and P has a finite ordinal dual D. In fact, the elements of D may be chosen as partitions of ordinals smaller than $m\omega$ for some integer $m > 0$.*

From Theorem 1.1, or more precisely from its proof, we obtain the following corollary, which was proved by Klazar [18, Theorem 3.1] in a different way. The inductive proof in [18] and the present proof both represent partitions by words. The present proof is more conceptual, as it derives the corollary from the Myhill–Nerode theorem on regular languages.

Corollary 1.2. *Let P be a hereditary class of partitions such that the number of blocks in $\pi \in P$ is bounded. Then the generating series of the growth function is rational, in fact*

$$f_P(x) = \sum_{n \geq 0} |P_n| x^n = \frac{p(x)}{\prod_{i=1}^{k}(1 - ix)^{a_i}}$$

where k is the maximum number of blocks in some $\pi \in P$, $p(x)$ is an integral polynomial and $a_i > 0$ are integers.

As we already mentioned, there are hereditary (even principal) classes of partitions admitting no finite ordinal dual. Specifically, we will prove the following result.

Theorem 1.3. *Let $\pi_0 = \{\{1,3\},\{2,4\},\{5\}\}$. For every finite family D of partitions of ordinals avoiding π_0 there is a finite normalized partition avoiding π_0 that is not contained in any member of D. Thus, the hereditary class of partitions P with the basis $F = \{\pi_0\}$ has no finite ordinal dual.*

The two previous theorems and corollary will be proved in Section 2.

In the second part of our paper, we consider more general duals, in which instead of partitions of ordinals, we allow infinite partitions of an arbitrary linearly ordered ground set. To handle this setting, it is convenient to represent partitions as relational structures, and to express partition containment as a special case of (induced) containment of relational structures. In fact, our results from the second part of the paper generalize to hereditary classes of relational structures of fixed (finite) signature.

Let us introduce the relevant definitions. A *relational structure* with *signature* $\delta = (\delta_1, \delta_2, \ldots, \delta_k)$ on the *ground set* V is a $(k+1)$-tuple $(V, R_1, R_2, \ldots, R_k)$, where R_i is a relation of arity δ_i on the ground set V, i.e., R_i is a set of ordered δ_i-tuples of (not necessarily distinct) elements of V. For example, relational structures of signature (2) correspond precisely to directed graphs, with loops, in which multiple edges are only allowed if they have opposite orientations. As with set partitions, the *size* of a relational structure is the size of its ground set.

If $\rho = (V, R_1, \ldots, R_k)$ and $\sigma = (W, S_1, \ldots, S_k)$ are two relational structures of the same signature δ, we say that σ *contains* ρ (or ρ is a *substructure* of σ, or $\rho \subseteq \sigma$), if there is an injective function $f\colon V \to W$ with the property that for any i and j and for any j-tuple $(v_1, \dots, v_j)$ of elements of V, this j-tuple belongs to R_i if and only if the j-tuple $(f(v_1), \dots, f(v_j))$ belongs to S_i. This notion of containment is a natural extension of the familiar notion of induced containment of graphs.

We call a relational structure *normalized* if its ground set is the set $[n]$ for some $n \geq 0$. Let us stress that the elements of $[n]$ here merely play a role of labels, and their standard linear order is not relevant for the notion of containment of relational structures. A *hereditary class of relational structures* is a set C of finite normalized relational structures, all of them sharing the same signature, with the property that for every structure $\pi \in C$, all the normalized substructures of π belong to C.

A set partition π of a linearly ordered set $(V, \triangleleft)$ may be represented as a relational structure (V, R_1, R_2) of signature $(2, 2)$, as follows: for a pair of vertices $i, j \in V$, we define $(i, j) \in R_1$ if $i \triangleleft j$, and we define $(i, j) \in R_2$ if i and j are in the same block of π. With this representation, the containment relation of set partitions corresponds precisely to the containment relation of relational structures.

Notice that any substructure of a relational structure representing a set partition is itself a representation of a set partition. Thus, a hereditary class P of partitions is naturally represented by a hereditary class C_P of relational structures of signature $(2, 2)$, where C_P is the class of all the normalized relational structures representing elements of P. Note, however, that a single set partition in P is typically represented by several (isomorphic) normalized relational structures in C_P. Thus, the growth rate of the class P is not the same function as the growth rate of C_P. This minor technical issue does not worry us, since we deal with structural results, rather than explicit enumerations.

Let us now focus on general hereditary classes of relational structures, without restricting ourselves to classes arising from set partitions. To introduce dualities in this setting, it is useful to start with the following fact, due to Fraïssé [10, 11].

Fact 1.4. *Let C be a hereditary class of relational structures of a finite signature δ. The following three conditions are equivalent:*

- *There is a (possibly infinite) relational structure Π of signature δ such that C is the set of all the finite normalized substructures of Π.*
- *C cannot be expressed as a union of two of its proper hereditary subclasses.*
- *For any two structures $\pi, \sigma \in C$, there is a structure $\tau \in C$ that contains both π and σ. (This condition is known as the* joint embedding property*.)*

A hereditary class satisfying the conditions of Fact 1.4 is called *atomic*. Recently, the concept of atomicity received attention in the study of hereditary permutation classes (see, e.g., Atkinson, Murphy and Ruškuc [2, 3], Murphy [20], or Vatter [24]). Besides that, a considerable amount of work has been devoted to the study of growth rates of atomic classes of general relational structures (a survey of this field has been presented by Pouzet [21]).

Let us stress that all relational structures considered in our paper have finite signature. Fact 1.4 does not directly generalize to classes of structures with infinite signature (see [9]).

We are mostly interested in the hereditary classes of relational structures that can be expressed as a union of finitely many atomic classes. We will call such classes *molecular*. In view of Fact 1.4, molecular classes are precisely those classes that can be described by a duality pair (F, D), where D contains finitely many (possibly infinite) relational structures.

