Graphical Abstract

Algorithmic coincidence classification of mesh patterns

Christian Bean, Bjarki Gudmundsson, Tomas Ken Magnusson, Henning Ulfarsson

Highlights

Algorithmic coincidence classification of mesh patterns

Christian Bean, Bjarki Gudmundsson, Tomas Ken Magnusson, Henning Ulfarsson

- We complete the coincidence classification of mesh patterns of length 3, meaning we are able to answer whether two such patterns are avoided by the same permutations or not.
- We show how techniques developed for the coincidence classification can be used to count the number of permutations avoiding classical permutation patterns.

Algorithmic coincidence classification of mesh patterns

Christian Bean[a], Bjarki Gudmundsson[a], Tomas Ken Magnusson[a], Henning Ulfarsson[a]

[a]*Department of Computer Science, Reykjavik University, Menntavegi 1, Reykjavik, 102, Iceland*

Abstract

Permutations have connections to other mathematical objects such as Schubert varieties, sorting networks, and genome rearrangements. Often the connection is described in terms of patterns that are absent from the permutations. There can be ambiguity in this description, in the sense that the same subset of permutations can be defined with two different patterns. This is called *coincidence* and we focus on the coincidence of mesh patterns, one of the most descriptive version of patterns in permutations. We review and extend previous results on coincidence of mesh patterns. We introduce the notion of a force on a permutation pattern and apply it to the coincidence classification of mesh patterns, completing the classification up to size three. We also show that this concept can be used to enumerate classical permutation classes.

Keywords: Permutation patterns, Wilf-classification, enumeration
2000 MSC: 05A05, 05A15

Permutation patterns have been studied since the beginning of the 20$^{\text{th}}$ century, starting with MacMahon [1], who considered the union of two decreasing sequences of points. Interest in modern day study was sparked by Simion and Schmidt [2]. Classical permutation patterns have been generalized to *vincular* patterns by Babson and Steingrímsson [3], to *bivincular* patterns by Bousquet-Mélou et al. [4] and to *barred* patterns by West [5]. The focus of this paper will be on *mesh* patterns, introduced by Brändén

Email addresses: `christianbean@ru.is` (Christian Bean), `bjarki.agust@gmail.com` (Bjarki Gudmundsson), `tomkenmag@gmail.com` (Tomas Ken Magnusson), `henningu@ru.is` (Henning Ulfarsson)

and Claesson [6], and subsume the prior generalizations. In particular, we study the classification of these patterns based on *coincidence*, an equivalence relation derived from the avoidance sets of the patterns. A related subject is the classification of patterns in terms of *Wilf-equivalence* where the size of the avoidance sets determines the relation. This line of inquiry for mesh patterns was started by Hilmarsson et al. [7], who provided sufficient conditions for the coincidence of mesh patterns with the so-called Shading lemma. This was further generalized by Claesson, Tenner, and Ulfarsson [8] in the Simultaneous Shading lemma. The relationship between the avoidance sets of mesh patterns and classical patterns has been studied by Tenner [9, 10], who determined which mesh patterns are coincident with classical patterns.

We will review these earlier results, and extend them using the notion of a *force* on a permutation pattern. This will culminate in the Shading Algorithm, which is powerful enough to coincidence classify the set of mesh patterns of size 3, except one case which we do by hand. We show how knowing the coincidence of mesh patterns can be used to enumerate permutation sets avoiding a classical pattern. Furthermore, we show how the concept of a force can be applied directly to that problem.

1. Preliminaries

A *permutation* on a set A is a bijection from A to itself. In this paper we always have $A = [1, n] = \{1, \ldots, n\}$ for some non-negative integer n. We denote the set of all permutations of size n as S_n and we write $\pi \in S_n$ as the word $\pi(1)\pi(2)\cdots\pi(n)$. Two sequences of integers $a_1a_2\cdots a_k$ and $b_1b_2\cdots b_k$ are *order isomorphic* if $a_i < a_j$ if and only if $b_i < b_j$ for all $i, j \in [1, k]$. The central definition in the study of permutation patterns is the following.

Definition 1.1. A permutation $\pi \in S_n$ *contains* a permutation $p \in S_k$ if there exists a sequence of indices $1 \leq i_1 < i_2 < \cdots < i_k \leq n$ such that the subsequence $c = \pi(i_1)\pi(i_2)\cdots\pi(i_k)$ of π is order isomorphic to p. The subsequence c is called an *occurrence* of p in π. If such a subsequence does not exist then π *avoids* p. In this context p is called a *classical* (*permutation*) *pattern*.

A permutation $\pi \in S_n$ can be represented graphically by plotting the points $\{(i, \pi(i)) \mid i \in [1, n]\}$ on a grid such that the final figure resembles a *mesh*. The elements $(i, \pi(i))$ are referred to as the *points* of the permutation π. An example is illustrated in the first subfigure of Figure 1.1.

An occurrence of a pattern in a permutation can be viewed in the graphical representation as scaling the indices and the values of the pattern to coincide with points of the permutation, while maintaining the same relative ordering. An example is depicted in Figure 1.1 with the graph of the permutation 42135 and one of its subsequences.

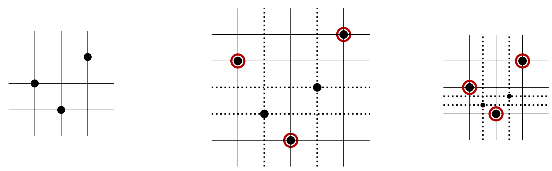

Figure 1.1: The graph of the pattern 213; an occurrence of this pattern within the permutation 42135; and a graph focusing on the occurrence, showing the locations of the remaining points in the permutation.

Using the graphical representation of permutations we recall the definition of *mesh patterns*, which are generalizations of permutation patterns with added restrictions. In the second subfigure of Figure 1.1, the points of the permutation not in the occurrence can be thought of as being mapped into the squares formed by the mesh. We can restrict when points are allowed to map into these squares by shading the mesh.

Definition 1.2 (Brändén and Claesson [6])**.** A *mesh pattern* is an ordered pair $p = (\tau, R)$ where $\tau \in S_k$, and R is a subset of the $(k+1)^2$ unit squares in $[0, k+1]^2$. The set R is the *mesh* (shading) and squares of the mesh are indexed by their lower-left corners; that is, $\lfloor i, j \rceil \in R$ refers to the square $[i, i+1] \times [j, j+1]$. The mesh pattern (τ, R) is depicted graphically by drawing τ as above, and shading the squares of the mesh R. The *size* of the mesh pattern p is k and denoted $|p|$.

Informally, a permutation $\pi \in S_n$ is said to contain a mesh pattern $p = (\tau, R)$ if π has an occurrence of the underlying classical pattern τ such that each of the shaded squares $\lfloor i, j \rceil \in R$ correspond to an empty region in the permutation; see Example 1.3 and [6].

Example 1.3. Consider the mesh pattern $p = (213, \{\lfloor 1, 2 \rceil, \lfloor 2, 2 \rceil, \lfloor 2, 3 \rceil\})$, shown in Figure 1.2. In the permutation 42135 the subsequence 415 is an occurrence of the classical pattern 213 and the squares in the mesh correspond to empty regions in the permutation, as can be seen in the second and third subfigures in Figure 1.2. Note that although the subsequence 215 is an occurrence of the pattern 213 it does not satisfy the requirements of the mesh, as the point 3 in the permutation is in the region corresponding to the shaded square $\lfloor 2,2 \rceil$.

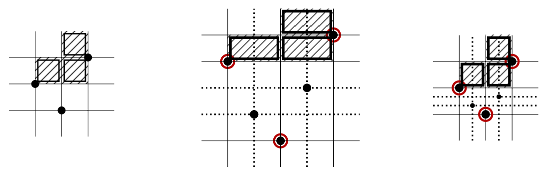

Figure 1.2: The graph of the mesh pattern $(213, \{\lfloor 1,2 \rceil, \lfloor 2,2 \rceil, \lfloor 2,3 \rceil\})$; an occurrence of this pattern within the permutation 42135; and a graph focusing on the occurrence.

The set of permutations of size n that avoid a pattern p is denoted $\mathrm{Av}_n(p)$. We also define $\mathrm{Av}(p) = \bigcup_{n=0}^{+\infty} \mathrm{Av}_n(p)$. Its complement, the set of permutations that contain the pattern p, is denoted $\mathrm{Co}(p)$. In more generality, if B is a set of patterns, we define $\mathrm{Av}_n(B)$ as the set of permutations of size n that avoid all the patterns in B, and $\mathrm{Av}(B) = \bigcup_{n=0}^{+\infty} \mathrm{Av}_n(B)$. If the set B is minimal, i.e., no $B' \subsetneq B$ with $\mathrm{Av}(B') = \mathrm{Av}(B)$ exists, then B is called a *basis*. When all the patterns in B are classical patterns the set $\mathrm{Av}(B)$ is called a *permutation class*.

Given two different classical patterns, p, q, it is never true that $\mathrm{Av}(p) = \mathrm{Av}(q)$. This property does not hold for mesh patterns. For example, the patterns 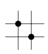 and 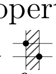 are different mesh patterns that have the same avoiding permutations. In fact, this is an equivalent way of stating that a permutation has an inversion if and only if it has a descent. This leads to the following definition.

Definition 1.4. Two patterns p and q are said to be *coincident* if $\mathrm{Av}(p) = \mathrm{Av}(q)$, denoted $p \asymp q$.

