

The undecidability of joint embedding for 3-dimensional permutation classes

Samuel Braunfeld

University of Maryland, College Park, USA

received 28th Feb. 2020, revised 25th Mar. 2021, 23rd July 2021, accepted 10th Aug. 2021.

As a step towards resolving a question of Ruškuc on the decidability of joint embedding for hereditary classes of permutations, which may be viewed as structures in a language of 2 linear orders, we show the corresponding problem is undecidable for hereditary classes of structures in a language of 3 linear orders.

Keywords: permutations, undecidable, joint embedding, atomic

1 Introduction

In Ruškuc (2005), Ruškuc posed several decision problems for finitely-constrained permutation classes, with the decidability of atomicity among them (and this question was recently re-posed in Jelínek (2017)). A permutation avoidance class is called *atomic* if it cannot be expressed as a union of two proper subclasses. A general hope is that understanding a permutation class can be reduced to understanding its atomic subclasses, as in the following lemma for calculating growth rates (see Vatter (2015) for a reference).

Lemma 1.1 *Suppose \mathcal{K} is a permutation class, with no infinite antichain in the containment order. Then \mathcal{K} can be expressed as a finite union of atomic subclasses. Furthermore, the upper growth rate of \mathcal{K} is equal to the maximum upper growth rate among its atomic subclasses.*

We may view permutations as structures in a language of two linear orders. Atomicity is then equivalent to the joint embedding property (see Vatter (2015)), a standard model-theoretic notion, so we may rephrase Ruškuc’s question.

Definition 1.2 *A class \mathcal{C} of structures has the joint embedding property (JEP) if, given $A, B \in \mathcal{C}$, there exists $C \in \mathcal{C}$ such that A, B embed into C .*

Question 1 *Is there an algorithm that, given finite set of forbidden permutations, decides whether the corresponding permutation class has the joint embedding property?*

This problem is known to be decidable in certain restricted classes of permutations, such as monotone grid classes Waton (2007). Also, whether a permutation class is a natural class, which is a strengthening of atomicity, is decidable Murphy (2003).

However, we believe there is a strong possibility Ruškuc’s problem is undecidable in general. We are not aware of many undecidability results in the permutation class literature, although Garrabrant and Pak (2015), using methods that seem quite different from ours, proves an undecidability result about comparing the parity of the number of permutations of size n in two permutation classes.

The author took a first step towards Ruškuc’s problem in Braunfeld (2019), proving the JEP is undecidable for hereditary graph classes. Although it is not yet clear whether that proof can be adapted to permutations, we here adapt it to *3-dimensional permutations*, i.e. structures in a language of 3 linear orders, proving the following theorem via a reduction from the string tiling problem.

Theorem 1.3 *There is no algorithm that, given a finite set of forbidden 3-dimensional permutations, decides whether the corresponding 3-dimensional permutation class has the JEP.*

A very rough sketch of the proof is as follows. We use a reduction from the (string) tiling problem, which asks whether we can tile an infinite grid using a given collection of tile types, subject to local constraints. The first two steps below ensure that the tiling problem is equivalent to whether we can jointly embed two particular 3-dimensional permutations, and the third step ensures that joint embedding for the class is equivalent to joint embedding for those two 3-dimensional permutations.

1. Construct two 3-dimensional permutations A^* , representing a grid, and B^* representing a suitable collection of tiles.
2. Choose a finite set of constraints to ensure that if C jointly embeds A^* and B^* , then it encodes a solution to the string tiling problem.
3. Show that if the string tiling problem admits a solution, then the chosen class has joint embedding.

2 Background

2.1 The (string) tiling problem

Rather than using a reduction from the halting problem to prove undecidability, we will use the string tiling problem, a variant of the tiling problem. The input to a tiling problem consists of a finite set $Tiles$ of tile types, as well as a set of rules of the form “Tiles of type i cannot be placed directly above tiles of type j ” and “Tiles of type k cannot be placed directly right of tiles of type ℓ ”. A solution to a tiling problem is a surjective function $\tau: \mathbb{N}^2 \rightarrow Tiles$, interpreted as placing tiles on a grid, that respects the tiling rules.

Theorem 2.1 (Berger (1966)) *There is no algorithm that, given a sets of tile types and tiling rules, decides whether the corresponding tiling problem has a solution.*

We will use a variant, called string tiling problems in Cherlin (2011). Here there are only two tile types, but there is some $D \in \mathbb{N}$ such that for every $d \leq D$, tiling rules may restrict which tiles are placed at distance d to the right of a given tile, or directly above a given tile. An encoding of tiling problems as string tiling problems is given in Lemma 7.6 of Cherlin (2011), the idea being to use several tiles in the string tiling problem to encode a single tile from the standard tiling problem. This proves the analogue of Theorem 2.1 for the string tiling problem.

As we will be reducing from the string tiling problem, which is co-recursively enumerable, we point out here that if \mathcal{C} is a hereditary class of finite structures in a finite relational language, then the JEP for \mathcal{C} is also co-recursively enumerable. To see this, consider $A, B \in \mathcal{C}$ that can be jointly embedded, as witnessed by

$C \in \mathcal{C}$ and embeddings $f: A \rightarrow C$ and $g: B \rightarrow C$. As \mathcal{C} is hereditary, the substructure of C induced on $f(A) \cup g(B)$ is also in \mathcal{C} . Thus, given $A, B \in \mathcal{C}$, $|A| + |B|$ is a bound on the size of the possible witnesses for joint embedding, and they can be exhaustively checked.

2.2 The argument for hereditary graph classes

We will now sketch the argument from Braunfeld (2019) for hereditary graph classes (in an expanded language with colored vertices and edges, and both directed and undirected edges), since our argument in this paper will attempt to re-encode it using 3-dimensional permutations. Although we are concerned with the JEP for finite structures in a hereditary class \mathcal{C} , the compactness theorem implies that the JEP for the finite members of \mathcal{C} is equivalent to the JEP for countable members of \mathcal{C} . Rather than work with families of increasingly large finite structures, we prefer to take our canonical models to be countable.




Fig. 1: A portion of the canonical models A_G^* and B_G^* , with the grid points in A_G^* tiled by tiles attached to grid points with the same coordinates in B_G^* . Path points are blue, with the origin a different shade. Grid points are red, their y -coordinate determined by an orange edge and their x -coordinate by a green edge. Tile points are purple. Points in 0-superscripted predicates have a black border, while points in 1-superscripted predicates do not.

This encodes a tiling of $(0,0)$ with tile-type 2, $(1,0)$ with tile-type 2, and $(2,0)$ with tile-type 1.

Suppose we are given a tiling problem \mathcal{T} . First, we describe graphs corresponding to A^* and B^* from the rough sketch in the introduction. A_G^* (see Figure 1) will contain a 1-way infinite directed path. To every pair of points in this path, we attach a point, representing a grid point with coordinates taken from the attached path points. Because we must distinguish between x and y -coordinates, we use the colored edges to attach each grid point to its coordinates. Furthermore the path points are colored distinctly from the grid points, and the origin of the path is also colored distinctly. B_G^* will look like a copy of A_G^* , although using a disjoint set of vertex colors. Furthermore, to each grid point in B_G^* , path of length t (where t is the number of tile types in the given tiling problem), using a new color for these points. These represent a full tile set available at each coordinate, with the different tile-types being distinguished by their distance from the corresponding grid point.

We then choose our constraints so that when we try to jointly embed A_G^* and B_G^* , the following is forced: for every grid point in A_G^* , with coordinates (x, y) , we must add an edge to one tile point attached to the grid point in B with the same coordinates. This is interpreted as tiling the point (x, y) by the corresponding

tile-type, and our constraints should further enforce the local tiling rules.