As the second main result of this paper, we provide a characterization of molecular classes of relational structures in terms of a joint embeddability of their elements. The result may in fact be stated in the (more general) language of partially ordered sets. For this, we need a few definitions. Let $(A, \leq)$ be a (typically infinite) poset. We say that two elements $x, y \in A$ are *compatible* in $(A, \leq)$, if they have a common upper bound $z \in A$ (i.e., if there is an element $z \in A$ satisfying $x \leq z$ and $y \leq z$). A set $I \subseteq A$ is an *ideal* if I is a down-set (i.e., $x \in I$ and $y \leq x$ implies $y \in I$), and I is up-directed (i.e., for every $x, y \in I$, the set I contains a common upper bound for x and y). We say that a poset $(A, \leq)$ is *covered* by ideals $I_1, \ldots, I_k$ if $A = I_1 \cup \cdots \cup I_k$.

Theorem 1.5. *Let $(A, \leq)$ be a (possibly infinite) poset and $k > 0$ be an integer. The poset $(A, \leq)$ can be covered by k ideals if and only if A does not contain $k+1$ pairwise incompatible elements. Moreover, if A cannot be covered by finitely many ideals, then it contains an infinite set of pairwise incompatible elements.*

The theorem can be easily applied to the containment order of (the isomorphism types of) relational structures in a given hereditary class. Let C be a hereditary class of relational structures. We say that two relational structures π and σ are *compatible in C* (or *C-compatible*), if there is a structure $\tau \in C$ that contains both π and σ; otherwise, π and σ are *C-incompatible*. A set of structures is *pairwise C-incompatible* if each two of its elements are C-incompatible. We now state the promised characterization of molecular classes.

Theorem 1.6. *Let C be a hereditary class of relational structures and let $k > 0$ be an integer. The class C can be expressed as a union of at most k atomic classes if and only if C does not have a pairwise C-incompatible subset of size $k+1$. Furthermore, C is molecular if and only if it has no infinite pairwise C-incompatible subset.*

Observe that the theorem implies that a class containing arbitrarily large finite pairwise incompatible subsets must also contain an infinite pairwise incompatible subset.

In Section 3, we prove Theorems 1.5 and 1.6. We then also point out that Theorem 1.6 implies, as a corollary, a result essentially due to Atkinson, Murphy and Ruškuc [2, Theorem 2.2] (see also [20, Proposition 188]), which states that every hereditary class that is partially well-quasi-ordered by inclusion is molecular.

Lastly, we provide an example of a finitely based class that is not molecular.

Let us now present the proofs of our results.

2 Finite ordinal duals

A partition σ of $[n]$ may be conveniently represented by a word $w = w_1 w_2 \cdots w_n$ with the property that $w_i = w_j$ if and only if i and j belong to the same block of σ. We say that two words w, w' are *isomorphic* (denoted by $w \sim w'$) if they represent the same partition. Thus, the word $AABAC$ is isomorphic to $XXYXA$, with both words representing the partition of [5] into the three blocks $\{1, 2, 4\}$, $\{3\}$ and $\{5\}$.

For two words w and w', we say that w is a *subword* of w' (denoted by $w \subseteq w'$) if w is obtained from w' by erasing some of its letters. We say that w is a *factor* of w' (denoted by $w \leq w'$) if w is a contiguous subword of w' (i.e. w' can be written as a concatenation of the form $w' = xwy$ for some words x, y).

Let π and σ be two normalized partitions, let p and s be two words representing π and σ, respectively. We have $\pi \subseteq \sigma$ if and only if there is a word p' isomorphic to p such that $p' \subseteq s$.

Let $\mathscr{A}$ be an alphabet. $\mathscr{A}^*$ denotes the set of all finite words over $\mathscr{A}$. We say that a language $\mathfrak{L} \subseteq \mathscr{A}^*$ is *hereditary*, if it is closed under taking subwords, we say that $\mathfrak{L}$ is *isomorphism-closed (with respect to the alphabet $\mathscr{A}$)* if for every two isomorphic words $w, w' \in \mathscr{A}^*$ we have $w \in \mathfrak{L} \iff w' \in \mathfrak{L}$.

We will also deal with infinite words over a given alphabet $\mathscr{A}$. The length of such a word is a countable ordinal α, and the word itself may be formally represented as a mapping from α to $\mathscr{A}$. The notions defined above extend to infinite words in the obvious way. We let ω denote the smallest infinite ordinal.

For a (typically infinite) word w, let $F(w)$ denote the language of all the finite factors of w; similarly, $S(w)$ denotes the language of all the finite subwords of w.

We say that a word u is *a universal word for the alphabet $\mathscr{A}$* if $F(u) = S(u) = \mathscr{A}^*$. The empty word is the only universal word of the empty alphabet. It is not hard to see that for any nonempty at most countable alphabet $\mathscr{A}$ there is a universal word of length ω.

Our proof of Theorem 1.1 is based on the following result.

Theorem 2.1. *Let $\mathscr{A}$ be a k-element alphabet and $\mathfrak{L}$ be a hereditary language over $\mathscr{A}$. Then there is a finite set F of words such that*

$$\mathfrak{L} = \{w \in \mathscr{A}^* \mid \text{ no subword of } w \text{ belongs to } F\}$$

and there are integers $m, n > 0$ and words $w_1, \ldots, w_n$, each of them of length at most $m\omega$, such that

$$\mathfrak{L} = \bigcup_{i=1}^{n} S(w_i).$$

The (typically infinite) words w_i have form

$$w_i = u_0 X_1 u_1 X_2 \cdots X_q u_q,$$

where each u_j is a universal word for a possibly empty alphabet $\mathscr{A}_j \subseteq \mathscr{A}$ and $X_1, \ldots, X_q$ are symbols of $\mathscr{A}$. The values of $\mathscr{A}_j$, X_j and q depend on i.

We first show how this theorem implies Theorem 1.1.

