

On Permutation Classes Defined by Token Passing
Networks, Gridding Matrices and Pictures:
Three Flavours of Involvement

Stephen D. Waton

PhD Thesis

February 11, 2007

Declarations

I Stephen Waton hereby certify that this thesis, which is approximately 40,000 words, has been written by me, that it is a record of work carried out by me, and that it has not been submitted in any previous application for a higher degree.

Signed:

Date:

I was admitted as a research student in September 2003 and as a candidate for the degree of PhD in September 2003; the higher investigation of which this is a record was carried out in the University of St Andrews between September 2003 and September 2006.

Signed:

Date:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Ph.D. in the University of St Andrews and that the candidate is qualified to submit this thesis in application to that degree.

Signature of Supervisor:

Date:

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and abstract will be published, and that a copy of the work may be supplied to any bona fide library or research worker, that my thesis will be electronically available for personal or research use, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration.

Signed:

Date:

Acknowledgements

I would like to thank my supervisor, Nik Ruškuc, for his support and encouragement. I would also like to thank:

- Michael Albert, Mike Atkinson, Robert Brignall, Fiona Brunk, Murray Elder, Sophie Huczynska, Steve Linton and Vince Vatter for their various and diverse contributions to my study of pattern classes;
- those I have shared an office with, who are, together with Robert, Andreas Disler, Nelson Silva and Morteza Jafarpour;
- those I have shared a flat with, Caroline McCarthy, Phil Rae, Gemma Dickenson, Mairi MacNaughton, Conor Parson, Jessica Forsythe and Anna Isherwood;
- my parents, Jill and Alan, for their contributions, financial, spiritual and practical;
- my brother, Andrew, for surviving countless “real doctor” jibes;
- Zoe-Lydia Tsivos, for her support, understanding and general great banter.

It has been a pleasure.

I acknowledge generous funding from the EPSRC.

Abstract

The study of pattern classes is the study of the involvement order on finite permutations. This order can be traced back to the work of Knuth. In recent years the area has attracted the attention of many combinatorialists and there have been many structural and enumerative developments. We consider permutations classes defined in three different ways and demonstrate that asking the same fixed questions in each case motivates a different view of involvement. Token passing networks encourage us to consider permutations as sequences of integers; grid classes encourage us to consider them as point sets; picture classes, which are developed for the first time in this thesis, encourage a purely geometrical approach. As we journey through each area we present several new results.

We begin by studying the basic definitions of a permutation. This is followed by a discussion of the questions one would wish to ask of permutation classes. We concentrate on four particular areas: partial well order, finite basis, atomicity and enumeration. Our third chapter asks these questions of token passing networks; we also develop the concept of completeness and show that it is decidable whether or not a particular network is complete. Next we move onto grid classes, our analysis using generic sets yields an algorithm for determining when a grid class is atomic; we also present a new and elegant proof which demonstrates that certain grid classes are partially well ordered.

The final chapter comprises the development and analysis of picture classes. We completely classify and enumerate those permutations which can be drawn from a circle, those which can be drawn from an X and those which

can be drawn from some convex polygon. We exhibit the first uncountable set of closed classes to be found in a natural setting; each class is drawn from three parallel lines. We present a permutation version of the famous ‘happy ending’ problem of Erdős and Szekeres. We conclude with a discussion of permutation classes in higher dimensional space.

Contents

1	Introduction and Preliminaries	1
1.1	Introduction	1
1.2	Permutations	2
1.3	The Involvement Order	3
1.3.1	A Point Set Definition of Involvement	5
1.3.2	A Geometric Definition of Involvement	6
1.4	Closed Classes of Permutations	7
1.5	Symmetries of Permutations under Involvement	8
2	The Algorithmic Problem Paradigm	13
2.1	Algorithmic Problems	13
2.2	Finite Closed Classes	14
2.3	Stack Sortable Permutations	15
2.4	The Basis Problem	18
2.5	The Atomicity Problem	20
2.6	The Enumeration Problem	22
2.7	The Partial Well Order Problem	26
2.8	Classes with a Single Basis Element of Length Three	32

3	Token Passing Networks	37
3.1	Definitions for Token Passing Networks	38
3.2	Some Examples of Token Passing Networks	46
3.3	Algorithmic Problems for Token Passing Networks	48
3.4	Incomplete Networks	50
3.4.1	Finite Token Passing Networks	50
3.4.2	Regular Languages, Finite State Automata and Ra- tional Generating Functions	52
3.4.3	Acyclic Networks	53
3.5	Complete Networks	54
3.5.1	Network Composition	59
3.5.2	Parallelizing Networks	62
3.5.3	Proof of the Main Theorem	66
3.6	Conclusions for Token Passing Networks	67
4	Grid Classes	69
4.1	Definitions for Grid Classes	69
4.2	Algorithmic Problems for Grid Classes	79
4.3	Early Results on Grid Classes	81
4.4	The Partial Well Order Problem for Grid Classes	83
4.5	The Atomicity Problem for Grid Classes	86
4.6	A Return to the Partial Well Order Problem	107
4.7	The Basis Problem for Grid Classes	111
4.7.1	Juxtaposition of Classes	112
4.7.2	A Larger Grid Class	113
4.8	The Enumeration Problem for Grid Classes	115

4.9	Conclusions for Grid Classes	119
5	Picture Classes	121
5.1	Definitions for Picture Classes	121
5.2	Algorithmic Problems for Picture Classes	123
5.3	Symmetries of the Plane which Respect Picture Classes . . .	124
5.4	The Atomicity Problem for Picture Classes	126
5.5	The Circle Class	126
5.5.1	The Basis of the Circle Class	128
5.5.2	Enumeration of the Circle Class	135
5.6	The X Class	140
5.7	Permutations Drawn from Parallel Lines	143
5.7.1	Permutations Drawn from a Single Line	144
5.7.2	Permutations Drawn from Two Infinite Parallel Lines	144
5.7.3	Permutations Drawn from Two Finite Parallel Lines .	146
5.7.4	Permutations Drawn from Three Infinite Parallel Lines	147
5.7.5	Permutations Drawn from Three Finite Parallel Lines	153
5.8	The Convex Class	154
5.8.1	Characterisation of the Convex Class	157
5.8.2	Enumeration of the Convex Class	161
5.8.3	Properties of the Convex Class	165
5.8.4	Permutations Drawn from Fixed Convex Polygons . .	168
5.8.5	A Combinatorial Problem for Permutations	171
5.9	An Extension to Higher Dimensions	175
5.10	Conclusions for Picture Classes	178

Chapter 1

Introduction and Preliminaries

1.1 Introduction

The study of pattern classes can be traced back to the work of Donald Knuth [45], who introduced the involvement ordering. His work was extended by Pratt [53] and Tarjan [59]. However, the area was largely ignored from around 1973 until an address by Herbert Wilf to the SIAM meeting on Discrete Mathematics in 1992. Since that date the field has grown rapidly. Atkinson, Albert and the Otago theory group, together with Ruškuc, Linton, Murphy and others from St Andrews have produced many beautiful results focusing on structural considerations. Much of this work is founded on the seminal theorems of Erdős and Szekeres [29]. Others, including West, Bóna, Zeilberger, Bousquet-Mélou, Sagan, Mansour and Steingrímsson have taken a more enumerative approach, for an introduction see Bóna [21]. The history of the structural approach to the study of pattern classes is the history of the search to find the terminology and methods which will enable us to describe and understand one particular poset, the set of permutations under the involvement ordering. The structure of this thesis in some sense mirrors this history. We begin, as we must, by defining our most basic objects.

Next we formalise the questions one may wish to ask about them. The main body of the thesis, Chapters 3, 4 and 5, consider permutation classes from different angles. There is, as the thesis progresses, a gentle descent into the world of discrete geometry. I believe this is both desirable and necessary, indeed, if this thesis has an underlying message it is that permutations are best considered neither as lists nor mappings, but instead as pictures we can draw, shrink and stretch.

1.2 Permutations

Definition 1.2.1. A *permutation* is a sequence of distinct integers from 1 to n with length n , that is any ordering of the integers 1 to n .

Definition 1.2.2. A *permutation* is a bijection from the set $\{1, \dots, n\}$ onto itself.

Definition 1.2.3. A *permutation* is an equivalence class, under order isomorphism, of sets of n distinct objects under two linear orderings.

It is easy to see that these three definitions yield the same objects, they correspond to viewing a permutation as a sequence, as an algebraic operation and as a picture (or relational structure). Each of them will be useful in different situations. In general we will write a permutation as a sequence according to the first definition. This is equivalent to listing the image points in the second definition, or to labeling our n points from the third definition and according to the first ordering and listing them according to the second ordering. It will also often be useful to plot a permutation, we do so by plotting the set $1, \dots, n$ (equivalently domain points or first linear order) along the x -axis and the sequence itself (image points, second linear order) along the y -axis.

Example 1.2.4. The permutation 134652 can be considered as the bijection from the set $\{1, 2, 3, 4, 5, 6\}$ on itself which maps 1 to 1, 2 to 3, 3 to 4, 4 to 6, 5 to 5 and 6 to 2. Alternatively we may think of it as the set of points

$\{1, 2, 3, 4, 5, 6\}$ under the following two linear orderings, $1 < 2 < 3 < 4 < 5 < 6$ and $1 \prec 3 \prec 4 \prec 6 \prec 5 \prec 2$. We can plot this permutation as shown in Figure 1.1.




Figure 1.1: A plot of the permutation 134652

1.3 The Involvement Order

Definition 1.3.1. Let $\Sigma = (\sigma_1, \dots, \sigma_k)$ and $\Pi = (\pi_1, \dots, \pi_k)$ be two sequences of distinct real numbers. We say that Σ and Π are *order isomorphic* if

$$\sigma_i < \sigma_j \Leftrightarrow \pi_i < \pi_j \text{ for all } i, j.$$

Definition 1.3.2. Let α and β be two permutations written as sequences of consecutive integers. We say that α is *involved* in β , written $\alpha \preceq \beta$ if there is a subsequence of β which is order isomorphic to α .

We will call this the subsequence definition of involvement. It is the oldest, and most common way to define involvement.

Example 1.3.3. The permutation 2341 is involved in the permutation 134652 as it is order isomorphic to the sequence 3462. This involvement is shown graphically in Figure 1.2.

Example 1.3.4. The permutation 2413 is not involved in the permutation 134652. There is no subsequence of 134652 which is order isomorphic to 2413, see Figure 1.3. We say that 134652 and 2413 are *incomparable* under involvement.




Figure 1.2: A plot of the permutation 2341 as a subpermutation of 134652.




Figure 1.3: A plots of the permutations 134652 and 2413 which are incomparable under involvement.

It is not surprising, however, that having defined a permutation in several ways we can also define involvement in several ways.

Definition 1.3.5. A set (permutation) $\sigma = (\{1, \dots, n\}, \leq_1^\sigma, \leq_2^\sigma)$, is said to be involved in a set $\tau = (\{1, \dots, m\}, \leq_1^\tau, \leq_2^\tau)$ if σ is a substructure of τ . That is, if there is some subset which is order isomorphic to σ .

It is clear that involvement is a partial ordering of the set of all permutations, that is, it is anti-symmetric, reflexive and transitive. The set of all permutations under involvement forms a very natural but incredibly complex structure with strong links to other areas of combinatorics.

We give two further, more radical definitions of involvement. These definitions are motivated by the concerns of Chapter 4, they come to motivate Chapter 5.

1.3.1 A Point Set Definition of Involvement

Definition 1.3.6. Following Felsner [32] we say that a finite set of points in the plane is *generic* if no two points are aligned either vertically or horizontally.

Definition 1.3.7. A generic point set $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ is said to be *involved* in a generic point set, $T = \{(z_1, w_1), \dots, (z_m, w_m)\}$, written $S \preceq T$ if there is a one to one mapping, f , from $\{1, \dots, n\}$ into $\{1, \dots, m\}$, satisfying the following conditions:

- If $x_i < x_j$ then $z_{f(i)} < z_{f(j)}$;
- If $y_i < y_j$ then $w_{f(i)} < w_{f(j)}$.

Involvement on generic point sets is a pre-order, it is reflexive and transitive.

Definition 1.3.8. We say that two generic point sets, S and T , are *order isomorphic* if and only if $T \preceq S$ and $S \preceq T$.

It is clear that order isomorphism is an equivalence relation. We factor the set of all generic point sets under involvement by this equivalence. We choose our class representatives to be those generic point sets with consecutive integer coordinates beginning with one. The partially ordered set we create is (isomorphic to) the set of permutations under involvement.

Example 1.3.9. The generic point sets shown in Figure 1.4 are isomorphic.

Definition 1.3.10. The *permutation image* of a generic set S , denoted $\Pi(S)$, is the permutation whose projection onto the plane is order isomorphic to S .

Lemma 1.3.11. *Given two generic sets $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ and $T = \{(z_1, w_1), \dots, (z_m, w_m)\}$ the following conditions are equivalent:*

1. S and T are order isomorphic.




Figure 1.4: Three point sets in the equivalence class of the permutation 52341

2. There exists a bijection g between the two sets $\{1, 2, \dots, n\}$ and $\{1, 2, \dots, m\}$ such that:

- $x_i < x_j$ if and only if $z_{g(i)} < z_{g(j)}$;
- $y_i < y_j$ if and only if $w_{g(i)} < w_{g(j)}$.

3. $\Pi(S) = \Pi(T)$.

Lemma 1.3.12. Given two generic sets S and T ,

$$S \preceq T \Leftrightarrow \Pi(S) \preceq \Pi(T)$$

1.3.2 A Geometric Definition of Involvement

Definition 1.3.13. A picture P is a set of points in the real plane.

Definition 1.3.14. Given two pictures $P = \{(x_1, y_1), \dots, (x_n, y_n)\}$ and Q we say that P is *involved* in Q , written $P \preceq Q$, if there exists a pair of order preserving bijections f and g from \mathbb{R} into \mathbb{R} such that the set $P_{f,g} = \{(f(x_1), g(y_1)), \dots, (f(x_n), g(y_n))\}$ is contained in Q .

Again this is a pre-order.

Definition 1.3.15. We say that two pictures P and Q are order isomorphic if $P \preceq Q$ and $Q \preceq P$.

Since we have not insured that our pictures are generic sets we will not get the set of permutations under involvement when we factor our pre-order

by this equivalence. Instead we will get the set of *full 0-1matrices* under the derived involvement order. A 0-1matrix is said to be full if it contains no row or column which is entirely zero. A full 0-1matrix M is said to be involved in a full 0-1matrix N if we can create M from N by deleting rows and columns and changing 1s to 0s. If we restrict ourselves to permutation matrices (or equivalently restrict pictures to generic sets) then we return to the set of permutations under involvement. There are close links between these two posets. Most famously the proof, by Marcus and Tardos, [49], of the Stanley-Wilf conjecture relies on an observation by Klazar, [42], that the Stanley-Wilf conjecture would follow from a proof of a conjecture by Füredi and Hajnal, [34], concerning 0-1 matrices. Marcus and Tardos in fact proved the Füredi-Hajnal conjecture.

The geometric definition of involvement corresponds precisely to the actions of stretching, squashing and sliding pictures (permutations) in the real plane. For this reason it is particularly easy to work with.

Just which definition of involvement is chosen depends entirely upon the setting we are in. As their name suggests, picture classes, Chapter 5, are easiest to work with using the geometric definition. In contrast, token passing networks, Chapter 3, demand the subsequence definition.

1.4 Closed Classes of Permutations

Definition 1.4.1. A *closed class* is a downset of the set of all permutations under involvement. That is, if C is a closed class and $\beta \in C$ with $\alpha \preceq \beta$ then $\alpha \in C$.

Closed classes occur naturally in a wide variety of settings. They form an extremely disparate collection of objects with incredibly diverse properties.

Example 1.4.2. The set of all permutations, S , forms a closed class under involvement. The beginnings of this poset are shown in Figure 1.5. There are $n!$ permutations of length n , so this poset becomes very wide very quickly.

Notice that even at short lengths the poset is becoming very varied, for example 1234 covers just one permutation, 123; 1243 covers two permutations; 1324 covers three permutations and 2413 covers four.

Example 1.4.3. We can see from Figure 1.5 that the set:

$$\{1, 12, 21, 132, 231, 213, 312, 2413\}$$

forms a closed class. Notice also that it has a single maximal element, 2413, in which every permutation is involved. We plot this closed class in Figure 1.6.

Example 1.4.4. We can see from Figure 1.5 that the set:

$$\{1, 12, 21, 123, 231, 321\}$$

forms a closed class. Notice that this class does not have a single maximal element. We plot this closed class in Figure 1.7.

There are many other finite closed classes which we can pick out from Figure 1.5. Indeed, there are infinitely many finite closed classes. Our next example, however, is of an infinite closed class.

Example 1.4.5. The set $\{1, 12, 123, \dots\}$ forms a closed class, the class of all increasing permutations. This is our first example of an infinite closed class. Notice from figure 1.5 that every permutation which is not in this class involves the permutation 21, so that this class may be characterised as precisely those permutations which do not involve 21.

1.5 Symmetries of Permutations under Involvement

For a permutation we define three symmetries. It is easy to see that these operations respect involvement and so we may talk about symmetries of closed classes.




Figure 1.5: The beginning of the poset of all permutations under involvement.




Figure 1.6: The closed class $\{1, 12, 21, 132, 231, 213, 312, 2413\}$.




Figure 1.7: The closed class $\{1, 12, 21, 123, 231, 321\}$.

Definition 1.5.1. Let α be a permutation given as a bijection from $1, \dots, n$ onto itself. The *inverse* of α is generated by swapping each point with its image.

Definition 1.5.2. Let α be a permutation given as a sequence of n distinct integers from 1 to n . The *reverse* of α is generated by reversing this sequence.

Definition 1.5.3. Let α be a permutation given as a sequence of n distinct integers from 1 to n . The *complement* of α is generated by replacing each integer i by the integer $n - i + 1$.

If we consider a permutation as a set of points under two orderings then these symmetries are even easier to see. Inverting the permutation corresponds to switching the two orderings, reversing to reversing the first ordering and complementing to reversing the second ordering. Furthermore it is easy to see that these symmetries correspond to isomorphisms of the plane, and so are easy to consider under the geometric definition of involvement. The inverse corresponds to a reflection in the line $y = x$. The reverse corresponds to a reflection in the y -axis. The complement corresponds to a reflection in the x -axis. Furthermore the inverse and either the reverse or complement together generate the dihedral group of order eight. In fact these are the only symmetries; the dihedral group of order eight is the automorphism group of the set of all permutations under involvement, see Smith [56].

Example 1.5.4. All the symmetries of the permutation 134652 are plotted in Figure 1.8. They are:

$$\{134652, 256134, 521346, 643125, 162354, 453261, 324516, 615423\}.$$




Figure 1.8: The symmetries of the permutation 134652

Chapter 2

The Algorithmic Problem Paradigm

The pattern class researcher sits in his office and waits. Sometimes a customer drops by, bringing with them some finite description of their latest pattern class. Naturally they wish to know more about their class. They have a particular property in mind, and wish to know whether this class possesses it. The researcher chews on his pen for a moment, offers a tentative fee and time scale, then begins his deliberations.

2.1 Algorithmic Problems

Almost every question we could wish to ask of pattern classes can be framed in the following way:

Algorithmic Problem 2.1.1.

Does there exist an algorithm which does the following?

Input: A closed class given by a finite description.

Output: TRUE if the class has the property P, FALSE otherwise.

We say that the property is *decidable* if such an algorithm exists, otherwise we say it is *undecidable*. It is worth noting that there are many different ways of defining a closed class, we shall see some of them shortly. We allow any finite description so that future approaches to closed classes may be studied in this framework. That said there are naturally some pathological definitions; in general we might expect a definition to allow membership testing, although we do not make this requirement explicit.

This approach has been championed by Ruškuc [54] in a talk to the Third International Conference on Permutation Patterns, 2005. We begin with some relatively easy examples which demonstrate its power.

2.2 Finite Closed Classes

In their seminal paper, A Combinatorial Problem in Geometry [29], Erdős and Szekeres proved a theorem that has become known as the Erdős-Szekeres Theorem. It states:

Theorem 2.2.1. *Given a set of n points in the plane, no two with the same x or y ordinates, it is possible to choose a set of at least \sqrt{n} points forming a monotonically increasing or monotonically decreasing sequence.*

This result allows us to characterise finite closed classes. Consider a closed class which does not contain the increasing permutation of length k , that is the permutation $12 \dots k$, nor the decreasing permutation of length l , that is $l \dots 21$. Since this is a closed class we know that it contains no increasing permutation of length greater than k and no decreasing permutation of length greater than l . Suppose without loss that $k \geq l$. Then, by the Erdős-Szekeres Theorem our class contains no permutations of length k^2 or greater. Conversely a class which contains every increasing or every decreasing permutation is clearly infinite. Thus we have the following decision theorem.

Theorem 2.2.2.

There exists an algorithm which does the following:

Input: A closed class given by some finite description.

Output: TRUE if the class is finite, FALSE otherwise.

2.3 Stack Sortable Permutations

Definition 2.3.1. A *stack* is a linearly ordered set S together with two operations. A *push* operation adds a new element to the top of the set S . A *pop* operation removes an element from the top of the set.

A stack is a device from theoretical computer science. It is a container which holds items of data. This data can only be accessed in a last in first out manner. New data is added to the top of the stack by a push operation. Old data can be removed from the top of the stack by a pop operation. Only the single data item at the top of the stack can be popped. In this way a stack resembles a stack of plates in a cafeteria.

We wish to consider the different orderings a stream of data may have when it leaves the stack, compared to when it enters. In particular we ask when can data be sorted using a stack. See Figure 2.1.




Figure 2.1: A stack sorting machine.

Definition 2.3.2. A permutation σ of length n , is said to be *stack sortable* if it is possible to pass the permutation through the stack so that the output is ordered from 1 to n .

We ask the obvious decision problem.

Algorithmic Problem 2.3.3.

Does there exist an algorithm which does the following?

Input: A Permutation given as a sequence.

Output: TRUE if the permutation is stack sortable,
FALSE otherwise.

Example 2.3.4. The permutation 1324756 is stack sortable. Figure 2.2 shows the sorting process.

It will be clear to the reader, from Figure 2.2, that stack sorting is deterministic; we only pop if the item on the top of the stack is the next item to be output, otherwise we push. Thus we may answer Problem 2.3.3.

Theorem 2.3.5.

There exists an algorithm which does the following:

Input: A Permutation given as a sequence.

Output: TRUE if the permutation is stack sortable,
FALSE otherwise.

We will leave a formal proof, together with a discussion of its implications, until later.

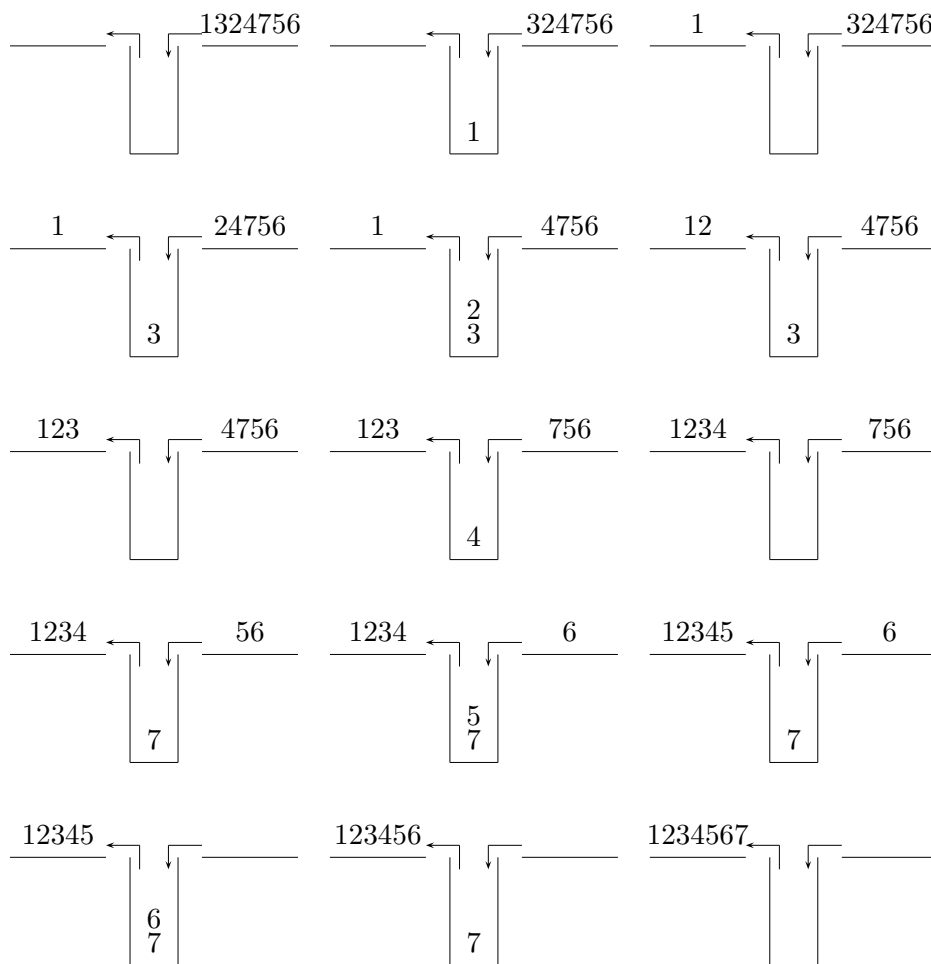


Figure 2.2: Sorting the permutation 1324756.

2.4 The Basis Problem

Definition 2.4.1. Let X be any set of permutations. The set of *avoiders* of X , denoted $\text{Av}(X)$ is the set of all permutations α with the property that if $\beta \in X$ then $\beta \not\preceq \alpha$.

Lemma 2.4.2. For any set of permutations X , $\text{Av}(X)$ is a closed class.

Proof. If $\alpha \in \text{Av}(X)$ then α avoids all $\beta \in X$. Then, since involvement is transitive, any $\gamma \preceq \alpha$ also avoids all β and so lies in $\text{Av}(X)$. \square

Example 2.4.3. The class of increasing permutations, see Example 1.4.5, is the set $\text{Av}(21)$. Note, however that it is also the set $\text{Av}(21, 321)$.

Definition 2.4.4. Let X be any set of permutations. We say that X is minimal if for any $\beta \in X$ there exists no $\alpha \in X$ with $\alpha \preceq \beta$.

Lemma 2.4.5. Given any set of permutations X , let $M(X)$ define the set of minimal elements in X . Then $\text{Av}(X) = \text{Av}(M(X))$.

Proof. It is clear that $\text{Av}(X) \subseteq \text{Av}(M(X))$, suppose now that $\alpha \in \text{Av}(M(X))$. For all $\beta \in X$ there exists some $\beta_1 \in M(X)$ such that $\beta_1 \preceq \beta$. Then since $\alpha \not\preceq \beta_1$ it follows that $\alpha \not\preceq \beta$. Thus $\text{Av}(M(X)) \subseteq \text{Av}(X)$. \square

Definition 2.4.6. Let C be a closed class of permutations. If D is the set of all permutations not in C then it is clear that $C = \text{Av}(D) = \text{Av}(M(D))$. We call $M(D)$ the *basis* of C , written $B(C)$. Thus $C = \text{Av}(B(C))$. The permutations in $B(C)$ satisfy the following equation:

$$\beta \in B(C) \Leftrightarrow \beta \notin C \ \& \ (\forall \alpha : (\alpha \preceq \beta) \ \& \ (\alpha \neq \beta) \Rightarrow \alpha \in C)$$

A basis gives an efficient way to characterise certain closed classes. It is an example of a finite description of a class. We have already seen another finite description, the stack sorting machine.

We have another, slightly different problem.

Algorithmic Problem 2.4.7.

Does there exist an algorithm which does the following?

Input: A class given by some finite description.

Output: A basis for the class.

It is worth noting that this formulation does not fit entirely within our original framework. A more precise decision problem would be:

Algorithmic Problem 2.4.8.

Does there exist an algorithm which does the following?

Input: A class given by some finite description and a basis B .

Output: TRUE if B is the basis of the class, FALSE otherwise.

However, rather than descend into a discussion of problems and certificates, we shall allow both formulations. Hence we have chosen to use the term algorithmic problem, rather than decision problem. Our motivation, after all, is to precisely define the questions one might wish to ask about pattern classes; we are not, at this stage, interested in actually implementing these algorithms.

We will demonstrate shortly that to find an algorithm which answers Problem 2.4.7 is a hopeless task, the set of all permutations under involvement is not partially well ordered and so bases may be infinite. Nonetheless a good portion of this thesis will be devoted to special cases where a solution can be found.

Definition 2.4.9. We say that a closed class is *finitely based* if the set of basis elements is finite.

Example 2.4.10. The set of all increasing permutations is finitely based. Its basis is the single permutation 21.

Lemma 2.4.11. *Every finite closed class is finitely based.*

Proof. It follows from the the Erdős-Szekeres Theorem 2.2.1 that every finite closed class has a basis element which is an increasing permutation and a basis element which is a decreasing permutation. Since the remaining basis elements must also avoid these permutations it follows that the class is finitely based. \square

A simple analysis of stack sorting yields the following result.