Note that if two mesh patterns p and q are coincident then they necessarily have the same underlying classical patterns. Recently Tannock and Ulfarsson [11, Definition 2.6] defined occurrences of mesh patterns in mesh patterns. Informally, the added restriction is that for a shaded square in the occurring pattern, the corresponding squares in the containing pattern must all be shaded; see Example 1.5 and [11].

Example 1.5. Consider the mesh pattern $p = (213, \{\lfloor 1,2 \rceil, \lfloor 2,2 \rceil, \lfloor 2,3 \rceil\})$, shown in Figure 1.3. In the mesh pattern $m = (42135, \{\lfloor 0,0 \rceil, \lfloor 0,1 \rceil, \lfloor 0,2 \rceil,$

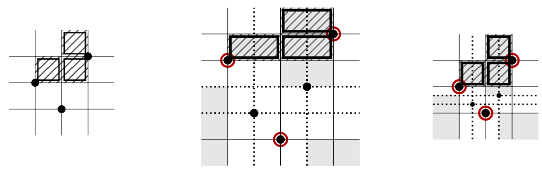

Figure 1.3: The graph of the mesh pattern $(213, \{\lfloor 1,2\rceil, \lfloor 2,2\rceil, \lfloor 2,3\rceil\})$; an occurrence of this pattern within the mesh pattern m, defined in Example 1.5; and a graph focusing on the occurrence.

$\lfloor 1,4\rceil$, $\lfloor 2,4\rceil$, $\lfloor 3,3\rceil$, $\lfloor 3,4\rceil$, $\lfloor 3,5\rceil$, $\lfloor 4,0\rceil$, $\lfloor 4,3\rceil$, $\lfloor 4,4\rceil$, $\lfloor 4,5\rceil$, $\lfloor 5,0\rceil\})$ the subsequence 415 is an occurrence of the classical pattern 213 and the squares in the mesh of p correspond to regions in m that are shaded and do not contain any points, as can be seen in the second and third subfigures in Figure 1.3.

The following remark follows from the previous definitions.

Remark 1.6. If a mesh pattern m contains a mesh pattern p and a permutation π contains m, then π also contains p.

To compare two occurrences of a mesh pattern in a permutation, or in another mesh pattern, we need the following definition.

Definition 1.7. Let $p = (\tau, R)$ and $q = (\sigma, T)$ be mesh patterns. If $u = \tau(u_1) \cdots \tau(u_k)$ and $v = \tau(v_1) \cdots \tau(v_k)$ are occurrences of q in p, then we say that

- u is *above* v with respect to the point $(i, \sigma(i))$ if $\tau(u_i) > \tau(v_i)$,
- u is *below* v with respect to the point $(i, \sigma(i))$ if $\tau(u_i) < \tau(v_i)$,
- u is *left of* v with respect to the point $(i, \sigma(i))$ if $u_i < v_i$, and,
- u is *right of* v with respect to the point $(i, \sigma(i))$ if $u_i > v_i$.

In Figure 1.4 are the occurrences $u = 415$ and $v = 435$ of the mesh pattern $(213, \{\lfloor 1,2\rceil, \lfloor 2,2\rceil, \lfloor 2,3\rceil\})$. The occurrence u is below v and u is left of v with respect to the point $(2, 1)$.

We will need to add points to mesh patterns as is formally defined in Tannock and Ulfarsson [11, Definition 3.2] and illustrated here in Figure 1.5. As in [11] we let $p = (\tau, R)$ be a mesh pattern such that $\lfloor i,j\rceil \notin R$, and $p^{\lfloor i,j\rceil} = (\tau', T')$ be the mesh pattern with a point inserted into the square

$\lfloor i,j \rceil$. We define the following four mesh patterns, which have the same underlying classical pattern as $p^{\lfloor i,j \rceil}$:

$$\begin{aligned} p^{\lfloor i,j \rceil \uparrow} &= (\tau', T' \cup \{\lfloor i, j+1 \rceil, \lfloor i+1, j+1 \rceil\}) \\ p^{\lfloor i,j \rceil \downarrow} &= (\tau', T' \cup \{\lfloor i, j \rceil, \lfloor i+1, j \rceil\}) \\ p^{\lfloor i,j \rceil \leftarrow} &= (\tau', T' \cup \{\lfloor i, j \rceil, \lfloor i, j+1 \rceil\}) \\ p^{\lfloor i,j \rceil \rightarrow} &= (\tau', T' \cup \{\lfloor i+1, j \rceil, \lfloor i+1, j+1 \rceil\}) \end{aligned}$$

Informally, these patterns are obtained by placing the highest, lowest, leftmost, or rightmost point in $\lfloor i,j \rceil$. We collect these mesh patterns in a set

$$p^{\lfloor i,j \rceil \star} = \{p^{\lfloor i,j \rceil \uparrow}, p^{\lfloor i,j \rceil \downarrow}, p^{\lfloor i,j \rceil \leftarrow}, p^{\lfloor i,j \rceil \rightarrow}\}$$

shown in Figure 1.6 for the case $\tau = 21$ and $\lfloor i,j \rceil = \lfloor 1,1 \rceil$.

We record the following easily proven statement for future reference.

Remark 1.8. Let $p = (\tau, R)$ be a mesh pattern such that $\lfloor i,j \rceil \notin R$. A permutation π that contains p either contains $(\tau, R \cup \lfloor i,j \rceil)$ or all of the patterns in $p^{\lfloor i,j \rceil \star}$.

Given a mesh pattern p it is clear that p has an occurrence in p. The indices of this occurrence, properly translated, give an occurrence of p in $p^{\lfloor i,j \rceil a}$, for any $a \in \{\uparrow, \downarrow, \leftarrow, \rightarrow\}$). We call this the *trivial occurrence*. Other occurrences of p in $p^{\lfloor i,j \rceil a}$ are called *non-trivial*. It follows that a non-trivial occurrence contains the inserted point $(i+1, j+1)$.

In Sections 2 and 3 we will review existing methods to prove that two mesh patterns are coincident and extend these methods. The central idea will be the definition of a *force* on a permutation pattern, given in Definition 2.8. This idea will be used to enumerate several permutation classes in Section 4.

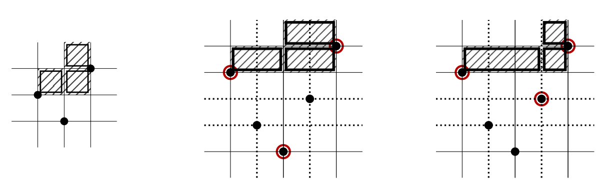

Figure 1.4: The mesh pattern $(213, \{\lfloor 1,2 \rceil, \lfloor 2,2 \rceil, \lfloor 2,3 \rceil\})$ and two occurrences of it in the permutation 42135.

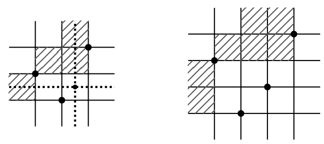

Figure 1.5: Adding a point to the mesh pattern $p = (213, \{\lfloor 0,1\rceil, \lfloor 1,2\rceil, \lfloor 2,2\rceil, \lfloor 2,3\rceil\})$, to create the mesh pattern $p^{\lfloor 2,1\rceil}$.

2. Extending the Shading Lemmas

Several proofs will make use of the following, easily proven, property.

Remark 2.1. Let $p = (\tau, R)$ and $q = (\tau, S)$ be two mesh patterns with the same underlying classical pattern. If $S \subseteq R$ and a permutation π contains the pattern p then it contains q, in other words $\mathrm{Av}(q) \subseteq \mathrm{Av}(p)$.

We start by recalling the Shading Lemma, from Hilmarsson et al. [7, Lemma 11], and give an alternative (sketch of a) proof.

Lemma 2.2 (Shading Lemma)**.** Let (τ, R) be a mesh pattern such that $\tau(i) = j$ and the square $\lfloor i,j\rceil \notin R$. If all of the following conditions are satisfied:

1. The square $\lfloor i-1, j-1\rceil$ is not in R;
2. At most one of the squares $\lfloor i, j-1\rceil$, $\lfloor i-1, j\rceil$ is in R;
3. If the square $\lfloor \ell, j-1\rceil$ is in R (with $\ell \notin \{i-1, i\}$) then the square $\lfloor \ell, j\rceil$ is also in R;
4. If the square $\lfloor i-1, \ell\rceil$ is in R (with $\ell \notin \{j-1, j\}$) then the square $\lfloor i, \ell\rceil$ is also in R;

then the patterns (τ, R) and $(\tau, R \cup \{\lfloor i,j\rceil\})$ are coincident. Symmetric conditions determine if other squares neighboring the point (i,j) can be added to R while preserving the coincidence of the corresponding patterns.

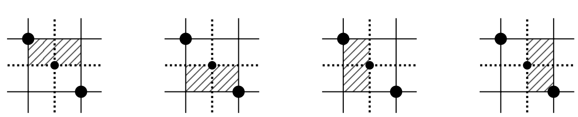

Figure 1.6: The set $p^{\lfloor i,j\rceil\star}$ for $p = 21$ and $\lfloor i,j\rceil = (1,1)$.

The conditions of the above lemma are satisfied for the mesh pattern $p = (12, \{\lfloor 0,2\rceil, \lfloor 1,0\rceil, \lfloor 2,0\rceil, \lfloor 2,1\rceil\})$, shown on the left below, and the square $\lfloor 2,2\rceil$. The lemma therefore implies the coincidence of p with the pattern q on the right.