Proof of Theorem 1.1. Let P be a hereditary class of partitions such that every partition in P has at most k blocks. Consider the language $\mathfrak{L}$ over a k-element alphabet, where $\mathfrak{L}$ contains all the words that represent the partitions in P. Clearly, $\mathfrak{L}$ is isomorphism closed and hereditary. By Theorem 2.1, $\mathfrak{L}$ can be characterized by a finite list F of forbidden subwords. Let G be the set of normalized partitions represented by the words in F. It is easy to see that G is the basis of P.

The second claim of follows equally easily from Theorem 2.1. We know that $\mathfrak{L}$ can be characterized as the language of all the finite subwords of a finite collection of words $w_1, \ldots, w_n$. Each such word w_i represents a partition σ_i of an ordinal not exceeding $m\omega$ for some integer $m > 0$. Let us check that the normalized partitions contained in σ_i are precisely the partitions in P. If $\rho \in P$ then $\mathfrak{L}$ contains a word w representing ρ, thus w is a subword of one of the words w_i and hence $\rho \subseteq \sigma_i$. Conversely, if a normalized partition ρ is contained in σ_i then the corresponding subword w of w_i represents ρ, thus $w \in \mathfrak{L}$ which implies $\rho \in P$. □

2.1 Proof of Theorem 2.1

A particular case of Higman's theorem from the wqo theory (see, e.g., Kruskal [19]) says that in an infinite set of words over a finite alphabet, two words must be comparable by subword relation. Hence each hereditary language $\mathfrak{L}$ over a finite alphabet $\mathscr{A}$ is characterized by a finite set F of forbidden subwords, namely the subword-minimal words over $\mathscr{A}$ not in $\mathfrak{L}$. This proves the first part of Theorem 2.1.

To construct the "dual" words w_i and thus to prove the second part we use the theory of regular languages. We introduce the necessary concepts and state certain classical results without proof.

A language $\mathfrak{L}$ over a k-element alphabet $\mathscr{A}$ is called *regular* if it is accepted by a finite automaton. Such an automaton G is a finite directed (multi-)graph with one initial vertex and a set of accepting vertices; it has the property that each vertex has exactly k outgoing edges, each of them labeled by a distinct symbol of $\mathscr{A}$. Every word $w \in \mathscr{A}^*$ corresponds to the walk in G which starts in the initial vertex and in its i-th step follows the edge labeled by the i-th letter of w. The word w is accepted, if and only if the corresponding walk in G ends in an accepting vertex. The vertices of an automaton are commonly called *states.*

For every language $\mathfrak{L} \subseteq \mathscr{A}^*$, we may define the equivalence relation $R_{\mathfrak{L}}$ on the set of words $\mathscr{A}^*$ as follows: $(x, y) \in R_{\mathfrak{L}}$ if for every word $w \in \mathscr{A}^*$ we have the equivalence $xw \in \mathfrak{L} \iff yw \in \mathfrak{L}$. For the following classical result see, for example, Hopcroft and Ullman [14].

Theorem 2.2 (Myhill–Nerode)**.** *A language $\mathfrak{L}$ is regular if and only if $R_{\mathfrak{L}}$ has finitely many equivalence classes. Furthermore, a regular language $\mathfrak{L}$ has a unique (up to isomorphism) accepting automaton G with the smallest number of states, and this automaton has the property that its number of states is equal* *to the number of equivalence classes of $R_{\mathfrak{L}}$, and two words belong to the same equivalence class of $R_{\mathfrak{L}}$ if and only if the two computations of G over these two words end in the same state.*

A language $\mathfrak{L}$ described by finitely many forbidden subwords is regular, because a language avoiding a single fixed subword is easily seen to be regular and regular languages are closed under finite intersections. The structure of the minimal automaton accepting a hereditary language is described by the following lemma.

Lemma 2.3. *Let A be the minimal automaton accepting a hereditary language $\mathfrak{L}$. The following holds:*

- *A has at most one non-accepting state. All the outgoing edges of this state are loops.*
- *Apart from possible loops, A has no other directed cycles.*

Proof. Both claims follow from the Myhill–Nerode theorem. For the first claim, assume that x and y are two words not belonging to $\mathfrak{L}$. Clearly, for any word w, neither xw nor yw can belong to $\mathfrak{L}$, which shows that $(x, y) \in R_{\mathfrak{L}}$, i.e., the computations over x and y both lead to the same non-accepting state. By the hereditary property, there can be no edge from a non-accepting state to an accepting state, implying that all the outgoing edges of the non-accepting state are loops.

Next, we will show that A has no cycles, apart from loops. By the previous claim, we know that any non-trivial cycle can only contain accepting states. Assume that A has a directed cycle C containing two distinct states a and b. Let w be a word corresponding to a path from the initial state to the state a, let x be the word corresponding to the path from a to b along the cycle C, let y be the word corresponding to the path from b to a along C. The computations of A over the word w and over the word wxy both end in the state a, while the computation over wx ends in b. We will now show that w and wx are $R_{\mathfrak{L}}$-equivalent. Let $z \in \mathscr{A}^*$ be any word. We need to show the equivalence $wz \in \mathfrak{L} \iff wxz \in \mathfrak{L}$. Clearly, if $wxz \in \mathfrak{L}$ then $wz \in \mathfrak{L}$ since $\mathfrak{L}$ is hereditary. On the other hand, if $wz \in \mathfrak{L}$, then $wxyz \in \mathfrak{L}$, since w and wxy are $R_{\mathfrak{L}}$-equivalent (their computations end in state a), and so $wxz \in \mathfrak{L}$ by hereditarity. Thus w and wx are $R_{\mathfrak{L}}$-equivalent, which is a contradiction, since their computations end in different states a and b. □

We are now ready to proceed with the proof of Theorem 2.1. Let $A = (V, E)$ be the minimal automaton accepting a hereditary language $\mathfrak{L}$. If A has no non-accepting state, then $\mathfrak{L} = \mathscr{A}^*$ and $\mathfrak{L}$ can be expressed as the set of subwords of a universal word for the alphabet $\mathscr{A}$. Thus we have $n = 1$ and $w_1 = u_0$ with $\mathscr{A}_0 = \mathscr{A}$.