Theorem 2.4.12 (Knuth [45]). *The set of stack sortable permutations is a closed class, with the single basis permutation 231. That is, a permutation is stack sortable if and only if it avoids the pattern 231.*

We leave the proof until later.

2.5 The Atomicity Problem

Definition 2.5.1. A closed class is said to be *atomic* if it cannot be expressed as a union of two proper subclasses.

Definition 2.5.2. Given a bijection f between two linearly ordered sets A and B the class $Sub(f)$ is the set of all permutations isomorphic to finite subsets of this bijection. For example, if $\exists i_1 < i_2 < i_3 < i_4 \in A$ with $f(i_1) < f(i_3) < f(i_4) < f(i_2) \in B$ then $1423 \in Sub(f)$.

Atomicity has been characterised in the permutation case by Atkinson, Murphy, and Ruškuc [14]; this is a special case of a theorem of Fraïssé [33] which applies to closed classes of all relational structures, see also Hodges [37, Section 7.1].

Theorem 2.5.3. *For a permutation class C the following are equivalent:*

1. C is atomic.

2. If α and β are permutations in C then there exists γ in C such that $\alpha \preceq \gamma$ and $\beta \preceq \gamma$. This is called the joint embedding property.
3. $C = \text{Sub}(f : A \rightarrow B)$ where A and B are linearly ordered sets and f is a bijection.

We again have an obvious decision problem.

Algorithmic Problem 2.5.4.

Does there exist an algorithm which does the following?

Input: A closed class given by some finite description.

Output: TRUE if the class is atomic, FALSE otherwise.

This problem, in such generality, is beyond the abilities of this author. However, much of this thesis will be directed toward answering it in specific cases. For the moment we settle for the following.

Example 2.5.5. The set of all increasing permutations is atomic. It can be defined as $\text{Sub}(f : \mathbb{N} \rightarrow \mathbb{N})$ where $f(x) = x$.

Example 2.5.6. The finite closed class $\{1, 12, 21, 132, 213, 231, 312, 2413\}$, see Example 1.4.3, is atomic. It can be defined as $\text{Sub}(2413)$. See Figure 1.6.

Example 2.5.7. The finite closed class $\{1, 12, 21, 123, 231, 321\}$, see Example 1.4.4, is not atomic. For example, the permutations 123 and 321 do not jointly embed. See Figure 1.7.

Theorem 2.5.8 (Knuth [45]). *The set of stack sortable permutations is atomic.*

Again the proof is delayed.

2.6 The Enumeration Problem

Definition 2.6.1. Given a set S of permutations, the number of permutations in S of length n is written $|S_n|$.

Definition 2.6.2. The *enumeration* of a closed class C is the sequence $(|C_n|)_n$, the sequence given by the number of permutations of each length, n , in the class.

Definition 2.6.3. A Wilfian formula for a sequence S_n is an algorithm which calculates the n^{th} element of the sequence in a time bounded by a polynomial in n .

Less formally, a Wilfian formula is any reasonable solution to an enumeration problem. For an elegant (and thought provoking) description of this concept see Wilf [61]. Two particularly important types of Wilfian formula we consider will be rational generating functions, see 3.4.2, and algebraic generating functions.

Algorithmic Problem 2.6.4.

Does there exist an algorithm which does the following?

Input: A closed class given by some finite description.

Output: A Wilfian formula for the sequence which enumerates it.

Again this is a problem which, at this level of generality, lies well beyond the range of this author's capabilities. Again much time will be spent on specific cases.

Example 2.6.5. The increasing permutations are enumerated by length by the Wilfian formula $|I_n| = 1$, where I_n denotes the set of increasing permutations of length n . Clearly there is just one increasing permutation of each length.

Definition 2.6.6. The *Catalan* numbers are the sequence of numbers:

$$\left(\binom{2n}{n} / (n+1) \right)_n.$$

Theorem 2.6.7 (Knuth [45]). *The stack sortable permutations are enumerated by the Catalan numbers.*

The proof is delayed, this time only for a moment.

Definition 2.6.8. A path in the plane from $(0, 0)$ to $(2n, 0)$ which uses only $(1, 1)$ and $(1, -1)$ as steps and never lies below the line $y = 0$ is called a *Dyck path*.

Example 2.6.9. Figure 2.3 shows a Dyck path.




Figure 2.3: A Dyck path

Lemma 2.6.10. *The number of Dyck paths of length $2n$ is the n^{th} Catalan number.*

Proof. We use André's reflection principle [7], see also Singmaster [55]. The total number of paths from $(0, 0)$ to $(2n, 0)$ is $\binom{2n}{n}$. Now consider those paths from $(0, 0)$ to $(2n, -2)$, there are $\binom{2n}{n-1}$ such paths. Any such path must cross the line $y = -1$, at the first such point reflect the remainder of the path in this line. This operation yields a path from $(0, 0)$ to $(0, 2n)$ which crosses the line $y = 0$, see figure 2.4. Furthermore, it is clear that we have

a bijection: every such path may be generated in this way. Thus the total number of Dyck paths is $\binom{2n}{n} - \binom{2n}{n-1} = \binom{2n}{n}/(n+1)$. \square




Figure 2.4: The path reflecting operation

Definition 2.6.11. A *stack word* is a word on the alphabet $\{a, b\}$ which if each a is interpreted as a push operation and each b is interpreted as a pop operation sorts some permutation σ .

Example 2.6.12. The permutation 1324756 is sorted by the stack word $abaabbabababb$. Figure 2.5 shows part of the sorting process.




Figure 2.5: Sorting the permutation 1324756; thus far the stack word prefix $abaabbaba$ has been applied.

We now set out to prove Theorems 2.3.5, 2.4.12, 2.5.8 and 2.6.7 by exploring the structure of stack sortable permutations further. The proof is essentially that given by Knuth [45].

Proof of the Theorems 2.3.5, 2.4.12, 2.5.8 and 2.6.7. The sorting algorithm for a stack is entirely deterministic. If the number on top of the stack is the next number to be output we pop, otherwise we push. Thus every stack sortable permutation has a unique stack word which sorts it. Furthermore it is clear that stack words are in one-to-one correspondence with Dyck paths, each a represents an up-step, each b a down-step, the condition that the path must never cross $y = 0$ is equivalent to saying that one cannot pop from an empty stack. Thus we have the enumeration result.

It is also clear that the sorting operation will fail if we see a 231 pattern, thus we have one half of Knuth's theorem. Let σ be a stack sortable permutation. Suppose that our sorting algorithm has failed. Return to the scene of the last pop operation. The next element to be output must lie in the stack, furthermore there must be some larger element above it on the stack. Thus if we take the first element to be popped, the next element to be output and the larger element above it on the stack we find a 231 pattern, completing out proof of Knuth's theorem.

We now delve deeper into the structure of σ . Clearly σ has a maximum element n , assume that n occurs in position $k + 1$ in the permutation. When σ is sorted n must go onto the bottom of the stack, since it must eventually leave the stack as the final pop operation. Thus the stack must be empty so that the numbers $1 \dots k$ must occur in positions $1 \dots k$. Thus the numbers $(k + 1) \dots (n - 1)$ must occur in positions $(k + 2) \dots n$. These two subpermutations must be stack sortable but are otherwise unrestricted, thus we have the joint embedding property, see figure 2.6.

In fact there is a far simpler way to see that the stack sortable permutations possess the joint embedding property. We simply pass two permutations through the stack, one after the other. \square




Figure 2.6: The decomposition of a stack sortable permutation.

2.7 The Partial Well Order Problem

Definition 2.7.1. A set of elements A belonging to a partially ordered set (S, \leq) is called an *antichain* if no pair of elements in A are comparable under \leq . In the case of permutations under involvement, if A is an antichain and α and β lie in A then neither $\alpha \preceq \beta$ nor $\beta \preceq \alpha$.

Notice immediately that the basis of any closed class must be an antichain. Basis elements must necessarily be incomparable under involvement.

Example 2.7.2. The set of all permutations of length k forms a finite antichain. Our first infinite antichain appears in Definition 2.8.5.

Definition 2.7.3. A set S with a partial order \leq is said to be *partially well ordered* if it contains no infinite antichains and no infinite descending chains.

Remark 2.7.4. For permutations under involvement an infinite descending chain is impossible, since if $\alpha \preceq \beta$ and $\alpha \neq \beta$ then α is necessarily shorter

than β . Thus a closed class is partially well ordered if it contains no infinite antichain.

We have another general decision problem.

Algorithmic Problem 2.7.5.

Does there exist an algorithm which does the following?

Input: A closed class given by some finite description.

Output: TRUE if the class is partially well ordered, FALSE otherwise.

Again the general case appears, at least to this author, hopeless. For now we settle for two simple examples followed by a more complex example.

Example 2.7.6. Every finite closed class is partially well ordered. There are only finitely many elements from which to construct any antichain.

Example 2.7.7. The class of all increasing permutations is partially well ordered, since every increasing permutation of length n or smaller is contained in the increasing permutation of length $n + 1$.

Next we will prove that the set of stack sortable permutations is partially well ordered, and in doing so prove a special case of Higman's theorem [36]. The proof is neither simple nor enlightening, thus all but the most diligent readers are advised to avoid it. They should return to the text at the Section 2.8.

Theorem 2.7.8 (Higman's Theorem (simplest case)).

The set of all words over a finite alphabet is partially well ordered.

Proof. The proof we give is due to Nash-Williams [52]. Let A be a finite alphabet of letters a_1, a_2, \dots, a_n . Let C be an infinite antichain of words under the subword ordering. Choose $C = c_1, c_2, \dots, c_k, \dots$ to be minimal

under the condition that $|c_i|$ is minimal such that c_1, c_2, \dots, c_i begins an infinite antichain.

Since C is infinite and A is finite there must be infinitely many words c_{k_1}, c_{k_2}, \dots which begin with the same first letter.

Let c'_k denote the word c_k with first letter removed.

Then $c_1, c_2, \dots, c_{k_1-1}, c'_{k_1}, c'_{k_2}, \dots$ is an infinite antichain contradicting the minimality of c_{k_1} . \square

This version of the theorem is sufficient to prove that certain closed classes are partially well ordered, for example all W -classes, see section 4.3. However, this approach fails for stack sortable permutations because, as we shall see, the operation which generates them is not associative.

Definition 2.7.9. An abstract algebra (A, M) is a set of elements A together with a set of operations M . Each operation μ in M is an n -ary operation for some positive integer n and maps each sequence (a_1, a_2, \dots, a_n) of n elements in A to some unique element $\mu(a_1, a_2, \dots, a_n)$. We will denote by M_n the set of all n -ary operations in M .

Definition 2.7.10. An abstract algebra (A, M) is said to be minimal if there exists no subset B of A such that (B, M) is an abstract algebra.

In what follows we shall recreate the notation of Higman [36] as far as possible, in order that this section may easily be compared with the original paper. In particular, in this setting, we do not allow abstract algebras to have generating sets, but instead allow 0-ary operations, or constants, in M .

Definition 2.7.11. Let (A, M) be an abstract algebra and let \leq be a partial order on the set of elements A . Let a and b be arbitrary elements of A . Let μ and λ be arbitrary r -ary operations in M . Let x, y, z be arbitrary sequences of elements in A , so that the expressions $\mu(x, a, y)$ and $\mu(z)$ make sense. Then (A, M) is said to be a *divisibility algebra* under \leq if the following conditions hold:

- $a \leq b \Rightarrow \mu(x, a, y) \leq \mu(x, b, y)$.

- $a \leq \mu(x, a, y)$.

In this case we say that \leq is a divisibility order on (A, M) . Furthermore, given partial orders on M_n we say that \leq is *compatible* with these partial orders if:

- $\lambda \leq \mu \Rightarrow \lambda(z) \leq \mu(z)$

Theorem 2.7.12 (Higman [36]).

Suppose that (A, M) is a minimal abstract algebra and that M_n , the set of n -ary operations in M , is partially well ordered for each n . Then (A, M) is partially well ordered under any divisibility ordering compatible with each M_n .

Theorem 2.7.13 (Knuth [45]). *The set of stack sortable permutations is partially well ordered.*

We will prove Theorem 2.7.13 and in doing so sketch the proof of Theorem 2.7.12, the complete proof is not substantially more complicated but neither is it substantially more interesting.

Proof. The set of permutations A avoiding 231 can be considered a minimal divisibility algebra by defining the following set of operations: $M = \{M_0 = \Omega, M_2 = x \hat{\oplus} y\}$.

We have chosen the notation $\hat{\oplus}$ rather than copying Higman's \oplus because the operation \oplus has come to have a specific meaning in the study of pattern classes, see Definition 2.7.14.

Consider Ω to be the empty permutation. If $|\alpha| = k$ and $|\beta| = l$ then $\alpha \hat{\oplus} \beta$ is the permutation $\alpha(k + l + 1)(\beta + k)$ where $\beta + k$ means to add k to each element of β . It is then clear that M generates A since every permutation avoiding 231 can be decomposed around its maximum, see figure 2.6. Thus (A, M) is a minimal algebra, that is, there is no proper subset of A which is closed under M .

When A is ordered under involvement it becomes a divisibility algebra (the operations respect the ordering) since:

- $x \preceq y \Rightarrow x \hat{\oplus} \alpha \preceq y \hat{\oplus} \alpha$
- $x \preceq y \Rightarrow \alpha \hat{\oplus} x \preceq \alpha \hat{\oplus} y$
- $x \preceq x \hat{\oplus} \alpha$
- $x \preceq \alpha \hat{\oplus} x$

For all α, x and y . Furthermore, involvement is trivially compatible with all operations, since there is just a single 2-ary operation and a single constant.

By way of a contradiction we will assume that A is not partially well ordered. Denote by $\text{cl}(\alpha)$ the upward closure of the element α , that is the set of all permutations which involve it. Let γ be of minimal length in A so that γ belongs to an infinite antichain. Then $A \setminus \text{cl}(\gamma)$ is not partially well ordered (since we can remove γ from the antichain). However, since γ cannot be the empty permutation, $\gamma = \alpha \hat{\oplus} \beta$ for some α and β and:

- $A \setminus \text{cl}(\alpha)$ is partially well ordered (since γ is minimal).
- $A \setminus \text{cl}(\beta)$ is partially well ordered (again since γ is minimal.)

Next we define two sets of unary operations:

- $f_{\alpha,a}(x) = a \hat{\oplus} x, \forall a \in A \setminus \text{cl}(\alpha)$
- $f_{\beta,b}(x) = x \hat{\oplus} b, \forall b \in A \setminus \text{cl}(\beta)$

We let $M' = M_0 \cup \{f_{\alpha,a} : \dots\} \cup \{f_{\beta,b} : \dots\}$. Then let (A', M') be the minimal sub-algebra of (A, M') (notice that we have reduced the “complexity” of our algebra by moving from binary to unary operations, this is the inductive step of Higman’s proof, repeat as required through the set of all operations).

We claim that A' is also not partially well ordered. We show, in fact, that $A = A' \cup \text{cl}(\gamma)$, hence $A \setminus \text{cl}(\gamma) \subseteq A'$. We have seen already that $A \setminus \text{cl}(\gamma)$ is not partially well ordered.

Notice first that $A' \cup \text{cl}(\gamma)$ is closed upwardly under M' since (A', M') is an algebra and $\text{cl}(\gamma)$ is an upwardly closed subset of A .

Let $d = x \hat{\oplus} y$ lie in A , furthermore let d be minimal outside $A' \cup \text{cl}(\gamma)$ so that x and y lie inside this set. If $\alpha \preceq x$ and $\beta \preceq y$ then $\gamma = \alpha \hat{\oplus} \beta \leq d$ so that $d \in \text{cl}(\gamma)$. Otherwise either $x \in A \setminus \text{cl}(\alpha)$ or $y \in A \setminus \text{cl}(\beta)$ which yields either $d = f(\alpha, x)y$ or $d = f(\beta, y)x$. Hence $d \in A' \cup \text{cl}(\gamma)$, and so $A \setminus \text{cl}(\gamma) \subseteq A'$.

Thus we have shown that if A is not partially well ordered then neither is (A', M') .

We complete our proof by demonstrating that (A', M') is, in fact, partially well ordered and so exhibiting a contradiction.


Let C be an infinite antichain chosen minimally as in 2.7.8. Each permutation in the antichain can be written $f_{\sigma_1, x_1}(f_{\sigma_2, x_2}(\dots(f_{\sigma_n, x_n})\dots))$ where $\sigma_i = \alpha$ or β and each $x_j \in A \setminus \text{cl}(\alpha) \cup A \setminus \text{cl}(\beta)$. Assume without loss that infinitely many begin f_{α, y_j} where the elements (y_j) form an increasing chain (we know that $(A \setminus \text{cl}(\alpha)) \cup (A \setminus \text{cl}(\beta))$ is a partially well ordered set). Performing the same trick as in the proof of Theorem 2.7.8 we consider the antichain formed by taking the first few members of C and then only those permutations with first elements from this chain and with their first elements removed. This new antichain contradicts the minimality of C . Hence A' is partially well ordered and our proof is complete. \square

Definition 2.7.14. The *direct sum* of two permutations is the permutation obtained by placing the second permutation above and to the right of the first. Thus if α has length k then the direct sum of α and β , written $\alpha \oplus \beta$ is the permutation $\alpha_1, \alpha_2, \dots, \alpha_k, \beta_1 + (k + 1), \beta_2 + (k + 1), \dots, \beta_n + (k + 1)$. See figure 2.8.

Dually we define the *skew sum* of two permutations, written $\alpha \ominus \beta$. See Figure 2.7.

Example 2.7.15. The direct sum of the permutations 2413 and 13245 is the permutation 241357689.

Definition 2.7.16. The permutations obtainable from 1 by repeated applications of \oplus and \ominus are said to be *separable*. Bose, Buss, and Lubiw [22] showed that the separable permutations are those that avoid 2413 and 3142.

Figure 2.7: The skew sum of α and β .


The class of separable permutations is also partially well ordered under involvement. In this case we have an algebra with a single constant, 1 and two binary operations, \oplus and \ominus . It is clear that we have a divisibility order under involvement. We make \oplus and \ominus incomparable so that involvement is compatible with their ordering, indeed in any situation where the number of operations is finite we may perform this trick.

2.8 Classes with a Single Basis Element of Length Three

The following is a consequence of the isomorphism laid out in Section 1.5.

Theorem 2.8.1. *The closed class of permutations C , defined as $\text{Av}(\alpha)$, where α is any single permutation from the set $\{231, 213, 132, 312\}$ has the following properties:*

- *C is in one to one correspondence with the stack sortable permutations.*

Figure 2.8: The direct sum of α and β .

- C is atomic.
- C is enumerated by the Catalan numbers.
- C is partially well ordered.

Proof. The proof follows immediately from the isomorphism operations, 312 is the inverse of 231, 132 is the reverse of 231 and 213 is the complement of 231. \square

Theorem 2.8.2. *The closed class of permutations D defined as $\text{Av}(\beta)$ where β is either 321 or 123 is also in one-to-one correspondence with stack sortable permutations.*

Proof. It is clear that 321 is the reverse of 123 so we consider only a single case. We present the standard bijection (see Bóna [21, Lemma 4.3]) which fixes left to right minima and yields a slightly stronger result.

Definition 2.8.3. An element i of a permutation σ is a *left-to-right minimum* if there is no element smaller than i which lies to the left of it in the

permutation. *Left-to-right maxima*, *right-to-left minima* and *right-to-left maxima* are defined mutatis mutandis.

Let σ be a permutation which avoids 123. Label the left-to-right minima of σ as (m_1, m_2, \dots, m_k) . Let S denote the sequence of elements of σ which are not left to right minima. It is immediately clear that S is a decreasing sequence. Form a new permutation τ by keeping the left to right minima (m_1, m_2, \dots, m_k) fixed and placing the elements of S from left to right at each step placing the smallest element which has not yet been placed but which is larger than the closest left to right minima on the left, see figure 2.9. We see immediately that τ is 132 avoiding and that τ is the only 132 avoider with this choice of left to right minima. Thus we have a bijection.

□




Figure 2.9: 123 and 132 avoiding permutations which are equivalent under our bijection

The proof of Theorem 2.8.2 extends naturally to a bijection between $\text{Av}(\sigma \oplus 21)$ and $\text{Av}(\sigma \oplus 12)$, see Babson and West [17]; Backelin, West and Xin [18] extend the proof further to a bijection between $\text{Av}(\sigma \oplus (k \dots 21))$ and $\text{Av}(\sigma \oplus (12 \dots k))$.

Theorem 2.8.4 (see Atkinson, Murphy and Ruškuc [14]). *The closed class of permutations D defined as $\text{Av}(\beta)$ where β is either 321 or 123 is atomic.*

Proof. It is clear that the direct sum of two permutations which avoid 321 also avoids 321. Thus the class of all permutations which avoid 321 posses the joint embedding property. \square

Definition 2.8.5. The set of permutations $\{u_1, u_2, u_3 \dots\}$ where $u_n = 235174 \dots (n-2)(n-5)(n-1)(n)(n-3)$ is called U , see Murphy [50].

Lemma 2.8.6. U is an infinite antichain.

Proof. Each member of U contains a copy of the pattern 2341 as its leftmost four elements and as its topmost four elements, furthermore there are no other copies of 2341. Thus any embedding of one member into another will have to match these two parts. It is immediately clear that the chain like structure of the remaining elements makes such an embedding impossible, see figure 2.10. These end patterns are known generally as *anchors*, the central structure is, in this case, called an *oscillation*. \square




Figure 2.10: Two members of the antichain U .

Theorem 2.8.7 (see Murphy [50]). The closed class of permutations D defined as $\text{Av}(\beta)$ where β is either 321 or 123 is not partially well ordered.

Proof. To prove that a class is not partially well ordered we have merely to exhibit an infinite antichain inside it. It is clear that every member of U avoids 321 thus the class $\text{Av}(321)$ is not partially well ordered. \square

The construction of just two more infinite antichains [50] is enough to yield the following:

Theorem 2.8.8.

There exists an algorithm which does the following:

Input: A closed class defined by a single basis element it avoids.

Output: TRUE if the class is partially well ordered, FALSE otherwise.

Proof. It is a result of Atkinson, Murphy and Ruškuc [14] that a closed class whose basis is a single permutation σ is partially well ordered if and only if $\sigma \in \{1, 12, 21, 132, 213, 231, 312\}$. \square

Chapter 3

Token Passing Networks

A Token passing network is a directed graph with a specified input node, or source, and a specified output node, or sink. The remaining vertices have one of six types, each of which stores data in a different way. Tokens travel through the network along the edges, starting at the input and finishing at the output, only one token may move at any time. We assume the tokens enter the network in ascending order. The order in which the tokens leave the network is the permutation generated by the network.

Finite token passing networks, those that can hold only finitely many tokens at any one time, have been studied at length. Such problems were first inspired by Knuth [45] who posed questions about systems of railway sidings. Atkinson and Tulley [11], Albert, Atkinson and Ruškuc [4] and finally Albert, Ruškuc and Linton [6] have all studied finite networks. We extend the concept, allowing networks to contain infinite components such as stacks and queues, in doing so we bring the work of Knuth [45], Pratt [53], Tarjan [59] and Atkinson, Murphy and Ruškuc [13] under the same definitional umbrella. We call such networks extended token passing networks. There have been many more variations of stack sorting and network sorting considered, for an overview see Bona [20].

3.1 Definitions for Token Passing Networks

Although the concept of a token passing network is an intuitive one, it will, nonetheless, require considerable effort to define formally. Readers familiar with these concepts are advised to skip immediately to Section 3.3.

Definition 3.1.1. An extended token passing network is a directed graph which satisfies the following conditions:

- There is a single specified input node, labelled I , which has zero in-degree.
- There is a single specified output node, labelled O , which has zero out-degree.
- The remaining nodes are either unlabelled or labelled D , ID , OD , S or Q .
- The edges are also labelled, each edge has a different label.

Nodes labelled D , ID , OD , S and Q represent, respectively, dequeues, input restricted dequeues, output restricted dequeues, stacks and queues. We do not need to worry about their exact operation until we have defined transitions in the network, but for now it is worth remembering that these nodes have unbounded capacity.




Figure 3.1: A Very Simple Network

So far we have only said what a token passing network looks like, we have not said anything about how tokens move through such networks or how tokens





Figure 3.2: Another Simple Network

are stored at nodes. We use an indirect approach, considering codewords before we consider permutations.

Definition 3.1.2. Given an extended token passing network N with edge set E we define an *alphabet* A as follows:

- If a is an edge which does not leave a node labelled D or ID and does not enter a node labelled D or OD then a is a member of A .
- If a is an edge which leaves a node labelled D or ID but does not enter a node labelled D or OD then a_1 and a_2 are members of A .
- If a is an edge which does not leave a node labelled D or ID but enters a node labelled D or OD then a^1 and a^2 are members of A .
- If a is an edge which leaves a node labelled D or ID and enters a node labelled D or OD then a_1^1 , a_2^1 , a_1^2 and a_2^2 are members of A .

As we shall see later on, the alphabet represents the set of all possible token moves. Thus if, for example, a particular node is a deque it requires two possible inputs and two possible outputs, which are represented by the subscripts and superscripts.



$$\text{Alphabet} = \{a, b, c^1, c^2, d_1, d_2\}.$$

Figure 3.3: A network and its alphabet.


For the network N we also define a set of states S which describe the number and position of tokens within the network.

Definition 3.1.3. A *state* is a set of pairs (v, i) where v is a node in N and i is a positive integer or zero, for an unlabelled node i is either 0 or 1, for a labelled node i can be any positive integer or 0. A state contains precisely one pair for each node in N except for the input and output nodes, we also define a special state e in which every i is 0. We call the value i the content of the node v .

For a network with only unlabelled nodes the set of states will be finite, otherwise it will be infinite.

There is a (partial) transition function F from a subset of $A \times S$ into S defined as follows.

Definition 3.1.4. F is defined on (a, s) if and only if the node at the start of the edge a does not have content 0 or is the input node and the node at the end of a is not an unlabelled node with content 1. If this is the case then $F(a, s) = t$ where t is identical to s except that the content of the node at



Possible states: $\{(X, 0), (Y, 0)\} = e$, $\{(X, 1), (Y, 5)\}$, $\{(X, 0), (Y, 1)\}$.

Figure 3.4: A network and some possible states.


the start of a is reduced by one unless it is the input node and the content of the node at the end of a is increased by one unless it is the output node.

Words on our alphabet allow us to move tokens through our network in the natural manner.

Definition 3.1.5. Let w be some word over the alphabet A . Let our initial state be e . Calculate a new state $F(a, e)$ where a is the first letter of w , call this state s_1 . Apply this repeatedly, thus if b is the m^{th} letter of w then the state s_m is $F(b, s_{m-1})$


If w is of length k and s_k is the state e then we say w is a *codeword* for N . If at any stage $F(b, s_{m-1})$ is undefined then w is an illegal word.

A transition describes a move of a single token across a single edge, from one node to another. Codewords describe sequences of transitions which move one or more tokens across the network, from source to sink. We are interested in the order in which the tokens reach the sink, relative to the order in which they left the source. In order to calculate this we need to know a bit more about the way in which our labelled nodes store tokens.



Let $S = \{(X, 0), (Y, 3)\}$.
 Then $F(S, c) = \{(X, 1), (Y, 3)\}$.
 Also $F(S, b) = \{(X, 0), (Y, 2)\}$.

Figure 3.5: A network and some possible transitions.



A legal codeword for the network above is $w = aacbabb$.

Figure 3.6: A network together with a legal codeword.

Each of our labelled nodes stores tokens in a different way, once we know how tokens are stored we can trace their paths across a network.

- A node labelled Q, called a queue, stores tokens in a first in first out

scheme.

- A node labelled S, called a stack, stores tokens in a first in last out scheme.
- A node labelled D, called a deque, stores tokens in a first or last in, first or last out scheme. An inward transition with superscript 1 sends a token to the start of the deque, an inward transition with superscript 2 sends a token to the end of the deque. An outward transition with a subscript 1 takes a token from the start of the deque, an outward transition with subscript 2 takes a token from the end of the deque.
- A node labelled OD, called an output restricted deque, stores tokens in a first or last in, first out scheme. Inward transitions work in the same way as the deque, outward transitions are always from the start.
- A node labelled ID, called an input restricted deque, stores tokens in a first in, first or last out scheme. Inward transitions are always to the start, outward transitions work in the same way as the deque.




Figure 3.7: A deque

Definition 3.1.6. Given a network N and a codeword w we say that w generates the permutation σ if the order in which the tokens arrive at the sink is precisely the order generated by applying the permutation σ to the order in which the tokens left the source.

In general it is useful to think of the tokens being labelled 1 to n and leaving the source in ascending order. However sometimes we will consider





Figure 3.8: The internal workings of a deque

the tokens arriving at the sink in ascending order having left the source in some other order, the tokens are then said to have been sorted. It is easy to see that if a codeword w generates σ then the same codeword will sort σ^{-1} . Also note that each transition will move a particular token, so that if w generates σ then we can generate any subpermutation of σ by deleting from w any letters which encode transitions moving tokens we wish to delete from σ .