As A_G^* and B_G^* will be in our hereditary class $\mathcal{C}_{\mathcal{T}}$, if $\mathcal{C}_{\mathcal{T}}$ has the JEP, then \mathcal{T} must have a solution, since we can read a valid tiling off the structure embedding A_G^* and B_G^* . We must then show that if \mathcal{T} has a solution $\tau: \mathbb{N}^2 \rightarrow Tiles$, then we may jointly embed any $A, B \in \mathcal{C}_{\mathcal{T}}$. For this, we add a variety of additional constraints ensuring that if we must add edges due to the constraints in the previous paragraph, and thus are attempting to encode a valid tiling, then A and B look approximately like one of our canonical models A_G^* and B_G^* . Crucially, we ensure that every grid point involved in our attempted tiling has unique coordinates (x, y) on a unique path; we thus have a well-defined input to give to τ , and add edges from grid points in A to tile points in B (or vice versa) as τ dictates.

The additional difficulties with (3-dimensional) permutations arise from the transitivity of the orders, which places severe limitations on how we may jointly embed a given pair of structures. Also, some concerns that are in common with the graph case shift in their difficulty. A key point in the graph case is that grid points and tile sets have unique coordinates. While that was simple to enforce in the graph case, it, and even the proper definition of coordinates, will be a significant concern here. On the other hand, the point of most concern in the graph case was ensuring that none of the configurations used to encode unary predicates were accidentally created by our joint embedding procedure, i.e. our method for constructing a third structure jointly embedding two given structures. Here this problem will be trivialized by taking advantage of the third linear order, but it returns to the fore when working with permutation classes.

3 The canonical models

3.1 Preliminary definitions

We first mention that the primary reason for using a third linear order is to obtain the first claim at the beginning of Lemma 6.3. The third order can largely be ignored otherwise, which may help in picturing the constructions.




Fig. 2: A typical antichain element from Bóna and Spielman (2000). The root is 13.

We choose an antichain \mathcal{A} of 3-dimensional permutations on which $<_1 = <_3^{opp}$, i.e. the opposite order of $<_3$, and containing at least 20 members (10 are used for the unary predicates in this section, and we will double that number later), which we will use to encode unary predicates, which may be thought of as coloring points. We also require that each element of \mathcal{A} have at least 5 points, and that the $<_1$ -greatest, $<_1$ -least, $<_2$ -greatest, and two $<_2$ -least points of each element are distinct, with the $<_1$ -greatest point

$<_2$ -below the $<_1$ -least point. For example, let \mathcal{A} be the infinite antichain from Bóna and Spielman (2000) (see Figure 2; in this figure and all others, $<_1$ is horizontal and $<_2$ is vertical), with the third order defined by $<_1 = <_3^{opp}$.

For $i \in \{0, 1\}$, select distinct antichain elements $E_X^i, E_Y^i, E_P^i, E_G^0, E_T^1$, and E_O^i , and let \mathcal{E} be the set of these members. Elements with a 0-superscript will be used to encode grid points, i.e. structures like the graph A_G^* and their coordinates, while those with a 1-superscript will encode tile sets and their coordinates, i.e. structures like the graph B_G^* .

If $E \in \mathcal{E}$, we say x is the *root* of E if it is the $<_1$ -least point. We will think of roots as actually representing points in the colored graph we are trying to encode, and almost all other points as being auxiliary to assist the encoding.

We also define the following unary predicates.

1. (path points) $x \in P^i$ if x is the root of a copy of E_P^i or E_O^i
2. (path origins) $x \in O^i$ if x is the root of a copy of E_O^i
3. (grid points) $x \in G^0$ if x is the root of a copy of E_G^0
4. (tile-type 1) $x \in T_1^1$ if x is the root of a copy of E_T^1
5. (tile-type 2) $x \in T_2^1$ if x is the $<_2$ -greatest point of a copy of E_T^1
6. (tile points) $T^1 = T_1^1 \cup T_2^1$

Note that we do not encode points representing tile-type 2 by roots. Since these will always be paired with points representing tile-type 1, we instead represent tiles by two different points in E_T^1 .

The antichain condition is meant to stop the following kind of situation. Suppose that E_P^i embedded into E_G^0 . Then whenever we encode a grid point by adding a copy of E_G^0 , we would also be adding a copy of E_P^i , and thus unintentionally be adding a path point.

In addition to encoding unary predicates, we will use elements of \mathcal{E} to encode edges between their roots and other points, using the following notion of *capture*.

Given a point x and E a copy of an element in $\mathcal{E} \setminus \{E_T^1\}$, we say x is *captured* by E if x is $<_2$ -between the two $<_2$ -least points of E , $E <_1 x$, and $E <_3 x$. This should be thought of as encoding a graph edge between x and the root of E .

An example of capture is shown in Figure 3. There are two copies of an element of \mathcal{E} projected onto $<_1$ and $<_2$, although the left copy should be taken $<_3$ -below the right copy. We view this as encoding a directed edge from the root of the left copy to the root of the right copy, and we may encode a directed path by continuing to daisy-chain such copies.

Before giving the remaining definitions, we discuss the stylized portion of the canonical models shown in Figure 4, which should be compared to the graph version in Figure 1. We have projected onto $<_1$ and $<_2$, and we will later see that $<_3$ will be determined by $<_1$. Since the precise structure of the elements of \mathcal{E} is unimportant, copies of those elements are represented by colored lines with two $<_2$ -least points marked to display capture. As in Figure 1, the blue points represent path points, with the origin a different shade. Grid points are in red and tile points in purple, with their x -coordinates determined by the path point captured by the copy of E_X^i denoted in green, and their y -coordinates determined by the path point captured by the copy of E_Y^i denoted in orange. Pairs of tiles at a given coordinate are connected by a line,




Fig. 3: Two copies of an element from \mathcal{E} . The left copy is capturing the root of the right copy.

with the tile of type 2 above and right of the tile of type 1. Finally, the copies of E_G^0 are the black lines connected to grid points, and they capture a tile from the pair of tiles at the corresponding coordinate. If the dividing line between $A_{<_1}^*$ and $B_{<_1}^*$ is ignored, this picture encodes a tiling of both $(0, 0)$ and $(1, 0)$ with tile-type 2. The dotted horizontal lines separate different regions of the picture, but are not part of the structure.

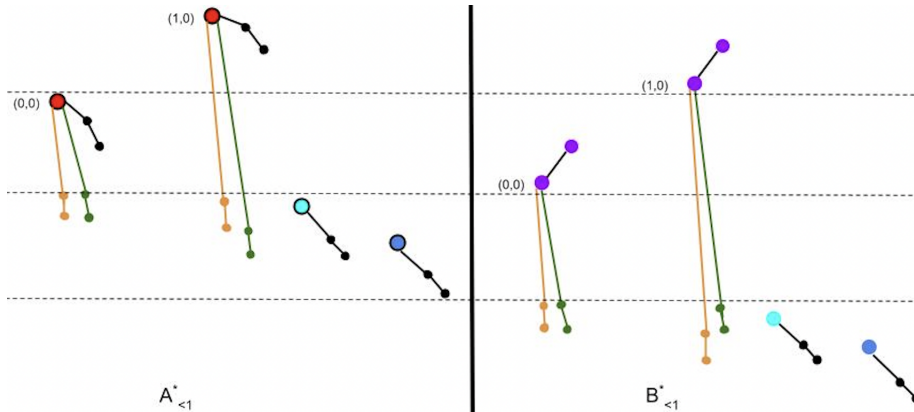


Fig. 4: A stylized projection onto $<_1$ and $<_2$ of a portion of the canonical models $A_{<_1}^*$ and $B_{<_1}^*$.

The following definitions may all be seen in Figure 4, with the exception that the P^i -paths shown only have length 2 instead of being infinite.

We say $x \in G^0$ is tiled by $y \in T^1$ if x is the root of a copy of E_G^0 that captures y .

We say $t_1 \in T_1^1$ and $t_2 \in T_2^1$ form a *tile set* if there exists E a copy of E_T^1 with root t_1 and $<_2$ -greatest point t_2 .