$p =$ 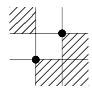 $\asymp$ 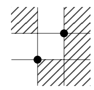 $= q$

The argument used in the original proof relies on replacing the point corresponding to 2 in an occurrence of the pattern with the rightmost (or highest) point in the region corresponding to the square $\lfloor 2,2\rceil$. See Figure 2.1, where we have an occurrence of p in the permutation 12536487 from which we produce an occurrence of q.

The motivation for Lemma 2.3 is an alternative argument for the coincidence of the two patterns. Consider again an occurrence of p in the same permutation, as in Figure 2.2. Out of all the occurrences of p consider the occurrence where the point corresponding to 2 is as far to the right as possible. If the square $\lfloor 2,2\rceil$ is not empty in this occurrence then taking the lowest point in it (for example) as a new 2 would give us another occurrence of p with the point corresponding to 2 further to the right, a contradiction. Therefore the square $\lfloor 2,2\rceil$ is empty and this occurrence of p must also be an occurrence of q.

Lemma 2.3. Let $p = (\tau, R)$ be a mesh pattern with $\lfloor i,j\rceil \notin R$. If a mesh pattern in the set $p^{\lfloor i,j\rceil\star}$ contains a non-trivial occurrence of p, then p and $q = (\tau, R \cup \{\lfloor i,j\rceil\})$ are coincident.

Proof. Take any $m \in p^{\lfloor i,j\rceil\star}$ such that m has a non-trivial occurrence of p, which exists by the premises. We consider the case where $m = p^{\lfloor i,j\rceil\leftarrow}$, as the

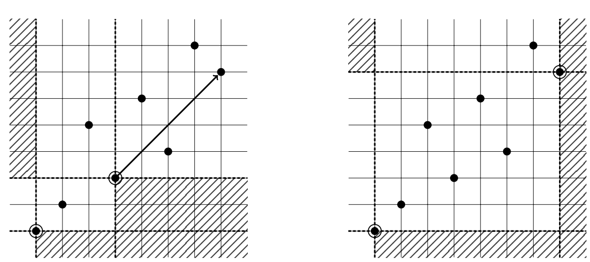

Figure 2.1: Choosing the rightmost point in a region, as in the original proof of the Shading Lemma.

other cases are symmetric to the following argument; see Figure 2.3. In a non-trivial occurrence of p in m, the inserted point $(i+1, j+1)$ corresponds to some point (k, τ_k) in p. The point corresponding to (k, τ_k) in the trivial occurrence is either left or right of $(i+1, j+1)$ in q. We consider the case where it is to the left of $(i+1, j+1)$, i.e., with a lower index, as the other case is analogous.

Take any permutation π that contains an occurrence of p, and consider the occurrence such that the point a corresponding to (k, τ_k) has the highest possible index in π, i.e., is as far to the right as possible. Assume that the region in π that corresponds to the square $\lfloor i, j \rceil$ in this occurrence of p is non-empty. The leftmost point in this region, denoted by b, along with the occurrence of p in π, gives us an occurrence of $m = p^{\lfloor i,j \rceil \leftarrow}$ in π; see Figure 2.4. This implies there is an occurrence of p in π with b corresponding to (k, τ_k). The point b is further to the right, i.e., has a higher index than a, which contradicts the choice of an occurrence of p in π. Hence, our assumption that the region corresponding to $\lfloor i, j \rceil$ was non-empty must be false. Therefore, the region is empty, and this occurrence of p is an occurrence of q as well.

We have shown that if a permutation contains p then it contains q. By Remark 2.1 it follows that if a permutation contains q then it contains p. Therefore p and q are coincident.

To show that the previous lemma implies any coincidence obtained with the Shading Lemma we need the following result.

Lemma 2.4. Let $p = (\tau, R)$ be a mesh pattern such that $\tau(i) = j$ and the square $\lfloor i, j \rceil \notin R$. If conditions (1)–(4) in Lemma 2.2 (Shading Lemma) are satisfied, then some $m \in p^{\lfloor i,j \rceil \star}$ contains a non-trivial occurrence of p.

Proof. We will only consider the case where $\lfloor i, j-1 \rceil$ is shaded, as the other cases are similar. Let $m = p^{\lfloor i,j \rceil \downarrow}$, as depicted in Figure 2.5. By swapping out

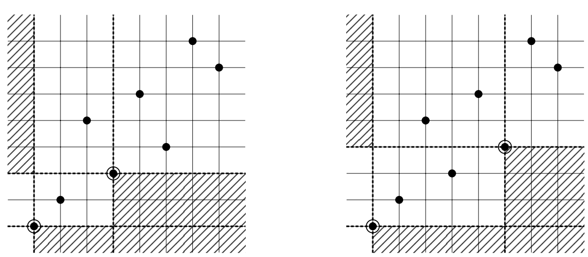

Figure 2.2: If the region is not empty, we can derive a contradiction.

the point (i, j) in the original occurrence of p in m with the inserted point, we get a non-trivial occurrence of the classical pattern τ in m. The conditions (3), (4) and the choice of $m = p^{\lfloor i,j \rceil \downarrow}$ guarantee that all the squares in R are still shaded in this new occurrence, making it a non-trivial occurrence of p in m.

The previous result implies the Shading Lemma (Lemma 2.2) is a consequence of Lemma 2.3. It is then natural to ask if these two lemmas are equivalent, in the sense that they can identify exactly the same coincidences of mesh patterns. In the following example we show that Lemma 2.3 is strictly stronger than the Shading Lemma.

Example 2.5. Consider the mesh patterns

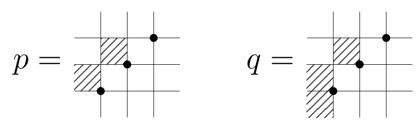

Then $p^{\lfloor 0,0 \rceil \uparrow}$ is the mesh pattern

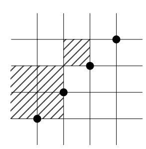

which contains a non-trivial occurrence of p, in the subsequence 123. By Lemma 2.3 this implies that p and q are coincident. However, the conditions of Lemma 2.2 are not satisfied for the square $\lfloor 0, 0 \rceil$ and, hence, the coincidence of these patterns does not follow from that lemma.

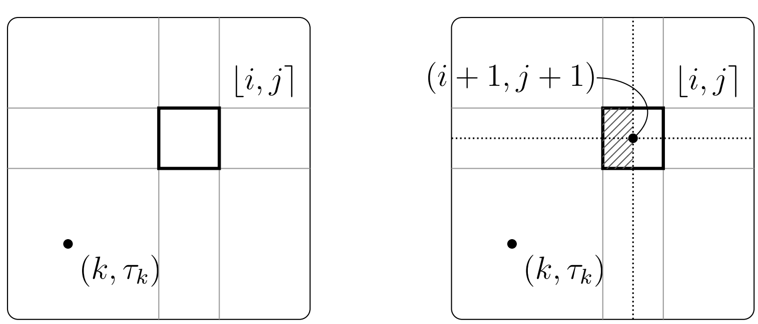

Figure 2.3: On the left is the pattern p. On the right is the pattern $m = p^{\lfloor i,j \rceil \leftarrow}$.

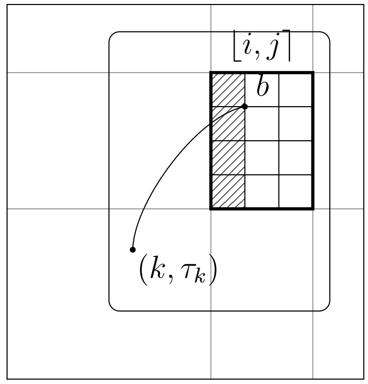

Figure 2.4: The region corresponding to the square $\lfloor i, j \rceil$ in a permutation π, containing a point b.

A previous strengthening of the Shading Lemma was given by Claesson, Tenner, and Ulfarsson [8, Lemma 7.6]:

Lemma 2.6 (Simultaneous Shading Lemma)**.** Let $p = (\tau, R)$ be a mesh pattern. Fix a subsequence G of τ and, for each $g \in G$, let U_g be a square or a pair of adjacent squares that are shadeable[1] from g. Then p is coincident with $(\tau, R \cup S)$, where $S = \bigcup_{g \in G} U_g$.

[1] As defined in Tenner, Claesson, and Ulfarsson [8], a square is *shadeable* if it satisfies the conditions of Lemma 2.2 or any of its symmetries, and a pair of adjacent squares are *shadeable* if they satisfy the conditions of Corollary 7.1 in [8] or any of its symmetries. The conditions of Corollary 7.1 are those required so that the two squares can be shaded by two applications of Lemma 2.2.

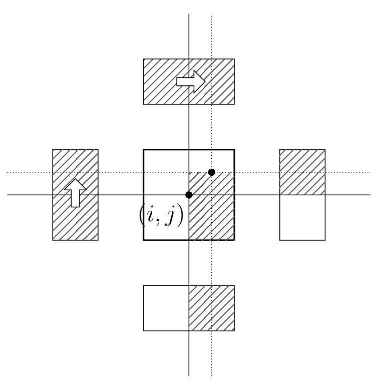

Figure 2.5: Obtaining a new occurrence of p.

An example of a coincidence following directly from the Simultaneous Shading Lemma is the following:

Example 2.7.

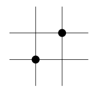 $\asymp$ 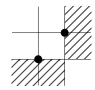

In the pattern on the left, the square pair $\{\lfloor 0, 0\rceil, \lfloor 1, 0\rceil\}$ is shadeable from the point 1 and the square pair $\{\lfloor 2, 1\rceil, \lfloor 2, 2\rceil\}$ is shadeable from the point 2. Thus the coincidence of the patterns follows from the Simultaneous Shading Lemma. This, however, does not follow from Lemma 2.3.