From a computational point of view, in order to calculate the permutation generated by a particular word w we need to store more information than we need to check if w is a codeword. We alter our set of states so that the content associated with each vertex is not an integer but an ordered set or list, describing the tokens stored at that vertex at that particular time. We also associate a content with the output node, which will eventually hold the permutation we are generating. When we perform a transition we pop a token from the appropriate end of one list and push it onto another list. Unlabelled nodes are only allowed to hold zero or one token.

Lemma 3.1.7. *Every extended token passing network generates a closed class of permutations.*

Proof. Subpermutations can be generated using the methods described above.



The codeword $w = aacbabdb$ generates the permutation 2431.

Figure 3.9: A network together with a codeword and the permutation it generates.

□

A further proof of this fact can be found in [16].

Remark 3.1.8. It is worth noting that in this setting the subsequence definition of involvement, Definition 1.3.2, is the most obvious and most natural. It is no surprise then, since the study of pattern classes began as the study of sorting machines, that the subsequence definition is so widely used.

Lemma 3.1.9. *Every extended token passing network generates a closed class which is atomic.*

Proof. A joint embedding of two permutations can be found simply by taking a codeword which generates the first followed by a codeword which generates the second. □




Figure 3.10: A stack token passing network

3.2 Some Examples of Token Passing Networks

Example 3.2.1. The stack sorting machine considered in Section 2.3 may be viewed as a token passing network, see Figure 3.10. We have seen that a stack can sort all those permutations avoiding 231. By running this sorting operation backwards we see that, if the input is $n(n-1)\dots 21$, then the output must avoid 132. This is the reverse symmetry. Taking complements shows that if the input is $12\dots(n-1)n$ then the output must avoid 312, thus a stack token passing network can generate any permutations which avoid 312. It happens that 312 is the inverse of 231, furthermore it is easy to see that a permutation can be generated by a stack if and only if its inverse can be sorted.

Example 3.2.2. Consider the token passing network shown in Figure 3.11. This network is sometimes called a 2-stack since it operates like a stack with maximum capacity 2. A simple analysis of possible fail states shows that this network can generate all permutations which avoid 312 and 321. There are 2^{n-1} such permutations of length n .

Example 3.2.3. As a final example consider the network shown in figure 3.12. This network can generate every permutation of length less than or equal to k . It is known as a k -buffer. It is clear that this network cannot generate any permutation of length $k+1$ whose first element is $k+1$. It can be shown that these are all the basis elements for the class generated by this network.





Figure 3.11: A 2-stack token passing network

Figure 3.12: A k -buffer.

3.3 Algorithmic Problems for Token Passing Networks

We have the following algorithmic problems:

Algorithmic Problem 3.3.1.

Does there exist an algorithm which does the following?

Input: An extended token passing network.

Output: A Wilfian formula which enumerates the class generated by this network.

Algorithmic Problem 3.3.2.

Does there exist an algorithm which does the following?

Input: An extended token passing network.

Output: A finite description for the basis of the class generated by this network.

Algorithmic Problem 3.3.3.

Does there exist an algorithm which does the following?

Input: An extended token passing network together with a permutation π .

Output: TRUE if π can be generated by this network, FALSE otherwise.

Algorithmic Problem 3.3.4.

Does there exist an algorithm which does the following?

3.3. ALGORITHMIC PROBLEMS FOR TOKEN PASSING NETWORKS49

Input: An extended token passing network.

Output: TRUE if the class generated by this network is partially well ordered, FALSE otherwise.

The question of partial well order remains entirely untouched for token passing networks. Certainly there exist networks which generate classes which are not partially well ordered, classifying them seems likely to prove immensely difficult.

Algorithmic Problem 3.3.5.

Does there exist an algorithm which does the following?

Input: Some finite description of a closed class C .

Output: An extended token passing network which generates this class if such a network exists.

Definition 3.3.6. An extended token passing network which can generate every permutation of any length is said to be *complete*. Otherwise a token passing network is said to be *incomplete*.

We will prove the following new result.

Theorem 3.3.7.

There exists an algorithm which does the following:

Input: An extended token passing network.

Output: TRUE if the class generated by this network is complete, FALSE otherwise.

3.4 Incomplete Networks

We first state known results about two special types of token passing networks, those without infinite components and those without cycles of directed edges, both of which are incomplete.

3.4.1 Finite Token Passing Networks

A finite token passing network is a network with no infinite components. Such networks have been widely studied in [11], [4] and [6], there are well understood methods for working with them. We have the following results.

Lemma 3.4.1. *Any finite token passing network is incomplete.*

Proof. A finite token passing network contains only finitely many nodes and no infinite components, thus it can hold only finitely many tokens at any one time. It is clear that such a network containing n nodes cannot generate any permutation which begins with $(n + 1)$ and hence is incomplete. \square

Definition 3.4.2. The *rank encoding* of a permutation is generated by replacing each element by its value relative to those elements which come after it.

Example 3.4.3. The permutation 2761453 has rank encoding 2651221.

Lemma 3.4.4. *Any permutation generated by a finite token passing network with n nodes has a rank encoding on the alphabet $\{1, \dots, n\}$.*

Proof. Follows from the proof of Lemma 3.4.1. \square

Following the work of Albert, Atkinson and Ruškuc [4] we define Ω_k to be the set of all permutations with rank encodings over the alphabet $\{1, \dots, k\}$. Let $E(\Omega_k)$ be the set of such encodings. It is easy to see that $E(\Omega_k)$ is a regular language. It is also easy to enumerate Ω_k , see that it has basis equal to the set of all permutations which begin $(k + 1)$ and see that Ω_k may be considered as the output of a k -buffer, see Figure 3.12.

Atkinson, Livesey and Tulley [11] proved that, under the rank encoding, the set of permutations generated by a finite token passing network is always a regular language.

Albert, Atkinson and Ruškuc [4] went further, giving an algorithm to construct an automaton accepting the basis of a closed class generated by a finite token passing network.

Finally Albert, Ruškuc and Linton [6] were able to show that relatively few classes can be generated by token passing networks. It should be mentioned however that the methods used are computationally expensive, only those networks with rank encodings on $\{1, 2, 3\}$ were completely classified, so these results remain more theoretical than practical.

These results allow us to give the following theorems.

Theorem 3.4.5.

There exists an algorithm which does the following:

Input: A finite token passing network and a permutation π .

Output: TRUE if π can be generated by this network, FALSE otherwise.

Theorem 3.4.6.

There exists an algorithm which does the following:

Input: A finite token passing network.

Output: A rational generating function which enumerates the class generated by this network.

Theorem 3.4.7.

There exists an algorithm which does the following:

Input: A finite token passing network.

Output: A regular expression for the basis elements
of the class generated by this network.

3.4.2 Regular Languages, Finite State Automata and Rational Generating Functions

Regular languages are very nice objects to work with combinatorially. They are in one to one correspondence with finite state automata and hence have linear membership tests, furthermore they have rational generating functions. For details see, for example, Hopcroft and Ullman [38].

Given a sequence $A_n = (a_0, a_1, a_2, \dots, a_i, \dots)$, its ordinary generating function is the formal power series:

$$G(A_n, x) = \sum_{n=0}^{\infty} a_n x^n$$

The analytical properties of this formal power series give specific information about the sequence. “A generating function is a clothesline on which we hang up a sequence of numbers for display.” Wilf [62].

A generating function is said to be rational if it can be expressed in the form:

$$G(A_n, x) = \sum_{n=0}^{\infty} a_n x^n = \frac{P(x)}{Q(x)}$$

where $P(x)$ and $Q(x)$ are polynomials in x .

The crucial theorem for rational generating functions is the following.

Theorem 3.4.8. *See Stanley [58, Theorem 4.1.1]*

Let $A_n = (a_1, a_2, \dots, a_i, \dots)$ be a sequence of complex numbers and let $\alpha_1, \alpha_2, \dots, \alpha_d$ be fixed complex numbers, $d \geq 1$, $\alpha_d \neq 0$. Then the following conditions are equivalent:

1.

$$G(A_n, x) = \sum_{n=0}^{\infty} a_n x^n = \frac{P(x)}{Q(x)}$$

where $Q(x) = 1 + \alpha_1 x + \alpha_2 x^2 + \dots + \alpha_d x^d$ and $P(x)$ is a polynomial in x of degree less than d .

2. For all $n \geq 0$,

$$a_{n+d} + \alpha_1 a_{n+d-1} + \alpha_2 a_{n+d-2} + \dots + \alpha_d a_n = 0$$

Essentially this theorem tells us that for a sequence with a rational generating function the denominator $Q(x)$ encodes a recurrence relation for the sequence whilst the numerator $P(x)$ encodes the initial conditions.

A second important result, again see Stanley [58, Theorem 4.1.1] is that the poles of the generating function determine the asymptotic behaviour of the sequence.

Algebraic generating functions are a generalisation of rational generating functions. Again it is possible to learn about the behaviour of a sequence with an algebraic generating function by studying analytical properties of the generating function.

3.4.3 Acyclic Networks

Definition 3.4.9. A network which contains no cycles of directed edges is said to be acyclic.

In such a network the length of any codeword will be restricted by the number of edges and the number of tokens being moved, since each token can use each edge at most once. This property allows us to prove the following lemma.

Lemma 3.4.10. *Every acyclic network is incomplete.*

Proof. The proof below is essentially the same as that given by Tarjan [59], we repeat it for completeness.

Let N be an acyclic network with E edges, the associated alphabet will have at most $4E$ letters. Consider all codewords which generate permutations of length l . Every such codeword has length at most $4El$. Then the number of permutations of length l that N can generate is less than or equal to the number of codewords of length $4El$ which is less than or equal to the number of words of length $4El$ on an alphabet of $4E$ letters.

The number of such words is $4E^{4El}$. The total number of permutations of length l is $l!$ which is greater than $4E^{4El}$ for large enough l . Hence there are some permutations N cannot generate and N is incomplete. \square

3.5 Complete Networks

There are two groups of simple networks which are complete. These are the infinite loop networks, which we will call IL and the dual stacks networks, which we will call DS . In some sense these are the only complete networks, all other complete networks contain them, as we will show. These networks are defined by Figures 3.13 and `refdsfig`.

Definition 3.5.1. Consider two networks M and N . M is said to be a *subnetwork* of N if N can be constructed from M by repeatedly applying any of the following steps.

- A new edge may be added, connecting two existing nodes.
- A new node may be added.
- An edge may be split into two by inserting a node halfway along it.
- Every unlabelled node may be labelled Q or S .
- Every node labelled Q or S may be relabelled ID or OD .
- Every node labelled ID or OD may be relabelled D .




Figure 3.13: The networks IL.

Under such a construction N and M will have the same input and output vertices.

Proposition 3.5.2. *Given two networks N and M with M a subnetwork of N together with a codeword w on M which generates the permutation σ then there exists a codeword v on N which also generates σ .*

Proof. It is clear that if two nodes are adjacent in M then there will be a path between them in N and that this path will not visit any of the nodes from M . Thus for each edge in M find a path in N which connects the appropriate vertices. Take the codeword w and replace each letter with the letters describing the path in N . This new codeword will then generate σ . \square

Corollary 3.5.3. *Any network which contains a complete subnetwork is complete.*

Theorem 3.5.4. *Every complete network contains a member of either IL or DS as a subnetwork.*

Before we can prove this theorem we need to prove several lemmas allowing us to extend and decompose networks, we do so in the next section.




Figure 3.14: The networks DS.






Figure 3.15: Constructing a network from a subnetwork.



Under the codeword $w = aababb$ the above network generates the permutation 231.



Under the extended codeword $w = ababcabcc$ the above network also generates 231.

Figure 3.16: Extending a codeword from a subnetwork into a network.

3.5.1 Network Composition

Networks can be composed in two essentially different ways, in series and in parallel.




Figure 3.17: Parallel composition of networks.

The following two lemmas are corollaries of Proposition 3.5.2.

Lemma 3.5.5. *The parallel composition of a complete network and any network is complete.*

Lemma 3.5.6. *The serial composition of a complete network and any network which contains a path from its input node to its output node is complete.*




Figure 3.18: Serial Composition of networks.

Notice that for parallel composition any network will suffice, whereas for serial composition we need to rule out certain pathological networks, dead-end networks which simply do not allow tokens to pass. We also prove two more lemmas on network composition.

Lemma 3.5.7. *The parallel composition of two incomplete networks is incomplete.*

The proof rests on the following construction.

Definition 3.5.8. The *wreath* product of two permutations α and β , written $\alpha \wr \beta$, is the permutation generated by replacing each point in α with a copy of β .

Example 3.5.9. The wreath product of the permutation 231 with the permutation 12, $231 \wr 12$, is the permutation 345612. The wreath product of 12 with 231 is 231564, note that the wreath product is not commutative.

Lemma 3.5.10. *If N is a token passing network which cannot generate the permutation $\alpha = \alpha_1\alpha_2 \dots \alpha_k$ and M is a token passing network which cannot generate the permutation $\beta = \beta_1\beta_2 \dots \beta_l$ then the network formed by the parallel composition of N and M cannot generate the permutation $\alpha \wr \beta$.*

Proof. Let N is a token passing network which cannot generate the permutation $\alpha = \alpha_1\alpha_2 \dots \alpha_k$. Let M is a token passing network which cannot generate the permutation $\beta = \beta_1\beta_2 \dots \beta_l$. Suppose further that the parallel composition of N and M can generate the permutation $\alpha\beta$. $\alpha\beta$ may be written as $\alpha_1\beta_1, \alpha_1\beta_2, \dots, \alpha_1\beta_l, \alpha_2\beta_1, \alpha_2\beta_2, \dots, \alpha_2\beta_l, \dots, \alpha_k\beta_1, \alpha_k\beta_2, \dots, \alpha_k\beta_l$. Call those elements corresponding to a single element in α a block of $\alpha\beta$. Clearly M cannot generate any single block, since each block is order isomorphic to β , thus at least one element of each block must pass through N , thus we can identify k elements which pass through as N , one from each block. The subpermutation formed by these elements is order isomorphic to α and has been generated by the network N , a contradiction. This proof can also be found in Atkinson and Beals [16]. \square

Clearly Lemma 3.5.7 follows as a corollary.

Lemma 3.5.11. *The serial composition of two incomplete networks is incomplete.*

Proof. If two networks, N and M are composed in series then it is easy to see that the class they generate is the composition of the two classes generated by the two networks. Hence the number of permutations which can be generated is at most the size of the two classes multiplied together.

Next we use the result of the Stanley-Wilf conjecture, proved by Marcus and Tardos [49], that for every pattern class there exists a constant C such that the number of permutations of length l in the class is less than or equal to C^l . Suppose then we compose the networks N and M in series with bounds C^l and D^l . Then we can generate at most $(CD)^l$ permutations, which will be less than $l!$ for some l .

Thus such a composition is incomplete. \square

Remark 3.5.12. Note that this is a far weaker result than the one we have for parallel composition, it offers no construction for a permutation which cannot be generated, it merely shows that such a permutation exists. It is my opinion that a method for constructing such permutations would be

a necessary first step toward answering the basis question for such compositions. That said, the basis question for parallel composition remains unsolved, indeed, Albert, Ruškuc and Linton [6], exhibit a network which is the parallel composition of two stacks of size two, yet is infinitely based.

3.5.2 Parallelizing Networks

Consider any network N . Consider all the paths through the network from I to O which do not return to a point they have already visited, i.e. do not follow a loop. There are obviously finitely many such paths. Construct a new network $p(N)$ where a copy of each such path runs from the input to the output without intersecting any other, i.e. all the paths are connected in parallel.




Figure 3.19: The network $p(N)$.

Definition 3.5.13. We say that one node is *strongly connected* to another if a token may leave the first node, travel through the network to the second and then return to the first, i.e. the two nodes are connected by one or several cycles.

Thus we may divide the network into a set of distinct strongly connected components, each consisting of one or more nodes.




Figure 3.20: A network divided into strongly connected components

For any path through the network we may specify the set of strongly connected components it passes through, together with an in-vertex and an out-vertex for each component. It is clear that such components will not overlap along the path, otherwise two such components would be strongly connected and hence a single component.


Along each path in $p(N)$ place a copy of each strongly connected component, so that the in and out-vertices of the component lie in the correct place on the path and the path passes through the component in the natural way.

$P(N)$ will be a group of paths connected in parallel, each path being a group of strongly connected components in series. We term $P(N)$ the *path-parallel expansion* of N .

We have the following results linking N and $P(N)$.

Lemma 3.5.14. *Each strongly connected component in $P(N)$, together with an input linked to the in-vertex and an output linked to the out-vertex, is a subnetwork of N .*


Proof. We construct the remainder of the path between the input and in-



Lemma 3.5.15. *A path in $P(N)$ is complete if and only if at least one strongly connected component subnetwork inside it is complete.*

Lemma 3.5.16. *$P(N)$ is complete if and only if some path through it is complete.*

Lemma 3.5.17. *If $P(N)$ is complete then N is complete.*

Figure 3.22: The general form of the network $P(N)$

Proof. This follows from Lemmas 3.5.14, 3.5.15 and 3.5.16. \square

Lemma 3.5.18. *If N can generate the permutation σ then so can $P(N)$.*

Proof. Let σ be generated on N by the codeword w . Each token will follow a particular path through N under w , there will be a copy of this path across $P(N)$. Construct a new codeword on $P(N)$ which moves the tokens along the correct paths moving the tokens in precisely the order that w moves them through N . Such a codeword will clearly generate σ on $P(N)$. \square

Corollary 3.5.19. *$P(N)$ is complete if and only if N is complete.*

These results mean we need only consider the strongly connected components of a network to establish if it is complete.

3.5.3 Proof of the Main Theorem

We are now ready to prove Theorem 3.5.4, that a network is complete if and only if it contains an infinite loop or a pair of strongly connected stacks.

Proof. By Corollary 3.5.19 we may consider the network $P(N)$ instead of N . By Lemmas 3.5.15 and 3.5.16 we may consider only the strongly connected components of a network. We consider only those strongly connected components containing two or more nodes, i.e. those which contain a cycle, since acyclic networks are incomplete. There are just four cases to consider: that the strongly connected component has no infinite nodes; that it contains a single stack but no other infinite nodes; that it contains two stacks and that it contains a queue. Note that a component containing any types of deque necessarily contains a queue as a subnetwork.

Case 1: No infinite components Such a network is finite, and hence incomplete.

Case 2: A single stack Suppose the strongly connected structure contains n single nodes and a single stack. Such a structure cannot generate the permutation $\sigma = (3n+2)1234...(3n+1)$, which we establish by the following argument. All the tokens must be placed in the network at the same time so that $(3n+2)$ may leave first. Since only n tokens may fit on the nodes at least one of the first $n+1$ tokens must be placed on the stack, call this token a . If more than n tokens greater than a are placed on the stack above a then it is clear that we cannot generate σ . However we still have $2n$ tokens greater than a to place in the network before the token $(3n+2)$ can pass through. At most n of these tokens can fit on the nodes and hence n of them must go on the stack above the token a , hence σ cannot be generated.

Case 3: Two stacks A strongly connected structure containing 2 stacks is clearly complete and contains a DS subnetwork.

Case 4: A queue A loop structure containing an infinite queue is clearly complete and contains an *IL* subnetwork.

□

There is a structural similarity between *IL* and *DS*, cases three and four, and Turing machines, the looping queue or the pair of stacks play the role of the Turing machine's infinite tape. We have seen that finite token passing networks, of which networks belonging to case one are examples, can be analysed using finite state automata. Case two, networks with a single stack, seem ideally constructed for analysis using push down automata since they contain a series of nodes and a single stack, further work in this area certainly seems worthwhile. An initial problem is the necessity to encode permutations over a finite alphabet, for finite networks a rank encoding can be used, but this does not work for those networks which contain a stack. One possible line of enquiry seems to be the insertion encoding developed recently by Albert, Ruškuc and Linton [5], however the question of whether such networks generate classes which can be encoded as context free languages is still an open one.

3.6 Conclusions for Token Passing Networks

We finish with several open questions, this author can offer no guarantee of either depth or difficulty.

- Does there exist an incomplete finitely based network containing an infinitely based subnetwork?
- If a network N is finitely based is its path-parallel expansion, $P(N)$, finitely based?

The concept of subnetworks generates a partial order on token passing networks. We can generate a second partial order by considering the power of

such networks. The network N is said to be at least as powerful as the network M if N can generate every permutation M can. Two networks which generate exactly the same set of permutations are said to be equivalent in power. There are some interesting questions that these concepts raise.

We call a network minimal if it contains no subnetwork of equivalent power.

- Do there exist two different minimal networks which are equivalent in power? (With the exception of *IL* and *DS*)
- Is there an algorithm to reduce a network to a minimal subnetwork of equivalent power?

Finally, and to motivate a move away from sorting machines, a quote from Pratt [53]: “from an abstract point of view, the [containment order] on permutations is even more interesting than the networks we were characterising.” The remainder of this thesis will attempt to support this assertion.

Chapter 4

Grid Classes

In all that follows we consider a permutation, first and foremost, as a pair of linear orders. Indeed, we might more properly define permutations as equivalence classes of sets of points in the real plane. Involvement is a very natural order in this setting, it corresponds precisely to deleting points.

4.1 Definitions for Grid Classes

We begin with the point set definition of involvement, this has already been seen in Section 1.3.1, however we repeat it here for completeness.

Definition 4.1.1. Following Felsner [32] we say that a set of points in the plane is *generic* if no two points are aligned either vertically or horizontally.

Almost every point set we use will be generic.

Definition 4.1.2. A generic point set $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ is said to be *involved* in a generic point set, $T = \{(z_1, w_1), \dots, (z_m, y_m)\}$, written $S \preceq T$ if there is a one to one mapping, f , from $\{1, \dots, n\}$ into $\{1, \dots, m\}$, satisfying the following conditions:

- If $x_i < x_j$ then $z_{f(i)} < z_{f(j)}$;

- If $y_i < y_j$ then $w_{f(i)} < w_{f(j)}$.

Involvement on generic point sets is a pre-order, it is reflexive and transitive.

Definition 4.1.3. We say that two generic point sets, S and T , are *order isomorphic* if and only if $T \preceq S$ and $S \preceq T$.

It is clear that order isomorphism is an equivalence relation. We factor the set of all generic point sets under involvement by this equivalence. We choose our class representatives to be those generic point sets with consecutive integer coordinates beginning with one. The partially ordered set we create is (isomorphic to) the set of permutations under involvement. Indeed we need only factor the set of all generic sets ordered as subsets by our equivalence to get the set of permutations under involvement, a natural ordering indeed.

Definition 4.1.4. The *permutation image* of a generic set S , denoted $\Pi(S)$, is the permutation whose projection onto the plane is order isomorphic to S .

Example 4.1.5. The permutation image of the generic set shown in Figure 4.1 is 13485276.




Figure 4.1: A generic set with permutation image 13485276

Lemma 4.1.6. Given two generic sets $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ and $T = \{(z_1, w_1), \dots, (z_m, y_m)\}$ the following conditions are equivalent:

1. S and T are order isomorphic.
2. There exist a bijection g between the two sets $\{1, 2, \dots, n\}$ and $\{1, 2, \dots, m\}$ such that
 - $x_i < x_j$ if and only if $z_{g(i)} < z_{g(j)}$;
 - $y_i < y_j$ if and only if $w_{g(i)} < w_{g(j)}$.
3. $\Pi(S) = \Pi(T)$.

Lemma 4.1.7. *If S is a generic set and $T \preceq S$ then $\Pi(T) \preceq \Pi(S)$; in particular if $T \subseteq S$ then $T \preceq S$ and $\Pi(T) \preceq \Pi(S)$. Furthermore if $\tau \preceq \sigma = \Pi(S)$ then there exists $T \subseteq S$ such that $\tau = \Pi(T)$.*

Definition 4.1.8. An $(r \times s)$ gridding is a pair of sequences of real numbers $((v_1, v_2, \dots, v_{r-1}), (h_1, h_2, \dots, h_{s-1}))$ where $v_1 < v_2 < \dots < v_{r-1}$ and $h_1 < h_2 < \dots < h_{s-1}$. For technical reasons we define $v_0 = h_0 = -\infty$ and $v_r = h_s = \infty$.

Intuitively we think of a gridding as a set of $(r - 1)$ distinct vertical lines, $x = v_1, x = v_2, \dots, x = v_{r-1}$, and $(s - 1)$ horizontal lines, $y = h_1, y = h_2, \dots, y = h_{s-1}$.

Definition 4.1.9. A *gridded set* is an ordered pair (S, G) where $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ is a generic set and $G = ((v_1, v_2, \dots, v_{r-1}), (h_1, h_2, \dots, h_{s-1}))$ is a gridding such that $x_i \neq v_j, y_i \neq h_k, (1 \leq i \leq n, 1 \leq j \leq r - 1, 1 \leq k \leq s - 1)$.

Given a gridded set, (S, G) , we will often be interested in various subsets of S which it defines. We define three types of subset: vertical strips, horizontal strips and cells.

Definition 4.1.10. Given a gridded set, (S, G) , where $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ is a generic set and $G = ((v_1, v_2, \dots, v_{r-1}), (h_1, h_2, \dots, h_{s-1}))$ is a gridding, the *vertical strip*, $V_i(S, G)$, is the set $\{(x_k, y_k) \in S : v_{i-1} < x_k < v_i\}$, (for $i = 1, \dots, r$).

Definition 4.1.11. Given a gridded set, (S, G) , where $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ and $G = ((v_1, v_2, \dots, v_{r-1}), (h_1, h_2, \dots, h_{s-1}))$, the *horizontal strip*, $H_j(S, G)$, is the set $\{(x_k, y_k) \in S : h_{j-1} < x_k < h_j\}$, (for $j = 1, \dots, s$).

Definition 4.1.12. Given a gridded set, (S, G) , where $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$ and $G = ((v_1, v_2, \dots, v_{r-1}), (h_1, h_2, \dots, h_{s-1}))$, the *cell*, $C_{i,j}(S, G)$ is the set $V_i(S, G) \cap H_j(S, G)$.

Definition 4.1.13. Given two gridded sets, (S, G) and (S', G') , where $S = \{(x_1, y_1), \dots, (x_n, y_n)\}$, $S' = \{(x'_1, y'_1), \dots, (x'_m, y'_m)\}$, $G = ((v_1, v_2, \dots, v_{r-1}), (h_1, h_2, \dots, h_{s-1}))$ and $G' = ((v'_1, v'_2, \dots, v'_{r-1}), (h'_1, h'_2, \dots, h'_{s-1}))$, we say that (S, G) and (S', G') are isomorphic if and only if S and S' are isomorphic and $|C_{i,j}(S, G)| = |C_{i,j}(S', G')|$, ($1 \leq i \leq r, 1 \leq j \leq s$).

Remark 4.1.14. Equivalently we may say that:

- $(S, G) \cong (S', G')$ if and only if $S \cong S'$ and $|V_i(S, G)| = |V_i(S', G')|$, $|H_j(S, G)| = |H_j(S', G')|$, ($1 \leq i \leq r, 1 \leq j \leq s$).
- (Harder) $(S, G) \cong (S', G')$ if and only if $V_i(S, G) \cong V_i(S', G')$ and $H_j(S, G) \cong H_j(S', G')$, ($1 \leq i \leq r, 1 \leq j \leq s$).

However cell isomorphism is not enough, $C_{i,j}(S, G) \cong C_{i,j}(S', G')$ does not imply $(S, G) \cong (S', G')$.

Example 4.1.15. Two isomorphic gridded sets, both with permutation image 13485276 are shown in Figure 4.2.

Lemma 4.1.16. *Given a gridded set, (S, G) , and a generic set, T , isomorphic to S , we can find a gridding, H , so that (T, H) is a gridded set isomorphic to (S, G) .*

Lemma 4.1.17. *Given an $(r \times s)$ gridded set, (S, G) , and an $(r \times s)$ gridding, H , then there exists a generic set, T , such that $(T, H) \cong (S, G)$.*

Proof. This follows from Lemma 4.1.7, we simply insert into the correct cell, respecting the overall ordering on the generic set as we proceed. \square




Figure 4.2: Two isomorphic gridded sets

Definition 4.1.18. An $(r \times s)$ *gridding matrix*, M , is an $(r \times s)$ matrix of 1s, -1 s or 0s.

We want our matrix entries to control cells in our griddings and so we choose, somewhat non-standardly, to index our matrices from the bottom left. Thus a (3×2) matrix M would be indexed $M = \begin{pmatrix} M_{1,2} & M_{2,2} & M_{3,2} \\ M_{1,1} & M_{2,1} & M_{3,1} \end{pmatrix}$.

Definition 4.1.19. Let M be an $(r \times s)$ gridding matrix. An M -*gridded set* is a gridded set (S, G) where G is an $(r \times s)$ gridding and (S, G) satisfies the following conditions:

- If $M_{i,j} = 0$ then $C_{i,j}(S, G) = \emptyset$.
- If $M_{i,j} = 1$ then $C_{i,j}(S, G)$ is strictly increasing.
- If $M_{i,j} = -1$ then $C_{i,j}(S, G)$ is strictly decreasing.

Definition 4.1.20. A generic set, S , is said to *admit* an M -gridding if there exists a gridding, G , for which (S, G) is an M -gridded set.

Definition 4.1.21. The *grid class* of the gridding matrix, M , denoted $\text{Grid}(M)$, is the set of all permutations, σ , which (when considered as generic sets) admit an M -gridding.




