

Open Research Online

The Open University's repository of research publications and other research outputs

On The Growth Of Permutation Classes

Thesis

How to cite:

Bevan, David Ian (2015). On The Growth Of Permutation Classes. PhD thesis The Open University.

For guidance on citations see [FAQs](#).

© 2015 The Author

Version: Version of Record

Link(s) to article on publisher's website:
<http://dx.doi.org/doi:10.21954/ou.ro.0000ab63>

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

ON THE GROWTH OF PERMUTATION CLASSES

DAVID IAN BEVAN

M.A. M.Sc. B.A.

a thesis submitted to
THE OPEN UNIVERSITY
for the degree of
DOCTOR OF PHILOSOPHY
in Mathematics

May 2015

Abstract

We study aspects of the enumeration of permutation classes, sets of permutations closed downwards under the subpermutation order.

First, we consider monotone grid classes of permutations. We present procedures for calculating the generating function of any class whose matrix has dimensions $m \times 1$ for some m , and of acyclic and unicyclic classes of gridded permutations. We show that almost all large permutations in a grid class have the same shape, and determine this limit shape.

We prove that the growth rate of a grid class is given by the square of the spectral radius of an associated graph and deduce some facts relating to the set of grid class growth rates. In the process, we establish a new result concerning tours on graphs. We also prove a similar result relating the growth rate of a geometric grid class to the matching polynomial of a graph, and determine the effect of edge subdivision on the matching polynomial. We characterise the growth rates of geometric grid classes in terms of the spectral radii of trees.

We then investigate the set of growth rates of permutation classes and establish a new upper bound on the value above which every real number is the growth rate of some permutation class. In the process, we prove new results concerning expansions of real numbers in non-integer bases in which the digits are drawn from sets of allowed values.

Finally, we introduce a new enumeration technique, based on associating a graph with each permutation, and determine the generating functions for some previously unenumerated classes. We conclude by using this approach to provide an improved lower bound on the growth rate of the class of permutations avoiding the pattern **1324**. In the process, we prove that, asymptotically, patterns in Łukasiewicz paths exhibit a concentrated Gaussian distribution.

Acknowledgements

*Jésus-Christ est l'objet de tout, et le centre où tout tend.
Qui le connaît connaît la raison de toutes choses.*¹

— Pascal, *Pensées* [135]

My journey (back) into mathematics research began with reading *Proofs from THE BOOK*, the delightful book by Aigner & Ziegler [1]. I am grateful to Günter Ziegler for his encouragement at that time, which resulted in my first published mathematics paper [35].

Lunchtimes at the Open University have been enlivened by numerous interesting stories from Phil Rippon. I am thankful to both him and Gwyneth Stallard for much helpful advice, and also for a couple of hundred cups of coffee! Thanks too, to other members of the department for friendship and support, including Robert, Toby, Ian, Ben, Tim and Jozef, and the other research students: Grahame, Rob, Matthew, Mairi, Rosie, Vasso and David.

I received a warm welcome and encouragement from many members of the permutation patterns community and other combinatorialists. Particular thanks are due to Vince Vatter for detailed feedback on some of my papers and for suggesting that it might be worthwhile investigating whether his upper bound on λ could be improved. I am also grateful to Mireille Bousquet-Mélou for valuable discussions relating to my (at that time, very woolly) ideas about using Hasse graphs for enumeration, during the Cardiff Workshop on Combinatorial Physics. Thanks are also due to Mike Atkinson, Michael Albert, Einar Steingrímsson, Jay Pantone, Dan Král', and Doron Zeilberger, among others.

Robert Brignall has been an excellent supervisor, guiding me with a light touch, and doing a first-rate job of teaching me how to write mathematics so that it can be understood by those trying to read it. Thank you!

¹Jesus Christ is the goal of everything, and the centre to which everything tends. Whoever knows Him knows the reason behind all things.

Michael Albert's *PermLab* software [6] was of particular benefit in helping to visualise and explore the structure of permutations in different permutation classes. I have also made extensive use of *Mathematica* [172] throughout my research.


Three years ago, my wife Penny encouraged me to pursue my dream, and has accompanied me on this latest adventure in our life together, bearing with good grace the (frequent) times when I'm in "maths world" and present only physically. She has also willingly taken on the primary bread-winning role during this season. Words are insufficient to express my gratitude.