Given a point $g \in G^0$ and $x, y \in P^0$ or $g \in T_1^1$ and $x, y \in P^1$, we say g is *coordinatized* by (x, y) if g is the root of a copy of E_X^0 that captures x and of E_Y^0 that captures y (or in the second case, we use E_X^1 and E_Y^1).

We say p' is a *path-successor* of p if $p, p' \in P^i$ and p is the root of a copy of E_P^i or E_O^i that captures p' .

We say h is a *horizontal successor* of g if $g, h \in G^0$ or $g, h \in T_1^1$ and there are $x, y, x' \in P^0$ (or P^1 in the second case) such that g is coordinatized by (x, y) , h is coordinatized by (x', y) , and x' is a path-successor of x . *Vertical successor* is defined similarly, but h is coordinatized by (x, y') and y' is a path-successor of y . *Horizontal predecessor* and *vertical predecessor* are defined conversely.

We define an *infinite one-way P^i -path* to be a copy of E_O^i with root p_0 and a sequence of copies of E_P^i , with roots (p_1, p_2, \dots) arranged such that p_{k+1} is captured by the copy of E_P^i (or E_O^i) rooted at p_k , and the copy of E_P^i (or E_O^i) rooted at p_k is $<_1$ -below that rooted at p_{k+1} . In this case, we say p_0 is the *path-origin* of the path.

We say $g \in G^0$ is a *grid-origin* if there is some $x \in O^0$ such that G is coordinatized by (x, x) . We say $t \in T_1^1$ is a *tile-origin* if there is some $x \in O^1$ such that G is coordinatized by (x, x) .

Let $g \in G^0 \cup T_1^1$ be coordinatized by (x, y) . We say g is *on the x -axis* if $y \in O^i$ (for the appropriate i), and g is *on the y -axis* if $x \in O^i$ (we will sometimes also refer to a tile set being on an axis if its first tile is). Note that a grid-origin or tile-origin is on both the x -axis and y -axis.

We define a *connector interval* to be the open $<_2$ -interval defined by the two $<_2$ -least points of a copy of E_G^0 . We define a *tile set interval* to be the open $<_2$ -interval defined by the $<_1$ -least point and the $<_2$ -greatest point (i.e. by the two tiles) of a copy of E_T^1 . Finally, we define a *special interval* to be either a connector interval or a tile set interval.

Given a special interval defined by some $E \in \mathcal{E}$, we call the $<_2$ -greater endpoint of the special interval its *top endpoint*, and the $<_2$ -lesser endpoint its *bottom endpoint*.

Although we defined a special interval as a $<_2$ -interval corresponding to a copy of an element of \mathcal{E} , we will often conflate the special interval with its corresponding copy of an element of \mathcal{E} . However, intersection of special intervals will always refer to intersection of the $<_2$ -intervals.

3.2 The canonical models

Our proof proceeds in two steps. First we prove the undecidability of the $<_1$ -JEP, defined below, following the rough sketch from the introduction. This will be finished by Section 6.1. Then, in Section 6.2, we quickly reduce from the $<_1$ -JEP to the JEP. In this section, we describe our canonical models for the first step.

Definition 3.1 *We say that a class of 3-dimensional permutations has the $<_1$ -JEP if it admits a joint embedding procedure in which, given factors labeled A and B , the procedure places $A <_1 B$.*

We now describe our canonical models $A_{<_1}^*$ and $B_{<_1}^*$ for the $<_1$ -JEP, corresponding to the graphs A_G^* and B_G^* from §2.2. We only describe $<_1$ and (sometimes) $<_2$, since $<_3$ will be determined as follows: if x, y are in the same copy of an element of \mathcal{E} that we specify below, then $x <_1 y \iff x >_3 y$. Otherwise, $x <_1 y \iff x <_3 y$. Note that this will ensure that the only copies of elements of \mathcal{E} appearing in either $A_{<_1}^*$ or $B_{<_1}^*$ will be those specified below. For suppose there is some further, unintended copy \widehat{E} of an $E \in \mathcal{E}$. By the antichain condition, it cannot embed into any single copy of an $E' \in \mathcal{E}$, so there must be

$x, y \in \widehat{E}$ occurring in two distinct specified copies of elements of \mathcal{E} . Suppose that $x <_1 y$, so then $x <_3 y$. But this contradicts that $<_1$ and $<_3$ disagree within E .

It would be helpful to keep in mind the image of part of $A_{<_1}^*$ and $B_{<_1}^*$ given in Figure 4 when reading the following construction. We start constructing $A_{<_1}^*$ by placing an infinite one-way P^0 -path with roots (p_0, p_1, \dots) . Then, $<_1$ -below and $<_2$ -above the path, we place a sequence of points indexed by \mathbb{N}^2 , increasing antilexicographically with respect to $<_1$ and $<_2$ (we say $(x, y) <_{antilex} (x', y')$ if $y < y'$ or $y = y'$ and $x < x'$); so, identifying a point with its indices, if $g <_{antilex} g'$ then $g <_{1,2} g'$. We now make each such point, which we will call grid points, the root of 3 different copies of elements of \mathcal{E} . Consider the point g indexed by (x, y) . We make g the root of a copy $E_{g,X}$ of E_X^0 , $E_{g,Y}$ of E_Y^0 , and $E_{g,G}$ of E_G^0 , satisfying the following.

1. $E_{g,X}$ captures p_x and $E_{g,Y}$ captures p_y .
2. $E_{g,G}$ is $<_2$ -above the path.
3. $E_{g,X} <_1 E_{g,Y} \setminus \{g\} <_1 E_{g,G} \setminus \{g\}$.
4. Let $g <_{antilex} g'$. Then for every $E, E' \in \mathcal{E}$ (not necessarily distinct), any copy of E rooted at g is $<_1$ -less than any copy of E' rooted at g' . Furthermore, $E_{g,G} <_2 E_{g',G}$.

We construct $B_{<_1}^*$ similarly, except using 1-superscripted elements of \mathcal{E} instead of 0-superscripted elements, and using copies of E_T^1 instead of E_G^0 .

As in the graph case, we will choose our constraints so that when performing the $<_1$ -JEP on $A_{<_1}^*$ and $B_{<_1}^*$, i.e. constructing a C that jointly embeds both, we will be forced to tile each grid point in $A_{<_1}^*$ by a tile from the corresponding tile set in $B_{<_1}^*$.

3.3 Picturing joint embedding

In this subsection, we refer back to Figure 4 to help picture the process of performing $<_1$ -joint embedding. This should explain some choices we made in constructing our canonical models and should help motivate some of the constraints in the next section. Also, although Figure 4 depicts the canonical models, what we say will apply to any 3-dimensional permutations we consider.

Suppose we want to $<_1$ -joint embed the structures $A = A_{<_1}^*$ and $B = B_{<_1}^*$ depicted in Figure 4. We must place $A <_1 B$. We will also choose to place $A <_3 B$ since, as in the construction of our canonical models, this will mean we don't create unintended copies of elements of \mathcal{E} . Given this choice, our only freedom is in determining the $<_2$ -relation between points in A and points in B . We may view our procedure as erasing the dividing line between the two structures, keeping the points in A fixed, and keeping the horizontal positions of the points in B fixed but being able to slide them around vertically while maintaining their vertical ordering. By sliding the points around, we may make certain special intervals in A and B intersect, and in particular make certain tiles fall into corresponding connector intervals.

Figure 4 was drawn so that minimal sliding is required. If we just erase the line and leave all the points fixed, this produces a tiling of both $(0, 0)$ and $(1, 0)$ by 2-tiles. If we instead wish to tile $(0, 0)$ by a 1-tile, we may slide the two tiles at $(0, 0)$ upwards while keeping all other points fixed until the 1-tile falls into the connector interval at $(0, 0)$ in A and the 2-tile falls out. (Since we are keeping the tiles at $(1, 0)$ fixed, the tiles at $(0, 0)$ will also be "compressed" closer together as we slide them upward. But this is an artifact of the picture, since in the actual structures there is no notion of distance, just relative ordering.) Our freedom

to move the tile set at $(0, 0)$ independently of the tile set at $(1, 0)$ stems from the fact that they form disjoint $<_2$ -intervals.