To give a common strengthening of Lemma 2.3 and the Simultaneous Shading Lemma we need some definitions. In Definition 1.7 we compared occurrences of mesh patterns using a point and a direction. We generalize this notion to include multiple points.

Definition 2.8. Given a mesh pattern $p = (\tau, R)$, with $\tau \in S_k$, we define a *force* on it as a tuple of pairs $F = ((\tau_{i_1}, a_1), (\tau_{i_2}, a_2), \dots, (\tau_{i_\ell}, a_\ell))$ where $\ell \in [0, k]$, the indices $i_j \in [1, k]$ are distinct, and $a_j \in \{\uparrow, \downarrow, \leftarrow, \rightarrow\}$ represents the direction we are *forcing* the point τ_{i_j} in. The *size* of the force F is ℓ and denoted $|F|$.

Let p be a mesh pattern with force F. If we have an occurrence $c = c_1 c_2 \cdots c_k$ of p in a permutation π, then for each (τ_{i_j}, a_j) we define the *strength* of the point c_{i_j} with respect to the force F as

$$\text{strength}_F(\pi, c, i_j) = \begin{cases} c_{i_j} & \text{if } a_j = \uparrow \\ -c_{i_j} & \text{if } a_j = \downarrow \\ -\pi^{-1}(c_{i_j}) & \text{if } a_j = \leftarrow \\ \pi^{-1}(c_{i_j}) & \text{if } a_j = \rightarrow \end{cases}$$

Finally, we define the *strength* of an occurrence c of p in a permutation π with respect to the force F as the tuple

$$\text{strength}_F(\pi, c) = (\text{strength}_F(\pi, c, i_1), \dots, \text{strength}_F(\pi, c, i_\ell)).$$

An occurrence c in π is *stronger* than another occurrence c' in π (*with respect to* F) if $\text{strength}_F(\pi, c) > \text{strength}_F(\pi, c')$, in the lexicographical order, otherwise it is *weaker*.

Example 2.9. Consider the pattern $\tau = (1342, \emptyset)$ along with the force $F = ((2, \uparrow), (3, \downarrow))$. In the permutation 2147563 the subsequence 2463 is an occurrence of τ with strength $(3, -4)$, while the subsequence 1563 has strength $(3, -5)$. The first occurrence is stronger with respect to this force.

Lemma 2.10. Let $p = (\tau, R)$ be a mesh pattern with force F, and assume $S = \{s_1, s_2, \dots, s_k\}$ where $S \cap R = \emptyset$. If all the sets $p_1 = (\tau, R)^{s_1\star}, p_2 = (\tau, R \cup \{s_1\})^{s_2\star}, \dots, p_k = (\tau, R \cup \{s_1, s_2, \dots, s_{k-1}\})^{s_k\star}$ contain an occurrence of p that is stronger than the trivial occurrence of p with respect to F, then p and $q = (\tau, R \cup S)$ are coincident.

Proof. Let π be a permutation and let c be an occurrence of p in π which has maximal strength with respect to the force F. Let $i \in \{1, 2, \dots, k\}$ and let c' be a non-trivial occurrence of p in some mesh pattern in p_i that is stronger than the trivial occurrence. The occurrence c' gives rise to an occurrence of p in π which is stronger than c, which is a contradiction. Hence, in the original occurrence c, the region corresponding to the square s_i is empty. Letting i range from 1 to k shows that the regions in π corresponding to all the squares in S are empty, i.e., c is an occurrence of q.

In the previous lemma, the special case where $|F| = 1$ and $k = 1$ is equivalent to Lemma 2.3. To show that Lemma 2.10 can prove any coincidence proven by the Simultaneous Shading Lemma we need the following result.

Lemma 2.11. Let $p = (\tau, R)$ be a mesh pattern. Fix a subsequence G of τ and, for each $g \in G$, let U_g be a square or pair of adjacent squares that are shadeable from g. Then there exists a force F such that $S = \bigcup_{g \in G} U_g$ satisfies the conditions of Lemma 2.10.

Proof. Let $k = |G|$. We define the force $F = ((g_1, a_1), (g_2, a_2), \dots, (g_k, a_k))$ as follows:[2]

$$a_i = \begin{cases} \uparrow & \text{if the square(s) } U_{g_i} \text{ are north of } g_i \\ \downarrow & \text{if the square(s) } U_{g_i} \text{ are south of } g_i \\ \leftarrow & \text{if the square(s) } U_{g_i} \text{ are west of } g_i \\ \rightarrow & \text{if the square(s) } U_{g_i} \text{ are east of } g_i \end{cases}$$

[2]Note that in some cases U_g satisfies many of the cases, in which case we can make an arbitrary choice.

It suffices to show that for each $s_i \in S$, some mesh pattern in $(\tau, R)^{s_i\star}$ contains a non-trivial occurrence of p that is stronger than the trivial occurrence of p. Let U_g be the square (or a pair of squares) corresponding to s_i. Since U_g is shadeable from g there is a mesh pattern in $(\tau, R)^{s_i\star}$ that contains a non-trivial occurrence of p that is stronger than the trivial occurrence of p.

The previous result implies the Simultaneous Shading Lemma (Lemma 2.6) is a consequence of Lemma 2.10. The following example shows that Lemma 2.10 is strictly stronger.

Example 2.12. Consider the two mesh patterns

$p =$ 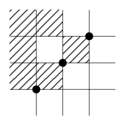 $q =$ 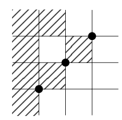

Then $p^{\lfloor 0,0 \rceil \uparrow}$ is the mesh pattern

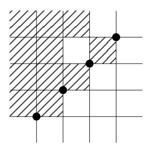

which contains p in the subsequence 123. This is a stronger occurrence of p than the trivial occurrence with respect to the force $F = ((1, \downarrow))$. By Lemma 2.10, p and q are coincident, but p does not satisfy the conditions of Simultaneous Shading Lemma for shading $\lfloor 0, 0 \rceil$.

We introduce one final strengthening of our lemmas, before presenting an algorithm that recursively applies them. Up to now all of the lemmas could be used to prove that two patterns are coincident, i.e., $\mathrm{Av}(q) = \mathrm{Av}(p)$. The next lemma, like Remark 2.1, can be used to prove results of the form $\mathrm{Av}(q) \subseteq \mathrm{Av}(p)$.

Lemma 2.13. Let $p = (\tau, R)$ be a mesh pattern with force F, and $q = (\tau, R')$ be another mesh pattern. Let $S = \{s_1, s_2, \ldots, s_k\}$ where $S = R' \setminus R$. If all the sets $p_1 = (\tau, R)^{s_1\star}, p_2 = (\tau, R \cup \{s_1\})^{s_2\star}, \ldots, p_k = (\tau, R \cup \{s_1, s_2, \ldots, s_{k-1}\})^{s_k\star}$ contain an occurrence of p that is stronger than the trivial occurrence of p, or an occurrence of a pattern that implies an occurrence of q,[3] then containment of p implies containment of q.

[3] For instance, a previous application of the lemma might have shown that a pattern p_i contains a pattern with shadings that form a superset of the shadings of q.

Proof. The proof is analogous to the proof of Lemma 2.10, with one exception. If a pattern p_i is ever reached such that containment of p_i implies the occurrence of q, then p implies the occurrence of p_i, which in turn implies the occurrence of q. ∎

Note that we can prove coincidence of two patterns p and q by applying the above lemma twice: Once to prove that containment of p implies containment of q, and then again in the other direction to prove that containment of q implies the containment of p.

We have already shown that the lemmas are ordered by implication as in Figure 2.6. To get a better idea of the power of these results, we compare

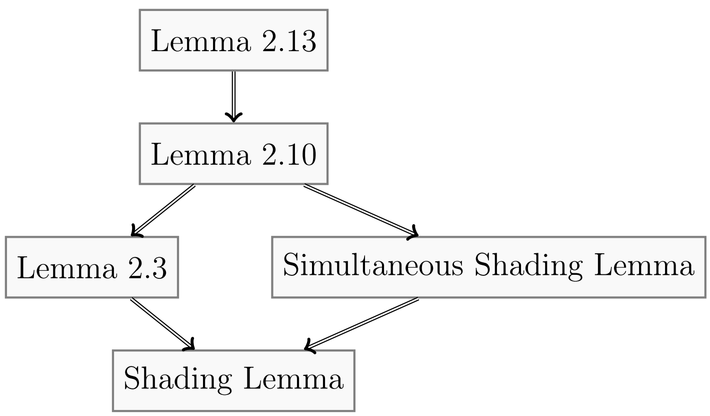

Figure 2.6: Comparison of the lemmas.

them across all mesh patterns with underlying classical patterns 1, 12, 123 and 132.[4] This is a total of 131.600 mesh patterns. Before we apply any of the lemmas we compute $\mathrm{Av}_k(p)$ for $k \leq 10$. This allows us to perform an experimental coincidence classification of these patterns: two patterns p, q are in the same *experimental class* if $\mathrm{Av}_k(p) = \mathrm{Av}_k(q)$ for $k \leq 10$. For any two patterns p and q in different experimental classes, the experimental classification finds a permutation that shows that p and q are not coincident. The number of experimental classes for each underlying pattern is given in Table 1. Some of these classes contain a single mesh pattern, and we call these *resolved* (res.) classes, since a pattern in such a class is not coincident

[4]The remaining patterns of size 2 and 3 are symmetries of these patterns.

to any other pattern. An example of such a class is the class which contains a fully shaded mesh pattern. Such a pattern is avoided by every permutation but the underlying pattern itself. Therefore it is coincident only to itself and the class is a singleton. The remaining classes are said to be *unresolved* (unr.). An unresolved class becomes resolved when we have shown that all the patterns in the class are coincident and therefore that the experimental class is in fact a coincidence class.