Assume now, that $\mathfrak{L} \neq \mathscr{A}^*$, and let f be the unique non-accepting state of A. Let $A' = (V', E')$ denote the directed graph obtained from A by the removal of the vertex f and all the edges incident to it. If f is the initial vertex of A, then $\mathcal{L}$ is the empty language, which can be expressed as the set of subwords of an empty collection of infinite words, satisfying Theorem 2.1 trivially. Let us assume that $\mathcal{L}$ is nonempty, and let a_0 denote the initial vertex of A. Lemma 2.3 implies that $a_0 \neq f$, i.e., $a_0 \in A'$.

We know that the words $w \in \mathcal{L}$ correspond bijectively to finite directed walks in A' starting in a_0. Since A' has no nontrivial cycles, the non-loop edges of any walk t in A' form a path; we will call this path *the backbone of t*. Since A' is finite, it has only finitely many maximal paths $p_1, p_2, \dots, p_k$ starting in a_0. For each such path p_i, we now construct an infinite word w_i of the form described in Theorem 2.1, with the property that for any walk t in A' whose backbone is an initial segment of the path p_i, the word w_i contains a subword w corresponding to the walk t. The word w_i is defined as follows. Let $a_0, a_1, \dots, a_m$ be the sequence of vertices of p_i, let e_j be the edge of p_i between a_{j-1} and a_j. Let $\mathscr{A}_j$ be the set of symbols that appear as labels of the loops at the vertex a_j, let u_j be a universal word of the alphabet $\mathscr{A}_j$ (with the empty word being universal for the empty alphabet), let X_j be the label of the edge e_j. Now we put $w_i = u_0 X_1 u_1 X_2 \cdots X_m u_m$. It is easy to check that any walk whose backbone is an initial segment of p_i corresponds to a subword of w_i (however, note that the converse is not necessarily true). This shows that any word of $\mathcal{L}$ is a subword of at least one of the words w_i, for $i = 1, \dots, k$.

To complete the proof of Theorem 2.1, we need to establish the following claim.

Lemma 2.4. *Any finite subword of any of the words w_i belongs to $\mathcal{L}$.*

Proof. Fix a word w_i, constructed from a path p_i as above. Let w be a finite subword of w_i. We will show that w is a subword of a finite word w' such that w' corresponds to a walk in A' with backbone p_i. This will imply that w' belongs to $\mathcal{L}$, and by the hereditarity of $\mathcal{L}$, we will know that w belongs to $\mathcal{L}$ as well.

Since $w \subseteq w_i = u_0 X_1 u_1 X_2 \cdots X_m u_m$, we may decompose w into a concatenation of the form $w = y_0 X_1^? y_1 X_2^? \cdots X_m^? y_m$, where each y_j is a subword of u_j, and $X_j^?$ is either the empty string or the symbol X_j. We now define $w' = y_0 X_1 y_1 X_2 \cdots X_m y_m$ and easily observe that w' has the required property. □

This completes the proof of Theorem 2.1.

2.2 Proof of Corollary 1.2

For the proof of Corollary 1.2, we use the results of the previous subsection. Let P be a hereditary class of partitions with at most k blocks and $\mathscr{A}$ be the alphabet $[k] = \{1, 2, \dots, k\}$. Let $\mathcal{L}$ be the language of all the words over $\mathscr{A}$ representing a partition in P.

A *rooted graph* is a directed graph with a special root vertex v_0 and with all edges labelled by letters of $\mathscr{A}$. We assume that no two edges leaving the same vertex have equal labels. A *rooted walk* in a rooted graph is a directed walk that starts in v_0. Each rooted walk represents a unique word formed by the labels of its edges in the order in which they were visited by the walk. Conversely, for each word, there is at most one rooted walk representing this word.

By the results of the previous subsection, we know that there is a rooted graph A' whose rooted walks represent precisely the words of $\mathcal{L}$. The problem is that a single partition of P is typically represented by several isomorphic words of $\mathcal{L}$, so we cannot use the enumeration of directed walks of A' directly.

Let us say that a word w over the alphabet $\mathscr{A}$ is *canonical* if w has the property that for each two letters $i, j \in \mathscr{A}$ such that $i < j$, if w contains the letter j, then w contains i as well, and furthermore, the first occurrence of i precedes the first occurrence of j. Note that each isomorphism class of words over $\mathscr{A}$ has exactly one canonical word. Let $\mathcal{L}_c$ be the set of all the canonical words in $\mathcal{L}$. There is a natural length-preserving bijection between P and $\mathcal{L}_c$. Our aim is to build a rooted graph A_c whose rooted walks represent the elements of $\mathcal{L}_c$. Furthermore, the graph A_c will contain no directed cycle except for possible loops. This will imply that its rooted walks are enumerated by generating function in the stated form.

Since the language $\mathcal{L}_c$ is not hereditary, we cannot use Lemma 2.3 directly. Instead, we apply a standard product construction used in the theory of finite automata. To describe the construction formally, we first construct a rooted graph B that generates precisely the canonical words. B has $k+1$ vertices $w_0, w_1, \ldots, w_k$ with w_0 being the root. For each $i \geq 1$, the vertex w_i has i loops labelled $1, \ldots, i$, and an in-going edge (w_{i-1}, w_i) labelled i. There are no other edges. It is easy to see that the rooted walks in B correspond to the canonical words over $\mathscr{A}$.

We now define the graph A_c. The vertex set of A_c is the Cartesian product of the vertex sets of A' and B. The edge set of A_c is determined by the following rule: A_c has a directed edge from (v, w) to (v', w') labelled i if and only if A' has a directed edge from v to v' labelled i and B has a directed edge from w to w' labelled i. The root of A_c is the vertex $x_0 = (v_0, w_0)$. It is clear that A_c has a rooted walk representing a word w if and only if both A' and B have rooted walks representing w. Since the walks in A' represent the words of $\mathcal{L}$ and the walks in B represent the canonical words, we conclude that the walks in A_c represent precisely the words from $\mathcal{L}_c$.