Figure 4.3: A gridded set whose permutation image is 4162357

Example 4.1.22. Figure 4.3 shows a $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ gridding of a generic set whose permutation image is 4162357, demonstrating that the permutation 4162357 lies in the grid class $\text{Grid} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$.

Definition 4.1.23. We say that a matrix, N , is a *submatrix* of a matrix, M , if N can be obtained from M by erasing some rows and columns and by changing any number of 1s or -1 s to 0s.

Lemma 4.1.24. Let M be an $(r \times s)$ gridding matrix, and let N be a submatrix of M . For every N -gridded set, (S, G) , there exists an $(r \times s)$ gridding, H , such that (S, H) is an M -gridded set. In particular $\text{Grid}(N) \subseteq \text{Grid}(M)$.

Proof. Suppose $N = M(i_k, j_l)_{p \times q}$. For each $x = 1, \dots, p - 1$ replace the x^{th} vertical line by $i_{x+1} - i_s$ vertical lines sufficiently close together that no point of the generic set lies between them. Perform the analogous operation on the $(q - 1)$ horizontal lines of G . \square

Lemma 4.1.25. Let M be a gridding matrix and S a generic set. If (S, G) is an M -gridded set and $T \subseteq S$ then (T, G) is an M -gridded set.

Proof. Removing points from any cell will not violate any on the conditions imposed by M . \square




Figure 4.4: White space in an increasing cell

Corollary 4.1.26. *Grid(M) is a closed class.*

Lemma 4.1.27. *If (S, G) is an M -gridded set and (T, H) is order isomorphic to (S, G) then (T, H) is an M -gridded set.*

Corollary 4.1.28. *If (S, G) is an M -gridded set and T is a generic set involved in S then there exists a gridding, H , such that (T, H) is an M -gridded set.*

Definition 4.1.29. Let M be a gridding matrix and (S, G) an M -gridded set. The *white space* of (S, G) is the set of all points (x, y) such that $(S \cup (x, y), G)$ is also an M -gridded set.

The white space of a M -gridded set is, in fact, a set of open rectangles, whose opposite corners are defined by adjacent points within a cell and by the corners of the cell itself. See Figure 4.4.

So, for any two adjacent points in a cell there is a white space rectangle between them, and a point may be inserted into it, yielding another M -gridded set with the same gridding. By a repeated application of this process we obtain:

Proposition 4.1.30. *Let (S, G) be an M -gridded set and let (T, G) be the gridded set obtained from (S, G) in the following way: for every white space*

rectangle $R \subseteq C_{ij}(S, G)$ insert an increasing set if $M_{ij} = 1$ or a decreasing set if $M_{ij} = -1$, then (T, G) is an M -gridded set.

Proof. It is clear that these insertions are in harmony with the conditions imposed by M . \square

Definition 4.1.31. A *Picture* is a set of points in the plane.

Generally we will consider a picture to be infinite set of points. We will be interested in its finite subsets.

Definition 4.1.32. The set of all permutations isomorphic to finite generic sets which are subsets of a picture, P , is called the *picture class* of P , denoted $Sub(P)$.

Picture classes have a beautiful structure of their own, but for the time being we will resist their temptations and make use of them only when absolutely necessary. A weak minded reader may find some relief scanning Chapter 5

Definition 4.1.33. A *gridded picture* is an ordered pair (P, G) where P is a picture and $G = ((v_1, v_2, \dots, v_{r-1}), (h_1, h_2, \dots, h_{s-1}))$. For technical reasons we allow points from dense subsets of P to lie on grid lines.

Definition 4.1.34. Given a gridded picture we define *cells* exactly as in Definition 4.1.12.

Definition 4.1.35. Two gridded pictures are said to be *isomorphic* precisely if they satisfy the conditions laid out in Definition 4.1.13.

Definition 4.1.36. Let M be an $(r \times s)$ gridding matrix. An *M -gridded picture* is a gridded picture, (P, G) , where G is an $(r \times s)$ gridding and (P, G) satisfies the following conditions:

- If $M_{i,j} = 0$ then $C_{i,j}(P, G) = \emptyset$.
- If $M_{i,j} = 1$ then $C_{i,j}(P, G)$ is strictly increasing.
- If $M_{i,j} = -1$ then $C_{i,j}(P, G)$ is strictly decreasing.

Definition 4.1.37. A picture P is said to *admit* an M -gridding if there exists a gridding, G , for which (P, G) is an M -gridded picture.

Lemma 4.1.38. *Let M be a gridding matrix and let P be an M -gridded picture. Then the picture class, $\text{Sub}(P)$, is a subclass of the grid class $\text{Grid}(M)$.*

Proof. Clearly if P admits an M gridding then so does every finite generic subset of P . \square

Definition 4.1.39. Given an M -gridded picture P , the $H_{i,j}$ -branch is the continuous function obtained by connecting points in the cell $C_{i,j}(P, G)$. If $M_{i,j} = 1$ then $H_{i,j}$ will be increasing, If $M_{i,j} = -1$ then $H_{i,j}$ will be decreasing.

Lemma 4.1.40. *Any M -gridded generic set, (S, G) can be extended to an infinite M -gridded picture (M, P) .*

Proof. Join up the points in each cell, to form piecewise linear functions which will necessarily be increasing or decreasing according to the conditions imposed by M . \square

Proposition 4.1.41. *Let M be a gridding matrix and let P be a picture which does not admit an M -gridding. Then there exists some finite permutation π in $\text{Sub}(P)$ which does not admit an M -gridding.*

Proof. Let M be an $(r \times s)$ gridding matrix and P be a picture. We will prove the Proposition by contradiction. That is, we will assume that every finite subpicture of P is M -griddable and prove that in this case P is also M -griddable.

Let $S_1 \subseteq S_2 \subseteq \dots$ be a sequence of finite subpictures of P whose union $\bigcup_{i \geq 1} (S_i) = S$ is dense in P , that is, whose closure, $\overline{S} = P$. That this is possible follows from a well known result of topology, that the real plane has a countable basis. We may assume without loss that all these pictures lie within the unit square, otherwise we may apply an order preserving mapping, such as \arctan followed by a linear map, to both axes.

Let F_i be an M -gridding of S_i . Clearly, for $i \leq j$, we see that F_j is also a gridding of S_i , since $S_i \subseteq S_j$.

The sequence (F_1, F_2, \dots) is a sequence in $\mathbb{R}^{((r-1)+(s-1))}$. Since this infinite sequence is bounded by the unit $((r-1)+(s-1))$ -cube it must contain an accumulation point. Call this accumulation point G . We will assume that S is dense everywhere it cuts a grid line of G , anywhere this does not hold it is clearly trivial to grid S , P and every S_i .

Let S'_i be the set S_i with any point which lies on a grid line of G removed. Given a particular S'_i there is clearly some δ which is less than the minimum distance, under the infinity norm, of any point in S'_i from any grid line (The infinity norm is simply the minimum of the horizontal and vertical distances). Now since G is an accumulation point there exists some S_j , $j \geq i$, such that F_j is less than δ , both horizontally and vertically from G (this is, every grid line of F_j is less than δ from the corresponding grid line of G). In this case (S'_i, F_j) and (S'_i, G) must be isomorphic M -gridded sets and so G grids every S'_i .

Let $\bigcup_{i \geq 1} (S'_i) = S'$. Since S is dense around the grid lines it follows that the closure of S' , $\overline{S'}$, is equal to \overline{S} . We have shown that G grids every S_i , and hence \overline{S} , thus all that remains is to show that G grids P .

Suppose this is not the case. There are two possibilities. Either a point of P lies in a cell corresponding to a matrix entry of M which is 0, or two points of P lie in the same cell, but in the wrong order (For example, two points form a decrease in a cell with matrix entry 1). In either case we know that these points do not lie in S and hence must lie in dense regions of P . Suppose first that there is a point, z , in a cell which is supposed to be empty. Let ϵ be the minimum distance under the infinity norm of this point from any grid line of G . Since P is dense and is the closure of S , there is some point from S a distance less than ϵ from z . Thus S was not gridded by G , a contradiction. Suppose next that two points lie in the same cell but in the wrong order. Call these points x and y . Let ϵ be the minimum distance under the infinity norm of the points x and y from any grid line of G and

from each other. Then there are points in S a distance at most $\epsilon/2$ from both x and y , these points too must lie in the wrong order. Again S could not be gridded, a contradiction.

□

This is an incredibly useful result. Rather than search for (possibly very large) finite permutations which lie outside a grid class we can consider instead pictures which lie outside the class. Often this will prove far easier.

Definition 4.1.42. Let M be an $(r \times s)$ gridding matrix, we define the *graph of the matrix*, $G(M)$, to be the bipartite graph with vertices $\{x_1, x_2, \dots, x_r\} \cup \{y_1, y_2, \dots, y_s\}$ and edges $\{(x_i, y_j) \text{ whenever } M_{i,j} = \pm 1\}$. Label the edges 1 or -1 depending on the value of the matrix entry.

We think of the graph $G(M)$ as the bipartite graph whose vertices are the rows and columns of our matrix and whose edges are the matrix entries.

Definition 4.1.43. Let M be a gridding matrix and let $G(M)$ be the graph of M . Let C be a cycle in M . The *sign* of C is the product of the edge labels of the cycle C .

Example 4.1.44. The graph of the gridding matrix $\begin{pmatrix} 1 & 1 & 1 \\ -1 & 1 & 1 \end{pmatrix}$ is shown in Figure 4.5, it contains cycles of both negative and positive sign.

4.2 Algorithmic Problems for Grid Classes

We ask our standard algorithmic questions for grid classes. We will take the same approach for picture classes in Chapter 5

Algorithmic Problem 4.2.1.

Does there exist an algorithm which does the following?




Figure 4.5: The graph of a gridding matrix.

Input: A gridding matrix M .

Output: TRUE if $\text{Grid}(M)$ is partially well ordered,
FALSE otherwise.

Algorithmic Problem 4.2.2.

Does there exist an algorithm which does the following?

Input: A gridding matrix M .

Output: TRUE if $\text{Grid}(M)$ is atomic, FALSE otherwise.

Algorithmic Problem 4.2.3.

Does there exist an algorithm which does the following?

Input: A gridding matrix M .

Output: The basis of the class $\text{Grid}(M)$.

Algorithmic Problem 4.2.4.

Does there exist an algorithm which does the following?

Input: A gridding matrix M .

Output: TRUE if $\text{Grid}(M)$ is finitely based, FALSE otherwise.

Algorithmic Problem 4.2.5.

Does there exist an algorithm which does the following?

Input: A gridding matrix M .

Output: A Wilfian formula to enumerate the class $\text{Grid}(M)$.

There is one problem for grid classes which we can answer immediately; the membership problem.

Theorem 4.2.6.

There exists an algorithm which does the following:

Input: A gridding matrix M and a permutation π .

Output: TRUE if $\pi \in \text{Grid}(M)$, FALSE otherwise.

Proof. Assume that π is a permutation of length n . Then there are at most $n + 1$ choices for different vertical grid lines and $n + 1$ choices for different horizontal grid lines, hence we can test all possible choices of grid lines. \square

4.3 Early Results on Grid Classes

Definition 4.3.1. If M is a one dimensional gridding matrix then $\text{Grid}(M)$ is called a *W-class*.

Proposition 4.3.2. *W-classes are partially well ordered, atomic, finitely based and have rational generating functions.*

W-classes were introduced by Atkinson, Murphy and Ruškuc, [12].

Proof. It is easy to see that a W-class is the vertical or horizontal juxtaposition of two smaller W-classes. A W-class can be put into one-to-one correspondence with a regular language and can be enumerated automatically using the insertion encoding, see Section 5.8.2. For the basis results see Section 4.7.1. The partial well order result follows from the proof in Section 4.6, indeed the graph of a one dimensional gridding matrix is precisely a star. \square

Definition 4.3.3. A permutation in the class $\text{Grid} \left(\begin{array}{cc} -1 & 1 \\ 1 & -1 \end{array} \right)$ is called a *skew merged* permutation.

Lemma 4.3.4. *The class of skew merged permutations has basis $\{2143, 3412\}$ and has an algebraic generating function.*

The basis was proved by Stankova [57] and then by Kézdy, Snevily and Wang [41]. Atkinson was able to enumerate the class [8]. The generating function is:

$$\frac{1 - 3x}{(1 - 2x)\sqrt{1 - 4x}}$$

In fact this is the only grid class for which the graph of the gridding matrix is not a forest that has been enumerated. This motivates the conjecture that a grid class is rationally enumerated if and only if the graph of the gridding matrix is a forest, see Huczynska and Vatter [39, Conjecture 2.8].

The grid classes we use here were first defined by Murphy and Vatter [51], under the name of profile classes, however it should be noted that they label matrices from top right and interchange the use of minus one and one in their definition. Grid classes have also been used by Huczynska and

Vatter [39], they provide a foundation for a simple proof of the Fibonacci dichotomy for permutations, see Kaiser and Klazar [40]. Huczynska and Vatter label matrices as we do and have the same interpretation of matrix entries, however their definitions still vary from ours constructively. We begin by gridding generic sets, then factor to get permutations, they simply grid permutations, so that generic sets do not play any role.

4.4 The Partial Well Order Problem for Grid Classes

Definition 4.4.1. The *identity* matrix of size n , I_n , is the gridding matrix which has ones along the main diagonal (bottom left to top right) and zeros elsewhere. The *reverse identity* matrix of size n , R_n , is the gridding matrix which has minus ones along the minor diagonal (top left to bottom right) and zeros elsewhere.

Remark 4.4.2. It is easy to see that $\text{Grid}(I_n)$ is equal to $\text{Grid}(1)$, the increasing permutations, and that $\text{Grid}(R_n)$ is equal to $\text{Grid}(-1)$, the decreasing permutations.

Definition 4.4.3. Given an $(r \times s)$ gridding matrix M , the matrix $M^{(i)}$ is the $(ir \times is)$ gridding matrix obtained by replacing each 1 in M with I_i , each -1 in M with R_i and each 0 with an $(i \times i)$ zero matrix.

Example 4.4.4.

$$\begin{pmatrix} & 1 & 1 \\ -1 & & 0 \end{pmatrix}^{(2)} = \begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix}$$

Lemma 4.4.5. $\text{Grid}(M^{(i)}) \subset \text{Grid}(M)$.

Proof. Given an $M^{(i)}$ -gridded set we may construct an M -gridded set simply by removing the appropriate grid lines. \square

It is important to notice, however, that the reverse inclusion does not hold in general, see Example 4.4.7. We do, however, have the following result.

Lemma 4.4.6. $\text{Grid}(M^{(i)}) = \text{Grid}(M)$ for odd i .

Proof. Suppose that i is odd, we can immediately pick out a submatrix of $M^{(i)}$ which is isomorphic to M as follows. Choose the $(i+1)/2$ th row and column in each block, that is, the middle cell in each block. These cells clearly form a submatrix isomorphic to M . Again see Example 4.4.7. \square

Example 4.4.7.

$$\begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}^{(2)} = \begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \\ 0 & -1 & 1 & 0 \end{pmatrix}$$

$$\begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}^{(3)} = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 1 \\ 0 & -1 & 0 & 0 & 1 & 0 \\ 0 & 0 & -1 & 1 & 0 & 0 \end{pmatrix}$$

Murphy and Vatter [51] proved the following theorem.

Theorem 4.4.8. *Given a gridding matrix M the class $\text{Grid}(M)$ is partially well ordered if and only if $G(M)$ contains no cycles.*

This, in turn, yields the following decision theorem as a corollary.

Theorem 4.4.9.

There exists an algorithm which does the following:

Input: A gridding matrix M .

Output: TRUE if $\text{Grid}(M)$ is partially well ordered,
FALSE otherwise.

Proof. Construct the graph $G(M)$ and check for cycles. \square

We will give some, but not all, of the details of Murphy and Vatter's construction. For the sake of completeness we also repeat the definition of partial well order.

Definition 4.4.10. A set S with a partial order \leq is said to be *partially well ordered* if it contains no infinite antichains and no infinite descending chains.

Observation 4.4.11. Let D be a subclass of a permutation class C . If D is not partially well ordered then neither is C .

Lemma 4.4.12. Let M be a gridding matrix whose graph $G(M)$ contains a cycle C . Then the graph of $M^{(2)}$, $G(M^{(2)})$ contains a cycle $C^{(2)}$ of positive sign.

Lemma 4.4.13. Let M be a gridding matrix whose graph $G(M)$ consists of just a single cycle of positive sign. Then we can construct an infinite antichain in $\text{Grid}(M)$.

We will not give the proofs, which can be found in [51] and are relatively technical. We will however give an example.

Example 4.4.14. Let M be the matrix $\begin{pmatrix} -1 & 1 & -1 \\ -1 & 1 & 1 \end{pmatrix}$ with graph shown in Figure 4.6. $G(M)$ contains several cycles, in particular it contains a cycle of negative sign from the submatrix $\begin{pmatrix} -1 & -1 \\ -1 & 1 \end{pmatrix}$, call this matrix N . Then $N^{(2)}$ is the matrix:

$$\begin{pmatrix} -1 & 0 & -1 & 0 \\ 0 & -1 & 0 & -1 \\ -1 & 0 & 0 & 1 \\ 0 & -1 & 1 & 0 \end{pmatrix}$$

We will exhibit an antichain inside $\text{Grid}(N^{(2)})$, since $\text{Grid}(N^{(2)}) \subseteq \text{Grid}(N) \subseteq \text{Grid}(M)$, this will show that $\text{Grid}(M)$ is not partially well ordered. A




Figure 4.6: The graph of a gridding matrix with several cycles.

typical member of this antichain is shown in Figure 4.7. Further members are constructed by taking different numbers of cycles around the matrix. The anchor points, ringed in the figure, prevent members of the chain from embedding into one another.

Vatter and Murphy's proof that $\text{Grid}(M)$ is partially well ordered if M is a forest is also somewhat technical. We present a new, simpler proof in Section 4.6. It rests on Higman's theorem, Theorem 2.7.12, and on building posets across a gridded permutation, the lack of cycles in the graph ensures these posets are consistent. We can also build an atomic representative, that is a bijection f between two linearly ordered sets with $\text{Grid}(M) = \text{Sub}(f)$.

4.5 The Atomicity Problem for Grid Classes

As we establish in this section, the atomicity of a grid class also depends on (cycle) properties of the graph of its gridding matrix. For the sake of




Figure 4.7: A typical antichain element.

completeness we repeat the definition of atomicity, Definition 2.5.3.

Definition 4.5.1. A closed class is said to be *atomic* if it cannot be expressed as a union of two proper subclasses.

Theorem 4.5.2. For a permutation class C the following are equivalent:

1. C is atomic.
2. If α and β are permutations in C then there exists γ in C such that $\alpha \preceq \gamma$ and $\beta \preceq \gamma$. This is called the joint embedding property.
3. $C = \text{Sub}(f : A \rightarrow B)$ where A and B are linearly ordered sets and f is a bijection.

Definition 4.5.3. Let $G(M)$ be the graph of a gridding matrix M . The *type* of each component in $G(M)$ is defined as follows:

Type 1 No cycles of negative sign.

Type 2 Precisely one cycle, this cycle being of negative sign.

Type 3 Not of type one or two, that is, contains a cycle of negative sign connected to another cycle.

We extend the definition of type to the entire graph as follows.

Type 1 Every component is of type one.

Type 2 Every component is of type one or two and at least one component is of type two.

Type 3 Not of type one or two, that is at least one component is of type three.

Theorem 4.5.4. *Let M be a gridding matrix. The following three conditions are equivalent.*

(I) *The grid class $\text{Grid}(M)$ is atomic.*

(II) *$\text{Grid}(M) = \text{Grid}(M^{(2)})$.*

(III) *The associated graph $G(M)$ is of type 1 or 2, i.e. every component of $G(M)$ which contains a cycle of negative sign contains no other cycle.*

This theorem yields the following decision theorem as an immediate corollary.

Theorem 4.5.5.

There exists an algorithm which does the following:

Input: A gridding matrix M .

Output: TRUE if $\text{Grid}(M)$ is atomic, FALSE otherwise.

The proof of Theorem 4.5.4 is in three parts. We prove first that condition (II) implies condition (I). Then that condition (III) implies condition (II). Finally we prove that the negation of condition (III) implies the negation of condition (I). Since the proof is lengthy we will divide it into three propositions, Propositions 4.5.6, 4.5.8 and 4.5.15.

Proposition 4.5.6. *Let M be a gridding matrix. If $\text{Grid}(M) = \text{Grid}(M^{(2)})$ then the grid class $\text{Grid}(M)$ is atomic.*

We need the following lemma.

Lemma 4.5.7. *Let A be a gridding matrix which contains disjoint copies of two gridding matrices, B and C . Let $\beta \in \text{Grid}(B)$ and $\gamma \in \text{Grid}(C)$. Then there exists $\alpha \in \text{Grid}(A)$ such that $\beta \preceq \alpha$ and $\gamma \preceq \alpha$.*

Proof. It is simple to construct such an α simply by mapping copies of β and γ into appropriate disjoint parts of a grid with the same dimensions as the matrix A . \square

We can now prove the Proposition.

Proof. Notice next that $M^{(3)}$ contains disjoint copies of M and $M^{(2)}$. Thus if $\text{Grid}(M) = \text{Grid}(M^{(2)})$ we have the joint embedding property and $\text{Grid}(M)$ is atomic. \square

Proposition 4.5.8. *Let M be a gridding matrix. If the associated graph $G(M)$ is of type 1 or 2, i.e. every component of $G(M)$ which contains a cycle of negative sign contains no other cycle then $\text{Grid}(M) = \text{Grid}(M^{(2)})$.*

Definition 4.5.9. An $(r \times s)$ gridding matrix M is a *partial multiplication table* if we can choose sequences of 1s and -1s, (C_1, \dots, C_r) and (R_1, \dots, R_s) , such that the matrix entry $M_{ij} = C_i R_j$ or 0.

We prove the second part of the main theorem in two parts, first dealing with those matrices which are partial multiplication tables and then those which are not.

Proposition 4.5.10. *Let M be a gridding matrix. M is a partial multiplication table if and only if the graph $G(M)$ is of type 1.*

Proof. (\Rightarrow)

Let M be an $(r \times s)$ gridding matrix with associated sequences (C_1, \dots, C_r) and (R_1, \dots, R_s) of 1s and -1s (so that $M_{ij} = C_i R_j$ or 0). The sequence (C_1, \dots, C_r) indexes the rows, the sequence (R_1, \dots, R_s) indexes the columns of M so that $/C_1, \dots, C_r, R_1, \dots, R_s/$ in one to one correspondence with the vertices of the graph $G(M)$. Let $(x_{i_1}, y_{j_1}, x_{i_2}, y_{j_2}, \dots, x_{i_n}, y_{j_n}, x_{i_{n+1}} = x_{i_1})$ be a cycle in $G(M)$. Its parity is:

$$M_{i_1 j_1} M_{i_2 j_1} M_{i_2 j_2} \dots M_{i_n j_n} M_{i_1 j_1},$$

which is equal to:

$$C_{i_1} R_{j_1} C_{i_2} R_{j_1} C_{i_2} R_{j_2} \dots C_{i_n} R_{j_n} C_{i_1} R_{j_n},$$

which, in turn, is equal to:

$$C_{i_1}^2 R_{j_1}^2 C_{i_2}^2 R_{j_2}^2 \dots C_{i_n}^2 R_{j_n}^2 = 1.$$

(\Leftarrow)

Let M be an $(r \times s)$ gridding matrix whose graph $G(M)$ contains only cycles of positive sign. Note that this condition ensures that any two paths between a pair of vertices have the same sign, since taken together they form a cycle. We will prove the case where $G(M)$ is connected; when this is not true the argument should be repeated for every connected component. We give a method for constructing the sequences (C_1, \dots, C_r) and (R_1, \dots, R_s) demonstrating that M is indeed a partial multiplication table. Let C_i be the sign of any (and hence every) path from x_1 to x_i . Let R_j be the sign of any (and hence every) path from x_1 to y_j .

CLAIM: If $M_{ij} \neq 0$ then $M_{ij} = C_i R_j$.

Proof (of Claim). Take any path from x_1 to x_i . Its sign is C_i . Extend this path to y_j by adding a single edge (x_i, y_j) . The sign of this new path is $R_j = C_i M_{ij}$. Then $C_i R_j = C_i^2 M_{ij} = M_{ij}$. \square

We now proceed to prove that if M is a partial multiplication table then $\text{Grid}(M)$ is equal to $\text{Grid}(M^{(2)})$.

Proof. Let $M = (M_{ij})_{m \times n}$ be a partial multiplication table.

We use the following mappings to unearth a copy of M inside $M^{(2)}$.

- $\rho(i) : \{1, \dots, m\} \rightarrow \{1, \dots, 2m\}$
- $\lambda(j) : \{1, \dots, n\} \rightarrow \{1, \dots, 2n\}$
- $\rho(i) = \begin{cases} 2i & \text{if } R_i = 1 \\ 2i - 1 & \text{if } R_i = -1 \end{cases}$
- $\lambda(j) = \begin{cases} 2j & \text{if } C_j = -1 \\ 2j - 1 & \text{if } C_j = 1 \end{cases}$

By considering the four possible pairs of values (C_i, R_j) it is easy to see that $M_{ij} = M_{\rho(i)\lambda(j)}^{(2)}$. \square

For matrices which are not partial multiplication tables, and so contain components of type 2, we begin by considering those matrices whose graph is a forest. We can build a type two component from a component whose graph is a tree by changing a single zero to a one or a minus one, as appropriate.

Definition 4.5.11. Let M and N be gridding matrices, with N a submatrix of M . Let (S, G) be an N -gridded set. A *refinement* of (S, G) into M is a pair of sequences $R = ((V_r), (H_r))$ such that $(S, G \cup R)$ is an M -gridded set.

We define refinements of picture in exactly the same manner, however, for technical reasons, we allow grid lines of the refinement to cut continuous branches of the picture.

Proposition 4.5.12. *Let M be a matrix such that $G(M)$ is a forest. Let (P, G) be an M -gridded picture. Let l be a horizontal or vertical line that intersects (P, G) properly (i.e. not a grid line). Then (P, G) can be refined to an $M^{(2)}$ picture containing l as a grid line.*

Proof. Assume without loss that l is horizontal. We proceed by induction on the number, z , of non-zero entries in M . If $z = 1$ and l passes through the only non-empty cell, C , of (P, G) then l intersects this H -branch, indeed by the mean value theorem and the fact that M -branches are either increasing or decreasing there is a single point of intersection, a . The desired refinement of (P, G) is obtained by taking l , the vertical line through a and arbitrary lines in any remaining rows or columns. If l does not intersect the non-empty cell then choose a second line l' which does and repeat, choosing l in the appropriate row or column. Suppose now that $z > 1$ and that the proposition is true for every matrix with fewer than z non-zero entries. Let $C_{i,j}$ be a cell corresponding to a leaf in $G(M)$, since $z > 1$ we may assume that $C_{i,j}$ is not the only non-zero cell through which l passes, for if it is we may instead take l to be vertical. Let M' be obtained by replacing M_{ij} by zero. Let (P', G) be the M' -gridded picture obtained by erasing the H_{ij} -branch. By induction there is an $M'^{(2)}$ refinement of P' containing l . Since C_{ij} corresponds to a leaf, and is not the only non-empty line through which l passes, there is a line k such that $k \neq l$, k intersects C_{ij} and k intersects no other cell. The other line through H_{ij} intersects the H_{ij} -branch at a single point b (again by the mean value theorem). We can now replace k by the line parallel to it through b , yielding an $M^{(2)}$ refinement of (P, G) . \square

We now complete the second part of the proof of Proposition 4.5.8, by proving that if M is a matrix of type 2 then $\text{Grid}(M) = \text{Grid}(M^{(2)})$.

Proof. We prove that every M -gridded picture can be refined to an $M^{(2)}$ -gridded picture. Without loss of generality we may assume that $G(M)$ is connected, for otherwise we may refine each component separately.

Let (P, G) be any M -gridded picture. Let C_{pq} be a cell belonging to the

only cycle of $G(M)$ and let $H_{p,q}$ be its associated H -branch. We may write the cycle as $H_{pq} \Rightarrow H_{p1} \Rightarrow \dots \Rightarrow H_{nq} \Rightarrow H_{pq}$. We now define the continuous function β to be the composition of the H -branches H_{p1} to H_{nq} around the cycle onto the cell C_{pq} . Since we have a cycle of negative sign it follows that if H_{pq} is increasing β will be decreasing and vice versa. Now let a be the point of intersection of $H_{p,q}$ and β (by the mean value theorem on $H_{p,q} - \beta$). Let l be the horizontal line through a . Let M' be obtained from M by replacing $M_{p,q}$ by zero. Let (P', G) be obtained from (P, G) by erasing $H_{p,q}$. By the previous proposition we may obtain an $M'^{(2)}$ refinement of (P', G) , using l . This will also be an $M^{(2)}$ refinement of (P, G) since the vertical line through $H_{p,q}$ must pass through a by the composition of branches. \square

The following Proposition 4.5.13 demonstrates the power of this construction.