To resolve an unresolved class we start by creating a directed graph with a vertex for each pattern in the class. The graph is completely disconnected at first. If two patterns p and q are shown to be coincident by the Shading Lemma, the Simultaneous Shading Lemma, Lemma 2.3 or Lemma 2.10, we add edges between the patterns in both directions. If Remark 2.1 or Lemma 2.13 shows that the containment of p implies containment of q then we add a directed edge from p to q. If the graph becomes strongly connected (i.e., is one strong component) then the class is resolved, as we have proven the coincidence of all the patterns in the class.[5]

Pattern	1		12		123		132	
	unr.	res.	unr.	res.	unr.	res.	unr.	res.
Exp.	1	7	59	161	9608	23908	10315	23035
Shading L.	0	8	2	218	205	33311	183	33167
Lemma 2.3	0	8	2	218	205	33311	183	33167
Sim. Shading L.	0	8	1	219	94	33422	145	33205
Lemma 2.10	0	8	1	219	94	33422	145	33205
Lemma 2.13	0	8	1	219	74	33442	121	33229

Table 1: The results of using the lemmas for coincidence classification of size 1, 2, and 3 mesh patterns.

As can be seen in Table 1, less than one percent of the classes remain unresolved after the Shading Lemma has been applied. We know from Lemma 2.4 and Example 2.5 that Lemma 2.3 is strictly stronger than the Shading Lemma. However, the lines for these two lemmas in Table 1 are identical, implying that (at least for mesh patterns with these underlying classical patterns), this extra strength is not enough to fully resolve any more classes

[5]The computations were performed on resources provided by the Icelandic High Performance Computing Centre at the University of Iceland.

than the Shading Lemma. The same phenomenon occurs for Lemma 2.10 and the Simultaneous Shading Lemma. After Lemma 2.13 has been applied, the remaining unresolved classes are 196. In the next section we further improve Lemma 2.13 and reduce this number down to one exceptional case, which we do by hand.

3. The Shading Algorithm

We define the Shading Algorithm (Algorithm 2) which iterates Lemma 2.13. The algorithm takes as input two mesh patterns $p = (\tau, R)$ and $q = (\tau, R')$, a force F on τ and a depth d. It outputs SUCCESS if it can show that containment of p implies containment of q. The core of the algorithm is the recursive function SA, which takes as input a mesh pattern $w = (\sigma, Y)$, an occurrence c of p in w and a depth d. The depth d serves as a maximum recursion depth of the function. The mesh pattern w represents a state and describes an occurrence of p in an arbitrary permutation. The algorithm uses w to explore the occurrence of p with maximum strength with respect to the force F, which in turn is used to infer on the shadings of p. Similar to the previous lemmas, the algorithm branches, depending on whether a square in w is empty or not. In the latter case it attempts to derive a contradiction by showing that the square can not contain a point, and is therefore empty. We start by giving an example of what the algorithm is meant to do.

Example 3.1. We want to show that an occurrence of the pattern $p = (\tau, R)$ implies an occurrence of the pattern $q = (\tau, R')$.

$p =$ 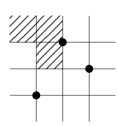 $q =$ 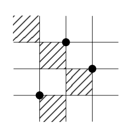

We think of p as an occurrence in a permutation. The goal is to show that $S = R' \setminus R = \{\lfloor 1, 0\rceil, \lfloor 2, 1\rceil\}$ can be shaded (or more precisely that there is an occurrence of p in the permutation where S is empty), or there is an occurrence of q. We choose the force $F = ((1, \rightarrow))$. Assume the occurrence of p in the permutation maximizes the strength with respect to the force F. Consider the case when $\lfloor 1, 0\rceil$ is not empty, and add the rightmost point in that square to p.

$w_1 = p^{\lfloor 1,0 \rceil \rightarrow} =$ 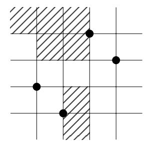 $q_1 =$ 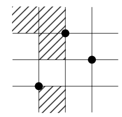

Consider the subsequence at indices 234 in w_1. This is an occurrence of q_1, which implies (since it has a superset of shadings) an occurrence of p, which is stronger than the original occurrence of p with respect to F. This is a contradiction. Thus $\lfloor 1,0 \rceil$ in p must have been empty. Consider the case when $\lfloor 2,1 \rceil$ is not empty in p, and add the leftmost point in that square to p.

$w_2 = (\tau, R \cup \lfloor 1,0 \rceil)^{\lfloor 2,1 \rceil \leftarrow} =$ 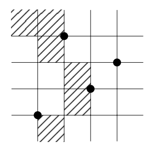 $q_2 =$ 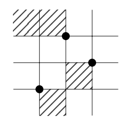

We take the subsequence at indices 123 in w_2. This is an occurrence of q_2. The square $\lfloor 1,2 \rceil$ is not shaded, and corresponds to $\lfloor 1,2 \rceil$ in w_2. Let us consider the rightmost point in $\lfloor 1,2 \rceil$ in w_2, giving:

$w_3 = w_2^{\lfloor 1,2 \rceil \rightarrow} =$ 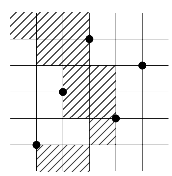 $q_3 =$ 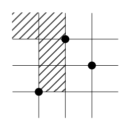

Here we take the subsequence at indices 235 in w_3 which is an occurrence of q_3, which implies a stronger occurrence of p with respect to F. Hence, $\lfloor 1,2 \rceil$ in w_2 is empty, which implies that $\lfloor 1,2 \rceil$ is empty in q_2. Therefore q_2 is a stronger occurrence of p with respect to F, thus $\lfloor 2,1 \rceil$ is empty in p. The original occurrence of p in the permutation is therefore an occurrence of q.

In Algorithm 1 we record two subfunctions we will need in the main algorithm.

Before proving our main result we prove a lemma about the case when the function SA in Algorithm 2 returns SUCCESS.

Lemma 3.2. For fixed mesh patterns p, q and a force F on p, if the function SA in Algorithm 2 returns SUCCESS for input $w = (\sigma, Y)$, c (an occurrence of p in w), F' and d, then at least one of the following is false:

1. In the strongest occurrence of w with respect to F' in any permutation π, the occurrence c of p in w corresponds to the strongest occurrence of p in π with respect to F.

2. The occurrence of w does not imply an occurrence of q.

Proof. We prove this by induction on the depth d. When $d = 0$, the base case of the procedure on line 4 checks whether statement (1) fails and the second base case of the procedure on line 7 checks whether statement (2) fails. Statement (1) fails when an occurrence of p that is stronger than c w.r.t. F is found in w. If we consider the strongest occurrence of w w.r.t. F' in some permutation, then this stronger occurrence of p in w is also a stronger occurrence of p in the permutation. The second statement fails when w contains an occurrence of q. Since the algorithm returned SUCCESS, either statement (1) or statement (2) is false.

Assume the lemma holds for d. We now prove the inductive case for $d+1$. Since the call to SA returned SUCCESS, there is some (τ, T) constructed on line 3 that resulted in SUCCESS. If one the two tests in line 4 or line 7 succeeds then we are done. We consider the case where the tests fail and the algorithm continues to line 10.

For each $i = 1, \ldots, k$ we claim that if there is some $a \in \{\uparrow, \downarrow, \leftarrow, \rightarrow\}$ such that $\mathrm{SA}((\sigma, Y \cup \{s_1, \ldots, s_{i-1}\})^{s_i a}, \mathrm{UO}(c, s_i), \mathrm{UF}(F', s_i), d-1)$ returns the value SUCCESS, then an occurrence of $(\sigma, Y \cup \{s_1, \ldots, s_{i-1}\})$ implies an occurrence of $(\sigma, Y \cup \{s_1, \ldots, s_i\})$. Consider the case when $i = 1$. By Remark 1.8 this occurrence of $w = (\sigma, Y)$ implies an occurrence of either $(\sigma, Y \cup \{s_1\})$ or the occurrence of $w^{\lfloor i,j \rceil a}$. The first case is trivial. For the second case we will show that the containment of $w^{\lfloor i,j \rceil a}$ leads to the same conclusion. The functions UO and UF update the occurrence c in w and the force F', respectively, such that they refer to the same points after the new point was inserted into $\lfloor i, j \rceil$. Since the strongest occurrence of p w.r.t. F is contained in the strongest occurrence of w w.r.t. F', it is contained in the strongest occurrence of $w^{\lfloor i,j \rceil a}$ w.r.t. $\mathrm{UF}(F', s_1)$. But since the recursive call on $w^{\lfloor i,j \rceil a}$ with the updated c and F' results in a contradiction, the strongest occurrence of w does not imply the strongest occurrence of $w^{\lfloor i,j \rceil a}$ and therefore implies the occurrence of $(\sigma, Y \cup \{s_1\})$. The same argument holds for every i and hence the claim holds.