Furthermore, we may easily check that A_c has no directed cycles except for loops. Indeed, the vertices of A' can be topologically ordered in such a way that all the directed edges are nondecreasing in this ordering, and the same is true for B. We may then order the vertices of A_c lexicographically using these two orderings, to obtain an ordering on A_c in which all the edges are nondecreasing. This shows that A_c has no directed cycles.

It remains to deduce the expression for the generating function from the structure of A_c. Let t be a rooted walk in A_c. Recall that the non-loop edges of t form a directed path p which we call the *backbone* of t. Let $x_0, x_1, \ldots, x_r$ be the vertices of p, and let l_i be the number of loops adjacent to the vertex x_i. Note that $l_i \leq k$, because the loops at a given vertex are labelled by distinct symbols. Let t_n be the number of walks of length n whose backbone is p, and let $g_p(x) = \sum_{n \geq 0} t_n x^n$ be the corresponding generating function. It is easy to obtain the expression

$$g_p(x) = \frac{x^r}{\prod_{i=0}^{r}(1 - l_i x)}.$$

Summing the functions $g_p(x)$ over the finitely many directed paths p of A_c starting in x_0, we obtain the generating function enumerating the rooted walks of A_c, which is equal to the generating function $f_P(x)$ enumerating P.

2.3 Proof of Theorem 1.3

To show that Theorem 1.1 cannot be directly extended to classes of partitions with an unbounded number of blocks, we now show that the class of partitions avoiding $\pi_0 = \{\{1,3\},\{2,4\},\{5\}\}$ has no finite ordinal dual.

We find it convenient to represent the partition π_0 by the word *ababc*. We prove that the partition *ababc* has no finite ordinal dual. In other words, we will prove that there is no finite collection D of partitions of ordinals with the property that every normalized partition avoiding *ababc* is contained in at least one member of D.

First we introduce some necessary notation and terminology. Let π be a partition of an ordinal α. We represent π as a word of length α over some (possibly infinite) alphabet $\mathscr{A}$ and make no distinction between the word and the partition it represents. We let π_i denote the i-th symbol of π (where $i \in \alpha$ is an ordinal index). Let $X, Y \in \mathscr{A}$ be two letters. We say that the letter X *crosses* the letter Y in π (or Y is *crossed by* X in π) if π contains the subword $XYXY$. More specifically, we say that X *crosses* Y *at* (i,j,k,l), if $i < j < k < l$ are four ordinal indices such that $\pi_i = \pi_k = X$ and $\pi_j = \pi_l = Y$.

The following proposition is the key to our proof.

Proposition 2.5. *Let π be an ababc-avoiding partition of an ordinal, represented as a word over $\mathscr{A}$. Then there is a (finite) number $M \in \mathbb{N}$ such that each letter $Y \in \mathscr{A}$ is crossed in π by at most M distinct letters.*

The proof of Proposition 2.5 is separated into three lemmas. In these lemmas, we assume that π is a fixed *ababc*-avoiding partition of an ordinal α, represented by a fixed word over $\mathscr{A}$.

Lemma 2.6. *Let X, X' and Y be three distinct letters of $\mathscr{A}$. Assume that X crosses Y at (i,j,k,l) and X' crosses Y at (i',j',k',l'), where $i < i'$. Then we have the following inequalities:*

$$i < i' < \min\{j, j'\} \le \max\{j, j'\} < k' < k < \min\{l, l'\} \le \max\{l, l'\}.$$

Hence, π contains the subword $XX'YX'XY$, and X' crosses Y at (i', j, k', l).

Proof. We will show that any other mutual arrangement of the eight indices $i, j, k, l, i', j', k', l'$ creates a pattern *ababc*. We have $i < i'$ by assumption, and $i < j < k < l$ and $i' < j' < k' < l'$ by definition.

Let us first show that k is greater than k': if $k < j'$, then the indices $i < j < k < j' < k'$ induce the subword $XYXYX'$, isomorphic to the forbidden pattern. Similarly, if $j' < k < k'$, then $i < i' < k < k' < l'$ corresponds to $XX'XX'Y$.

Next, we observe that $k < l'$, otherwise $i' < j' < k' < l' < k$ produces $X'YX'YX$.

It remains to show that j is between i' and k': if $j < i'$, then $j < i' < j' < k' < k$ gives $YX'YX'X$, and if $k' < j$ then $i' < j' < k' < j < k$ gives $X'YX'YX$.

These inequalities imply the lemma. $\square$

Let $\mathscr{B}(Y)$ denote the set of all the letters that cross the letter Y in π.

Lemma 2.7. *For each $Y \in \mathscr{A}$, the set $\mathscr{B}(Y)$ is finite.*

Proof. Choose an arbitrary $Y \in \mathscr{A}$. We introduce an ordering $\prec$ on the set $\mathscr{B}(Y)$ in the following way: for two distinct letters $X, X' \in \mathscr{B}(Y)$, we write $X \prec X'$ if the first occurrence of X in π precedes the first occurrence of X'. Assume that $X_1 \prec X_2 \prec X_3 \prec \cdots$ is the ascending chain of all the elements of $\mathscr{B}(Y)$. For contradiction, assume that this chain is infinite. We will now find an infinite descending chain in the ordinal α ordered by $\leq$, which is a contradiction, since ordinals do not contain infinite descending chains.