Also note that it is crucial that given an enumeration of the coordinates, the connector intervals in A and the tile sets in B occur in the same $<_2$ -order. For if A were unchanged but the $(0, 0)$ -tile set appeared $<_2$ -above the $(1, 0)$ -tile set in B , then it would be impossible for both the connector intervals in A to capture tiles from the corresponding tile set. In the canonical models, the disjointness and ordering are ensured by point (4), requiring that the connector intervals and tile set intervals are $<_2$ -increasing when ordered antilexicographically.

4 Constraints

For a given string tiling problem \mathcal{T} , we now describe the forbidden substructures in our class $\mathcal{P}_{\mathcal{T}}$, which will impose constraints on arbitrary 3-dimensional permutations in our class and how they may be jointly embedded. We will give some discussion of the constraints before listing them, dividing them into three groups.

The first group of constraints is meant to ensure that when we attempt to perform $<_1$ -joint embedding on the canonical models, we must produce a solution to \mathcal{T} . This includes Constraints 7, 8, and 12. In §2.2, we noted that we would wish our constraints to force a grid point to be tiled using a tile from a tile set with the same coordinates. However, as we are forbidding a *finite* number of finite structures, our constraints must have a *local* character. Since determining the coordinates of a grid point requires walking back to the origin, and thus looking at an unbounded number of vertices, we cannot use our constraints to directly say that a grid point with given coordinates should be tiled using the tile set with the same coordinates. Instead, we will start the tiling at the origin (Constraint 7), and then propagate it by local constraints (Constraint 8). Finally, Constraint 12 ensures this tiling gives a valid solution of \mathcal{T} .

The next group of constraints is meant to ensure that elements of $\mathcal{P}_{\mathcal{T}}$ look like the canonical models in certain ways. This includes Constraints 1–6. These ensure that the origin, path, and grid points/tile sets actually encode something grid-like. They also enforce some conventions we chose for the canonical models, such as that the grid points/tile sets should be antilexicographically increasing in $<_2$. We would like to demand even more from the structures in $\mathcal{P}_{\mathcal{T}}$: for example, we would like every path to have an origin point, or every grid point to have coordinates from a path. However, we cannot enforce such “totality” conditions since $\mathcal{P}_{\mathcal{T}}$ must be closed under substructure, so we must allow for partial structures.

In addition to allowing for partial grids, we don’t impose any constraints to ensure the uniqueness of grids or paths, so a structure in $\mathcal{P}_{\mathcal{T}}$ can have multiple distinct paths each with its own grid attached. While we could enforce uniqueness by forbidding substructures, this would require us to make points in distinct factors equal when performing joint embedding (for example, if we demand a unique path origin, then if there is a path origin in each factor, they must be identified with each other). While this may be viable, it seemed as though it would greatly increase the complexity of the arguments.

Our final group of constraints is concerned with the intersection of special intervals, primarily to control the interaction between multiple copies of a grid, some of which may be partial grids. For example, if we have multiple grids of connector intervals in a single structure, all connector intervals with given coordinates (x, y) should intersect, otherwise we could not insert a given tile into all of them when performing joint embedding. This group, which will play a more technical role, includes Constraints 7–11.

The reason for the overlap between the first and last groups is that forcing tilings is essentially a special case of forcing the intersection of special intervals. If I_G is a connector interval and I_T a tile interval, then

the grid point of I_G is tiled by a tile from I_T if and only if I_G and I_T intersect and $I_G <_{1,3} I_T$.

Given a string tiling problem \mathcal{T} , we now define a class $\mathcal{P}_{\mathcal{T}}$ of 3-dimensional permutations by forbidding substructures to enforce the constraints below. Some constraints are initially described informally, with the formal description of the forbidden substructures nested below. For a few other constraints, we describe the forbidden substructures after the entire list.

1. Path points have at most 1 predecessor and at most 1 successor.
2. Path origins have no predecessor.
3. Special intervals are coordinatized by at most one pair of points.
4. Path points, and their associated copies of E_P^i (or E_O^i), are $<_2$ -below all copies of E_G^0 and E_T^1 .
5. Special intervals are antilexicographically increasing in $<_2$ with respect to their coordinates.
 - (a) Let I, I' be a pair of special intervals, with I' a horizontal or vertical successor of I . Then $I <_2 I'$.
 - (b) Let I, I' be a pair of special intervals. Suppose that I' is on the y -axis, and I has a horizontal predecessor I_{hp} with $I_{hp} <_2 I'$. Then $I <_2 I'$.
6. No point can belong to a copy of both a 0-superscripted and 1-superscripted element of \mathcal{E}
7. All special intervals corresponding to grid-origins or tile-origins intersect. Furthermore, if I^0 corresponds to a grid-origin and I^1 to a tile-origin and $I^0 <_1 I^1$, then $I^0 <_3 I^1$.
8. Two special intervals must intersect if their respective predecessors intersect.

Let I, I' be special intervals.

- (a) Suppose I is on neither the x nor y -axis. Suppose I has horizontal predecessor I_{hp} and vertical predecessor I_{vp} , and I' has horizontal predecessor I'_{hp} and vertical predecessor I'_{vp} . If I_{hp} intersects I'_{hp} and I_{vp} intersects I'_{vp} , then I must intersect I' .
- (b) Suppose I is on the x -axis. Suppose I has horizontal predecessor I_{hp} and I' has horizontal predecessor I'_{hp} . If I_{hp} intersects I'_{hp} , then I must intersect I' .
- (c) Suppose I is on the y -axis. Suppose I has vertical predecessor I_{vp} and I' has vertical predecessor I'_{vp} . If I_{vp} intersects I'_{vp} , then I must intersect I' .

Furthermore, if $I <_1 I'$ in any of the above cases, then $I <_3 I'$.

9. If two special intervals intersect, then their respective predecessors must intersect.

Let I, I' be special intervals (allowing $I = I'$).

- (a) Suppose I is on neither the x nor y -axis. Suppose I has horizontal predecessor I_{hp} and vertical predecessor I_{vp} , and I' has horizontal predecessor I'_{hp} and vertical predecessor I'_{vp} . If I intersects I' , then I_{hp} must intersect I'_{hp} and I_{vp} must intersect I'_{vp} .
- (b) Suppose I is on the x -axis. Suppose I has horizontal predecessor I_{hp} and I' has horizontal predecessor I'_{hp} . If I intersects I' , then I_{hp} must intersect I'_{hp} .

- (c) Suppose I is on the y -axis. Suppose I has vertical predecessor I_{vp} and I' has vertical predecessor I'_{vp} . If I intersects I' , then I_{vp} must intersect I'_{vp} .
10. If I_1, I_2 , and I_3 are special intervals, and I_1 and I_2 intersect I_3 , then I_1 and I_2 intersect.
11. Let I and I' be special intervals that intersect, and suppose I is on the x -axis (resp. y -axis). If I' is coordinatized, then I' is on the x -axis (resp. y -axis).
12. The tiling rules of \mathcal{T} are respected.

We now describe the forbidden substructures corresponding to Constraints 2, 3, 11, and 12.

For Constraint 2, we first forbid any structure consisting of a copy of E_O^i with its root captured by a copy of E_P^i . Note that this corresponds to several forbidden substructures, since there are many ways the points of the copy of E_O^i can interleave with those of the copy of E_P^i in $<_2$, while still having the root of E_O^i be captured. This will be true in all other constraints, i.e. the forbidden configurations we describe will correspond to many forbidden substructures, depending on how the points of the copies of the elements of \mathcal{E} interleave. For Constraint 2, we also forbid any structure consisting of a copy of E_O^i with its root captured by another copy of E_O^i .

For Constraint 3, we forbid any structure consisting of two copies of E_X^i (or of E_Y^i) intersecting at their root. We also forbid any copy of E_X^i or E_Y^i capturing two points that are in P^i .