We have shown that for every $i = 1, \ldots, k$, an occurrence of the mesh pattern $(\sigma, Y \cup \{s_1, \ldots, s_{i-1}\})$ implies an occurrence of $(\sigma, Y \cup \{s_1, \ldots, s_{i-1}, s_i\})$, and hence that w implies the occurrence of $(\sigma, Y \cup \{s_1, \ldots, s_{i-1}, s_i\})$. This pattern contains an occurrence of q, thus w implies an occurrence of q, concluding our proof.

We now state our main result.

Theorem 3.3. If Algorithm 2 returns SUCCESS for p, q and some choice of F, then an occurrence of p implies an occurrence of q.

Proof. Algorithm 2 calls the function SA with p, τ, F and d. By Lemma 3.2, the algorithm returns SUCCESS when either in the strongest occurrence of p w.r.t. F, the occurrence τ in p does not correspond to the strongest occurrence of p w.r.t. F, or an occurrence of p implies the occurrence of q. Since the first case would be a contradiction, it must be the second case.

When we run the algorithm with depth $d = 2$ we are able to automatically classify mesh patterns of size 2, showing that the total number of coincidence classes is 220. The results from running Algorithm 2 with $d = 1, \ldots, 6$ on mesh patterns of size 3 are given in Table 2. The algorithm fully classifies the coincidence of mesh patterns with the underlying pattern 123 at $d = 4$, while two classes remain unresolved for 132 at $d = 6$. The implementation is available on GitHub [12], and further description of the repository is given in Section Appendix A.

Pattern	$d=1$	$d=2$	$d=3$	$d=4$	$d=5$	$d=6$
123	74	8	6	0	0	0
132	121	32	13	6	2	2

Table 2: The number of unresolved classes after running the Shading Algorithm at depths $d = 1, \ldots, 6$.

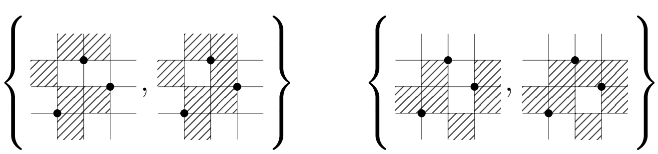

Figure 3.1: The two unresolved classes with underlying pattern 132 after Algorithm 2 has been applied.

The two remaining unresolved classes are symmetries of each other as can be seen in Figure 3.1. We will therefore only prove the coincidence of the first class since the arguments will be identical for the second class. We start by proving that the the pattern 21 is coincident with a *decorated pattern* (See Ulfarsson [13] for the definition).

Proposition 3.4. The patterns p and q, shown below, are coincident. The second pattern is a decorated pattern that contains a (possibly empty) increasing sequence in the union of the squares $\lfloor 0,0 \rceil$ and $\lfloor 0,1 \rceil$, denoted by the diagonal line.

$p =$ 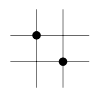 $\asymp$ 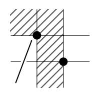 $= q$

Proof. Let π be a permutation. If π contains q then it contains p. Assume that π contains p. Since an occurrence of p is an inversion, π must have a descent. Consider the first descent in π, ab, with $a > b$.

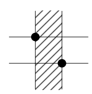

Consider the points in the region left of the point a in p, the squares $\lfloor 0,0 \rceil$, $\lfloor 0,1 \rceil$ and $\lfloor 0,2 \rceil$. If this region contains an inversion, then it must contain a descent. As we picked the leftmost occurrence of p already, this region must avoid any 21 and can only contain an increasing sequence. Furthermore, the square $\lfloor 0,2 \rceil$ must be empty, since the rightmost point in that region would form a descent with its right adjacent point. Thus, this descent ab is an occurrence of q.

Proposition 3.5. The following mesh patterns are coincident.

$p =$ 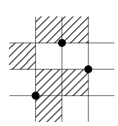 $\asymp$ 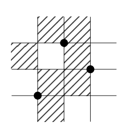 $= q$

Proof. By Remark 2.1 it suffices to show that an occurrence of p implies an occurrence of q. Let π be a permutation that contains p. Either the points in the region corresponding to the square $\lfloor 1,2 \rceil$ form an increasing sequence (possibly empty) or contain an inversion. We will show in both cases, which are depicted below, that the occurrence implies an occurrence of q.

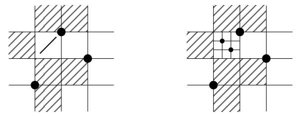

In the first case, when the square $\lfloor 1, 2\rceil$ contains an increasing sequence, the square $\lfloor 2, 2\rceil$ either contains a point or is empty. If the square is empty, we are done, since this occurrence will then be an occurrence of q. Otherwise we pick the leftmost point in the square and place it into the occurrence of the pattern p.

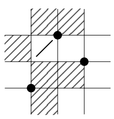

with $\lfloor 2, 2\rceil$ non-empty implies the occurrence of

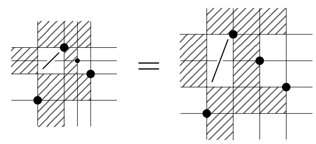

The square $\lfloor 1, 2\rceil$ is either empty or contains a point. If it is empty, the occurrence implies an occurrence of q with the subsequence 143. Otherwise, the square contains a point. Pick the rightmost point, which is also the highest point.

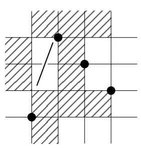

with $\lfloor 1, 2\rceil$ non-empty implies the occurrence of

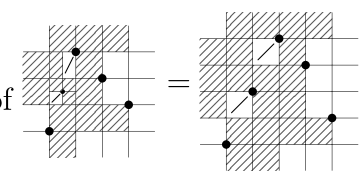

This occurrence of p with the inferred points forms an occurrence of q, namely the subsequence 354.

Consider the case when the square $\lfloor 1, 2\rceil$ in p contains an inversion. By Proposition 3.4 an inversion is coincident with the decorated pattern in that proposition, which we place instead of the inversion into the occurrence of p.

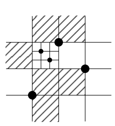

implies the occurrence of

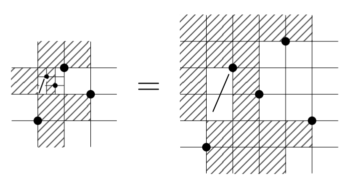

The square $\lfloor 1, 2\rceil$ in the occurrence is either empty or contains a point. If the square is empty, the occurrence forms an occurrence of q with the subsequence

132. We assume it contains a point and place the rightmost point in the square into the pattern of the occurrence.

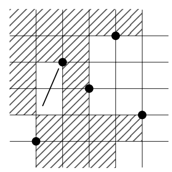 implies the occurrence of 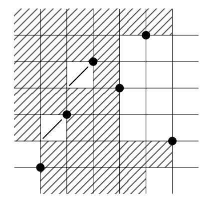

The subsequence 354 forms an occurrence of q, and we conclude that an occurrence of p implies an occurrence of q.

This completes the coincidence classification of mesh patterns of sizes 1, 2, and 3. There are 8 coincidence classes of mesh patterns of size 1. The number of classes for longer patterns are shown in Table 3. The number of

Pattern	12	123	132
Number of mesh patterns	512	65536	65536
Coincidence classes	220	33516	33350
Coincidence classes of size 1	161	23908	23035
Coincidence classes of size 2	37	6116	6598
Coincidence classes of size 3	2	132	286
Coincidence classes of size 4	11	1961	2182
Coincidence classes of size 5	0	16	46
Coincidence classes of size 6	0	172	164
Coincidence classes of size 7	0	0	0
Coincidence classes of size ≥ 8	9	1211	1039

Table 3: Number of coincidence classes of mesh patterns with underlying patterns 12, 123 and 132.

singleton classes in Table 3 shows how effective the experimental classification is, since roughly a third of the patterns are not coincident with any other mesh pattern. From the table we can also see that for 123 and 132 more than 90% of the coincidence classes contain at most four mesh patterns, which greatly reduces the number of comparisons in contrast to running the algorithm on every pair of patterns.

4. Applications of mesh patterns and the force to enumeration

Knowing coincidences of mesh patterns can be used to enumerate permutation classes, as the following example shows.

Example 4.1. Let $B = \{1234, 1243, 1324, 1342, 1423, 2314, 2341, 3124, 4123\}$ and consider the permutation class $\mathrm{Av}(B)$. Since 123 appears as a subpattern in every pattern in B we can write $\mathrm{Av}(B)$ as the disjoint union

$$\mathrm{Av}(B) = \mathrm{Av}(123) \sqcup (\mathrm{Av}(B) \cap \mathrm{Co}(123)).$$

The enumeration of $\mathrm{Av}(123)$ is well known to be given by the Catalan numbers, which have the generating function

$$C(x) = \frac{1 - \sqrt{1 - 4x}}{2x}.$$

Therefore we only need to consider $\mathrm{Av}(B) \cap \mathrm{Co}(123)$. Using any of the lemmas in Table 1 we can show that 123 is coincident with the pattern

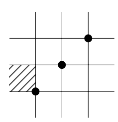

A permutation in $\mathrm{Av}(B)$ contains the previous pattern if and only if it contains the following pattern.