By Lemma 2.6, we know that there are two indices $j, l \in \mathscr{A}$, such that every letter $X \in \mathscr{B}(Y)$ crosses Y at (i, j, k, l), for some values of i and k. In fact, if $i(X)$ denotes the index of the first occurrence of X in π, and if $k(X)$ is the smallest index greater than j such that $\pi_{k(X)} = X$, then X crosses Y at $(i(X), j, k(X), l)$. If $X \prec X'$ are two letters from $\mathscr{B}(Y)$, we know that $i(X) < i(X')$, and by Lemma 2.6, we also have $k(X) > k(X')$. In particular, $k(X_1) > k(X_2) > \cdots$ is an infinite decreasing sequence of indices, which is a contradiction, as mentioned above. $\square$

Lemma 2.8. *There is at most one letter $Y \in \mathscr{A}$ for which $|\mathscr{B}(Y)| > 1$.*

Proof. Assume that there is a letter Y such that $|\mathscr{B}(Y)| > 1$. Choose distinct letters $X, X' \in \mathscr{B}(Y)$. We may assume, by Lemma 2.6, that there are two indices $j, l \in \alpha$ such that X crosses Y at (i, j, k, l) for some i, k, and X' crosses Y at (i', j, k', l), for some i', k'. We claim that for each $m \geq l$, we have $\pi_m = Y$. Indeed, the indices $i < j < k < l < m$ induce the word $XYXY\pi_m$ and the indices $i' < j < k' < l < m$ induce the word $X'YX'Y\pi_m$. If $\pi_m \neq Y$, then at least one of these words is isomorphic to the pattern *ababc* (while the other might be isomorphic to *ababa*).

We conclude that if $|\mathscr{B}(Y)| > 1$, then there is a value $l \in \alpha$ such that $\pi_m = Y$ for each $m \geq l$. Clearly, there can be at most one Y with this property. $\square$

Proof of Proposition 2.5. By Lemma 2.7, we know that $\mathscr{B}(Y)$ is finite for every Y, and by lemma 2.8, we see that $|\mathscr{B}(Y)| \leq 1$ for all Y except one. Thus, we may define $M = \max\{|\mathscr{B}(Y)|; Y \in \mathscr{A}\}$, knowing that M is finite. $\square$

We complete the proof of Theorem 1.3. Let C be a finite collection of ordinal partitions avoiding *ababc*. Let us choose M such that no letter $Y \in \mathscr{A}$

is crossed more than M times in any partition from C. We take the *ababc*-avoiding partition $\rho = X_1X_2\cdots X_{M+1}YX_{M+1}X_M\cdots X_1Y$. In ρ, the letter Y is crossed $M+1$ times, so ρ is not contained in any member of C.

3 Molecular classes

We now focus on general classes of relational structures. Our first aim is to prove Theorems 1.5 and 1.6 announced in the introduction.

3.1 Proofs

Let us begin with the proof of Theorem 1.5. Let $(A,\le)$ be a poset. A subset of A whose elements are pairwise incompatible in A will be called a *pie set* of A.

It is clear that if A has $k+1$ pairwise incompatible elements, then it cannot be covered by k ideals. To prove the theorem, it suffices to show that either A has an infinite pie set, or there is an integer k such that A can be covered by k ideals and it contains a pie set of size k.

Assume first that there is an integer $k \in \mathbb{N}$ such that the largest pie set of A has cardinality k. Let $P = \{\alpha_1,\dots,\alpha_k\}$ be a pie set of size k. We define k sets $C_1, C_2, \dots, C_k$, where C_i is the set of all the elements of A that are compatible with α_i. We claim that these k sets are ideals that cover A.

It is easy to see that see that the sets $C_1,\dots,C_k$ cover A. If there were an element $\beta \in A$ not belonging to any of the sets $C_1,\dots,C_k$, then β would be incompatible with all the elements of P, yielding a pie set of size $k+1$.

We now show that each C_i is an ideal. Choose $i \in [k]$ arbitrarily. Clearly C_i is a down-set. Let β and β' be two elements of C_i. Our goal is to show that β and β' have a common upper bound in C_i. By definition of C_i, the elements β and β' are compatible with α_i, so A contains a common upper bound γ of α_i and β, and a common upper bound γ' of α_i and β'. Since γ and γ' are greater than α_i, they both belong to C_i. Note that γ and γ' are incompatible with all the elements of P except α_i, because if γ or γ' were compatible with some α_j, then α_i would be compatible with α_j as well. We conclude that γ and γ' must be compatible in A, otherwise $P\cup\{\gamma,\gamma'\}\setminus\{\alpha_i\}$ would be a pie set of size $k+1$. Thus, γ and γ' have a common upper bound δ, which necessarily belongs to C_i. Since δ is also a common upper bound of β and β', we see that C_i is an ideal, as claimed.

To complete the proof of Theorem 1.5, we need to show that if the poset A has arbitrarily large finite pie sets, then it has an infinite pie set as well. Assume for contradiction that A has arbitrarily large finite pie sets, but no infinite pie set. We will inductively construct a sequence $\beta_1,\beta_2,\dots$ of elements of A and a sequence $A_0, A_1, A_2, \dots$ of subposets of A, with these properties:

- For each $n \in \mathbb{N}\cup\{0\}$, the poset A_n is an up-set of A. In particular, any two elements that are incompatible in A_n are also incompatible in A, and any pie set of A_n is also a pie set of A.

- For each $n \in \mathbb{N} \cup \{0\}$, the poset A_n contains arbitrarily large pie sets.
- For each $n \in \mathbb{N}$, A_n is a subposet of A_{n-1}.
- For each $n \in \mathbb{N}$, β_n belongs to A_{n-1}, and β_n is incompatible with all the elements of A_n. Note that this implies that $\{\beta_n;\ n \in \mathbb{N}\}$ is a pie set of A.

As the base step of the induction, define $A_0 = A$. For the induction step, let n be greater than 0, and assume that the poset A_{n-1} has already been defined and that it satisfies the conditions above. Note that if $P_1 \subseteq P_2 \subseteq P_3 \subseteq \cdots$ is a chain of pie sets of A_{n-1}, then the union $\bigcup_{i \geq 1} P_i$ is also a pie set of A_{n-1}. By Zorn's lemma, A_{n-1} has an $\subseteq$-maximal pie set P. Since every pie set of A_{n-1} is also a pie set of A, the set P is finite. Since A_{n-1} contains arbitrarily large pie sets, we may assume that P has at least two elements. Let $\alpha_1, \ldots, \alpha_k$ be the elements of P. Let C_i be the set of all the elements of A_{n-1} that are compatible with α_i, and let G_i be the set of all the elements of A_{n-1} that are greater than or equal to α_i. Obviously, G_i is a subset of C_i, and any element of G_i is incompatible with any element of G_j whenever $i \neq j$.