Proposition 4.5.13. *Grid(M) where $M = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ is atomic.*

Proof. We will prove that every M -gridded set (S, G) can be refined to an $M^{(2)}$ -gridded set. Without loss of generality assume that $S \in (0, 2) \times (0, 2)$ and $G = ((1), (1))$. For each cell $C_{ij}(S, G)$ define a piecewise linear function λ_{ij} to be a subset of the square $[i - 1, i] \times [j - 1, j]$ obtained by connecting successive points to each other, and connecting the first and last points to the appropriate corners of the square. We are simply constructing H -branches in each cell. We remark that with the exception of the points $C_{ij}(S, G)$ and two corner points, λ_{ij} is entirely contained in the white space of (S, G) . Now we use $\lambda_{22}, \lambda_{12}$ and λ_{11} to construct a function $\mu \in [1, 2] \times [0, 1]$ as follows:

$$\mu = \{(x_2, y_1) : (\exists x_1, y_2)((x_2, y_2) \in \lambda_{22}, (x_1, y_2) \in \lambda_{12}, (x_1, y_1) \in \lambda_{11})\}$$
Note that μ is in fact the composition $\lambda_{22} \cdot \lambda_{12}^{-1} \cdot \lambda_{11}$; therefore it is also piecewise linear. Next note that $(1, 0), (2, 1) \in \mu$ while $(1, 1), (2, 0) \in \lambda_{21}$, and so μ and λ_{21} intersect in a point $(p_2, q_1) \in \mu \cap \lambda_{21}$. By the definition of μ there exist (p_1, q_2) such that $(p_i, q_j) \in \lambda_{ij}$ for all $i, j \in \{1, 2\}$. By the definition of a generic set at most one of the points (p_i, q_j) is in S . If S contains no such points then $(S, G, ((p_1, p_2), (q_1, q_2)))$ is a refinement to an




Figure 4.8: A refinement of a gridded picture.

$M^{(2)}$ -gridding. Suppose now that $(p_r, q_s) \in S$. Let ε be smaller than the minimum non-zero horizontal or vertical distance from a point (p_i, q_j) to a point in S . If $(r, s) \in \{(1, 1), (1, 2), (2, 2)\}$ the triple $(S, G, ((p_1 + \varepsilon, p_2 + \varepsilon), (q_1 + \varepsilon, q_2 + \varepsilon)))$ is a refinement. Otherwise $(S, G, ((p_1 - \varepsilon, p_2 - \varepsilon), (q_1 + \varepsilon, q_2 + \varepsilon)))$ is a refinement. See Figure 4.8. \square

Before we consider Proposition 4.5.15 we give an example of a grid class which is not atomic.

Proposition 4.5.14. *The class $\text{Grid} \begin{pmatrix} 1 & 1 & 1 \\ -1 & 1 & 1 \end{pmatrix}$ is not atomic.*

Proof. We claim that the following holds:

$$\text{Grid} \begin{pmatrix} 1 & 1 & 1 \\ -1 & 1 & 1 \end{pmatrix} = \left(\text{Grid} \begin{pmatrix} 1 & 1 & 1 \\ -1 & 1 & 1 \end{pmatrix} \cap \text{Grid} \begin{pmatrix} 0 & 1 & 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 & 0 & 1 \\ 0 & -1 & 1 & 0 & 1 & 0 \end{pmatrix} \right) \cup$$

$$\left(\text{Grid} \begin{pmatrix} 1 & 1 & 1 \\ -1 & 1 & 1 \end{pmatrix} \cap \text{Grid} \begin{pmatrix} 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 1 \\ -1 & 0 & 0 & 1 & 0 & 1 \\ 0 & -1 & 1 & 0 & 1 & 0 \end{pmatrix} \right)$$

We will show that we can refine every M -gridded picture into either an

$$\begin{pmatrix} 0 & 1 & 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 & 0 & 1 \\ 0 & -1 & 1 & 0 & 1 & 0 \end{pmatrix} = T$$

or an

$$\begin{pmatrix} 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 1 \\ -1 & 0 & 0 & 1 & 0 & 1 \\ 0 & -1 & 1 & 0 & 1 & 0 \end{pmatrix} = B$$

gridded picture.

Clearly we can refine any $\begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$ gridded picture into a

$$\begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \\ 0 & -1 & 1 & 0 \end{pmatrix}$$

picture by the method given in the proof Proposition 4.5.8. We do this to the left hand part of our picture. We then extend the lower horizontal grid line from the refinement until it crosses the increase in cell $(3,1)$. At this crossing point we draw a vertical grid line. This grid line must cross the upper horizontal grid line of the refinement at some point x . If x lies below the increase in cell $(3,2)$ then we have gridded our picture into T , otherwise we have gridded into B .




Figure 4.9: Pictures which cannot jointly embed in $\text{Grid}(M)$.

Finally we exhibit two pictures, the first in $T \cup M$ but not $B \cup M$, the second in $B \cup M$ but not $T \cup M$, Lemma 4.1.41 assures us that there are finite permutations which witness this fact. Thus we have divided $\text{Grid}(M)$ into two proper subclasses. See Figure 4.9. \square

We now proceed to prove the final part of Theorem 4.5.4.

Proposition 4.5.15. *Let M be a gridding matrix. If the associated graph $G(M)$ is of not of type 1 or 2, i.e. there exists some component of $G(M)$ which contains a cycle of negative sign connected to another cycle, then the grid class $\text{Grid}(M)$ is not atomic.*

We first consider a type three matrix whose graph is connected, that is, consists of a single component. The proof is a generalisation of Proposition 4.5.14.

Definition 4.5.16. Given a gridding matrix, M , whose graph, $G(M)$, consists of a single type three component we identify a *target* cell, $C_{i,j}$ and a set of *free cells* $\{C_{a_1,b_1}, \dots, C_{a_n,b_n}\}$ as follows. First identify a cycle, C^- , of negative sign in $G(M)$, next construct a spanning tree which connects C^- to the remaining vertices of $G(M)$. This spanning tree will not contain every edge in $G(M)$, of the cells in M corresponding to the remaining edges choose one to be the target cell, the remainder become the free cells. Note that the edges corresponding to the target cell and free cells must all lie outside C^- but inside some cycle in $G(M)$.

Definition 4.5.17. Given an $(r \times s)$ gridding matrix, M , whose graph, $G(M)$, consists of a single type three component and a target cell, $C_{p,q}$, we

construct a *top heavy* M -gridded picture using the following method. First divide the positive $(r \times s)$ rectangle into unit squares, these divisions will be our gridding. In each cell, $C_{i,j}$, except the target cell, $C_{p,q}$, draw an increasing diagonal if $M_{i,j} = 1$, a decreasing diagonal if $M_{i,j} = -1$, or leave blank if $M_{i,j} = 0$. In the target cell draw an increasing concave curve if $M_{i,j} = 1$ or a decreasing concave curve if $M_{i,j} = -1$. A concave curve must pass strictly above the center point of the cell.

Definition 4.5.18. Dually we define a *bottom heavy* M -gridded picture, this time using convex curves which pass strictly below the centre point of the target cell.

These are the natural generalisations of the pictures in Figure 4.9.

Definition 4.5.19. Given an $(r \times s)$ gridding matrix, M , whose graph, $G(M)$, consists of a single type three component, a target cell, $C_{p,q}$ and a set of free cells, $\{C_{a_1,b_1}, \dots, C_{a_n,b_n}\}$, we define the *top heavy* matrix, T , to be the matrix constructed by taking a copy of $M^{(2)}$, replacing each (2×2) block in correspondence with a free cell by the matrix $\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ if $M_{a_i,b_i} = 1$ or by $\begin{pmatrix} -1 & -1 \\ -1 & -1 \end{pmatrix}$ if $M_{a_i,b_i} = -1$ by replacing the (2×2) block in correspondence with the target cell, $C_{p,q}$, with the matrix $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ if $M_{p,q} = 1$ or by $\begin{pmatrix} -1 & -1 \\ 0 & -1 \end{pmatrix}$ if $M_{p,q} = -1$.

Definition 4.5.20. Dually we define the *bottom heavy* matrix B , this time replacing the target cell block by $\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$ if $M_{p,q} = 1$ or by $\begin{pmatrix} -1 & 0 \\ -1 & -1 \end{pmatrix}$ if $M_{p,q} = -1$.

These are the generalisations of the matrices defined in the proof of Proposition 4.5.14, we have had to add target cells because of the possibility of there being several cycles outside C^- .

Lemma 4.5.21. *Given a gridding matrix, M , whose graph, $G(M)$, consists of a single type three component, every M -gridded set, (S, G) , can be refined to either a T -gridded set or a B -gridded set.*

The proof follows the same pattern as that of Proposition 4.5.14.

Proof. We begin by refining the cycle C^- . We follow the method given in the proof of Lemma 4.5.12 and the proof of Proposition 4.5.8. Each refinement we construct is then allowed to propagate around the rest of the gridded set, following the spanning tree constructed in Definition 4.5.16. Finally we must grid those points corresponding to the free cells and, as in the proof of Proposition 4.5.14, those points corresponding to the target cell. It is clear that the H -branch in each free cell can be gridded. The H -branch in the target cell must either lie above or below the intersection of the propagated grid lines. If it lies above this point we have gridded into T , otherwise we have gridded into B . \square

Corollary 4.5.22.

$$\text{Grid}(M) = (\text{Grid}(T) \cap \text{Grid}(M)) \cup (\text{Grid}(B) \cap \text{Grid}(M)).$$

Proposition 4.5.23. *If M is a gridding matrix, whose graph, $G(M)$, consists of a single type three component then $\text{Grid}(M)$ is not atomic.*

Proof. Finally we observe that our top heavy picture lies in $\text{Grid}(T) \cap \text{Grid}(M)$ but not $\text{Grid}(B) \cap \text{Grid}(M)$ and that our bottom heavy picture lies in $\text{Grid}(B) \cap \text{Grid}(M)$ but not $\text{Grid}(T) \cap \text{Grid}(M)$. Thus we have a disjoint union and by Lemma 4.1.41 there are finite permutations which bear witness to this fact. We conclude that $\text{Grid}(M)$ is not atomic. \square

Finally we consider the case where M consists of several components, at least one of which is of type three.

Lemma 4.5.24. *Let P^T be a top heavy picture constructed from a gridding matrix whose graph consists of a single type 3 component. Any gridding of P^T will correspond to a gridding matrix of type three.*

Proof. The connected nature of P^T ensures it can only be embedded into a single component. We walk around the cycles C^- and C^+ , demonstrating that these ensure we are in a type three component. Assume without loss that the target cell $C_{p,q}$ is increasing, that is that the part of the picture P^T corresponding to the target cell is an increasing convex curve. Choose any branch within the cycle of negative sign, C^- , we will assume, without loss, that this branch has positive sign. For any gridding of P^T it must be the case that in the top right hand corner of this branch there are infinitely many points which will lie in the same cell. Following the cycle C^- around the picture will lead us to the bottom left hand corner of this branch, by a path of negative sign. Following the cycle C^+ will bring us back to a subset of our starting point, by a path of positive sign. Finally, start at any point on the branch, follow the cycle C^+ , since the curve in the target cell is convex this will lead us to a point on the branch above and to the left of our starting point by a path of positive sign. Thus our starting point, which we took to be a cell in our new gridding is the intersection of a cycle of negative sign and a cycle of positive sign, ensuring that we are indeed in a type three component. See Figure 4.10. \square

The dual result for bottom heavy pictures follows by symmetry.

Lemma 4.5.25. *Two top heavy pictures, P_1^T and P_2^T , constructed from gridding matrices whose graphs consist of single type 3 components cannot be gridded so that they lie in the same type 3 component.*

Proof. The proof follows from the proof of Proposition 4.5.8. We see immediately that for the pictures to embed into the negative cycle of the type three component one would have to lie inside the other, contradicting their disjoint nature. \square




Figure 4.10: Gridding a top heavy picture

We are now ready to complete the proof of Proposition 4.5.15.

Proof. Given an $(r \times s)$ gridding matrix, M , whose graph, $G(M)$ consists of several components, at least one of which is of type three, we shall construct pictures and matrices which demonstrate that $Grid(M)$ can be written as a disjoint union.

Begin by ordering the components of $G(M)$ according to the maximum number of cells between any member of the component and the bottom of the matrix, we term this number the height of the component. Construct a new top heavy picture P^T by dividing an $(r \times s)$ rectangle into unit squares and by placing the top heavy picture corresponding to each component into the correct squares. Construct a new bottom heavy picture P^B by placing the bottom heavy picture corresponding to each component into the correct squares. Leave the squares corresponding to remaining components empty. Next construct analogues of the top and bottom heavy matrices. First construct the matrix T by constructing $M^{(2)}$ and replacing the appropriate blocks in the highest type three component blocks with blocks from the top heavy matrix of this component. For any other type three component perform the analogous operation, but treat the target cell as a free cell. Construct B analogously, but replace blocks in the highest type three component with blocks from the bottom heavy matrix of that component. That every M -gridded set can be gridded into T or B follows from the proofs of Lemmas 4.5.12 and 4.5.21. Thus we get our equation:

$$Grid(M) = (Grid(T) \cap Grid(M)) \cup (Grid(B) \cap Grid(M)).$$

It follows from Lemmas 4.5.24 and 4.5.25 that P^T lies in $Grid(T) \cap Grid(M)$ but not $Grid(B) \cap Grid(M)$ and that P^B lies in $Grid(B) \cap Grid(M)$ but not $Grid(T) \cap Grid(M)$. Thus we have an independent union and $Grid(M)$ is not atomic. \square

Our decision theorem, Theorem 4.5.5, follows as an immediate corollary.

The following proposition illustrates these methods.




Figure 4.11: The graph of the matrix M , unlabelled edges have positive sign.

Proposition 4.5.26. *Let M be the matrix:*

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & -1 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & -1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

The grid class $\text{Grid}(M)$ is not atomic.

We begin by drawing the associated graph $G(M)$, we notice that it contains two disjoint components, both of type 3. See Figure 4.11.

Next we choose two negative cycles, one from each component, and construct a spanning tree linking these cycles to the remainder of their components. See Figure 4.12. We choose a target cell in the highest type 3 component, in this case we choose cell $(6,6)$. The remaining cells which are not in the spanning tree are cells $(3,5)$ and $(1,3)$, these become free cells, notice that there is no target cell in the second component.

The matrix B is:


$$\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & -1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$


We can certainly grid any member of $\text{Grid}(M)$ into one of these matrices, the results of Lemmas 4.5.12 and 4.5.21 yield this result, we grid the negative cycles first then propagate these griddings around the tree. The free cells and target cell give us the freedom we need to complete the griddings.

Top and bottom heavy pictures which embed only into their respective matrices are shown in Figures 4.13 and 4.14, the reader is invited to attempt to grid them. We conclude that

$$\text{Grid}(M) = (\text{Grid}(T) \cap \text{Grid}(M)) \cup (\text{Grid}(B) \cap \text{Grid}(M))$$

This union is independent, thus $\text{Grid}(M)$ is not atomic.

Figure 4.13: A top heavy picture for the matrix M .

Figure 4.14: A bottom heavy picture for the matrix M .

4.6 A Return to the Partial Well Order Problem

We present a new proof that $\text{Grid}(M)$ is partially well ordered if the graph of M , $G(M)$ is a forest. We construct posets using M -gridded sets. Underlying this proof, though in the background, is the fact that we can build a simple atomic representative of the class.

Definition 4.6.1. Let M be an $(r \times s)$ gridding matrix containing k non-zero entries whose graph, $G(M)$, is a forest. Label the non-zero entries of M with the labels $\{1, \dots, k\}$. Let $T = ((C_1, \dots, C_r), (R_1, \dots, R_s))$ be the partial multiplication table associated with M . Let (S, G) be an M gridded set. For each vertical strip, $V_i(S, G)$, we build a linear ordering, \leq_{c_i} , of the points in that strip by reading from left to right if $C_i = 1$ and from right to left if $C_i = -1$. For each horizontal strip, $H_j(S, G)$, we build a linear ordering, \leq_{r_j} of the points in that strip by reading from bottom to top if $R_j = 1$ and from top to bottom if $R_j = -1$.

Definition 4.6.2. A set of linear orders is said to be *consistent* if their union forms a partially ordered set, that is there are no cycles in the union.

Lemma 4.6.3. *Any set of linear orderings constructed by the method given in Definition 4.6.1 is consistent.*

Proof. It follows from the definition of a partial multiplication table that we cannot have $a <_{c_i} b$ and $b <_{r_j} a$. Thus we need only worry about cycles. If $a \leq_{c_i} b \leq_{r_j} c$ then $M_{i,j}$ is a non-zero entry in our matrix. Thus any minimal cycle in $P(G)$ corresponds to a cycle in the graph $G(M)$ contradicting the fact that $G(M)$ is a forest. \square

Definition 4.6.4. The poset $P_{(S,G)}$ is the union of these linear orderings.

Definition 4.6.5. A word $W_{(S,G)}$ is generated by taking any linear extension L of $P_{(S,G)}$ and writing down the matrix entry labels of points in the order given by L . Thus there will be several words for each gridded set.

Lemma 4.6.6. *Every word $W_{(S,G)}$ encodes the gridded permutation isomorphic to (S, G) .*

Proof. Let $W_{(S,G)} = w_1 w_2 \dots w_n$ be such a word. Create points P_1, P_2, \dots, P_n corresponding to each letter. It suffice to give the left to right and bottom to top orderings of these points. Let $\text{col}(w_i)$ denote the column of M containing the entry with label w_i , $\text{row}(w_i)$ is the row of M containing the entry with label w_i . Let $\text{colsign}(w_i)$ and $\text{rowsign}(w_i)$ denote the corresponding entries in the multiplication table.

$$P_i \text{ left of } P_j \Leftrightarrow \begin{cases} \text{col}(w_i) < \text{col}(w_j) \text{ or} \\ \text{col}(w_i) = \text{col}(w_j) \text{ and } (i < j) \text{ and } \text{colsign}(w_i) = 1 \text{ or} \\ \text{col}(w_i) = \text{col}(w_j) \text{ and } (j < i) \text{ and } \text{colsign}(w_i) = -1. \end{cases}$$

$$P_i \text{ below } P_j \Leftrightarrow \begin{cases} \text{row}(w_i) < \text{row}(w_j) \text{ or} \\ \text{row}(w_i) = \text{row}(w_j) \text{ and } (i < j) \text{ and } \text{rowsign}(w_i) = 1 \text{ or} \\ \text{row}(w_i) = \text{row}(w_j) \text{ and } (j < i) \text{ and } \text{rowsign}(w_i) = -1. \end{cases}$$

□

Lemma 4.6.7. *Let $U = u_1 u_2 \dots u_m$ be a subword of $W = w_1 w_2 \dots w_n$. The the permutation determined by u is involved in the permutation determined by w .*

Proof. The result follows from the proof of Lemma 4.6.6. □

Theorem 4.6.8. *Let M be a gridding matrix whose graph is a forest. Then $\text{Grid}(M)$ is partially well ordered.*

Proof. Observe that every word on the alphabet $\{1, \dots, k\}$ can be formed in the manner described above and that by Higman's Theorem, 2.7.12, this set is partially well ordered under the subword ordering. Since $\text{Grid}(M)$ can be put into one to one correspondence with a subset of this set, and since the subword ordering implies involvement, we see that $\text{Grid}(M)$ is partially well ordered. □

Example 4.6.9. Let M be the matrix in Figure 4.15, shown as a partial multiplication table and with cell labels as subscripts. Let (G, S) be the M -

$$\left(\begin{array}{c|ccc} -1 & -1_2 & 1_3 & 0 \\ 1 & 0 & -1_4 & 0 \\ 1 & 1_1 & 0 & 1_5 \\ \hline & 1 & -1 & 1 \end{array} \right)$$

Figure 4.15: A matrix whose graph is a tree, written as a partial multiplication table with cell labels.

gridded set, isomorphic to the permutation 1984675(10)32, shown in Figure 4.16, with the direction of the row and column orders from the partial multiplication table shown as arrows. The linear row and column orders are:

- $1 <_{c_1} 9 <_{c_1} 8 <_{c_1} 4$
- $10 <_{c_2} 5 <_{c_2} 7 <_{c_2} 6$
- $2 <_{c_3} 3$
- $1 <_{r_1} 2 <_{r_1} 3 <_{r_1} 4$
- $5 <_{r_2} 6$
- $10 <_{r_3} 9 <_{r_3} 8 <_{r_3} 7$

The poset $P_{(S,G)}$ is shown in Figure 4.17. A linear extension is $1 < 10 < 2 < 9 < 5 < 3 < 8 < 4 < 7 < 6$. The word associated with this extension is 1352452134. If we associate the points $a, b, c, d, e, f, g, h, i, j$ with the letters of the word we can construct the left to right and bottom to top orders. The left to right order is $a < d < g < h \parallel j < i < e < b \parallel f < c$ where \parallel indicates a column break. The bottom to top order is $a < c < f < h \parallel e < j \parallel i < j < d < b$ where \parallel indicates a row break. We can recover the permutation by numbering from bottom to top then reading from left to right, so that $a = 1, c = 2, f = 3, h = 4, e = 5, j = 6, i = 7, g = 8, d = 9, b = 10$ and the permutation is 1984675(10)32 as expected.




Figure 4.16: A gridded set isomorphic to 1984675(10)32 with column and row directions shown.




Figure 4.17: The poset of a gridded set.

We conclude by building an atomic representative for these sets, choosing particularly nice linear extensions of our posets.


Definition 4.6.10. Let M be an $(r \times s)$ gridding matrix containing k non-zero entries whose graph $G(M)$ is forest. Let T be its associated multiplication table. First label the non-zero entries in M with the labels $\{0, \dots, k-1\}$. Next take an $(r \times s)$ grid and in each box $B_{i,j}$ which corresponds with a non-zero entry in M place an increasing or decreasing sequence which is in one to one correspondence with the natural numbers \mathbb{N} , so that the first natural number, 1, lies on the left hand side of the box if the multiplication table entry $C_i = 1$ and on the right hand side of the box if the multiplication table entry $C_i = -1$, and so that the first natural number, 1, lies at the bottom of the box if the multiplication table entry $R_j = 1$ and at the top of the box if $R_j = -1$. Notice that this guarantees the sequence will be increasing in $M_{ij} = 1$ and decreasing if $M_{ij} = -1$. Finally for each box $B_{i,j}$ erase all the points except those of the form $nk + l$ where l is the label of the matrix entry $M_{i,j}$. Call the picture created P .

P is an atomic representative of $\text{Grid}(M)$, that is $\text{Sub}(P) = \text{Grid}(M)$ and P is a bijection. It is clear that any permutation in $\text{Grid}(M)$ can be embedded into P , this is simply a particular case of the previous methods, the cell labels of each elements are the remainders modulo k . It is worth noting that these methods fail if $G(M)$ is not a forest, since Lemma 4.6.3 does not hold, indeed it is easy to construct gridded sets which are explicit counter-examples to the lemma if $G(M)$ contains a cycle.

Example 4.6.11. Let M be the gridding matrix shown in Figure 4.15. An atomic representative of $\text{Grid}(M)$ is shown in Figure 4.18.

4.7 The Basis Problem for Grid Classes

Basis results for grid classes are relatively scarce. For the sake of completeness we give a proof that W -classes are finitely based. We then consider larger grid classes.

Figure 4.18: An atomic representative for $\text{Grid}(M)$.

4.7.1 Juxtaposition of Classes

Definition 4.7.1. Let A and B be closed classes. A generic set S lies in the *horizontal juxtaposition* of A and B , denoted $A \mid B$, if it is possible to draw a vertical line through S which does not contain any points of S , so that the generic set to the left of the line belongs to A and the generic set to the right of the line belongs to B . Dually we define the *vertical juxtaposition*.

Lemma 4.7.2. Let A and B be finitely based closed classes, then both the horizontal and vertical juxtapositions of A and B are finitely based.

Proof. Let α be a basis element of the horizontal juxtaposition $A \mid B$. Draw a vertical line through α so that the generic set to the left of the line belongs to A and any line further right does not have this property. Clearly this is possible, otherwise α lies in the juxtaposition, we can identify a basis element

of A which determines the position of this line. Now draw a line so that the generic set to the right of it belongs to B but any line further left does not have this property. Again we can identify a basis element of B . Since α is a basis element it must be equal to the union of these two basis elements. Thus if A is finitely based with longest basis element of length k and B is finitely based with longest basis element of length l then $A \mid B$ is finitely based with longest basis element of length at most $k + l$. The proof for vertical juxtaposition is the dual. \square

Remark 4.7.3. The above proof also yields an algorithm for calculating the basis of $A \mid B$ given the bases of A and B .

Proposition 4.7.4. *Let M be a gridding matrix consisting of a single row (or dually a single column), that is a W -class. Then $\text{Grid}(M)$ is finitely based.*

Proof. It is clear that the basis of the grid class $\text{Grid}(1)$ is the single permutation 21 and the basis of the grid class $\text{Grid}(-1)$ is the single permutation 12. The result then follows from the lemma above. It is also clear that the basis can be easily calculated. \square

4.7.2 A Larger Grid Class

Larger grid classes cannot be constructed simply by juxtaposition. Instead they are subclasses of classes constructed by juxtaposition. To see this consider the permutation 13245768 which lies inside the vertical juxtaposition of $\text{Grid}(1, 1)$ with itself but certainly does not lie inside $\text{Grid}\left(\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}\right)$.

Theorem 4.7.5. *The grid class, $\text{Grid}(M)$, where $M = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ is finitely based.*

Proof. We first note that the vertical juxtaposition of the classes $\text{Grid}(1, 1)$ and $\text{Grid}(1, 1)$ is finitely based. We also note that $\text{Grid}(M)$ lies inside this

juxtaposition. Thus we need only prove that there is no infinite antichain of basis elements within this juxtaposition. Let α be a basis element of $\text{Grid}(M)$ that lies within the juxtaposition. Then there is at least one horizontal line that demonstrates this fact. Let H_t be the highest such horizontal line, its position is determined by a basis element of $\text{Grid}(1, 1)$, that is every higher horizontal line has a basis element of $\text{Grid}(1, 1)$ below it. Above and below this line there must be genuine members of $\text{Grid}(1, 1)$, since we are outside $\text{Grid}(M)$. Let H_b be the lowest such horizontal line, again we can identify a basis element of $\text{Grid}(1, 1)$ which determines this, again we must have genuine members of $\text{Grid}(1, 1)$ both above and below the line. Using H_t we can identify the following four points: point 1 which lies at the top of the first increase above the line, point 2 which lies at the bottom of the second increase above the line, point 3 which lies at the top of the first increase below the line and point 6 which lies at the bottom of the second increase below the line. It is clear that the two inversions determined by points 1 and 2 and points 3 and 6 cannot overlap since we are outside $\text{Grid}(M)$. Next consider the line H_b , we can identify the following four points: point 1 at the top of the first increase above the line, point 4 at the bottom of the second increase above the line, point 5 at the top of the first increase below the line and point 6 at the bottom of the second increase below the line. It is clear that points 1 and 6 are the same for both lines and that the inversions determined by points 1 and 4 and points 5 and 6 cannot overlap. Since the inversion determined by the points 3 and 6 is contained by the inversion determined by the points 5 and 6 it follows that the inversions determined by the points 1 and 2 and the points 5 and 6 cannot overlap. Thus these six points, together with the two basis elements already identified demonstrate that α is not in $\text{Grid}(M)$ and since α is a basis element it contains at most these points. Hence $\text{Grid}(M)$ is finitely based. See Figure 4.19. \square




Figure 4.19: The structure of a basis element

4.8 The Enumeration Problem for Grid Classes

If basis results for grid classes are scarce enumeration results are almost non-existent. Exceptions are the W -classes and the skew merged permutations, see Proposition 4.3.2 and Lemma 4.3.4. We give a method for enumerating the class $\text{Grid} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$.

Theorem 4.8.1. *The class $\text{Grid} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ is rationally enumerated.*

Before we prove this result we require some definitions.

Definition 4.8.2. A simple permutation is a permutation σ which does not map any nontrivial interval onto an interval.

Example 4.8.3. 2413 is a simple permutation. 3412 is not simple, it maps $\{1, 2\}$ onto $\{3, 4\}$ and $\{3, 4\}$ onto $\{1, 2\}$.

The study of simple permutations provides a powerful technique for the enumeration of pattern classes, see Albert and Atkinson, [1], and Albert, Atkinson and Klazar, [2], for more details. A further treatment of simple permutations can be found in Brignall, Huczynska and Vatter, [24, 23] and Brignall, Ruškuc and Vatter, [25].




Figure 4.20: The structure of an irreducible permutation.

We will not use the full power of these techniques. The relatively simple nature of the class we are studying makes it unnecessary. Instead we will consider those permutations which are irreducible, see Atkinson, [9], and Atkinson and Stitt, [15].

Definition 4.8.4. An *irreducible* permutation, π , contains no pair $i, i + 1$ with $\pi(i + 1) = \pi(i) + 1$.

Lemma 4.8.5. If $g(x)$ is the generating function for the irreducible permutations in $\text{Grid} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$, then the generating function for all permutations in $\text{Grid} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ is $g \left(\frac{x}{1-x} \right)$.