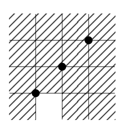

This is because the existence of a point in any square (except $\lfloor 0, 1 \rceil$) would imply an occurrence of one of the basis elements in B. In such a permutation the region corresponding to the square $\lfloor 1, 0 \rceil$ must avoid 12, or else an occurrence of 3124 would be realized. This implies that every permutation in $\mathrm{Av}(B) \cap \mathrm{Co}(123)$ has a unique occurrence of 123 that is an occurrence of the pattern above. Moreover, any decreasing sequence of points can be placed in this region without creating a basis element. Every permutation in $\mathrm{Av}(B) \cap \mathrm{Co}(123)$ can therefore be constructed by starting with the permutation 123 and placing a 12 avoiding permutation in the region corresponding to $\lfloor 1, 0 \rceil$. Hence, we obtain the following generating function $F_B(x)$ of the permutation class $\mathrm{Av}(B)$:

$$\begin{aligned} F_B(x) &= F_{123}(x) + x^3 \cdot F_{12}(x) \\ &= C(x) + \frac{x^3}{1-x}. \end{aligned}$$

This example motivates the following definitions.

Definition 4.2. Let $\mathcal{C}$ be a permutation class. Two patterns p and q such that $\mathrm{Av}(p) \cap \mathcal{C} = \mathrm{Av}(q) \cap \mathcal{C}$ are said to *coincident with respect to* $\mathcal{C}$, denoted $p \overset{\mathcal{C}}{\asymp} q$.

Definition 4.3. Define $\mathrm{occ}_p(\pi)$ as the number of occurrences of a pattern p in a permutation π. A pattern p is *binary* if $\mathrm{occ}_p(\pi) \leq 1$ for every permutation π. A pattern p is *binary with respect to a permutation class* $\mathcal{C}$ if $\mathrm{occ}_p(\pi) \leq 1$ for every $\pi \in \mathcal{C}$.

Example 4.1 gives an approach to enumerating permutation classes $\mathcal{C}$: First choose a classical pattern $p \in \mathcal{C}$, and find a coincident pattern $q \asymp p$ with the Shading Algorithm. Add in the shadings implied by the basis of $\mathcal{C}$, obtaining a pattern $q' \overset{\mathcal{C}}{\asymp} p$. If q' is binary with respect to $\mathcal{C}$ it can be used to find structural information about $\mathcal{C} \cap \mathrm{Co}(p)$.

It is not clear what is a good choice for the pattern $p \in \mathcal{C}$, or for the coincident pattern $q \asymp p$. In more generality, how can it be determined when a pattern can only occur at most once in any permutation?

Proposition 4.4. For every classical permutation pattern p (except ϵ, the pattern of length 0) and every $i > 0$ there exists a permutation π such that $\mathrm{occ}_p(\pi) \geq i$.

Proof. Let p be a classical permutation pattern of size n. Insert the element $n+1$ to the left of the element n. This new permutation is of size $n+1$ and has at least two occurrences of p. The occurrence where the element n corresponds to the element n in p and the occurrence where $n+1$ corresponds to the element n in p. This process can be repeated and the element $n+2$ added, giving a permutation with at least one more occurrence of p than the previous permutation. This process can be continued to obtain a permutation with as many occurrences of p as desired.

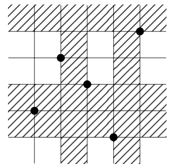

Figure 4.1: An example of an anchored pattern. The point with value 5 is anchored to the boundary as the highest point and the remaining points are anchored to it through each other.

This proposition does not hold for mesh patterns in general. Take for example the mesh pattern $m =$ . An occurrence of m in a permutation corresponds to the leftmost point in the permutation. Every permutation has at most one leftmost point, hence every permutation has at most one occurrence of m. There are also bigger and more complex mesh patterns that exhibit this behavior, i.e., occur exactly once or never in any permutation.

Although Proposition 4.4 implies that no classical permutation pattern is binary, the pattern m shows that mesh patterns can be binary. The pattern contains a single point that is *forced*, in a similar manner to what was discussed in Definition 2.8. Larger patterns of this type exist, such as the pattern $q =$ 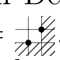. The 1 in an occurrence of q in π corresponds to 1. Similarly 2 must be the rightmost point in π. A point in a mesh pattern of size n is *anchored to a boundary* if it is an occurrence of at least one of the mesh patterns , , , . Furthermore, a point is *anchored* to another point if together they are an occurrence of at least one of the mesh patterns 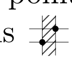, 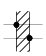, 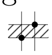, 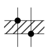.

Definition 4.5. A mesh pattern p is *anchored* if for every point p_{i_1} of the pattern there exists a sequence of points $p_{i_1}p_{i_2}\dots p_{i_n}$ such that p_{i_ℓ} is anchored to $p_{i_{\ell+1}}$ and p_{i_n} is anchored to a boundary.

An example of an anchored pattern is given in Figure 4.1 in which the point $(5,5)$ is anchored to the top boundary. The point $(4,1)$ is anchored to $(5,5)$, $(1,2)$ to $(4,1)$, $(3,3)$ to $(1,2)$ and $(2,4)$ to $(3,3)$, hence each point is anchored through a sequence of points to a boundary-anchored point.

Proposition 4.6. Every anchored mesh pattern is a binary mesh pattern.

Proof. Let p be an anchored mesh pattern and let π be a permutation that contains p. Any occurrence of p in π will use the same point in π for a boundary-anchored point in p, e.g., if p contains a point anchored to the bottom boundary (fully shaded bottom row), $(i, 1)$ then the corresponding point in π must be the lowest point in π. Any point anchored to $(i, 1)$ is therefore uniquely determined in any occurrence of p in π, since it must have an adjacent index, $i + 1$, or value, 2. This argument can be iterated on the points anchored to the points anchored to i and so on. Since every point in p is anchored to a boundary-anchored point through a sequence of points, each point is uniquely determined in every occurrence of p in π. Hence, there can only be one occurrence of p in π.

The previous result implies:

Corollary 4.7. There are infinitely many binary mesh patterns.

Proposition 4.6 shows that anchored patterns are binary. However, non-anchored binary patterns do exist, such as the pattern in Figure 4.2.[6]

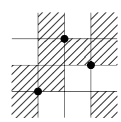

Figure 4.2: Example of a non-anchored binary mesh pattern.

While several permutation classes can be enumerated using the method in Example 4.1 we now turn our attention to a more powerful method, which is including a force with the pattern.

Definition 4.8. A *forced pattern* is a tuple (p, F) of a pattern p and a force F. An *occurrence* of a forced pattern (p, F) in a permutation π is an occurrence c of p in π such that

$$\text{strength}_F(\pi, c) = \max_{\text{occ. } c' \text{ of } p \text{ in } \pi} \text{strength}_F(\pi, c')$$

where the force strengths are compared in the lexicographical order.

From the definition it follows that a permutation π contains a forced pattern (p, F) if and only if it contains its underlying pattern p. We extend Definition 4.3 to also apply to these new patterns.

[6]This pattern can be shown to be binary with Lemma 4.11

Remark 4.9. We note that although forced patterns are similar to anchored mesh patterns in some aspects, these are different in general. We leave it to the reader to show that the forced pattern $(12, ((1, \leftarrow)))$ does not have the same occurrences as the mesh pattern 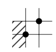, or the mesh pattern 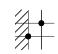 in a permutation.

Proposition 4.10. A forced pattern (p, F) with $|p| = |F|$ is binary.

Proof. Assume a permutation σ has two occurrences of p, $c_1 = \sigma_{i_1} \dots \sigma_{i_n}$ and $c_2 = \sigma_{j_1} \dots \sigma_{j_n}$ that have equal strength with respect to F. Then for every $k \in [1, n]$ we have either $i_k = j_k$ or $\sigma_{i_k} = \sigma_{j_k}$ from which it follows that the two occurrences are equal.

Proposition 4.10 shows that any (classical) pattern can be made binary by adding a force to it. The more points that are forced, the fewer occurrences there are of the pattern. If all the points are forced, then there is a unique occurrence of the pattern with maximum strength.

Sometimes it may be preferable to use a force of smaller size. This can be important if one wants to apply the technique of Example 4.12 to several permutation classes. Then for each potential pattern of length k it might be too computationally hard to check if a force of length k gives a good description of the subclass of the permutation class that contains the pattern.

Lemma 4.11. Let p be a pattern that is not binary with respect to the permutation class $\mathcal{C}$. If p has size n, then there exists a permutation $\pi \in \mathcal{C}$ of size at most $2n$ such that $\mathrm{occ}_p(\pi) > 1$.

Proof. Since p is not binary with respect to $\mathcal{C}$ there exists a permutation $\sigma \in \mathcal{C}$ such that $\mathrm{occ}_p(\sigma) > 1$. Let $c_1 = \sigma_{i_1} \dots \sigma_{i_n}$ and $c_2 = \sigma_{j_1} \dots \sigma_{j_n}$ be two distinct occurrences of p. Let π be the permutation that is order-isomorphic to the union of the occurrences c_1 and c_2. Then π has size at most $2n$, contains at least two occurrences of p, and is a member of $\mathcal{C}$.

By Lemma 4.11 we only need to check permutations up to size $2n$ to verify that a forced pattern of size n is binary. Given a classical pattern, this allows us to discover, in a brute force manner, a small force that makes the pattern binary. We start with the pattern with the empty force. If the pattern is binary we are done, otherwise we pick a point and direction and add this to the force. If this forced pattern is now binary, we are done, else we repeat this process. By Proposition 4.10, this will eventually result in a binary forced pattern.