From the maximality of P, we know that the sets $C_1, \ldots, C_k$ cover A_{n-1}. In particular, at least one of these sets contains arbitrarily large pie sets of A_{n-1}. Assume, without loss of generality, that C_1 contains arbitrarily large pie sets. It follows that G_1 contains arbitrarily large pie sets as well, because every element of C_1 has an upper bound in G_1, and when we replace an element of a pie set with its upper bound, we get a new pie set of the same cardinality.

We now define $A_n = G_1$ and $\beta_n = \alpha_2$. It is clear that A_n and β_n have all the required properties. In particular, $\{\beta_n;\ n \in \mathbb{N}\}$ is an infinite pie set of A, contradicting our assumptions. This contradiction completes the proof of Theorem 1.5.

The proof of Theorem 1.6 is now easy. Assume that C is a hereditary class of relational structures. The containment relation $\subseteq$ on the class C is a quasi-order. To be able to apply Theorem 1.5 directly, we choose one representative from every isomorphism class of structures in C, and we let C_I be the set of these representatives. Then $(C_I, \subseteq)$ is a poset. It is clear that the poset $(C_I, \subseteq)$ can be covered by k ideals if and only if the class C can be expressed as a union of k atomic classes, and that the poset $(C_I, \subseteq)$ has a subset S_I of pairwise incompatible elements, if and only if the class C has a set S of pairwise incompatible elements of the same cardinality. Applying Theorem 1.5 for the poset $(C_I, \subseteq)$, we obtain Theorem 1.6.

A direct consequence of Theorem 1.6 is the following corollary, which we already hinted at in the introduction. A quasi-order relation is called a *well-quasi-order* if it has no infinite antichains and no strictly decreasing infinite chains. In the context of permutation classes, the following result appears in Murphy's thesis [20], and a slightly weaker statement is given in Atkinson et al. [2, Theorem 2.2]. Their proofs generalize immediately to the setting of relational structures and are only slightly longer than the proof given below.

Corollary 3.1. *A hereditary class C of relational structures that is well-quasi-ordered by containment is molecular.*

Proof. Since C is well-quasi-ordered by inclusion, it has no infinite $\subseteq$-antichain, hence no infinite pie set. By Theorem 1.6, it is molecular. □

The converse of Corollary 3.1 does not hold: it is easy to come up with examples of molecular (even atomic) classes that are not well-quasi-ordered by inclusion. For instance, the class of all finite permutations is atomic but not well-quasi-ordered.

3.2 Examples of non-molecular finitely based classes

Recall that a hereditary class of relational structures is called *finitely based*, if its set of minimal forbidden patterns is finite. We now present a finitely based class of ordered graphs that is not molecular. Then we will show how this example may be extended to similar classes of unordered graphs, permutations, and set partitions. Our construction is related to previously known constructions of infinite antichains of permutations [23].

An *ordered graph* is a (simple undirected) graph whose vertex set is linearly ordered. We shall represent ordered graphs as relational structures with two binary relations, where the first relation is a linear ordering of the ground set, and the second relation is a symmetric relation that represents edge adjacency. Note that in this representation, set partitions are represented as ordered graphs whose every connected component is a clique.

Let C be the class of all ordered graphs that have no vertex of degree greater than three, and at most two vertices of degree three. It is easily observed that C is a finitely based hereditary class. We now present an infinite sequence of elements of C that are incompatible in C, thus showing that C is not molecular.

For an even number $n \geq 8$, let G_n be an ordered graph on the vertex set $[n]$, which is a union of two triangles on the vertices $1, 2, 3$ and $(n-2), (n-1), n$ and a path connecting the vertices $1, 5, 4, 7, 6, 9, 8, \dots, 2i+1, 2i, \dots, n-5, n-6, n-3, n-4, n$. Note that G_n has exactly two vertices of degree three (1 and n), while the remaining vertices have degree two. In particular, $G_n \in C$. Note that each graph G_n consists simply of a pair of triangles joined by a path. The purpose of the slightly peculiar vertex ordering is to ensure that G_n represents a permutation, in the sense defined below.

It is not difficult to check that for any two distinct even numbers $m, n \geq 8$, the graphs G_m and G_n are incompatible in the class C. This shows that C is not molecular.

This construction works in the same way within the realm of unordered graphs, and it can be easily extended to hereditary classes of permutations: indeed, let π_n be the permutation

$$3, 2, 5, 1, 7, 4, 9, 6, \dots, 2i+1, 2i-2, \dots, n-3, n-6, n, n-4, n-1, n-2.$$

We say that an ordered graph G on the vertex set $[n]$ represents a permutation π if ij is an edge of G if and only if the two elements i, j have their order inverted by the permutation. With this representation, the containment of permutations is translated into induced containment of ordered graphs. The graphs that represent a permutation are called *permutation graphs*.

The permutation π_n is represented by the graph G_n. Let Π be the class containing all the permutations whose representing graphs belong to C. It follows from the previous discussion that the permutations π_n form an infinite sequence of permutations which are pairwise incompatible in Π. It is not difficult to see that the minimal forbidden permutations of Π are exactly those permutations that are represented by the permutation graphs among the minimal forbidden induced graphs of C. We conclude that Π is a non-molecular finitely based class of permutations.

We remark that this construction can be further adapted to the poset of set partitions. Note that a permutation π of order n can be represented by a partition of the set $[2n]$ into two-element sets $\{i, n+\pi(i)\}$, where the containment relation of these special set partitions corresponds to the containment relation among permutations. Thus, we might use the above construction to obtain a finitely based hereditary class of set partitions that is not molecular. We omit the details of the argument.

4 Conclusions and open problems

We have shown that classes of set partitions with a bounded number of blocks admit a simple description in terms of finite ordinal duals, and that this description is useful in deriving enumeration results. While these results do not generalize to arbitrary classes of partitions, there are classes of unbounded number of blocks that admit finite ordinal duals as well. It is at this point an open problem to characterize the classes that have a finite ordinal dual.