Proof. Any single point in an irreducible permutation can be expanded into an increase of arbitrary length to yield a unique permutation in $\text{Grid} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$. \square

We proceed to prove Theorem 4.8.1.

Proof. The structure of any large irreducible permutation in $\text{Grid} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$




Figure 4.21: The four forms of irreducible permutations.

is shown in Figure 4.20. We make the following structural remarks:

- An irreducible permutation in $\text{Grid} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ consists of three interleaved increases.
- Lines h and v are determined by the endpoint of the increase labelled 1.
- There can be at most one point in boxes a, b, and c.

These conditions are sufficient to show that for any large irreducible permutation this decomposition is unique.

Given an irreducible permutation in $\text{Grid} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ we can add new maximum elements to create longer irreducible permutations, furthermore every irreducible may be created in this way. Thus, for the purposes of enumeration, we can ignore the contents of box a. An irreducible permutation then has one of the four forms shown in Figure 4.21.




Figure 4.22: The two non-irreducible permutations which may be created.

We consider removing maximum elements from permutations of each form. These processes are reversible, we have already noted that we may construct longer irreducibles from shorter ones. Thus we will be able to construct recurrence relations and so enumerate the irreducible permutations.

For a permutation of form U the maximum element lies on the left, thus removing it will yield an irreducible permutation of form V , R or S , giving the recurrence $U_n = V_{n-1} + R_{n-1} + S_{n-1}$.

For a permutation of form V , removing the maximal element will either yield a permutation of form S or one of the two non-irreducible permutations shown in figure 4.22. Combining the increasing pair will then yield a permutation of form V or S . The recurrence is $V_n = S_{n-1} + V_{n-2} + S_{n-2}$.

Permutations of forms S and R can only yield permutations of forms V and U respectively when the maximum is removed. Thus $R_n = U_{n-1}$ and $S_n = V_{n-1}$.

We set $\Sigma_n = U_n + V_n + R_n + S_n$, the number of irreducibles of length n . Solving for Σ yields $\Sigma_n = 2\Sigma_{n-2} + 2\Sigma_{n-3} - 2\Sigma_{n-5} - \Sigma_{n-6}$. From here it is a simple exercise to find the short irreducible elements of the class and hence enumerate. The number of irreducible elements given by length is 1, 1, 2, 4, 9, 17, 30, 51, 85, 140, 229, ... Note that it is necessary to calculate these values by hand because for very short irreducible permutations our decomposition is not unique. Using these values we can solve the recurrence

relation to yield the generating function:

$$\frac{x(x^5 + 2x^4 + x^3 - x + 1)}{(x^3 - 2x + 1)}$$

From this we can construct the generating function for the grid class itself, it is:

$$\frac{-(2x^5 - 9x^4 + 15x^3 + 6x - 14x^2 - 1)x}{(2x^3 + 5x - 7x^2 - 1)(-1 + x)^3}$$

The first few terms of the sequence are:

$$1, 2, 5, 14, 42, 128, 384, 1123, 3204.$$

□

4.9 Conclusions for Grid Classes

Two of the four main decision problems for grid classes have been solved, the partial well order and atomicity problems. The remaining two, finite basis and enumeration, are open, and for grid classes with complicated graphs seem almost intractable.

The basis problem is particularly frustrating. It is very natural to conjecture that every grid class is finitely based, see, for example, Huczynska and Vatter [39, Conjecture 2.3]. Indeed, the natural bound on the length of a basis element for a grid class whose matrix contains k non-zero entries would seem to be $2k + 1$. Nonetheless, a proof is not only elusive, even an approach that hints at the beginnings of a proof has not been found.

Enumeration appears even more difficult. However an attack on those grid classes whose graph is a forest, using the underlying bijections hinted at in Section 4.6, might yield results. The method used to prove Theorem 4.8.1 does not appear to generalise easily.

Chapter 5

Picture Classes

Geomancy (noun):

- Divination by means of lines or figures formed by jotting down on paper a number of dots at random.

[Oxford English Dictionary]

5.1 Definitions for Picture Classes

We begin with a geometric definition of involvement, Definition 1.3.2, which we repeat for completeness.

Definition 5.1.1. A *picture* is a set of points in the real plane.

Definition 5.1.2. Given two pictures $P = \{(x_1, y_1), \dots, (x_n, y_n)\}$ and Q we say that P is *involved* in Q , written $P \preceq Q$, if there exists a pair of order preserving mappings f and g from \mathbb{R} into \mathbb{R} such that the set $P_{f,g} = \{(f(x_1), g(y_1)), \dots, (f(x_n), g(y_n))\}$ is contained in Q .

This is a pre-order on the set of all finite pictures, it is reflexive and transitive.

Definition 5.1.3. We say that two pictures P and Q are order isomorphic if $P \preceq Q$ and $Q \preceq P$.

Clearly order isomorphism of pictures is an equivalence relation. If we restrict ourselves to generic sets, and factor the pre-order on generic sets by this equivalence we obtain (a poset isomorphic to) the set of all permutations under involvement. Permutations are class representatives for equivalence classes of generic sets under these stretching, squashing and sliding operations (axis parallel order preserving mappings).

Pictures can also play a second role; we can obtain permutations from them. Given any picture we can obtain a permutation by choosing any set of n points from the picture, no two lying in the same vertical or horizontal line, this will be a generic set, isomorphic to some permutation.

Definition 5.1.4. A permutation is said to be *drawn from* (or, *on*) a picture, if it can be constructed by the above method.

Lemma 5.1.5. *The set of all permutations which can be drawn from a picture P forms a closed class. We denote this class by $D(P)$.*

Proof. If a permutation π can be drawn from a picture P then so can any subpermutation of π ; simply take those points corresponding to elements of the subpermutation. \square

Example 5.1.6. The set of permutations which can be drawn from a single line of positive gradient is simply the set of increasing permutations, that is, the closed class of permutations avoiding the single basis element 21.

Example 5.1.7. The set of permutations which can be drawn from the picture in Figure 5.1 is the class of permutations avoiding the basis elements 213 and 312. It may also be thought of as the set of all juxtapositions of an increase and a decrease or as the grid class $\text{Grid}(1, -1)$.




Figure 5.1: A picture defining the closed class of permutations avoiding 213 and 312.

5.2 Algorithmic Problems for Picture Classes

Since a picture is simply a set of points in the plane, it is only reasonable to ask questions about pictures for those pictures which admit some finite description. An infinite picture which can only be described by listing its points will make for a very long question. We ask the following algorithmic problems:

Algorithmic Problem 5.2.1.

Does there exist an algorithm which does the following?

Input: A finite description of a picture P .

Output: TRUE if the class $D(P)$ is atomic, FALSE otherwise.

Algorithmic Problem 5.2.2.

Does there exist an algorithm which does the following?

Input: A finite description of a picture P .

Output: TRUE if the class $D(P)$ is finitely based, FALSE otherwise.

Algorithmic Problem 5.2.3.

Does there exist an algorithm which does the following?

Input: A finite description of a picture P and a list of basis elements B .

Output: TRUE if $D(P) = \text{Av}(B)$, FALSE otherwise.

Algorithmic Problem 5.2.4.

Does there exist an algorithm which does the following?

Input: A finite description of a picture P and a Wilfian formula.

Output: TRUE if the class $D(P)$ is enumerated by this formula, FALSE otherwise.

These questions are, unsurprisingly, incredibly difficult. Before considering them, however, we ask the following:

Algorithmic Problem 5.2.5.

Does there exist an algorithm which does the following?

Input: A finite description of two pictures P and Q .

Output: TRUE if $D(P) = D(Q)$, FALSE otherwise.

5.3 Symmetries of the Plane which Respect Picture Classes

Since permutations are class representatives of generic sets we immediately see that we may also stretch, squash and slide any pictures which define a class. More formally, given a picture, P , any order preserving mapping




Figure 5.2: The equality of the circle and C -classes, (not to scale).

applied to either the x or y axis will not affect the set of permutations which can be drawn from P . This observation has some surprising consequences, for example, we see immediately that every permutation can be drawn from an infinite spiral and that every permutation can be drawn from any “solid” region, no matter how small. Indeed pictures that look very different can turn out to define the same class. The following example demonstrates that the circle class, Section 5.5, could equally well be defined as the diamond class, the C -class or even the egg class.

Example 5.3.1. Consider the following two classes defined by pictures. The circle class is the set of permutations which can be drawn from the unit circle. The C -class is the set of permutations which can be drawn from the unit circle with all points within 45 degrees of the positive x -axis removed. Clearly the C -class is contained within the circle class, since its defining picture is a subset of the unit circle. However, by applying order preserving mappings to parts of the x and y -axes, we can demonstrate that there is a defining picture for the circle class contained in the C -class picture. First stretch the point $x = 0$ until it has size $6 + 4\sqrt{2}$, effectively splitting the circle in two vertically. Next stretch the x -axis by a factor of $2\sqrt{2}(\sqrt{2} + 1)$ about the point $y = 0$. Finally shrink the whole picture by a factor of $2\sqrt{2}(\sqrt{2} + 1)$ in both axes, about the origin. What remains is a subset of the C -class picture. See Figure 5.2.

5.4 The Atomicity Problem for Picture Classes

Recall that a class is atomic if it has the joint embedding property, or equivalently if it can be expressed as $Sub(B)$ where B is an infinite bijection between two linearly ordered sets. For picture classes the latter is often easier to prove, we find some underlying bijection which is similar to our picture. Indeed we might expect every picture class to be atomic but this is not the case.

Lemma 5.4.1. *Not every picture class is atomic.*

Proof. Let P be the finite picture consisting of the three points $(0, 0)$, $(1, 1)$, $(2, 0)$, then $D(P) = \{1, 12, 21\} = \{1, 12\} \cup \{1, 21\}$. \square

One possible approach would be to demand that every picture be a bijection, although that seems unnecessarily limiting; we would then be studying atomic classes that in some sense “look nice”.

In fact, it is often easy to see that a picture contains an underlying bijection. For example it is clear that the circle class does. We simply divide into quadrants, and restrict each quadrant to a different continuous subset of the real plane. This is easiest to envisage if we first stretch and squash our circle to form a diamond centred on the origin with height and width one. We then assign each quadrant a different prime number p_i and restrict the range and domain of that quadrant to those numbers in \mathbb{Q} whose denominator is a power of p_i .


Notwithstanding the depth of this problem we will not give it much more consideration. Unfortunately it is simply too far from the combinatorial stomping ground of this author.

5.5 The Circle Class


The circle class (or C, diamond or egg class) provides a very elegant introduction to picture classes, it is far from trivial, yet we can answer all of our

main questions. We will demonstrate that it is finitely based and partially well ordered and we will determine its enumeration. This section consists of joint work with Vince Vatter.

Choose n points from a circle — no two with the same x -coordinate or y -coordinate (i.e. choose a generic set) — label them 1– n by height, reading bottom-to-top, and record these labels reading left-to-right. This operation produces a permutation; for example, the set of points shown on the left below gives the permutation 45312, a plot of which is on the right.



We say that such permutations π can be *drawn on (or, from) a circle*, and we refer to the points of the circle as a *circle drawing* of π . Not all permutations can be drawn on a circle. For example and sake of contradiction suppose that 1324 can be drawn on a circle, take such a drawing, and connect the points with lines. This gives a quadrilateral such as the one below.



Because we have assumed that these points lie on a circle, this quadrilateral is a *cyclic quadrilateral*. Moreover, Proposition 22 of Book III of Euclid's *Elements*, see Heath [35], states that the sum of opposite angles in a cyclic quadrilateral equals 180° , but $\angle 213$ and $\angle 342$ must both be less than 90° ,

since, for example, both 2 and 3 must lie above and to the right of 1. We conclude that 1324 (and, of course, its reverse 4231) cannot be drawn on a circle.

5.5.1 The Basis of the Circle Class

For the rest of the section we will label the leftmost point of a permutation π as L , the rightmost as R , the topmost as T , and the bottommost as B . Note that these labels can coincide. We begin by constructing some canonical drawings.

Proposition 5.5.1. *If the permutation π of length at least 2 can be drawn on a circle then there is a circle drawing of π in which either L or R lies on the x -axis and either T or B lies on the y -axis.*

Proof. Consider any circle drawing of π . We may assume by symmetry that T lies closer to the y -axis than B . Note that there cannot be any points closer to the y -axis than T ; if such a point lay on the lower half circle then it would be lower than B , while if it lay on the upper half circle it would be greater than T . Thus we can shift T to the y -axis.

We now consider this new circle drawing of π . We assume, again by symmetry, that L lies closer to the x -axis than R . Note that $L \neq T$, since no point lies as far away from the x -axis than T and π has length at least 2. By essentially the same argument as in the T/B case it follows that there are no points closer to the x -axis than L , and thus we may also draw L on the x -axis. \square

Once L or R is placed on the x -axis and T or B is placed on the y -axis, the quadrants of the other points are fixed. Thus Proposition 5.5.1 shows that we may restrict our attention to four types of circle drawings. We refer to these (when they exist) as LT , RT , LB , and RB drawings, based on the points that lie on the axes.

If π can be drawn with L on the x -axis and a and b are points in such a

drawing, both below L , then a is closer to the x -axis if and only if a lies above b (in either the drawing or the permutation, both are equivalent). Similarly, if a and b both lie above L , a is closer to the x -axis if and only if a lies below b . Motivated by these observations, we define the linear orders

$$a <_L^- b \iff a \text{ lies above } b$$

on L and the points below L , and

$$a <_L^+ b \iff a \text{ lies below } b$$

on L and the points above L . The same holds with L replaced by R , and we define $<_R^-$ and $<_R^+$ analogously.

Similarly, if π can be drawn with T on the y -axis and a and b lie to the left of T in such a drawing, a will lie closer to the x -axis if and only if a lies to the left of b . We thus define

$$a <_T^- b \iff a \text{ lies to the left of } b$$

on T and the points to the left of T , and

$$a <_T^+ b \iff a \text{ lies to the right of } b$$

on T and the points to the right of T . The orders $<_B^-$ and $<_B^+$ are defined by replacing T by B . Note that we can construct these orders for any permutation, not just circle permutations. It is the interactions of these orders, once constructed, which allow us to determine whether or not a permutation can be drawn on a circle.

Thus we have four linear orders on (subsets of) the points of π for each of the four types of circle drawings we are considering. For example, the *LT orders* are $\{<_L^-, <_L^+, <_T^-, <_T^+\}$. To continue this example, from an *LT* drawing of π in which no two points lie the same distance from the x -axis (which can always be achieved by shifting the points by a very small distance ε) we can

build a linear order $<_x$ on the points of π by defining

$$a <_x b \iff a \text{ lies closer to the } x\text{-axis than } b.$$

The importance of this order is that each of the LT orders agrees with $<_x$; i.e., if $a < b$ for some order in $\{<_L^-, <_L^+, <_T^-, <_T^+\}$, then $a <_x b$. In other words, if π has an LT drawing then the LT orders are *consistent*: there does not exist a cycle

$$a_1 <_1 a_2 <_2 \cdots <_k a_1$$

where $<_i \in \{<_L^-, <_L^+, <_T^-, <_T^+\}$ for all $i = 1, \dots, k$. For this reason we are able to define the LT poset as the transitive closure of the union of these orders, i.e., as the poset on the points of π in which

$$a <_{LT} b \iff a = a_1 <_1 a_2 <_2 \cdots <_k a_k = b$$

for some points a_2, \dots, a_{k-1} and $<_i \in \{<_L^-, <_L^+, <_T^-, <_T^+\}$. In this language, our previous comment was that $<_x$ is a linear extension of $<_{LT}$. Conversely, any linear extension of this poset naturally produces a circle drawing:

Proposition 5.5.2. *If the LT (resp. RT , LB , RB) orders for π are consistent, then π has an LT (resp. RT , LB , RB) drawing.*

Proof. By symmetry we may assume that the LT orders for π are consistent. Consider a linear extension $a_1 < \cdots < a_n$ of the LT poset of π , and note that a_1 must be L and a_n must be T . We begin by placing a_1 on the x -axis. Now set $\theta = \pi/2(n-1)$ and for $j = 2, \dots, n-1$ place a_j at $(\cos(j\theta), \sin(j\theta))$ if a_j lies above L and to the right of T , $(\cos(\pi/2 - j\theta), \sin(\pi/2 - j\theta))$ if a_j lies above L and to the left of T , $(\cos(\pi/2 + j\theta), \sin(\pi/2 + j\theta))$ if a_j lies below L to the left of T , and $(\cos(-j\theta), \sin(-j\theta))$ if a_j lies below L to the right of T . We finish by placing $a_n = T$ on the y -axis. Notice that our choice of θ guarantees that every point may be placed, that our use of \cos and \sin guarantees that every point lies on the unit circle and that the properties of the \cos and \sin curves guarantee that the four linear orders are respected. See Figure 5.3. \square




Figure 5.3: The construction of an LT drawing of the permutation 462891735. In this case we took the linear extension $4 < 5 < 6 < 3 < 7 < 2 < 8 < 1 < 9$.

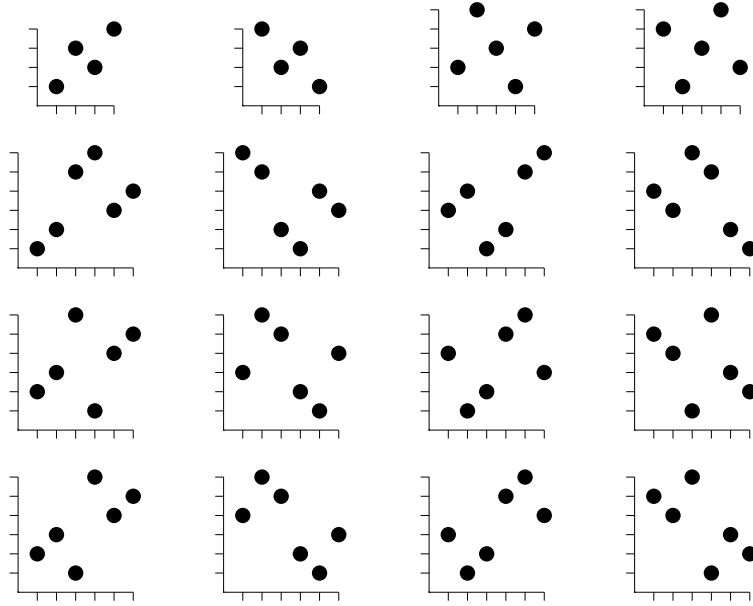


Figure 5.4: The basis for the class of permutations that can be drawn on a circle.

We have the following algorithmic theorem as a corollary.

Theorem 5.5.3.

There exists an algorithm which does the following:

Input: A Permutation π .

Output: A circle drawing of π if π can be drawn on a circle, FALSE otherwise.

Our previous proposition allows us to finish the characterisation of the circle class.

Theorem 5.5.4. *The basis elements for the circle class are*

1324, 4231,
 25314, 41352,
 125634, 231645, 236145, 312564,
 341256, 365214, 412563, 436521,
 465213, 541632, 546132, 652143.

These basis elements are plotted in Figure 5.4.

Proof. Suppose that π cannot be drawn on a circle. Thus by Proposition 5.5.2, each of the four sets of linear orders is inconsistent. Suppose, without loss of generality, that the LT orders for π are inconsistent, or in other words, that they contain a cycle, and consider a cycle of minimal length, say


$$a_1 <_1 a_2 <_2 \cdots <_k a_1,$$

where each $<_i$ is one of $\{<_L^-, <_L^+, <_T^-, <_T^+\}$. Clearly each a_i must be distinct, as otherwise the cycle could be shortened. Now suppose that some symbol, say $<_L^-$ occurs twice in a minimal cycle. Thus we have a cycle of the form $\dots < u <_L^- v < \dots < w <_L^- t < \dots$. Since $<_L^-$ is a linear order we may compare u and w . If $u <_L^- w$ then $u <_L^- t$ and $\dots < u <_L^- t < \dots$ forms a shorter cycle. Otherwise $w <_L^- u$ in which case $w <_L^- v$ and $v < \dots < w$ forms a shorter cycle. Thus each of $<_L^-, <_L^+, <_T^-$, and $<_T^+$ can occur at most once in a minimal cycle, so the length of the cycle is at most 4.

Moreover, each point (except L and T , but these points clearly cannot participate in a cycle) can participate in precisely one of $\{<_L^-, <_L^+\}$ and precisely one of $\{<_T^-, <_T^+\}$. Thus cycles of length 3 are impossible. Now consider a minimal cycle of length 4. There are two cases. If this cycle is of the form

$$a_1 <_L^+ a_2 <_T^+ a_3 <_L^- a_4 <_T^- a_1,$$

then it is easy to check that these points are order isomorphic to 1324:



As observed earlier 1324 cannot be drawn on a circle, it is easily seen to be a basis element for this class, and, indeed, lies in Figure 5.4. In the other case, $a_1 >_L^+ a_2 >_T^+ a_3 >_L^- a_4 >_T^- a_1$ and it is easy to check that the points are order isomorphic to 4231, which is also a basis element and is listed in Figure 5.4. The same situations occur if the RT , LB , or RB orders contain cycles of length 4, so we may assume that this does not occur.

Thus we may assume that each of the four sets of orders contains a cycle of length 2. Choose one such cycle for each set of orders. These eight points, which may not all be distinct, together with L , R , T , and B have inconsistent LT , RT , LB , and RB orders, and thus cannot be drawn on a circle. Therefore the length of a basis element for this class is at most 12. A routine computer check finishes the proof, having observed that the permutations given in the Theorem are indeed basis elements we simply construct all permutations of length at most 12 which avoid them and check that each has at least one set of consistent orders. \square

Recall Helly's theorem, which states that, given a finite collection of convex sets in d -dimensional Euclidean space, if every $d+1$ of them have non-empty intersection then the entire collection has non-empty intersection; $d+1$ is called the *Helly number* of d -dimensional Euclidean space. We can state Theorem 5.5.4 in the following way, which demonstrates its similarity to this classic theorem of combinatorial geometry.

Corollary 5.5.5. *Given a permutation σ , if every subpermutation of length 6 can be drawn on a circle then σ can be drawn on a circle.*




Figure 5.5: Three circle drawings of the permutation 46187235, giving the words 21244311, 13244311, and 13244113, from left to right. The third of these words, 13244113, happens to be the minimal word for this permutation.

5.5.2 Enumeration of the Circle Class

We count the permutations that can be drawn on a circle by putting them in bijection with the words of a regular language. From any circle drawing of π in which no two points lie the same distance from the x -axis, we order the points according to $<_x$ (the distance from the x -axis order) and record which quadrant each point lies in. In doing so we take quadrant 1 as $\{(x, y) : x, y \geq 0\}$, quadrant 2 as $\{(x, y) : x < 0, y \geq 0\}$, quadrant 3 as $\{(x, y) : x \leq 0, y < 0\}$, and quadrant 4 as $\{(x, y) : x > 0, y < 0\}$. This gives a word over the alphabet $A = \{1, 2, 3, 4\}$ that corresponds to π ; see Figure 5.5 for an example.

As a permutation may correspond to many different words, we need to choose a canonical word for each permutation and then characterise the language of canonical words. Our choice for the canonical word to associate to a permutation will be the lexicographically minimal word (henceforth shortened to *minimal word*). We begin with some simple observations

Proposition 5.5.6. *No minimal word is of the form $3u$, $4u$, $u2$, or $u4$, where u is any word over the alphabet $\{1, 2, 3, 4\}$.*

Proof. It follows from the proof of Proposition 5.5.1 and our labelling of quadrants that the words $2u$ and $3u$ correspond to the same permutation,

as do the words $1u$ and $4u$, $u1$ and $u2$, and $u3$ and $u4$. \square

Because of Proposition 5.5.6, we need only consider words that come from LT , RT , LB , or RB drawings. Indeed, each such word comes from a linear extension of one of these drawing posets. Our next observation yields a canonical linear extension for each poset.

Proposition 5.5.7. *No minimal word is of the form $u31u$ or $u42u$, where u is any word over the alphabet $\{1, 2, 3, 4\}$.*

Proof. The words $u13v$ and $u31v$ encode the same permutation, as do $u24v$ and $u42v$. \square

Notice that having chosen which points will lie on the axes, that is having chosen a poset, we have fixed the quadrant of every point. Since each linear order compares points in adjacent quadrants our only freedom comes from diagonally opposite quadrants. Proposition 5.5.7 eliminates this freedom and so eliminates all but one word that arises from each of the four posets. Recall that for any alphabet A the set A^* is the set of all words (including the empty word) over that alphabet. We have seen that every permutation which can be drawn on a circle corresponds to at most four words in the language:

$$A^* \setminus (3A^* \cup 4A^* \cup A^*2 \cup A^*4 \cup A^*42A^* \cup A^*31).$$

Next we introduce terms for the permutations that correspond to more than one word in the language above, i.e., the permutations that have more than one drawing poset.

- We say that π has an *ambiguous beginning* if both L and R can be drawn on the x -axis.
- We say that π has an *ambiguous ending* if both T and B can be drawn on the y -axis.




Figure 5.6: An epicene permutation drawn on a circle in two radically different ways.




Figure 5.7: A second epicene permutation.

We must pay particular attention to a particular kind of ambiguity; we say that a permutation is *epicene* if it consists of three distinct monotonic subpermutations in one of two configurations. Either an increase, above and to the left of a decrease, which is in turn above and to the left of an increase, or a decrease, below and to the left of an increase, which is in turn below and to the left of a decrease. Epicene permutations admit two radically different drawings — see Figure 5.6 — and so need extra consideration.

We now introduce an extension of Proposition 5.5.1. We give only the L/R version, the T/B version follows by symmetry.

Proposition 5.5.8. *Let π be a permutation which can be drawn on a circle with either L or R on the x -axis. There are points vertically between L and R if and only if π is an epicene permutation.*

Proof. Suppose that π satisfies these hypotheses and assume without loss

that L lies above R . In any L drawing of π (i.e., any drawing with L on the x -axis), the points lying vertically between L and R must lie in quadrant 3 (because they cannot lie to the right of R) and thus form a decreasing subsequence. These points must similarly lie in quadrant 1 in any R drawing of π . Let us call this decreasing subsequence D . By considering the position of D in an L and an R drawing of π , it is apparent that no points may lie amongst D either horizontally or vertically.

If there are any points above L then they must lie in quadrant 2 in both drawings and thus form an increase — say I_1 — above and to the left of D ; again, no points may lie amongst this increasing subsequence either horizontally or vertically. Finally, any and all points below R must lie in quadrant 4 in both drawings and so must form an increasing subsequence I_2 below and to the right of D , and no points may lie amongst I_2 . Thus π consists of an increasing subsequence (I_1 together with L) above and to the left of a decreasing subsequence (D), which lies above and to the left of an increasing subsequence (I_2 together with R) and is therefore epicene. \square

To eliminate the remaining non-minimal words we make the following remarks:

- No minimal word is of the form 23^*1u or 2^+4u , where u is any word over the alphabet $\{1, 2, 3, 4\}$, because these are non-minimal words for non-epicene permutations with ambiguous beginnings. The minimal word for such permutations are obtained from R drawings.
- No minimal word is of the form $u31^+$ or $u42^*1$, where u is any word over the alphabet $\{1, 2, 3, 4\}$, because these are non-minimal words for non-epicene permutations with ambiguous endings. The minimal words for such permutations are obtained from B drawings. (Although note that these words are already eliminated by Proposition 5.5.7.)
- No minimal word lies in the sets 2^+1^* , 2^+3^* , $23^*2^*1^+$, or $2^+3^*4^*3$, because these are sets of non-minimal words for epicene permutations.

The minimal words for such permutations are obtained from R drawings.

From the regular language we have obtained, it is routine to enumerate the permutations that can be drawn on a circle. We begin by constructing a finite state automaton from the language, this can be done by hand, however the automaton we construct is relatively large (14 states) and so it is easier to use a computer. Having obtained the automaton we use the transition matrix to compute the generating function. Solving the equations by hand would again be relatively difficult, as the transition matrix is a 14×14 matrix, so again it is easier to use a computer. The details of these methods can be found in Hopcroft and Ullman [38].

Theorem 5.5.9. *The generating function (by length) for permutations that can be drawn on a circle is:*

$$\frac{1 - 6x + 12x^2 - 10x^3 + 5x^4 + 2x^5 - 2x^6}{(1 - 4x + 2x^2)(1 - x)^3}$$

This sequence begins:

1, 2, 6, 22, 84, 308, 1090, 3782, 13000, 44504, 152102, 519506, 1773948, 6056932.

Proposition 5.5.10. *The circle class is partially well ordered.*

Proof. It is clear that the set of circle words is partially well ordered under the subword ordering, by Nash-Williams' [52] proof of Higman's theorem for words, Theorem 2.7.8. Now suppose that u and w are circle words and that w is a subword of u . It is clear that the permutation corresponding to w is involved in the permutation corresponding to u , we simply plot u on a circle then remove those points which correspond to letters which are not in w , what remains is a circle drawing of the permutation corresponding to w . Thus, since that class is partially well ordered under a relation which implies involvement it is also partially well ordered under involvement. \square

A class which is partially well ordered and finitely based is said to be *strongly*

finitely based, see Atkinson, Murphy, Ruškuc [12]. In this case every subclass is also finitely based. Our enumeration method then allows any subclass to be enumerated, we eliminate from the language all words corresponding to extra basis elements and perform the same automatic computations. Thus we have the following algorithmic theorem.