For Constraint 11, we forbid any structure consisting of the intervals I and I' as described, where the x -coordinate (resp. y -coordinate) of I is in O^i and the x -coordinate (resp. y -coordinate) of I' is in $P^i \setminus O^i$.

For Constraint 12, suppose \mathcal{T} forbids a tile of type j at distance d to the right of (resp. above) a tile of type i . Then we forbid any substructure consisting of $g, g' \in G^0$ such that g' is a d -fold horizontal (resp. direct vertical) successor of g (note that this implies the existence of suitable copies of E_G^0, E_X^0, E_Y^0 , and E_P^0), $t \in T_i^1, t' \in T_j^1$ such that t' is the d -fold horizontal (resp. direct vertical) successor of t , and finally g is tiled by t and g' is tiled by t' .

5 Weak coordinates

When we perform joint embedding on two structures A and B , where A contains a G^0 -grid, and B a grid of tile-sets, Constraints 7 and 8 will force that the connector intervals in the G^0 -grid in A are tiled using points from B , since tiling is a particular case of special intervals intersecting, as discussed before the constraint list. However, connector intervals may be forced to capture tiles for other reasons.

Consider the following scenario. There is a connector interval $I \subset A$ that is part of the G^0 -grid and another connector interval $I' \subset A$ that is part of another G^0 -grid that is missing a grid-origin. Constraints 7 and 8 will not force us to tile I' . However, it may be that the endpoints of I are $<_2$ -between the endpoint of I' , so by tiling I we must also inadvertently tile I' . Then, if I' has successors in its own partial G^0 -grid, Constraint 8 takes effect and we may be forced to tile them as well.

We see that in addition to the tiling of a connector interval being forced by the usual propagation along coordinate paths, the tiling can also be forced due to intersection properties, and then propagate as usual. Thus, in addition to considering a special interval to have coordinates (x, y) if it is coordinatized by the x^{th} and y^{th} points on a path with a path-origin, we will also want to consider all special intervals that intersect such intervals to have coordinates (x, y) .

Definition 5.1 Given a special interval I , we say I is weakly coordinatized by $(x, y) \in \mathbb{N}^2$ if one of the following cases holds.

1. $(x, y) = (0, 0)$: I is, or intersects, a grid-origin or tile-origin.
2. $x = 0, y \neq 0$: I has a vertical predecessor weakly coordinatized by $(0, y - 1)$ or intersects some special interval I' with such a predecessor.
3. $x \neq 0, y = 0$: I has a horizontal predecessor weakly coordinatized by $(x - 1, 0)$ or intersects some special interval I' with such a predecessor.
4. $x, y \neq 0$: I has a horizontal predecessor weakly coordinatized by $(x - 1, y)$ and a vertical predecessor weakly coordinatized by $(x, y - 1)$, or intersects some special interval I' with such predecessors.

We say a point is weakly coordinatized by (x, y) if it is an endpoint of some special interval, i.e. either part of a tile set or one of the two $<_2$ -least points of E_G^0 , weakly coordinatized by (x, y) .

In the definition above, instead of specifying that I has weak coordinates (x, y) if it intersects an appropriate I' , we might have only required that there is a chain of special intervals ($I_0 = I, I_1, \dots, I_n = I'$) with each I_i intersecting I_{i+1} . But then by Constraint 10, we would have I intersects I' , so this does not yield anything new.

We will now show that several properties enforced by our constraints for our earlier notion of coordinates will also hold for weak coordinates.

Lemma 5.2 *The weak coordinates of special intervals are unique.*

Proof: Suppose I has weak coordinates (x, y) and (x', y') . First, suppose $(x, y) = (0, 0)$. Then I must intersect a grid origin or path origin J (allowing $I = J$), and also I intersects a special interval J' (allowing $I = J'$) such that J' has predecessor(s) with weak coordinates $(x' - 1, y')$ and/or $(x', y' - 1)$, so J' has coordinates on a path. By Constraint 10, J and J' intersect. By Constraint 11, the x and y -coordinates of J' must be path origins, and so cannot have predecessors, which is a contradiction.

Now suppose $(x, y), (x', y') \neq (0, 0)$, with $(x, y) <_{antilex} (x', y')$. We will further suppose $x, x' \neq 0, y, y' \neq 0$, although we will return to these cases afterward. By induction, we may assume all special intervals with weak coordinates antilexicographically less than (x, y) have unique weak coordinates.

Then we may find special intervals J and J' (possibly equal to I) such that the following hold.

- (i) J and J' intersect I , and thus intersect each other.
- (ii) J has a horizontal predecessor weakly coordinatized by $(x - 1, y)$ and a vertical predecessor weakly coordinatized by $(x, y - 1)$.
- (iii) J' has a horizontal predecessor weakly coordinatized by $(x' - 1, y')$ and a vertical predecessor weakly coordinatized by $(x', y' - 1)$.

As J and J' intersect, by Constraint 9 the horizontal predecessor of J must intersect that of J' , and similarly for vertical predecessors. By induction, we may assume the predecessors of J and J' have unique weak coordinates. Thus we have $x = x', y = y'$.

In the case $y = 0$ (the case $x = 0$ is similar), we must also have that $y' = 0$ by Constraint 11. We then only get horizontal predecessors for J and J' , but we may still finish as in the previous case. \square

Lemma 5.3 *All special intervals weakly coordinatized by (x, y) intersect.*

Proof: We proceed by antilexicographic induction on (x, y) . If $(x, y) = (0, 0)$, then this is immediate from Constraint 7.

Otherwise, assume $x, y \neq 0$ (as in Lemma 5.2, these cases just require using Constraint 11 and a single predecessor), and let I_1, I_2 have weak coordinates (x, y) . Then I_1 intersects a special interval I'_1 such that I'_1 has a horizontal predecessor weakly coordinatized by $(x - 1, y)$ and a vertical predecessor weakly coordinatized by $(x, y - 1)$, and I_2 similarly intersects some interval I'_2 . By induction, the respective predecessors intersect. Thus by Constraint 8, I'_1 and I'_2 intersect, and so I_1 and I_2 intersect by Constraint 10. \square

Corollary 5.4 (1) *Suppose a is the endpoint of a special interval I and is $<_2$ -between 2 points weakly coordinatized by (x, y) . Then a is weakly coordinatized by (x, y) .*

(2) *All 1-tiles weakly coordinatized by (x, y) are $<_2$ all 2-tiles weakly coordinatized by (x, y) .*

(3) *Suppose I is weakly coordinatized by (x, y) . If $x, y \neq 0$ and I_{hp} and I_{vp} are horizontal and vertical predecessors of I , then I_{hp} is weakly coordinatized by $(x - 1, y)$ and I_{vp} by $(x, y - 1)$. If $y = 0$ (resp. $x = 0$), the same holds, but only with I_{hp} (resp. I_{vp}).*

Proof: (1) Suppose a is $<_2$ -between b, c weakly coordinatized by (x, y) . If b, c belong to the same copy of an element of \mathcal{E} , then that copy intersects I , and we are done by Lemma 5.3. If b, c belong to different copies of elements of \mathcal{E} , their respective special intervals intersect each other by Lemma 5.3, and so intersect I , and we are again done.

(2) If not, there would be a pair of non-intersecting tile-intervals weakly coordinatized by (x, y) , contradicting Lemma 5.3.

(3) If not, the weak coordinates of I would not be unique, contradicting Lemma 5.2. \square

Lemma 5.5 *Suppose I is weakly coordinatized by (x, y) , I' is weakly coordinatized by (x', y') , and $(x, y) <_{antilex} (x', y')$. Then $I <_2 I'$.*

Proof: Fix I with weak coordinates (x, y) and I' with weak coordinates (x', y') . By induction, it is sufficient to consider the cases $(x', y') = (x + 1, y)$ and $(x', y') = (x', y + 1)$.