Example 4.12. Let $B' = \{1324, 1342, 1423, 2143, 2413, 3142\}$ and consider the permutation class $\mathrm{Av}(B')$. Similarly as in Example 4.1 we get

$$\mathrm{Av}(B') = \mathrm{Av}(132) \sqcup (\mathrm{Av}(B') \cap \mathrm{Co}(132)).$$

An occurrence of 132, in a permutation in $\mathrm{Av}(B')$, will be an occurrence of the mesh pattern

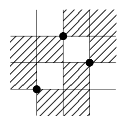

It is possible to check that there is no mesh pattern $q \overset{\mathrm{Av}(B')}{\asymp} 132$ that is binary with respect to $\mathrm{Av}(B')$. However, 132 is coincident to the forced pattern $(132, ((3, \uparrow), (1, \downarrow), (2, \downarrow)))$, which is binary by Proposition 4.10. An occurrence of this pattern, in a permutation in $\mathrm{Av}(B')$, will be of the form

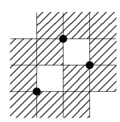

Because of the force the region corresponding to the square $\lfloor 0, 3 \rceil$ avoids the pattern 132. The restrictions from the basis B' imply that the regions corresponding to the squares $\lfloor 1, 1 \rceil$, $\lfloor 2, 2 \rceil$, $\lfloor 3, 0 \rceil$ avoid 21, 12, B', respectively. It is easy to check that there are no other conditions causing these regions to be dependent. We therefore obtain the following equation satisfied by the generating function $F_{B'}(x)$ for the permutation class $\mathrm{Av}(B')$:

$$\begin{aligned} F_{B'}(x) &= F_{132}(x) + F_{132}(x) \cdot F_{21}(x) \cdot F_{12}(x) \cdot F_{B'}(x) \cdot x^3 \\ &= C(x) + \frac{x^3 C(x) F_{B'}(x)}{(1-x)^2}, \end{aligned}$$

where we have used the fact that $F_{132}(x) = C(x)$, the generating function of the Catalan numbers. Solving this equation gives

$$\begin{aligned} F_{B'}(x) &= \frac{(1-x)^2 C(x)}{(1-x)^2 - x^3 C(x)} \\ &= 1 + x + 2x^2 + 6x^3 + 18x^4 + 54x^5 + 167x^6 + 534x^7 \\ &\quad + 1755x^8 + 5896x^9 + 20167x^{10} + \cdots . \end{aligned}$$

By applying the process used to enumerate $\text{Av}(B')$ in the previous example to non-insertion encodable[7] permutation classes with bases consisting of size 4 classical patterns we are able to enumerate 316 classes. No permutation class with fewer than six patterns in its basis is successful. Detailed results are in Table 4.

Size of basis	Nr. of classes	Successes	%
12	1	1	100.0
11	10	10	100.0
10	48	39	81.3
9	151	86	57.0
8	337	106	31.5
7	547	62	11.3
6	659	12	1.8
5	578	0	0.0

Table 4: Number of permutation classes with bases consisting of size 4 patterns that can be enumerated with the process used in Example 4.12.

5. Future work

The experimental classification is sufficient to coincidence classify (without proof) the mesh patterns of sizes 0, 1, 2, 3 correctly by considering the permutations of size 1, 3, 5, 10, respectively. This leads to the following conjecture.

Conjecture 5.1. The mesh patterns of size n can be coincidence classified by experimental classification with permutations of size $(n+1)^2 + n$.

The intuition behind the expression $(n+1)^2 + n$ is similar to the proof of Lemma 4.11. Consider two mesh patterns p, q with the same underlying classical pattern and a permutation π that contains p but not q. Then in every occurrence of p in π, there must be a point in a region that corresponds to a shaded region in q. Since there are $(n+1)^2$ squares and n points in a mesh pattern of size n, we obtain the term in Conjecture 5.1.

[7] Albert, Linton, and Ruškuc [14] studied permutation classes with a regular language for their insertion encoding. Vatter[15] provided an algorithm for automatically computing the generating function of such classes.

One of the obvious next steps with the Shading Algorithm would be the classification of size 4 mesh patterns. The main issue is the number of mesh patterns to classify since the size of the underlying pattern increases and the number of different shadings increases to 2^{25}. The experimental classification becomes even more vital in this case, but the size of the permutations to consider also increases dramatically and the task becomes computationally impractical.

We end with two open enumerative problems:

Problem 5.2. Given a classical pattern p, determine the number of anchored mesh patterns (p, R).

Problem 5.3. Given a classical pattern p, determine the number of binary mesh patterns (p, R).

Acknowledgements

We thank a reviewer for a very careful review.

Appendix A. Implementation of the Shading Algorithm

The Python implementations of the Shading Algorithm, Shading Lemma and The Simultaneous Shading Lemma are available at Bean et al. [12]. The core of the implementations is in the file `tsa5_knowledge.py` under a directory called `the_shading_algorithm`. The script `classify.py` reads in experimental classes with known coincidence relations of the patterns and calls the Shading Algorithm on pairs of patterns to decide their coincidence. The full classification of the mesh patterns with underlying classes 12, 123 and 231 is given in the files located in the directory `results/final_results`.

Each file in the `results/final_results` of the GitHub repository [12], contains multiple lines, where each line represents a coincidence class. We represent the mesh of the patterns with an integer such that the binary representation of the integer describes the shadings. Starting with the least significant bit, the i-th bit is set to 1 if $\lfloor \lfloor i/(n+1) \rfloor, i \mod (n+1) \rceil$ is shaded.

References

[1] P. A. MacMahon, Combinatory analysis, Vol. 137, American Mathematical Soc., 2001.

[2] R. Simion, F. W. Schmidt, Restricted permutations, European Journal of Combinatorics 6 (4) (1985) 383–406.

[3] E. Babson, E. Steingrímsson, Generalized permutation patterns and a classification of the mahonian statistics, Sém. Lothar. Combin 44 (B44b) (2000) 547–548.

[4] M. Bousquet-Mélou, A. Claesson, M. Dukes, S. Kitaev, (2+ 2)-free posets, ascent sequences and pattern avoiding permutations, Journal of Combinatorial Theory, Series A 117 (7) (2010) 884–909.

[5] J. West, Sorting twice through a stack, Theoretical Computer Science 117 (1-2) (1993) 303–313.

[6] P. Brändén, A. Claesson, Mesh patterns and the expansion of permutation statistics as sums of permutation patterns, The Electronic Journal of Combinatorics 18 (2) (2011) P5.

[7] Í. Hilmarsson, I. Jónsdóttir, S. Sigurðardóttir, L. Viðarsdóttir, H. Ulfarsson, Wilf-classification of mesh patterns of short length, The Electronic Journal of Combinatorics 22 (4) (2015).

[8] A. Claesson, B. E. Tenner, H. Ulfarsson, Coincidence among families of mesh patterns, Australasian Journal of Combinatorics 63 (2015) 88–106.

[9] B. E. Tenner, Coincidental pattern avoidance, Journal of Combinatorics 4 (3) (2013).

[10] B. E. Tenner, Mesh patterns with superfluous mesh, Advances in Applied Mathematics 51 (5) (2013) 606–618.

[11] M. Tannock, H. Ulfarsson, Equivalence classes of mesh patterns with a dominating pattern, Discrete Mathematics & Theoretical Computer Science 19 (2) (2018).

[12] C. Bean, B. A. Gudmundsson, T. K. Magnusson, H. Ulfarsson, The shading algorithm implementation, `https://github.com/PermutaTriangle/TheShading` (2015).

[13] H. Ulfarsson, Describing west-3-stack-sortable permutations with permutation patterns, Séminaire Lotharingien de Combinatoire 67 (2012) B67d.

[14] M. H. Albert, S. Linton, N. Ruškuc, The insertion encoding of permutations, The Electronic Journal of Combinatorics 12 (1) (2005) R47.

[15] V. Vatter, Finding regular insertion encodings for permutation classes, Journal of Symbolic Computation 47 (3) (2012) 259–265.

Algorithm 1 Subfunctions of the Shading Algorithm.

Updating an occurrence after the insertion of a point

function UO($c = c_1 \cdots c_n$, $\lfloor x, y \rceil$)
    for all c_i **do**
        if $c_i \geq y$ **then** $c'_i \leftarrow c_i + 1$ **else** $c'_i \leftarrow c_i$
    end for
    return $c'_1 c'_2 \cdots c'_n$
end function

Updating a force after the insertion of a point

function UF($F' = ((t_1, a_1), \ldots, (t_k, a_k))$, $\lfloor x, y \rceil$)
    for all t_i **do**
        if $t_i \geq y$ **then** $t'_i \leftarrow t_i + 1$ **else** $t'_i \leftarrow t_i$
    end for
    return $((t'_1, a_1), \ldots, (t'_k, a_k))$
end function

Algorithm 2 The Shading Algorithm.

```
function SA(w = (σ, Y), c, F′, d)
    for each occurrence c′ of the classical pattern τ in w do
        Let T be the maximal shading so c′ is an occurrence of (τ, T) in w
        if R is a subset of T and strength_F(c′) > strength_F(c) then
            return Success
        end if
        if R′ is a subset of T then
            return Success
        end if
        if d > 0 then
            OK ← True
            Let S = {s_1, s_2, ..., s_k} be squares in w corresponding to R′\T
            for i ← 1 to k do
                if for all a ∈ {↑, ↓, ←, →},
                SA((σ, Y ∪ {s_1, s_2, ..., s_{i−1}})^{s_i a}, UO(c, s_i), UF(F′, s_i), d−1)
                returns Failure then
                    OK ← False
                end if
            end for
            if OK then
                return Success
            end if
        end if
    end for
    return Failure
end function

return SA(p, τ, F, d)
```