Theorem 5.5.11.

There exists an algorithm which does the following:


Input: Any subclass of the circle class, given as
extra basis elements.

Output: A rational generating function enumerating
that subclass.

Finally note that the circle class is a partially well ordered subclass of the grid class $\text{Grid} \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$, sometimes called the diamond grid class. Unfortunately for this grid class, unlike for circles, we cannot choose a fixed finite set of griddings for any permutation. The number of possible griddings grows as the length of the permutations grow. For this reason the methods used here cannot be applied to this grid class.

5.6 The X Class

A second easy, and indeed very similar example of a picture class, comes from drawing permutations on an X. As an added advantage we have (with circles, Xs and grid classes) all the machinery necessary for a permutational game of noughts and crosses, or, for American readers, Tic-Tac-Toe.



The set of permutations that can be drawn in this manner happens to be the intersection of two well-studied permutation classes: the separable and skew-merged permutations. Skew merged permutations were discussed in Section 4.3.

Recall the definitions of direct and skew sums, Definition 2.7.14 and its dual, and the definition of separable permutations, Definition 2.7.16.

Proposition 5.6.1. *A permutation can be drawn on an X if and only if it is both separable and skew-merged, i.e., if and only if it avoids 2143, 2413, 3142, and 3412.*

Proof. Clearly every permutation that can be drawn on an X is both separable and skew-merged, so it suffices to show that any separable skew-merged permutation, say π , can be drawn on an X. We prove this by induction on the length of π . Since π is separable, it can be decomposed as either a direct sum or a skew sum; without loss suppose that $\pi = \sigma \oplus \tau$ for non-empty σ and τ . By induction, both σ and τ can be drawn on an X.


If σ contains a nontrivial decreasing subsequence (one of length at least 2) then it is clear that τ must be entirely increasing, and we can draw π on an X by first drawing σ on an X and then drawing τ above and to the right of this σ -drawing. Similarly, if τ contains a nontrivial decreasing subsequence then π can be drawn on the X by first drawing τ on the X and then drawing the necessarily increasing σ below and to the left of this τ -drawing, completing the proof. \square

Considerations similar to (but simpler than) those for circle case show that the permutations drawable on an X can be encoded with the regular lan-

guage:

$$A^* \setminus (A^*2 \cup A^*3 \cup A^*4 \cup A^*31 \cup A^*41 \cup A^*31A^* \cup A^*42A^*).$$

We do not give the full details here, but instead sketch the construction on a canonical word. We begin by dividing the X into quadrants. We order the points by distance from the center of the X and demand that the point closest to the center lies in quadrant 1. Again we have freedom in choosing where diagonally opposite points lie, and so eliminate words containing a 31 or 42. Finally we deal with those cases where the points closest to the centre of a the X form a long increase or a long decrease; we encode these permutations as $1 \dots 1$ for an increase and $2 \dots 21$ for a decrease, thus we avoid words ending 41 and 31.



The above picture shows non-canonical and canonical drawings of the permutation 7253461. The word corresponding to the canonical drawing is 2413211.

The generating function for these permutations is:

$$\frac{1 - 2x}{1 - 4x + 2x^2}.$$

Curiously, this generating function also arises in Knuth [46, Section 5.4.8, Exercise 8], where it counts permutations π such that for every i there is




Figure 5.8: Knuth's exercise.

at most one $j < i$ with $\pi(j) > i$; see Figure 5.8. These permutations can be shown to be the avoiders of 3412, 3421, 4312, and 4321, and thus this equinumerosity result follows from the work of Atkinson [10]. A bijection between these two sets has been discovered by Elizalde [private communication]. Finally Atkinson [Private Communication] has demonstrated that the X class is precisely the set of permutations which can be constructed from a single point using $1 \oplus \tau$, $1 \ominus \tau$, $\tau \oplus 1$, $\tau \ominus 1$, where τ has already been constructed. However, permutations $1 \oplus \sigma \oplus 1$ and $1 \ominus \sigma \ominus 1$ appear twice among these forms. The enumeration result follows immediately from this observation.

5.7 Permutations Drawn from Parallel Lines

In this section we take our picture to be a set of k parallel lines in the plane. For any particular set of k parallel lines we see that the set of all permutations which can be drawn from them forms a closed class. Furthermore, it is clear that vertical and horizontal lines will yield only trivial results. Finally, observe that if a permutation can be drawn from k increasing parallel lines then its reverse can be drawn from k corresponding decreasing parallel lines. Thus we shall consider only lines of positive gradient.

For $k = 1, 2, 3$ we ask the following questions:

- Which permutations can be drawn from k parallel lines?
- How does the configuration of the lines affect the permutations which can be drawn?

The results for $k = 3$ are both striking and intricate. In particular we characterise the first uncountable set of independent closed classes to be found in a natural setting. That such uncountable sets exist was well known, however, all previous examples had been constructed by taking subsets of an infinite antichain either as basis elements or as part of a Sub construction.

5.7.1 Permutations Drawn from a Single Line

It is immediately apparent that only one permutation of each length n can be drawn from a single line of positive gradient. This is the increasing permutation of length n . The set of all such permutations forms a closed class, the class of increasing permutations.

5.7.2 Permutations Drawn from Two Infinite Parallel Lines

The observations on symmetries of the plane, Section 5.3, demonstrate that there is just one class of permutations drawn from two infinite parallel lines of positive gradient, since neither the angle of the lines, nor the distance between them is relevant.

Lemma 5.7.1. *Those permutations which can be drawn from two infinite parallel lines of positive gradient are precisely those that do not involve the permutation 321.*

Thus 321 is the basis of this class and a permutation can be drawn on two infinite parallel lines of positive gradient if every subpermutation of length three can.

Proof. It is apparent that any permutation which can be drawn from two parallel lines avoids 321 thus we need only prove that those permutations

which cannot be drawn must contain it. Let α be a permutation which cannot be drawn on two parallel lines. Furthermore suppose α is minimal in this respect under involvement, that is, every permutation involved in α can be drawn. We remove the minimum point of α and construct a 2 parallel line drawing. We do this greedily so that every point which can be drawn on the top line is; points lie on the bottom line only if they play the role of a 1 in a 21 pattern. It is clear that the minimum point cannot be added to this drawing on one of the lines, otherwise α is drawable. If there is a point above and to the left of this minimum point on the bottom line then we immediately have a 321 pattern. Otherwise the only points above and to the left of the minimum lie on the top line. These points can then be shifted up the top line until the minimum point can be placed. \square

Those permutations avoiding 321 have been studied at length, see Theorem 2.8.2. They are enumerated by the Catalan numbers. They are not partially well ordered.

Lemma 5.7.2. *The infinite two line class is equal to the grid class of the following infinite matrix.*

$$\begin{pmatrix} 0 & \dots & \dots & \dots & \dots & \dots & \dots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 1 & 1 & \ddots & \ddots & \vdots \\ 0 & 0 & 1 & 1 & \ddots & \ddots & \ddots & \vdots \\ 0 & 1 & 1 & 0 & \ddots & \ddots & \ddots & \vdots \\ 1 & 1 & 0 & 0 & \dots & \dots & \dots & 0 \end{pmatrix}$$

This infinite gridding matrix is a mapping from $\mathbb{N} \times \mathbb{N}$ into $\{0, 1\}$. It is 0 everywhere except along the diagonal and just below the diagonal. That this is a reasonable definition for an infinite matrix remains a moot point, however, the definitions from Chapter 4 extend naturally to this setting.

Proof. It is clear that this infinite grid class avoids 321, thus we need only show that every infinite 2 parallel line picture can be gridded in this manner. We first notice that for any finite permutation we may choose an arbitrary horizontal line below which we do not place any points, we call this line the base line. We begin on the base line. At the base of the second parallel line draw vertical and horizontal grid lines; where the vertical grid line crosses the first line draw a horizontal grid line; at the next crossing point draw a vertical grid line; repeat this process. See Figure 5.9. \square

5.7.3 Permutations Drawn from Two Finite Parallel Lines

It is surprising, given earlier results, to discover that two parallel lines of finite length in the plane yield a variety of smaller classes. In contrast to the infinite case these classes are all partially well ordered.

The results from section 5.3 demonstrate that the precise distance between two lines is unimportant. Instead we are interested in something closer to the ratio of the length of the lines to the distance between them. We will consider this concept in more detail later on.

Lemma 5.7.3. *Any finite two line class can be embedded into a finite grid class whose gridding matrix is a finite submatrix of the infinite matrix from Lemma 5.7.2.*

Proof. The finite length of the lines guarantees that the line drawing process from the proof of Lemma 5.7.2 will terminate. Thus the matrix will be finite. See Figure 5.9. \square

Lemma 5.7.4. *The class of permutations drawn from any pair of finite parallel lines is partially well ordered.*

Proof. It is a result of Vatter and Murphy [51] that the grid class of any matrix whose associated graph is a forest is partially well ordered, see Section 4.6. \square




Figure 5.9: Dividing two parallel lines into a grid class.

It is a strong conjecture of this author that the class of permutations drawn from any pair of finite parallel lines is finitely based. There are two possible approaches, the first would be to demonstrate that the grid class of any finite matrix of the staircase form (Lemma 5.7.2) is finitely based. Hence such a class is strongly finitely based, see Atkinson, Murphy, Ruškuc [12], and so every subclass is finitely based. A more elegant approach would be to demonstrate that the only infinite antichains in $\text{Av}(321)$ are of the same form as the antichain U , Definition 2.8.5. Clearly only oscillations, see Lemma 2.8.6, of a finite length can be drawn from a finite pair of parallel lines, so that all but finitely many members of this antichain would be removed by a single basis element.

5.7.4 Permutations Drawn from Three Infinite Parallel Lines

It follows from the results in section 5.3 that the angle of the lines and the absolute distances between them are irrelevant. Thus we are left to consider only the ratio of the distances between the lines, see Figure 5.10.

In fact this ratio is a somewhat misleading quantity. It is useful for parallel lines of infinite length at a forty-five degree angle but what we are actually seeking is something more geometrical.




Figure 5.10: Three parallel lines in the plane.

Definition 5.7.5. A rectangular walk on a three line picture is a path which is always horizontal or vertical, which changes direction only when it touches a line and which never repeats itself.

Definition 5.7.6. A three line picture is said to be *rational* if every rectangular walk which stays within a bounded area of the picture is finite.

Example 5.7.7. Consider three infinite lines of unit gradient. The first begins at the origin, the second at the point $(3, 0)$, the third at the point $(5, 0)$. This three line picture is rational. Every rectangular walk which stays within a bounded area is finite. See Figure 5.11.

Theorem 5.7.8. *There are uncountably many different classes of permutations which can be drawn from three parallel lines, one for every positive real number.*

Proof. Given fixed integers a and b and any small real number ϵ we consider three different infinite three line pictures, one of rational ratio a/b , one of ratio $a/(b + \epsilon)$ and one of ratio $a/(b - \epsilon)$. We exhibit a permutation which lies in each class but in neither of the others. Notice that we do not demand ϵ be irrational, the proof is sufficient to demonstrate the difference between any pair of three line classes.

- a/b

We begin by constructing a rectangular walk from the middle line which consists of b up then right steps followed by a down then left




Figure 5.11: A rational three line picture with a rectangular walk shown

steps. See Figure 5.11. the geometry of this situation demands that we return to our starting point.

We next place pairs of points around the vertices of this walk. We begin at the starting point of our walk on the middle line, the points must be sufficiently close together that any pair of points vertically between them on either the top or bottom line is distinct from any pair of points between them horizontally on either the top or bottom line. Label the lower point $x_{1,1}$ and the higher point $y_{1,1}$, next place a pair of points vertically and between these two points, on the top line, around the next vertex of the walk.. Label these points $x_{2,1}$ and $y_{2,1}$. Continue to place points contained in the previous pair either horizontally or vertically, around vertices of the walk, until points $x_{2b+2a,1}$ and $y_{2b+2a,1}$ are placed on the middle line. We see immediately that $x_{2b+2a,1}$ and $y_{2b+2a,1}$ lie horizontally between $x_{1,1}$ and $y_{1,1}$. We may place points $x_{1,2}$ and $y_{1,2}$ on the middle line, again around the

first vertex of our walk, and so repeat this loop as often as we like. Notice the telescoping structure which the pairs of points form. All permutations constructed in this manner will lie inside this picture class. See Figure 5.13.

- $a/(b + \epsilon)$

We attempt to place those permutations constructed on the a/b picture onto this picture.

We begin again by following the walk they are to follow. The geometry of the situation demands that after the $(2a + 2b)$ sets the walk must lie a distance $(a\epsilon)$ below the starting point.

As we attempt to place points we must again follow this walk to ensure the telescoping structure. Clearly we can place all those points until $x_{2b+2a,1}$ and $y_{2b+2a,1}$, which must lie vertically between $x_{2b+2a-1,1}$ and $y_{2b+2a-1,1}$ and horizontally between $x_{1,1}$ and $y_{1,1}$. However the geometry of this picture demands that $y_{2b+2a,1}$ lies a distance $a\epsilon$ below $y_{1,1}$. Since the maximum distance between the first pair $x_{1,1}$ and $y_{1,1}$ is bounded by the condition that points contained vertically and horizontally are distinct, we may simply repeat our looping construction until we run out of space. Thus if the maximum distance between $x_{1,1}$ and $y_{1,1}$ is z we choose n such that $na\epsilon > z$. Then we must place points $x_{2b+2a,n}$ and $y_{2b+2a,n}$ below the point $x_{1,1}$ and hence violate the horizontal containment condition. Not only have we failed to draw our intended permutation, but we have constructed a permutation which cannot be drawn on the a/b picture. See Figure 5.12.

- $a/(b - \epsilon)$

We attempt to place those permutations constructed on the a/b picture onto this picture.

We begin again by following the walk they are to follow. The geometry of the situation demands that after the $(2a + 2b)$ sets the walk must lie a distance $(a\epsilon)$ above the starting point.

As we attempt to place points we must again follow this walk to ensure




Figure 5.12: Three parallel line pictures of rational and irrational ratio

the telescoping structure. Clearly we can place all those points until $x_{2b+2a,1}$ and $y_{2b+2a,1}$, which must lie vertically between $x_{2b+2a-1,1}$ and $y_{2b+2a-1,1}$ and horizontally between $x_{1,1}$ and $y_{1,1}$. However the geometry of this picture demands that $x_{2b+2a,1}$ lie a distance $a\epsilon$ above $x_{1,1}$. Since the maximum distance between the first pair $x_{1,1}$ and $y_{1,1}$ is bounded by the condition that points contained vertically and horizontally are distinct, we may simply repeat our looping construction until we run out of space. Thus if the maximum distance between $x_{1,1}$ and $y_{1,1}$ is z we choose n such that $na\epsilon > z$. Then we must place points $x_{2b+2a,n}$ and $y_{2b+2a,n}$ above the point $y_{1,1}$ and hence violate the horizontal containment condition. Not only have we failed to draw our intended permutation, but we have constructed a permutation which cannot be drawn on the a/b or $a/(b + \epsilon)$ picture. It is also clear that our $a/(b + \epsilon)$ permutation cannot be drawn here. See Figure 5.12.

□

As with two lines we may attempt to express these classes as grid classes of infinite matrices. The results are striking.

Lemma 5.7.9. *The class of points drawn from three infinite parallel lines of rational ratio is a subclass of the grid class of an infinite matrix which avoids 4321.*

Proof. Any small area can be gridded into a finite matrix simply by drawing




Figure 5.13: Constructing permutations inside a rational ratio three line class

enough lines. The rational ratio guarantees that the line drawing process from Lemma 5.7.2 terminates. \square

We can, however, exhibit a permutation for each grid class which belongs to the grid class but not to the corresponding three line class. Together these permutations form an infinite antichain which we call the Escher antichain see Lemma 5.7.15.

Lemma 5.7.10. *The class of points drawn from three infinite parallel lines of irrational ratio is not a subclass of a 4321 avoiding grid class for any infinite matrix.*

Proof. Consider an irrational three line picture. Any rectangular walk must move from cell to cell, two vertices of such a walk may not lie in the same cell. Hence we require infinitely many cells to grid a bounded area of such a picture. Lemma 4.1.41 demonstrates that any attempt at a finite gridding can be defeated by a finite permutation. \square

Remark 5.7.11. This is something of a philosophical result, we demand that an infinite gridding matrix be a mapping from $\mathbb{N} \times \mathbb{N}$ into $\{0, 1, -1\}$. That this is a reasonable definition for an infinite matrix remains a moot point.

5.7.5 Permutations Drawn from Three Finite Parallel Lines

Here the distinction between rational and irrational ratios is again apparent.

Lemma 5.7.12. *A rational three finite parallel line picture may be gridded into a finite grid class which avoids 4321.*

On this occasion however we cannot guarantee that the associated graph is a forest, indeed on almost every occasion we would expect the graph to contain a cycle. Thus we do not necessarily expect these classes to be partially well ordered. However, the antichains constructed by Murphy and Vatter in [51] cannot be drawn on three parallel lines. Thus the question remains open.

Lemma 5.7.13. *An irrational three finite parallel lines picture cannot be gridded into a finite 4321 avoiding grid class.*

These results follow from the infinite cases, Lemmas 5.7.9 and 5.7.10.

It is a conjecture of this author that the grid class of any finite matrix is finitely based.

Definition 5.7.14. The *three line class*, T , is the set of all permutations which can be drawn from some three parallel line picture. That is, it is the union, over all possible ratios, of three line classes.

Lemma 5.7.15. *The three line class, T , is not finitely based.*

Proof. We exhibit an infinite antichain of basis elements. First we choose a rational three line picture and construct a finite walk around it. We place points around the vertices of this walk exactly as in the proof of 5.7.8, however instead of telescoping these points we allow each to be contained in the previous and contain the next, that is, when we come to place points $x_{2b+2a,1}$ and $y_{2b+2a,1}$, which our three lines force to lie horizontally between $x_{1,1}$ and $y_{1,1}$, we place them instead so that they contain $x_{1,1}$ and $y_{1,1}$ horizontally. It is necessary to add extra points to the top and bottom of the permutation, so that every point lies in a 321 pattern. It is clear that such a permutation cannot be drawn on three lines, and further that it is a basis element. Furthermore, we can construct a new permutation for every pair of integers a, b with a and b coprime. Clearly these permutations form an antichain. See Figure 5.14. We call this antichain the Escher antichain, because of its similarity to the picture Ascending and Descending, by M.C. Escher, which may be found in [31]. \square

5.8 The Convex Class

We have already considered those permutations which can be drawn from a particular convex shape, the circle (or diamond), in Section 5.5. Here we




Figure 5.14: The permutation $(7, 2, 1, 5, 8, 3, 6, 10, 9, 4)$, the shortest member of the Escher antichain, with cycle structure shown




Figure 5.15: The permutation 1243, convex and standard drawings.

consider those permutations which may be drawn from any convex shape. This class will be the union, over all convex shapes, of those permutations which may be drawn from each shape. This section contains joint work with Michael Albert, Steve Linton, Nik Ruškuc and Vince Vatter.

Definition 5.8.1. A permutation is said to be *convex* if it is isomorphic to some set of points in the plane all of which lie on the edge of some convex polygon.

Example 5.8.2. The permutation 1243 is convex, see Figure 5.15. Note that its “standard” drawing, the drawing with integer coordinates, is not convex. There are far fewer permutations whose standard drawing is convex than there are convex permutations. This should come as no surprise, permutations in the circle class do not generally have standard drawings which are circular.

It is clear that the set of all convex permutations forms a closed class, we will refer to this class as the convex class, \mathcal{C} . It should be noted immediately that the convex class is not a picture class. Instead it is a union of infinitely many picture classes. A characterisation which describes precisely which pictures are necessary would be very interesting. This author suspects that a set of rectangles with one side of variable length would suffice.




Figure 5.16: Plots of the basis elements of the class of a convex permutations.

5.8.1 Characterisation of the Convex Class

Our first task is to characterise the convex permutations. We do so by identifying the basis elements of the class.

Theorem 5.8.3. *The convex class is finitely based with basis elements: $\{52341, 52314, 51342, 51324, 42351, 42315, 41352, 41325, 25341, 25314, 15342, 15324, 24351, 24315, 14352, 14325\}$.*

We can also rephrase Theorem 5.8.3 as follows:

Corollary 5.8.4. *A permutation can be drawn on a convex shape if every five point subpermutation can.*

We have already defined left-to-right minima, see Definition 2.8.3, however we repeat the definition here for the sake of completeness.

Definition 5.8.5. In a permutation π , a point p is said to be a *left-to-right minimum* if there is no point to the left of p which is smaller than it.

Similarly we define *left-to-right maxima*, *right-to-left minima* and *right-to-left maxima*.

Definition 5.8.6. In a permutation π , a point p is said to be *extremal* if it is either a left-to-right maximum, a left-to-right minimum, a right-to-left maximum or a right-to-left minimum.

Notice that the basis elements listed in Theorem 5.8.3 and shown in Figure 5.16 all consist of a central element surrounded by four extremal elements, one of each type.

Definition 5.8.7. A permutation, π , is said to be *extremal* if every point, p , in π is extremal.

Lemma 5.8.8. *The set of all extremal permutations forms a closed class with basis in Theorem 5.8.3.*

Proof. It is clear that an extremal permutation avoids these patterns. Furthermore it is clear that any element which does not play the role of a ‘3’ in one of the above patterns is an extremal point, so that any permutation which avoids these patterns is extremal. \square

Lemma 5.8.9. *Every convex permutation is extremal.*

Proof. This is clear from any convex drawing of the permutation. \square

We can now proceed to prove the main theorem.

Proof of Theorem 5.8.3. We will prove that every extremal permutation can be drawn on a convex polygon. We require four additional lemmas.

Lemma 5.8.10. *Any permutation which avoids 213 and 312 can be drawn on a convex cap using only vertical stretching operations. See Figure 5.17*

Proof. A permutation which avoids both 213 and 312 can be characterised as the horizontal juxtaposition of an increase and a decrease. Given a point




Figure 5.17: Forming a convex cap by vertical stretches

set in the plane, order isomorphic to such a permutation, we join the points from left to right to form a cap. It is clear that using only vertical stretches we can make this cap convex, and that doing so will produce a point set which is isomorphic to our initial point set. \square

Lemma 5.8.11. *Any permutation which avoids 132 and 231 can be drawn on a convex cup using only vertical stretches.*

Proof. The proof of lemma 5.8.10 goes through, mutatis mutandis. \square

Notice that having constructed a convex cap or a convex cup we may scale the entire point set vertically without destroying its convex nature.

Lemma 5.8.12. *Any permutation which avoids 321 can be drawn on a pair of parallel lines of any positive gradient.*

Proof. We have already proved the slightly stronger result, that the class of permutations which can be drawn on two parallel lines of positive gradient is precisely $\text{Av}(321)$, see Lemma 5.7.1. \square

By symmetry we also have:

Lemma 5.8.13. *Any permutation which avoids 123 can be drawn on a pair of parallel lines of any negative gradient.*

We now return to the proof of the main theorem.

For any extremal permutation we can identify topmost, bottommost, rightmost and leftmost points, which we shall label T , B , R and L respectively. We shall consider the possible patterns of these elements. There are eleven cases in all, $\{1, 12, 21, 132, 213, 231, 312, 2143, 2413, 3142, 3412\}$, after symmetry these reduce to five cases:

- 1 The simplest case, such a permutation has size one and so is trivially convex.
- 12 Such a permutation must avoid 321 since it is also extremal. We may then draw it on two parallel lines.
- 132 We divide such permutations into two parts vertically. The set of all points below R , which plays the role of ‘2’, must avoid 321 and so can be drawn on two parallel lines. The set of points above R must avoid 312 and 213 and so can be drawn on a convex cap. Appropriate vertical scaling of this top segment will ensure a convex shape.
- 2143 We divide such permutations into three parts vertically. The set of all points between L and R must avoid 321 and so can be drawn on two parallel lines. The set of points below L must avoid 132 and 231 and so can be drawn on a convex cup. The set of points above R must avoid 213 and 312 and so can be drawn on a convex cap. Appropriate vertical scaling of the cup and cap will ensure a convex shape.
- 2413 We divide such permutations into three parts vertically. The set of all points between L and R must avoid 321 and so can be drawn on two parallel lines. The set of points below L must avoid 132 and 231 and so can be drawn on a convex cup. The set of points above R must avoid 213 and 312 and so can be drawn on a convex cap. Appropriate vertical scaling of the cup and cap will ensure a convex shape.

□

5.8.2 Enumeration of the Convex Class

To count the convex permutations we use the insertion encoding, introduced by Albert, Linton and Ruškuc [5]. The insertion encoding is a correspondence between permutation classes and languages. This correspondence equates each permutation with a word describing the way in which it “evolved”. The permutation is built up in a step by step process. At each step a new maximum element is inserted into an open *slot* (represented by a \diamond), until the permutation is complete, at which point no open slots remain. The insertion can occur in one of four ways:

- The slot can be filled (replacing \diamond by n).
- The new maximum can be placed on the left of the slot (replacing \diamond by $n\diamond$).
- The new maximum can be placed on the right of the slot (replacing \diamond by $\diamond n$).
- The new maximum can be placed in the middle of the slot (replacing \diamond by $\diamond n \diamond$).

These operations are called fills (denoted **f**), lefts (denoted **l**), rights (denoted **r**) and middles (denoted **m**) respectively. In addition these symbols are subscripted by the slot they apply to (read from left to right). For example, the permutation 31254 has insertion encoding $\mathbf{m}_1\mathbf{l}_2\mathbf{f}_1\mathbf{r}_1\mathbf{f}_1$ because its evolution is:

$$\begin{array}{c} \diamond \\ \diamond 1 \diamond \\ \diamond 12 \diamond \\ 312 \diamond \\ 312 \diamond 4 \\ 31254 \end{array}$$

Each intermediate step in this process is called a configuration. The insertion encoding can be used to enumerate some regular classes automatically. However, its full power can only be brought to bear when the number of slots is unlimited, and in some such cases we can derive algebraic generating functions.

Theorem 5.8.14. *The generating function for the convex class is:*

$$\frac{t(1 - 6t - 10t^2 + 4t^2\sqrt{1 - 4t})}{(1 - 4t)^2}$$

The Taylor series begins

$$t + 2t^2 + 6t^3 + 24t^4 + 104t^5 + 464t^6 + 2088t^7 + 9392t^8 + 42064t^9 + \dots$$

Proof. Using the insertion encoding to build a convex permutation, consider the insertion of an arbitrary element in a configuration. If, before the insertion, there were smaller elements on both sides of it and, after the insertion, there are slots on both sides of it, then the element inserted will inevitably become a ‘3’ in one of the basis elements of \mathcal{C} . Thus, no such insertions can be permitted. On the other hand, if every element is inserted in such a way that either there are smaller elements on at most one side of it, or slots remain on at most one side of it after insertion, then no inserted element can play the role of a ‘3’ in one of the basis elements of \mathcal{C} and hence the final permutation produced will belong to \mathcal{C} . Summarizing these conditions:

- if there are three or more slots, then no insertions are possible in any but the outer two slots;
- if there are elements at both ends of the configuration, then the only permitted operations are at the extremes of the configuration (fills at either end, lefts at the left hand end, or rights at the right hand end);
- if there are elements at a single end of the configuration, say the left hand end, then at that end only lefts and fills are possible, while at the other end all operations are possible;

- if there are elements at neither end of the configuration then all operations are possible in the two (or one) outermost slots.

The enumeration strategy follows this summary. Specifically, we introduce six generating functions. Each counts (using the coefficient of t^k) the number of words of length k from a configuration to the final configuration. Three correspond to the configurations having a single slot with small elements at neither end, a single end, or both ends. The remaining three deal with the configurations having two or more slots – in these, a second variable z is used to tag the number of slots (minus two, for convenience). The six generating functions will be called: $B_1(t)$, $B(t, z)$ (for small elements at both ends), $S_1(t)$, $S(t, z)$ (for small elements at a single end), and $N_1(t)$, $N(t, z)$ (for small elements at neither end).