Claim 1 *Let I, I', J, J' be special intervals. Suppose I intersects I' , J intersects J' , and $I' <_2 J'$. Then $I <_2 J$.*

Proof of Claim: Suppose not. Then I must intersect J . But then by Constraint 10, I' must intersect J' . \diamond

First assume $(x', y') = (x + 1, y)$. Then I' intersects some interval J' with a horizontal predecessor J'_{hp} weakly coordinatized by (x, y) , which in turn intersects I by Lemma 5.3. By Constraint 5(a), we have $J'_{hp} <_2 J'$, and so $I <_2 I'$ by the Claim. By induction, we get the same result for $(x', y') = (x + i, y), i > 0$.

The case $(x', y') = (x', y + 1)$ is similar, though more involved. Then, I' intersects some interval J' with a vertical predecessor J'_{vp} weakly coordinatized by (x', y) . If $x < x'$, then by the previous case, $I <_2 J'_{vp}$, and if $x = x'$ then I intersects J'_{vp} by Lemma 5.3. As $J'_{vp} <_2 J'$ by Constraint 5(a), the Claim gives $J'_{vp} <_2 I'$, and so $I <_2 I'$.

So suppose $x = x' + i$, $i \geq 0$. It suffices to consider the case $x' = 0$, since increasing x' only increases the $<_2$ -position of I' , by the first case. We proceed by induction on i , with the case $i = 0$ handled above. We get J'_{vp} as above, and similarly get that I intersects some interval J with a horizontal predecessor J_{hp} weakly coordinatized by $(x' + (i - 1), y)$. By induction, $J_{hp} <_2 I'$. Then by Constraint 5(b), $J <_2 I'$, so the Claim gives $I <_2 I'$. \square

6 Reductions

6.1 Reductions with the $<_1$ -JEP

We begin with the easy direction, that if $\mathcal{P}_{\mathcal{T}}$ has the $<_1$ -JEP then performing joint embedding on our canonical models will encode a solution to \mathcal{T} .

Lemma 6.1 *Let \mathcal{T} be a string tiling problem, and $\mathcal{P}_{\mathcal{T}}$ the corresponding 3-dimensional permutation class. If $\mathcal{P}_{\mathcal{T}}$ has the $<_1$ -JEP, then \mathcal{T} has a solution.*

Proof: Let $A_{<_1}^*$ and $B_{<_1}^*$ be the canonical models from §3.2. Then $A_{<_1}^*, B_{<_1}^* \in \mathcal{P}_{\mathcal{T}}$, so we can apply the $<_1$ -JEP yielding $C_{<_1}^*$. As $A_{<_1}^* <_1 B_{<_1}^*$ there can be no identifications of points between the factors, so we may assume $C_{<_1}^*$ has $A_{<_1}^* \sqcup B_{<_1}^*$ as a base set. Furthermore, by Constraint 7, the grid-origin in $A_{<_1}^*$ must capture some tile from the tile-origin in $B_{<_1}^*$. This then propagates to a tiling of the entire grid in $A_{<_1}^*$ by tiles from the grid in $B_{<_1}^*$ by Constraint 8, while respecting the rules of the tiling problem by Constraint 12. We thus associate to $C_{<_1}^*$ the tiling $\theta(x, y) = i$ if the connector interval associated to the G^0 -point with coordinates (x, y) captures a tile of type i (if it captures tiles of both types, we may pick either). \square

Now we must assume \mathcal{T} has a solution and show $\mathcal{P}_{\mathcal{T}}$ has the $<_1$ -JEP. The difficulties in defining a joint embedding procedure arise from the way in which arbitrary $A, B \in \mathcal{P}_{\mathcal{T}}$ can deviate from the canonical models. We have already mentioned that an element of $\mathcal{P}_{\mathcal{T}}$ can have multiple grids and partial grids, which led us to consider weak coordinates. Another possible issue that arises from this is that a structure can contain an incorrect tiling. For a simple example, consider a tiling problem where the only constraints are that a 2-tile cannot be adjacent to either a 1-tile or a 2-tile, so the unique solution is to tile the whole grid with 1-tiles. However, $A \in \mathcal{P}_{\mathcal{T}}$ can contain a full grid of connector intervals, and a single tile-set at the origin, such that the connector interval at the origin captures the 2-tile. Then if $B \in \mathcal{P}_{\mathcal{T}}$ contains a full grid of tile sets, when we joint-embed A and B , we will have to arrange so that the connector interval at the origin of A captures the 1-tile from the grid of tile-sets in B , while still capturing the 2-tile from the partial grid of tile sets in A , and making the tile origins in A and B intersect. This will be possible, but relies on having only two types of tile, which is why we imposed this restriction.

From Lemma 5.5, we see that the special intervals weakly coordinatized by (x, y) are contained in a particular $<_2$ -interval, separated from all other special intervals with distinct weak coordinates, and that the intervals are antilexicographically increasing in $<_2$. Thus an arbitrary element of $\mathcal{P}_{\mathcal{T}}$ recovers some of the structure of Figure 4, although with possibly many special intervals at a given weak coordinate. As is already done Figure 4, our first step when performing joint embedding will be to put the intervals in the

two factors with the same weak coordinates at roughly the same $<_2$ -level. This is done in the following definition, where we would like to simply say that the bottom coordinates of the $<_2$ -intervals I_A and I_B are set equal with respect to $<_2$, as are the top coordinates. However, as distinct points cannot be equal with respect to $<_2$, the definition is more convoluted.

Definition 6.2 *Let C be a structure equipped with a partial order $<$, and let $A, B \subset C$ be totally $<$ -ordered. Let I_A, I_B be closed $<$ -intervals in A, B . Extending $<$ such that $b_1 < I_A < b_2$ for any $b_1 < I_B < b_2$, and such that $a_1 < I_B < a_2$ for any $a_1 < I_A < a_2$, will be called $<$ -aligning I_A with I_B . Note, this may not be possible, depending on the initial $<$ -configuration.*

Given A, B , we will use the definition in our joint embedding procedure as follows. After taking the disjoint union $C = A \sqcup B$, $<_2$ will be a partial order on C . We will partition A into $<_2$ -intervals $I_{A,i}$ for $i \in \mathbb{N}$, with the condition that if $i < j$ then the $I_{A,i} <_2 I_{A,j}$, and similarly partition B into $<_2$ -intervals $I_{B,i}$. For each i , we will then align $I_{A,i}$ with $I_{B,i}$. This yields a sequence of disjoint increasing $<_2$ -intervals in C , and we will then complete $<_2$ to a linear order on C by completing it on each such interval separately. By the disjointness, the completion in any one interval can be done independently of the completion on other intervals.

Before beginning our next lemma, we repeat that the claim at the beginning of its proof is the reason we use a third linear order in this paper.

Lemma 6.3 *Let \mathcal{T} be a string tiling problem, and $\mathcal{P}_{\mathcal{T}}$ the corresponding 3-dimensional permutation class. If \mathcal{T} has a solution, then $\mathcal{P}_{\mathcal{T}}$ has the $<_1$ -JEP.*

Proof: Let $A, B \in \mathcal{P}_{\mathcal{T}}$, and begin by defining C to be the disjoint union of A and B , so the order relations are only partially defined. Next, complete $<_1$ and $<_3$ in C so that $A <_1 B$ and $A <_3 B$. Thus, all that remains is to determine how points in A are $<_2$ -related to points in B , while satisfying the constraints.

Claim 1 *For any $<_2$ -completion of C , and any $E \subset C$ a copy of some element of \mathcal{E} , either $E \subset A$ or $E \subset B$.*

Proof of Claim: Suppose $a \in E$ and $a \in A$. As $A <_1 B$, all points $a' \in E$ with $a' <_1 a$ are in A . As $<_1 = <_3^{opp}$ on E , for any $a' \in E$ such that $a' >_1 a$ we have $a' <_3 a$; as $B >_3 A$, such a' are also in A . \diamond

We now extend $<_2$ in C so that all copies of E_O^i and E_P^i in B are $<_2$ -below all points in A , and similarly so that all copies of E_O^i and E_P^i in A are $<_2$ -below all copies of E_G^0 and E_T^1 in B (here we use Constraint 4). In Figure 4, this is satisfied if we erase the dividing line. Given this extension, we prove the next claim.