In the setting of relational structures, we have introduced the concept of molecular classes, and have provided a characterization of molecular classes in terms of joint embeddability of their elements.

As we mentioned in the introduction, many researchers have recently studied hereditary classes of various structures, with the goal of characterizing their possible speeds. It appears, though, that many results in this area are of a 'negative' nature, i.e., they show that the speeds of hereditary classes can behave in a rather arbitrary fashion. For instance, Balogh, Bollobás and Weinreich [6] have constructed, for any $c > 1$ and any $\varepsilon > 1/c$, hereditary classes of graphs with speeds oscillating infinitely often between $n^{(c+o(1))n}$ and $2^{n^{2-\varepsilon}}$.

In view of such results, it might be desirable to find more 'well-behaved' subfamily of hereditary classes, in which such arbitrary behavior does not appear. Molecular classes might provide such a convenient subfamily. Indeed, the 'oscillating' classes constructed by Balogh et al. are non-molecular (they in fact contain infinitely many graphs that are not proper subgraphs of any other graph in the class). In contrast, in the setting of graphs [5] as well as in the setting of permutations [17, 24], there are 'positive' results which show that classes of sufficiently small speeds are 'well-behaved', in the sense that their speeds are tightly constrained. These results are based on structural descriptions of small classes, which actually easily imply that these small classes are molecular (or even well-quasi-ordered). We may thus hope that by restricting our study to molecular classes of ordered structures, as opposed to general hereditary classes, we might recover some of the nice behavior that we observe in small speeds.

5 Acknowledgements

We are grateful for the numerous suggestions given by the anonymous referees. In particular, the proof of Theorem 1.5 given in this paper has been suggested by a referee, and is much simpler than our original argument.

References

[1] Michael H. Albert and Steve A. Linton. Growing at a perfect speed. *Combinatorics, Probability and Computing*, 18(3):301–308, 2009.

[2] Mike D. Atkinson, Maximillian M. Murphy, and Nikola Ruškuc. Partially well-ordered closed sets of permutations. *Order*, 19(2):101–113, 2002.

[3] Mike D. Atkinson, Maximillian M. Murphy, and Nikola Ruškuc. Pattern avoidance classes and subpermutations. *Electronic Journal of Combinatorics*, 12, 2005.

[4] József Balogh, Béla Bollobás, and Robert Morris. Hereditary properties of partitions, ordered graphs and ordered hypergraphs. *European Journal of Combinatorics*, 27(8):1263–1281, 2006.

[5] József Balogh, Béla Bollobás, and David Weinreich. The speed of hereditary properties of graphs. *Journal of Combinatorial Theory, Series B*, 79(2):131–156, 2000.

[6] József Balogh, Béla Bollobás, and David Weinreich. The penultimate rate of growth for graph properties. *European Journal of Combinatorics*, 22(3):277–289, 2001.

[7] Béla Bollobás. Hereditary and monotone properties of combinatorial structures. In Anthony Hilton and John Talbot, editors, *Surveys in Combinatorics 2007*, volume 346 of *London Mathematical Society Lecture Note Series*, pages 1–40. Cambridge University Press, 2007.

[8] Anna de Mier. On the symmetry of the distribution of k-crossings and k-nestings in graphs. *Electronic Journal of Combinatorics*, 13(1), 2006.

[9] Christian Delhommé, Maurice Pouzet, Gábor Sági, and Norbert Sauer. Representation of ideals of relational structures. *Discrete Mathematics*, 309(6):1374 – 1384, 2009.

[10] Roland Fraïssé. Sur l'extension aux relations de quelques propriétés des ordres. *Annales Scientifiques de l'École Normale Supérieure*, 71(4):363–388, 1954.

[11] Roland Fraïssé. *Theory of Relations*. North-Holland, second edition, 2000.

[12] Adam M. Goyt. Avoidance of partitions of a three-element set. *Advances in Applied Mathematics*, 41(1):95–114, 2008.

[13] Pavol Hell and Jaroslav Nešetřil. *Graphs and Homomorphisms*. Oxford Lecture Series in Mathematics and Its Applications, 28. Oxford University Press, 2004.

[14] John E. Hopcroft and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley Series in Computer Science. Addison-Wesley, first edition, 1979.

[15] Sophie Huczynska and Nik Ruškuc. Pattern classes of permutations via bijections between linearly ordered sets. *European Journal of Combinatorics*, 29(1):118–139, 2008.

[16] Vít Jelínek and Toufik Mansour. On pattern-avoiding partitions. *Electronic Journal of Combinatorics*, 15, 2008.

[17] Tomáš Kaiser and Martin Klazar. On growth rates of closed permutation classes. *Electronic Journal of Combinatorics*, 9(2), 2003.

[18] Martin Klazar. Counting pattern-free set partitions I: a generalization of Stirling numbers of the second kind. *European Journal of Combinatorics*, 21(3):367–378, 2000.

[19] Joseph B. Kruskal. The theory of well-quasi-ordering: A frequently discovered concept. *Journal of Combinatorial Theory, Series A*, 13:297–305, 1972.

[20] Maximillian M. Murphy. *Restricted permutations, antichains, atomic classes, and stack sorting*. PhD thesis, University of St. Andrews, 2002.

[21] Maurice Pouzet. The profile of relations. *Global Journal of Pure and Applied Mathematics*, 2(3):237–272, 2006.

[22] Bruce E. Sagan. Pattern avoidance in set partitions. arXiv:math/0604292; to appear in Ars Combinatoria.

[23] Daniel A. Spielmann and Miklós Bóna. An infinite antichain of pemutations. *Electronic Journal of Combinatorics*, 7(1), 2000.

[24] Vincent Vatter. Small permutation classes. arXiv:0712.4006, 2007.

[25] Vincent Vatter. Permutation classes of every growth rate above 2.48188. *Mathematika*, 56:182–192, 2010.