We will now compute these generating functions, beginning with the B 's. These are very straightforward. For B_1 , the only allowed operations are \mathbf{l}_1 , \mathbf{r}_1 and \mathbf{f}_1 (which terminates the encoding). Thus:

$$\begin{aligned} B_1(t) &= 2tB_1(t) + t \\ B_1(t) &= \frac{t}{1-2t}. \end{aligned}$$

If there are $s \geq 2$ slots and small elements at both ends, then the allowed operations are \mathbf{l}_1 , \mathbf{r}_{-1} and $\mathbf{f}_{\pm 1}$ (note that we use the subscript -1 to refer the rightmost open slot). Thus:

$$\begin{aligned} B(t, z) &= 2tB(t, z) + 2tzB(t, z) + 2tB_1(t) \\ B(t, z) &= \frac{2t^2}{(1-2t)(1-2t-2tz)}. \end{aligned}$$

The situation for the S generating functions is a little more complex. For simplicity we assume that there are small elements at the right hand end of the configuration (and hence not at the left). The allowed operations are then: \mathbf{f}_1 and \mathbf{l}_1 (leading to B configurations), \mathbf{m}_1 and \mathbf{r}_1 (leading to S configurations) and \mathbf{f}_{-1} and \mathbf{r}_{-1} (also leading to S configurations). These

are all distinct if there are at least two slots, but only the operations carrying a 1 subscript can be applied when there is a single slot. Thus:

$$\begin{aligned} S_1(t) &= t + tB_1(t) + tS(t, 0) + S_1(t) \\ S(t, z) &= tS_1(t) + tzS(t, z) + t(S(t, z) - S(t, 0))/z + \\ &\quad 2tS(t, z) + (t + tz)B(t, z) + tB_1(t). \end{aligned}$$

The $S(t, z)$ terms in the second equation can be collected together in order to apply the kernel method:

$$(1 - 2t - tz - t/z) S(t, z) = tS_1(t) - tS(t, 0)/z + (t + tz)B(t, z) + tB_1(t).$$

Setting the first factor on the left to 0, specifically

$$z = \frac{1 - 2t - \sqrt{1 - 4t}}{2t}$$

and substituting in the right hand side allows us to produce a second equation connecting $S_1(t)$ and $S(t, 0)$. Solving this system yields:

$$\begin{aligned} S_1(t) &= \frac{t}{\sqrt{1 - 4t}} \\ S(t, 0) &= \frac{(2t^2 - 3t + 1)\sqrt{1 - 4t} - 4t^2 + 5t - 1}{(1 - 2t)(1 - 4t)}. \end{aligned}$$

In turn, we can now substitute these values in the original equation and solve for $S(t, z)$. The resulting complex expression is of interest to us only as an ingredient in the next step of the computation so we shall not reproduce it here. It can be found using the Maple code in the appendix of [3], this code can also be used to verify all of the computations found in this proof.

We repeat the procedure to produce the N generating functions. This time the original equations are:

$$\begin{aligned} N_1(t) &= t + tN(t, 0) + 2tS_1(t) \\ N(t, z) &= 2tN(t, z) + 2t(N(t, z) - N(t, 0))/z + (2t + 2tz)B(t, z) + 2tB_1(t). \end{aligned}$$

The kernel in this case is simply $z = 2t/(1 - 2t)$ and substitution and simplification yields:

$$N_1(t) = \frac{t(1 - 6t - 10t^2 + 4t^2\sqrt{1 - 4t})}{(1 - 4t)^2}.$$

which, as the state with 1 slot and no small elements on either side is the initial configuration for the insertion encoding, and hence is the generating function of \mathcal{C} . \square

The kernel method, used above, is a standard but some what technical tool. It was originally developed by Donald Knuth [45]. He used it to enumerate the Schröder numbers in connection with his work on restricted dequeues. Indeed, the kernel method has become an import part of the study of pattern classes, see, for example, [19].

5.8.3 Properties of the Convex Class

Definition 5.8.15. A rectangular walk inside a polygon is a path which is always horizontal or vertical, which changes direction only when it touches an edge of the polygon and which never repeats itself.

Definition 5.8.16. A polygon is said to be *rational* if every rectangular walk is finite.

These definitions are the natural extension of Definition 5.7.5 which considered three parallel lines. This author suspects there is a radical difference between those classes drawn from polygons which are rational and those drawn from polygons which are not.

Example 5.8.17. Figure 5.18 shows rectangular walks on two tilted squares, one at forty five degrees and one at thirty degrees. For the square at forty five degrees (diamond) every walk is finite, for the square at thirty degrees almost every walk is infinite, although all walks tend to the unique inscribed square.




Figure 5.18: Rectangular walks on tilted squares.

Recall the definition of atomicity, that a class is atomic if it cannot be expressed as the union of two proper subclasses, and that a class is atomic if every pair of permutations possess the joint embedding property, see Theorem 2.5.3.

Theorem 5.8.18. *The convex class is not atomic, indeed it cannot be written as the union of finitely many atomic classes.*

Proof. Consider the following infinite sequence of permutations.




Figure 5.19: Permutations which do not join in the convex class.

45817236
 4581(10)23(11)6(12)97
 4581(10)23(11)6(13)(16)7(15)9(12)(14)
 ⋮
 ⋮
 4581(10)23(11) ... (n-6)n(n-3)(n-5)
 4581(10)23(11) ... (n-1)(n-7)(n-4)(n-2)
 ⋮
 ⋮

These permutations are constructed from the convex figures shown in Figure 5.19 by following the same telescoping methods as in the proof of Theorem 5.7.8.

To see that joint embedding is impossible consider the rectangular walks from which these permutations are constructed. It is easy to see that no convex shape can admit any pair of these walks. Hence these permutations do not jointly embed. \square

Since the permutations in the proof of Theorem 5.8.18 do not jointly embed

inside the convex class, they are incomparable, i.e. they form an infinite antichain, and so the convex class is not partially well ordered. Another way to see this is to consider the antichains constructed by Murphy and Vatter in [51], many of which lie inside the convex class.

5.8.4 Permutations Drawn from Fixed Convex Polygons

The convex class may be considered as the union, over all convex polygons, of the class of permutations drawn from each polygon. In this section we consider those subclasses drawn from individual polygons.

Lemma 5.8.19. *There exists a convex polygon from which we can draw a closed class which is not partially well ordered.*

Proof. We construct an infinite antichain of permutations drawn from a tilted square. In doing so we defeat a conjecture of Murphy [50], that “the closure of a tilted square does not contain an infinite antichain, no matter what the angle of tilt.” A schematic drawing of a typical member of this antichain is shown in Figure 5.20. To construct the permutation place a set of four decreasing points at a . Place a further four decreasing points at b so that all four lie horizontally between the middle pair at a . At the next vertex place a pair of points vertically between the middle pair at b . At each further vertex place a pair of points either horizontally or vertically between the previous pair as appropriate. Finally place a set of three points at z between the previous pair. Further members can be constructed by taking different numbers of steps around the square before placing the set of three points which act as the end anchor. Notice that we are again telescoping points around a rectangular walk, in this case moving clockwise around the square. It is clear that these permutations do indeed form an antichain. \square

Conjecture 5.8.20. *The set of all permutations drawn from a fixed convex polygon is partially well ordered if and only if the polygon is rational.*

Lemma 5.8.21. *The set of all permutations drawn from any fixed convex polygon with no edges which are either vertical or horizontal is atomic.*




Figure 5.20: A schematic drawing of of typical member of an antichain constructed on a tilted square.




Figure 5.21: A rational quadrilateral, here a/b is $2/3$.

Proof. Follows from Theorem 2.5.3. We need only note that we can find an underlying bijection f from some subset of \mathbb{R} into some subset of \mathbb{R} such that $Sub(f)$ is equal to the class drawn from our polygon. \square

Lemma 5.8.22. *There are uncountably many closed classes which may be drawn from convex quadrilaterals.*

The proof is a simple extension of a Theorem 5.7.8, which demonstrates that uncountably many closed classes may be drawn from three parallel lines in the plane. It rests on the observation that for a rectangle rotated forty five degrees every rectangular path is finite (i.e. the picture is rational) if and only if the ratio of the lengths of the sides is rational. See Figure 5.21.

Corollary 5.8.23. *Not every class drawn from a convex quadrilateral is finitely based.*

Proof. There are uncountably many different classes, but only countably many finite antichains to act as bases. \square

Conjecture 5.8.24. *A closed class drawn from a fixed convex polygon is finitely based if and only if that polygon is rational.*

Conjecture 5.8.25. *The set of all permutations drawn from any fixed convex polygon is rationally enumerated if and only if that polygon is rational.*

5.8.5 A Combinatorial Problem for Permutations

In [29] Erdős and Szekeres give the following problem:

“Can we find for a given number n a number $N(n)$ such that from any set containing at least N points it is possible to select n points forming a convex polygon?”

It is termed the “Happy Ending Problem”.

Erdős and Szekeres proved that $N(n)$ exists for all n and gave the following upper and lower bounds on its value:

$$2^{n+1} + 1 \leq N(n) \leq \binom{2n-4}{n-2} + 2$$

The lower bound, $2^{n+1}+1$, is sharp for $n = 2, 3, 4, 5$ and has been conjectured to be sharp for all n .

The upper bound has been gradually improved over time, first by Erdős and Szekeres themselves in [30] to $\binom{2n-4}{n-2} + 1$. Then by Chung and Graham in [28] to $\binom{2n-4}{n-2}$. Then by Kleitman and Pachter in [44] to $\binom{2n-4}{n-2} + 7 - 2n$. The current tightest bound was set by Tóth and Valtr in [60] at $\binom{2n-5}{n-2} + 2$.

We consider a similar problem, but for permutations.

Definition 5.8.26. For every integer n let $f(n)$ be the smallest integer such that every permutation of length $f(n)$ contains a convex permutation of length n .

Computing $f(n)$ is a “Happy Ending Problem” for permutations.

In [27] Chung considers a similar problem, she proves that in any sequence of n distinct real numbers (permutation) there is a unimodal subsequence

(cup or cap) of length at least $(3n - 3/4)^{1/2} - 1/2$ and that this bound is tight.

Proposition 5.8.27. *There exists a permutation of length $\frac{(n-1)(n+3)}{8}$ whose longest convex subpermutation is of length $n - 1$*

A *layered* permutation is constructed by replacing each point in an increasing permutation with a decreasing sequence of points, i.e. a layered permutation is a sequence of decreases, each lying above and to the right of the previous one.

Proof. We construct a layered permutation with this property. It consists of m layers. The first $m/2$ layers have lengths which increase by 2 each time from 2 up to m . The remaining $m/2$ layers have lengths which decrease by 2 each time from m to 2. It is clear that the longest convex subpermutation has length $2m$ and that the permutation contains $m(m+2)$ points. Setting $2m = n - 1$ yields our bound. See Figure 5.22. \square

Proposition 5.8.28. *Every permutation of length $\frac{n^4}{4(n-2)^2}$ contains a convex subpermutation of length n .*

Proof. We use a double application of the Erdős-Szekeres theorem [29], see Section 2.2. First let π be any permutation of length k . Let a be the length of the longest increasing subsequence of π . Let b be the length of the longest decreasing subsequence. Assume without loss that $a \geq b$. The proof of the Erdős-Szekeres theorem implies $ab \geq k$. There are three cases:

1. $a \geq 2\sqrt{k}$

The length of the longest convex subpermutation is at least the length of this increase. Thus π contains a convex subpermutation of length $2\sqrt{k}$. If $k = \frac{n^4}{4(n-2)^2}$, then $2\sqrt{k} > n$ as required.

2. $a < 2\sqrt{k}$ and $b \leq \sqrt{k}$

We will use the result of Lemma 5.8.12. Remove the a points of




Figure 5.22: A permutation of length $\frac{(n-1)(n+3)}{8}$ with longest convex sub-permutation is of length $n - 1$.

our longest increasing subsequence from π . The permutation that remains has a longest decreasing subsequence of length $b \leq \sqrt{k}$ and a longest increasing subsequence on length c . It contains at least $k - 2\sqrt{k}$ points. Apply the Erdős-Szekeres theorem again: $bc \geq k - 2\sqrt{k}$ hence $c \geq \frac{k-2\sqrt{k}}{b}$. Since $b \leq \sqrt{k}$ we have $c \geq \frac{k-2\sqrt{k}}{\sqrt{k}}$. Together these two increasing subpermutations will form a convex subpermutation of length $a + c = \frac{2k-\sqrt{k}}{\sqrt{k}}$. If $k = \frac{n^4}{4(n-2)^2}$, then $\frac{2k-\sqrt{k}}{\sqrt{k}} > n$ as required.

3. $a < 2\sqrt{k}$ and $b > \sqrt{k}$

Here we use Lemma 5.8.12 or its dual. This time remove the a points of the longest increasing subsequence and the b points of the longest decreasing subsequence. The subpermutation which remains has length at least $k - 4\sqrt{k}$. It has a longest increase of length c and a longest decrease of length d . Apply Erdős-Szekeres again to get $cd \geq k - 4\sqrt{k}$. Without loss let us assume that $c \geq \sqrt{k - 4\sqrt{k}}$. Take the two increases with a and c points to get a convex sub-permutation of length $a + c \geq \sqrt{k} + \sqrt{k - 4\sqrt{k}}$. If $k = \frac{n^4}{4(n-2)^2}$, then $\sqrt{k} + \sqrt{k - 4\sqrt{k}} = n$. If d had been greater than c we would have chosen the two decreases and performed the same analysis.

□

Finally, and for the sake of clarity we will reverse the problem and take a limit. We wish to find the limit, as n tends to infinity, of the length of the longest convex subpermutation which we can guarantee in any permutation of length n . Let us call this limit l . Our results now become:

$$2\sqrt{n} \leq l \leq 2\sqrt{2n}.$$

Tightening these bounds is an open problem.

5.9 An Extension to Higher Dimensions

Permutations can be drawn from sets of points in the plane precisely because permutations can be defined as point sets under two linear orders and points in the plane can be defined as pairs of real numbers. There is an obvious extension to higher dimensions, see, for example, Cameron [26]. Indeed higher dimensional permutations have also been studied in the setting of token passing networks and sorting machines, see Atkinson, Walker and Linton [48]. However, we will stick to our picture based world: here the natural symmetries are especially apparent.

Definition 5.9.1. An *r-picture* is a set of points in r -dimensional Euclidean space, that is, a set of r -tuples of real numbers or a set of vectors in \mathbb{R}^r .

Definition 5.9.2. An *r-generic set* is a set of points in \mathbb{R}^r such that no pair share the same k th ordinate for any k in $\{1 \dots r\}$.

Thus the generic sets considered in Chapter 4 become 2-generic sets.

Definition 5.9.3. Given two r -generic sets S and T with

$$S = \{(s_{1,1}, s_{1,2}, \dots, s_{1,r}), \dots, (s_{n,1}, s_{n,2}, \dots, s_{n,r})\}$$

we say that S is *involved* in T , written $S \preceq T$ if there exist a set of r order preserving mappings $\{f_1, \dots, f_r\}$, from \mathbb{R} into \mathbb{R} , such that the set

$$S_{f_1, \dots, f_r} = \{(f_1(s_{1,1}), f_2(s_{1,2}), \dots, f_r(s_{1,r})), \dots, (f_1(s_{n,1}), f_2(s_{n,2}), \dots, f_r(s_{n,r}))\}$$

is contained in T .

Just as in the geometric definition of involvement this is a pre-order.

Definition 5.9.4. Two r -generic sets S and T are said to be order isomorphic if $S \preceq T$ and $T \preceq S$.

Clearly this is an equivalence relation. If we order r -generic sets by involvement, then factor this partially ordered set by our equivalence we obtain the

r -dimensional analogue of involvement on permutations. If we were to allow finite r -dimensional pictures, rather than generic sets, we would obtain an r -dimensional analogue of involvement on 0-1 matrices, see Section 1.3.2.

As with ordinary permutations we can also construct other definitions of r -dimensional permutations. We will give only the relational definition together with a further definition of involvement (although the temptation of the trijection [sic] is hard to resist).

Definition 5.9.5. An r -dimensional permutation of length n is a set of n points under r linear orders, $\leq_1, \leq_2, \dots, \leq_r$.

Definition 5.9.6. The r -dimensional permutation image of an r -generic set, S , $\Pi(S)$ is the r -dimensional permutation which is isomorphic to S .

Definition 5.9.7. Given two r -dimensional permutation σ and τ we say that σ is *involved* in τ , written $\sigma \preceq \tau$ if and only if σ is a subset of some r -generic set S whose permutational image is τ .

Definition 5.9.8. A set of r -dimensional permutations which is closed under involvement is called an r -dimensional pattern class.

Clearly these classes can be characterised by the minimum r -dimensional permutations they avoid. This set will be termed the *basis* of the class. Many of the results and conjectures for (two dimensional) permutation patterns translate directly into the higher dimensional setting. We begin with one of the most famous, the r -dimensional analogue of the Stanley-Wilf conjecture. Marcus and Klazar, [43], extended Marcus and Tardos's proof of the Stanley-Wilf conjecture, [49], to relational structures of arbitrary dimension.

Theorem 5.9.9 (Marcus and Klazar).

The number of r -dimensional permutations of length n which avoid a particular r -dimensional permutation is less than $(C^n)^{r-1}$ for some constant C .

The Erdős-Szekeres theorem also translates directly into this setting. We define an r -dimensional permutation to be monotone if it is monotone with

respect to every pair of orderings, thus there are 2^{r-1} directions of monotonicity. The result is due to De Bruijn, see Kruskal [47].

Lemma 5.9.10. *Every r -dimensional permutation of length $n^{2^{r-1}} + 1$ contains a monotone subpermutation of length $n + 1$.*

The proof is simply a repeated application of the proof of the Erdős-Szekeres theorem [29].

Our theories for picture classes also translate, thus it is reasonable to ask which 3-dimensional permutations can be drawn from a sphere, which r -dimensional permutations can be drawn on a hypersphere, which r -dimensional permutations can be drawn from a pair of parallel hyperplanes or which r -dimensional permutations can be drawn on an r -dimensional convex hull.

Lemma 5.9.11. *The class of r -dimensional permutations which can be drawn on an r -dimensional convex hull is finitely based, with basis elements of length $2^r + 1$.*

We can rephrase this theorem: An r -dimensional permutation of length n can be drawn on a convex hull if every subpermutation of length $2^r + 1$ can.

Proof. The proof in the two dimensional setting carries through, [3], however in r dimensions there are 2^r types of extremal point. \square

Lemma 5.9.12. *The class of all r -dimensional permutations which can be drawn on an $(r - 1)$ -dimensional hyperplane is finitely based with basis elements of length $r + 1$.*

Remark 5.9.13. $r + 1$ is the Helly number for r -dimensional Euclidean space.

Proof. An r -generic set of size r defines a $(r - 1)$ -dimensional hyperplane. By way of contradiction consider a large basis element. We choose r points from this basis element and consider all $r + 1$ point subsets which contain these points. Clearly they can all be drawn on a hyperplane. Appropriate

axis parallel scaling will allow us to line up these hyperplanes, where upon appropriate scaling within the hyperplane will allow us to align our control points. Finally we may overlay all our hyperplanes to produce a planar representation of our basis element. \square

Remark 5.9.14. This provides an alternative proof that the union of the class of increasing permutations with the class of decreasing permutations is finitely based with basis elements of length 3.

5.10 Conclusions for Picture Classes

Picture classes are an innovation in the field of pattern class research. They give a new way of defining closed classes, using the geometrical ideas of stretching and squashing to emphasise the relational structure of permutations. The study of simple examples such as the circle and the X has shown that this approach can lead to elegant and complete results. Furthermore, it is highly unlikely that anyone would study the circle class without approaching it from this viewpoint.

The more complicated examples, considering three parallel lines and fixed convex polygons have served to further illuminate the depth and intricacy that pattern classes may possess. In particular the exposition of uncountably many closed classes drawn from three lines and from particular fixed polygons demonstrates the huge complexity that even simple constructions can offer.

Finally the prospect of extension to higher dimensions offers many intriguing possibilities of its own. Although many of the two dimensional problems are incredibly difficult, perhaps intractable, these generalisations may throw up new insights which can be applied to the two dimensional case.

Bibliography

- [1] M. H. Albert and M. D. Atkinson. Simple permutations and pattern restricted permutations. *Discrete Math.*, 300(1-3):1–15, 2005.
- [2] M. H. Albert, M. D. Atkinson, and M. Klazar. The enumeration of simple permutations. *J. Integer Seq.*, 6(4):Article 03.4.4, 18 pp. (electronic), 2003.
- [3] M. H. Albert, M. D. Atkinson, S. Linton, N. Ruškuc, Vatter V., and S. Waton. On convex permutations. In preparation.
- [4] M. H. Albert, M. D. Atkinson, and N. Ruškuc. Regular closed sets of permutations. *Theoret. Comput. Sci.*, 306(1-3):85–100, 2003.
- [5] M. H. Albert, S. Linton, and N. Ruškuc. The insertion encoding of permutations. *Electron. J. Combin.*, 12(1):Research paper 47, 31 pp. (electronic), 2005.
- [6] M. H. Albert, N. Ruškuc, and S. Linton. On the permutational power of token passing networks. In preparation.
- [7] D. André. Sur les permutations alternées. *J. Math. Pures Appl.*, 7:167–184, 1881.
- [8] M. D. Atkinson. Permutations which are the union of an increasing and a decreasing subsequence. *Electron. J. Combin.*, 5:Research paper 6, 13 pp. (electronic), 1998.

- [9] M. D. Atkinson. Restricted permutations. *Discrete Math.*, 195(1-3):27–38, 1999.
- [10] M. D. Atkinson. Some equinumerous pattern-avoiding classes of permutations. *Discrete Math. Theor. Comput. Sci.*, 7(1):71–73 (electronic), 2005.
- [11] M. D. Atkinson, M. J. Livesey, and D. Tulley. Permutations generated by token passing in graphs. *Theoret. Comput. Sci.*, 178(1-2):103–118, 1997.
- [12] M. D. Atkinson, M. M. Murphy, and N. Ruškuc. Partially well-ordered closed sets of permutations. *Order*, 19(2):101–113, 2002.
- [13] M. D. Atkinson, M. M. Murphy, and N. Ruškuc. Sorting with two ordered stacks in series. *Theoret. Comput. Sci.*, 289(1):205–223, 2002.
- [14] M. D. Atkinson, M. M. Murphy, and N. Ruškuc. Pattern avoidance classes and subpermutations. *Electron. J. Combin.*, 12(1):Research paper 60, 18 pp. (electronic), 2005.
- [15] M. D. Atkinson and T. Stitt. Restricted permutations and the wreath product. *Discrete Math.*, 259(1-3):19–36, 2002.
- [16] Michael D. Atkinson and Robert Beals. Permuting mechanisms and closed classes of permutations. In *Combinatorics, computation & logic '99 (Auckland)*, volume 21 of *Aust. Comput. Sci. Commun.*, pages 117–127. Springer, Singapore, 1999.
- [17] Eric Babson and Julian West. The permutations $123p_4 \cdots p_m$ and $321p_4 \cdots p_m$ are Wilf-equivalent. *Graphs Combin.*, 16(4):373–380, 2000.
- [18] J. Backelin, Julian West, and Guoce Xin. Wilf-equivalence for singleton classes. Preprint.
- [19] C. Banderier, M. Bousquet-Mlou, A. Denise, Ph. Flajolet, D. Gardy, and D. Gouyou-Beauchamps. On generating functions of generating

- trees. In *Proceedings of 11th Conference on Formal Power Series and Algebraic Combinatorics (FPSAC'99)*, pages 40–52, 1999.
- [20] Miklós Bóna. A survey of stack-sorting disciplines. *Electron. J. Combin.*, 9(2):Article 1, 16 pp. (electronic), 2003.
- [21] Miklós Bóna. *Combinatorics of permutations*. Discrete Mathematics and its Applications (Boca Raton). Chapman & Hall/CRC, Boca Raton, FL, 2004. With a foreword by Richard Stanley.
- [22] Prosenjit Bose, Jonathan F. Buss, and Anna Lubiw. Pattern matching for permutations. *Inform. Process. Lett.*, 65(5):277–283, 1998.
- [23] Robert Brignall, Sophie Huczynska, and Vincent Vatter. Decomposing simple permutations, with enumerative consequences. In preparation.
- [24] Robert Brignall, Sophie Huczynska, and Vincent Vatter. Simple permutations and algebraic generating functions. In preparation.
- [25] Robert Brignall, Nik Ruškuc, and Vincent Vatter. Simple permutations: decidability and unavoidable substructures. In preparation.
- [26] Peter J. Cameron. Infinite permutation groups in enumeration and model theory. In *Proceedings of the International Congress of Mathematicians, Vol. I, II (Kyoto, 1990)*, pages 1431–1441, Tokyo, 1991. Math. Soc. Japan.
- [27] F. R. K. Chung. On unimodal subsequences. *J. Combin. Theory Ser. A*, 29(3):267–279, 1980.
- [28] F. R. K. Chung and R. L. Graham. Forced convex n -gons in the plane. *Discrete Comput. Geom.*, 19(3, Special Issue):367–371, 1998. Dedicated to the memory of Paul Erdős.
- [29] Paul Erdős and George Szekeres. A combinatorial problem in geometry. *Compos. Math.*, 2:463–470, 1935.

- [30] P. Erdős and G. Szekeres. On some extremum problems in elementary geometry. *Ann. Univ. Sci. Budapest. Eötvös Sect. Math.*, 3–4:53–62, 1960/1961.
- [31] M. C. Escher. *The Graphical Work of M. C. Escher*. Pan Books Ltd, London, sixth edition, 1978.
- [32] Stefan Felsner. Empty rectangles and graph dimension, 2006.
- [33] Roland Fraïssé. Sur l’extension aux relations de quelques propriétés des ordres. *Ann. Sci. Ecole Norm. Sup. (3)*, 71:363–388, 1954.
- [34] Zoltán Füredi and Péter Hajnal. Davenport-Schinzel theory of matrices. *Discrete Math.*, 103(3):233–251, 1992.
- [35] Sir Thomas Heath, editor. *The thirteen books of Euclid’s elements (3 vols)*. Dover, 2 edition, 1956.
- [36] Graham Higman. Ordering by divisibility in abstract algebras. *Proc. London Math. Soc. (3)*, 2:326–336, 1952.
- [37] Wilfrid Hodges. *Model theory*, volume 42 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1993.
- [38] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to automata theory, languages, and computation*. Addison-Wesley Publishing Co., Reading, Mass., 2nd edition, 2001.
- [39] Sophie Huczynska and Vincent Vatter. Grid classes and the Fibonacci dichotomy for restricted permutations. *Electron. J. Combin.*, 13:Research paper 54, 14 pp. (electronic), 2006.
- [40] Tomáš Kaiser and Martin Klazar. On growth rates of closed permutation classes. *Electron. J. Combin.*, 9(2):Research paper 10, 20 pp. (electronic), 2003.

- [41] André E. Kézdy, Hunter S. Snevily, and Chi Wang. Partitioning permutations into increasing and decreasing subsequences. *J. Combin. Theory Ser. A*, 73(2):353–359, 1996.
- [42] Martin Klazar. The Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture. In *Formal power series and algebraic combinatorics (Moscow, 2000)*, pages 250–255. Springer, Berlin, 2000.
- [43] Martin Klazar and Adam Marcus. Extensions of the linear bound in the Füredi-Hajnal conjecture, 2005.
- [44] D. Kleitman and L. Pachter. Finding convex sets among points in the plane. *Discrete Comput. Geom.*, 19(3, Special Issue):405–410, 1998. Dedicated to the memory of Paul Erdős.
- [45] Donald E. Knuth. *The art of computer programming. Vol. 1: Fundamental algorithms*. Addison-Wesley Publishing Co., Reading, Mass., 1969.
- [46] Donald E. Knuth. *The art of computer programming. Volume 3*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1973. Sorting and searching, Addison-Wesley Series in Computer Science and Information Processing.
- [47] Joseph B. Kruskal, Jr. Monotonic subsequences. *Proc. Amer. Math. Soc.*, 4:264–274, 1953.
- [48] Louise Walker M. D. Atkinson and S. A. Linton. Priority queues and multisets. *Electron. J. Combin.*, 2:Paper R24, (electronic), 1995.
- [49] Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Combin. Theory Ser. A*, 107(1):153–160, 2004.
- [50] Maximillian M. Murphy. *Restricted permutations, antichains, atomic classes, and stack sorting*. PhD thesis, Univ. of St Andrews, 2002.

- [51] Maximillian M. Murphy and Vincent Vatter. Profile classes and partial well-order for permutations. *Electron. J. Combin.*, 9(2):Research paper 17, 30 pp. (electronic), 2003.
- [52] C. St. J. A. Nash-Williams. On well-quasi-ordering finite trees. *Proc. Cambridge Philos. Soc.*, 59:833–835, 1963.
- [53] Vaughan R. Pratt. Computing permutations with double-ended queues, parallel stacks and parallel queues. In *STOC '73: Proceedings of the fifth annual ACM symposium on Theory of computing*, pages 268–277, New York, NY, USA, 1973. ACM Press.
- [54] N. Ruškuc. Decidability questions for pattern avoidance classes of permutations, 2005.
- [55] David Singmaster. An elementary evaluation of the Catalan numbers. *Amer. Math. Monthly*, 85(5):366–368, 1978.
- [56] Rebecca Smith. Permutation reconstruction. *Electron. J. Combin.*, 13:Note 11, (electronic), 2006.
- [57] Zvezdelina E. Stankova. Forbidden subsequences. *Discrete Math.*, 132(1-3):291–316, 1994.
- [58] Richard P. Stanley. *Enumerative combinatorics. Vol. 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1997.
- [59] Robert Tarjan. Sorting using networks of queues and stacks. *J. Assoc. Comput. Mach.*, 19:341–346, 1972.
- [60] G. Tóth and P. Valtr. Note on the Erdős-Szekeres theorem. *Discrete Comput. Geom.*, 19(3, Special Issue):457–459, 1998. Dedicated to the memory of Paul Erdős.
- [61] Herbert S. Wilf. What is an answer? *Amer. Math. Monthly*, 89(5):289–292, 1982.

- [62] Herbert S. Wilf. *generatingfunctionology*. A K Peters Ltd., Wellesley, MA, third edition, 2006.