Claim 2 *For any $<_2$ -completion of C , let $E \subset C$ be a copy of some element of \mathcal{E} in one factor. Then E captures no P^i -points in the other factor.*

Proof of Claim: If $E \subset B$, then it captures no points in A , as $A <_1 B$. If $E \subset A$, it captures no P^i -points in B , as all such points are $<_2$ -below all points in A . \diamond

Constraints 1–3 and 5(a) follow immediately from the claims above and the fact that the constraints hold in each factor. Constraint 4 holds by the paragraph before Claim 2, and Constraint 6 holds as we have identified no points.

The remaining constraints concern the relations between special intervals. For each $(x, y) \in \mathbb{N}^2$, we may consider the closed $<_2$ -interval $I_{x,y}^A$, whose endpoints are the $<_2$ -least and greatest points weakly




Fig. 5: An example of joint embedding at (x, y) with $\theta(x, y) = 1$, projected onto \langle_1, \langle_2 . The lines with black endpoints represents connector intervals, while the purple points represent tiles.

coordinatized by (x, y) in A , and similarly $I_{x,y}^B$. By Lemma 5.5, in each factor these intervals are non-overlapping and antilexicographically increasing with respect to \langle_2 . We may thus \langle_2 -align each $I_{x,y}^A$ with $I_{x,y}^B$, and set $I_{x,y}^X \langle_2 I_{x',y'}^Y$ for $X, Y \in \{A, B\}$ and $(x, y) \langle_{antilex} (x', y')$ (as in Figure 4). From this, it follows that Constraint 5(b) is satisfied. Thus the Constraints 1–6 that we have discussed so far will remain satisfied for any completion of \langle_2 .

We now consider each coordinate-pair (x, y) one at a time, and determine the \langle_2 -order of the points weakly coordinatized by (x, y) independently of what we do at other weak coordinates.

Let $\theta: \mathbb{N}^2 \rightarrow \{1, 2\}$ be a valid tiling. For now, we assume there is a connector interval in A and tile set in B , each weakly coordinatized by (x, y) .

Suppose $\theta(x, y) = 1$. We will work entirely in $I_{x,y}^A$ and $I_{x,y}^B$ (and by Corollary 5.4(1), all special interval endpoints in these intervals are weakly coordinatized by (x, y)). Figure 5 shows an example of the joint embedding procedure at a coordinate (x, y) with $\theta(x, y) = 1$, where we have zoomed in on the special intervals. In Figure 5, in A there is a connector interval capturing a 2-tile from a tile set and in B there is a connector interval capturing a 1-tile. In C , the connector interval from A captures the correct tile from B and all the special intervals intersect.

We now describe the general procedure when $\theta(x, y) = 1$. Let I_A be the intersection of all special intervals in $I_{x,y}^A$, and I_B for $I_{x,y}^B$ (these are non-empty by Lemma 5.3, and are shown in Figure 5). We first set all points from A \langle_2 -below all the 2-tiles from B . Note the bottom endpoint of I_B must be \langle_2 -below all the 2-tiles in B , as must all the 1-tiles in B by Corollary 5.4(2). Thus we may set I_A to contain all the 1-tiles from B as well as the bottom endpoint of I_B . This makes all the connector intervals in $I_{x,y}^A$ capture all 1-tiles in $I_{x,y}^B$, and makes I_A intersect I_B (this intersection is shown in Figure 5), so all special intervals in $I_{x,y}^A$ intersect those in $I_{x,y}^B$. Finally, we arbitrarily complete \langle_2 to a linear order, when restricted to points in $I_{x,y}^A \cup I_{x,y}^B$.

The case $\theta(x, y) = 2$ is similar.

If there is no connector interval in A and tile set in B , each weakly coordinatized by (x, y) , the process is simpler. We just intersect I_A with I_B to ensure all the special intervals in A weakly coordinatized by (x, y) intersect all those in B weakly coordinatized by (x, y) .

Because we have made every special interval from $I_{x,y}^A$ intersect every special interval $I_{x,y}^B$, we will satisfy Constraints 7–10, with Constraint 9 additionally using Corollary 5.4(3). We have also tiled every G^0 -point in A weakly coordinatized by (x, y) according to $\theta(x, y)$, and not tiled any G^0 -point in B , and

so will satisfy Constraint 12. Because we have only intersected special intervals on a given axis with those on the same axis, Constraint 11 holds as well.

To finish, we must complete $<_2$ to a linear order while still satisfying Constraints 7–12. By inspection, they will still be satisfied if we do not intersect any further special intervals during the completion. Given a special interval I in one factor and I' in the other that do not yet intersect, we will be forced to intersect I and I' only if either there is another special interval J intersecting both I and I' , or there are special intervals J and J' both intersecting I' with $I <_2$ -between them (or the same configuration with the roles of I and I' reversed). Recall we only intersected special intervals with the same weak coordinates. Since we also intersected all special intervals at a given set of weak coordinates, the first configuration cannot appear. Since our partially-defined structure satisfies Constraint 10, J and J' must intersect each other in the second configuration, and thus I cannot be $<_2$ -between them.

Thus there is some way to complete $<_2$ to a linear order while not intersecting any further special intervals, and any such completion will suffice to finish defining C . \square

6.2 From the $<_1$ -JEP to the JEP

We first describe why we initially restricted ourselves to the $<_1$ -JEP. Note that our definition of capture and the final parts of Constraint 7 and 8 are asymmetric with respect to $<_1$. When jointly embedding our canonical models $A_{<_1}^*$ and $B_{<_1}^*$, if we were not forced to put $A_{<_1}^* <_1 B_{<_1}^*$, we could trivially jointly embed them by putting $A_{<_1}^* >_1 B_{<_1}^*$. But then no connector intervals in $A_{<_1}^*$ would capture any tiles in $B_{<_1}^*$, and so this would not encode a solution to the tiling problem. One could try to remove these asymmetries, but then the JEP will fail.

In order to remove the requirement of $<_1$ -JEP from Lemma 6.1, we slightly adjust the class $\mathcal{P}_{\mathcal{T}}$ we are working in. For each 0-superscripted element of \mathcal{E} , we introduce a corresponding 2-superscripted element to \mathcal{E} from \mathcal{A} , and for each 1-superscripted element of \mathcal{E} we introduce a corresponding 3-superscripted element to \mathcal{E} from \mathcal{A} , thus doubling the size of \mathcal{E} . We define the corresponding unary predicates as before.

The idea is that 2-superscripted elements should behave like 0-superscripted ones, and 3-superscripted elements like 1-superscripted ones, with the exception that 0-superscripted grids should be tiled by 1-superscripted tiles while 2-superscripted grids should be tiled by 3-superscripted tiles. We will also use $<_2$ to separate the 0, 1-superscripted elements from 2, 3-superscripted elements.

Thus, given a string tiling problem \mathcal{T} , we define a 3-dimensional permutation class $\mathcal{Q}_{\mathcal{T}}$ as follows. We use all the constraints from $\mathcal{P}_{\mathcal{T}}$, and then duplicate those constraints replacing 0-superscripted and 1-superscripted predicates with 2-superscripted and 3-superscripted predicates, respectively.

We also add the following constraints.

- (6*) Constraint 6 is replaced by a constraint forbidding the identification of any points from 2 distinctly-superscripted elements of \mathcal{E} .
- (13) All copies of $\{0, 1\}$ -superscripted elements of \mathcal{E} must be $<_2$ -below all copies of $\{2, 3\}$ -superscripted elements of \mathcal{E} .

Lemma 6.4 *Let \mathcal{T} be a string tiling problem, and $\mathcal{Q}_{\mathcal{T}}$ the corresponding 3-dimensional permutation class. If \mathcal{T} has a solution, then $\mathcal{Q}_{\mathcal{T}}$ has the JEP.*

Proof: Fix a tiling $\theta: \mathbb{N}^2 \rightarrow \{1, 2\}$. Given A, B in our new class, split both into 2 $<_2$ -intervals so that the lesser interval contains all copies of $\{0, 1\}$ -superscripted elements of \mathcal{E} , and the greater interval contains

all copies of $\{2, 3\}$ -superscripted elements of \mathcal{E} . We may then apply the joint embedding procedure of Lemma 6.3 separately to the pair of $<_2$ -lesser intervals and the pair of $<_2$ -greater intervals. \square

In the following lemma, we weaken the $<_1$ -JEP from earlier to simply the JEP. This is done by adjusting the canonical models so that we must perform the $<_1$ -JEP with either a copy of our earlier canonical models, or with a copy of the earlier canonical models using $\{2, 3\}$ -superscripted elements instead of $\{0, 1\}$ -superscripted elements.




Fig. 6: The canonical models in $\mathcal{Q}_{\mathcal{T}}$, projected onto $<_1, <_2$.

Lemma 6.5 *Let \mathcal{T} be a string tiling problem, and $\mathcal{Q}_{\mathcal{T}}$ the corresponding 3-dimensional permutation class. If $\mathcal{Q}_{\mathcal{T}}$ has the JEP, then \mathcal{T} has a solution.*

Proof: We describe our new canonical models, which are pictured in Figure 6. Let A_0 be as $A_{<_1}^*$ in Lemma 6.1 and B_3 be as $B_{<_1}^*$ in Lemma 6.1 but with 3-superscripted elements of \mathcal{E} instead of 1-superscripted elements of \mathcal{E} . Let $A^* = A_0 \sqcup B_3$, with $A_0 <_{1,2,3} B_3$.

Let A_2 be as $A_{<_1}^*$ in Lemma 6.1 but with 2-superscripted elements of \mathcal{E} instead of 0-superscripted elements of \mathcal{E} and B_1 be as $B_{<_1}^*$ in Lemma 6.1. Let $B^* = A_2 \sqcup B_1$, with $A_2 <_{1,3} B_1$ and $B_1 <_2 A_2$.

In A^* , as $<_1$ and $<_3$ agree between A_0 and B_3 , any copy of $E \in \mathcal{E}$ that occurs must be contained either in A_0 or in B_3 . Similarly in B^* , any copy of $E \in \mathcal{E}$ that occurs must be contained either in A_2 or in B_1 .

As $A_{<_1}^*, B_{<_1}^*$ in Lemma 6.1 were in $\mathcal{P}_{\mathcal{T}}$, A^*, B^* will be in $\mathcal{Q}_{\mathcal{T}}$. If $\mathcal{Q}_{\mathcal{T}}$ has the JEP, there is some C^* embedding A^*, B^* .

By Constraint 6*, C^* must contain $A^* \sqcup B^*$. Suppose in C^* that $A_0 <_1 B_1$. Then as in Lemma 6.1, we must produce a tiling. If we don't have $A_0 <_1 B_1$ in C^* , then it must be that $A_2 <_1 B_3$, and again we must produce a tiling as in Lemma 6.1. \square

Theorem 6.6 *There is no algorithm that, given a finite set of forbidden 3-dimensional permutations, decides whether the corresponding 3-dimensional permutation class has the JEP.*

Proof: Given a string tiling problem \mathcal{T} , consider the 3-dimensional permutation class $\mathcal{Q}_{\mathcal{T}}$. By Lemmas 6.4 and 6.5, \mathcal{T} has a solution if and only if $\mathcal{Q}_{\mathcal{T}}$ has the JEP. As the string tiling problem is undecidable, we are finished. \square

Corollary 6.7 *The JEP is undecidable for n -dimensional permutation classes with finitely many forbidden n -dimensional permutations, for $n \geq 3$*

Proof: We have already shown this for $n = 3$, so fix $n > 3$. To any 3-dimensional pattern class \mathcal{C} , we can associate an n -dimensional permutation class $L(\mathcal{C})$ whose constraints are all expansions of the constraints from \mathcal{C} to n orders. Also, given any n -dimensional permutation, we may consider its reduct to the first 3 orders. Both of these operations preserve that there are only finitely many forbidden n -dimensional permutations.

We claim that \mathcal{C} has the JEP if and only if $L(\mathcal{C})$ has the JEP. Suppose $L(\mathcal{C})$ has the JEP. Given $A, B \in \mathcal{C}$, we may expand them to structures in $L(\mathcal{C})$, jointly embed the expansions, and then take the reduct, giving a joint embedding of A, B . Now suppose \mathcal{C} has the JEP. Given $A, B \in L(\mathcal{C})$ we may jointly embed their reducts, and any expansion of the result will give a joint embedding of A and B . \square

7 Concluding Remarks

We finish by discussing the obstructions to adapting this proof to permutation classes. As mentioned before, the main issue is the loss of an easy proof for the claim at the beginning of Lemma 6.3. If simply taking the projection of our 3-dimensional joint embedding procedure to the first 2 orders, transitivity will force us to produce many configurations we do not intend to. For example, consider the following situation. Let F be a forbidden permutation, and suppose $F = F_1 \sqcup F_2$ with $F_1 <_{1,2} F_2$ (more elaborate constructions can remove this requirement). Suppose we are performing the $<_1$ -JEP on A, B , and there are $a \in A$ and $b \in B$ such that we must set $a <_2 b$. We may instead consider the structure A' formed from A by placing $F_1 <_2 a$ and B' formed from B by placing $b <_2 F_2$. If A' and B' are still in our permutation class, then when jointly embedding them, transitivity will force $F_1 <_{1,2} F_2$, and so we will create a copy of F .

A perhaps more basic manifestation of the difficulty in ruling out unintended configurations in permutation classes is the question of how to represent arbitrarily large grids, and thus the canonical models, in a class $\mathcal{P}_{\mathcal{T}}$ such that $\mathcal{P}_{\mathcal{T}}$ is not the class of all permutations.

7.1 Acknowledgments

I thank Gregory Cherlin for many discussions on the material in this paper, and the referee for several corrections and suggestions for improving the presentation.

References

- R. Berger. *The undecidability of the domino problem*. Number 66 in Memoirs of the American Mathematical Society. American Mathematical Society, 1966.
- M. Bóna and D. A. Spielman. An infinite antichain of permutations. *The Electronic Journal of Combinatorics*, 7:Paper N2, 2000.
- S. Braunfeld. The undecidability of joint embedding and joint homomorphism for hereditary graph classes. *Discrete Mathematics & Theoretical Computer Science*, 21, 2019.
- G. Cherlin. Forbidden substructures and combinatorial dichotomies: WQO and universality. *Discrete Mathematics*, 311(15):1543–1584, 2011.

- S. Garrabrant and I. Pak. Pattern avoidance is not P-recursive. *arXiv preprint arXiv:1505.06508*, 2015.
- V. Jelínek. Ramsey-type and amalgamation-type properties of permutations. In *Surveys in Combinatorics 2017*, pages 272–311. Cambridge University Press, 2017.
- M. M. Murphy. *Restricted permutations, antichains, atomic classes and stack sorting*. PhD thesis, University of St Andrews, 2003.
- N. Ruškuc. Decidability questions for pattern avoidance classes of permutations. In *Third International Conference on Permutation Patterns, Gainesville, Fla, 2005*. Available at www-groups.mcs.st-and.ac.uk/~nik/talks/cambridge04.pdf.
- V. Vatter. Permutation classes. In *Handbook of Enumerative Combinatorics*, pages 777–858. Chapman and Hall/CRC, 2015.
- S. D. Waton. *On permutation classes defined by token passing networks, gridding matrices and pictures: three flavours of involvement*. PhD thesis, University of St Andrews, 2007.