Above all else, like Blaise Pascal, I am convinced that the crux to understanding life, the universe and everything — including the ontology, epistemology and "unreasonable effectiveness"² of mathematics — lies in the life, death and resurrection of a certain itinerant Jewish rabbi and miracle-worker. I thus conclude:

*To the only wise God
be glory forever
through Jesus Christ!*³

²See Wigner [168]

³Romans 16:27



Permutation class art

The permutations of length 7 in $\text{Av}(21\bar{3}54, 4\bar{5}\bar{3}12)$
that are both sum and skew indecomposable

Contents

Abstract	iii
Acknowledgements	v
1 Enumerating permutation classes	1
1.1 Enumeration	1
1.2 Classes of permutations	11
1.3 Historical overview	14
1.4 Outline and list of main results	19
I GRID CLASSES	23
2 Introducing grid classes	25
2.1 Grid classes and griddings	25
2.2 Outline of Part I	27
3 Skinny grid classes	29
4 Gridded permutations	37
4.1 Acyclic classes	40
4.2 Unicyclic classes	42
4.3 Exploitation	43
5 Growth rates of grid classes	47
5.1 Introduction	47
5.2 Tours on graphs	48
5.3 Grid classes	65
5.4 Implications	68

6	Grid class limit shapes	75
7	Geometric grid classes	81
7.1	Introduction	81
7.2	Proof of growth rate theorem	83
7.3	Consequences	87
II	THE SET OF GROWTH RATES	97
8	Intervals of growth rates	99
8.1	Introduction	99
8.2	Non-integer bases and generalised digits	101
8.3	Sum-closed permutation classes	105
8.4	Constructions that yield intervals of growth rates	108
III	HASSE GRAPHS	117
9	Introducing Hasse graphs	119
9.1	Building Hasse graphs	119
9.2	Avoiding 1324	121
9.3	Outline of Part III	122
10	Avoiding 1324 and 2314	125
11	Avoiding 1324 and 1432	129
12	Forest-like permutations	135
13	Plane permutations	139
14	Avoiding 1234 and 2341	145
15	Avoiding 1243 and 2314	155

16 Avoiding 1324	161
16.1 Introduction	161
16.2 Permutations with a regular structure	165
16.3 Concentration of distributions	171
16.4 Patterns in Łukasiewicz paths	177
16.5 Summing up	180

Chapter 1

Enumerating permutation classes

Doron Zeilberger begins his expository article [177] on enumerative and algebraic combinatorics in *The Princeton Companion to Mathematics* by observing that

“Enumeration, otherwise known as counting, is the oldest mathematical subject”.

Subsequently, he declares,

“The fundamental theorem of enumeration, independently discovered by several anonymous cave dwellers, states that

$$|A| = \sum_{\alpha \in A} 1. \tag{1}$$

... While this formula is still useful after all these years, enumerating specific finite sets is no longer considered mathematics.”

In this study, we participate in the millennia-old pursuit of counting, and contribute to the growth of mathematical knowledge, by addressing the question of the enumeration of certain infinite sets, called permutation classes.

In the rest of this introductory section, we first consider what it means to enumerate a class of combinatorial objects and describe the techniques we use. Then we introduce permutation classes and give an overview of research concerning their enumeration. Finally, we present an outline of the contents of the three parts of this thesis.

1.1 Enumeration

A *combinatorial class* \mathcal{A} is a countable (i.e. finite or countably infinite) set endowed with a *size* function, such that

- the size of each element of \mathcal{A} is a non-negative integer, and
- the number of elements of \mathcal{A} of any given size is finite.

The size of an element $\alpha \in \mathcal{A}$ is denoted by $|\alpha|$. The finite set of elements in \mathcal{A} whose size is n is written \mathcal{A}_n . We consistently use calligraphic uppercase letters, e.g. $\mathcal{R}, \mathcal{S}, \mathcal{T}$, for combinatorial classes and lowercase Greek letters, e.g. ρ, σ, τ , for their members.

Given some combinatorial class, or family of combinatorial classes, the goal of enumerative combinatorics is to determine the number of elements of each size in the class or classes. As discussed by Wilf [170], there are a number of possible ways of answering the question, “How many things are there?”. While claiming that (1), applied to each \mathcal{A}_n , provides a simple formula that “answers” all such questions at once, Wilf rejects such an answer since it is just a restatement of the question.

Generating functions

For the most part, the answers we give make use of *generatingfunctionology*. According to Wilf, who coined this neologism as the title of his classic book [171],

“A generating function is a clothesline on which we hang up a sequence of numbers for display.”

The (ordinary univariate) *generating function* for a combinatorial class \mathcal{A} is defined to be the formal power series

$$A(z) = \sum_{n \geq 0} |\mathcal{A}_n| z^n = \sum_{\alpha \in \mathcal{A}} z^{|\alpha|}.$$

Thus, each element $\alpha \in \mathcal{A}$ makes a contribution of $z^{|\alpha|}$, the result being that, for each n , the coefficient of z^n is the number of elements of size n . In generating functions, we consistently use the variable z to mark the size of objects.

As an elementary example, the number of distinct ways of tossing a coin n times is clearly 2^n , hence the generating function for sequences of coin tosses, in which the size of a sequence is the number of tosses, is

$$S(z) = 2z + 4z^2 + 8z^3 + \dots = \frac{2z}{1 - 2z}.$$

Given the generating function $F(z)$ for a combinatorial class, we use the notation $[z^n]F(z)$ to denote the operation of extracting the coefficient of z^n (i.e. the number of elements of the class with size n) from the formal power series $F(z)$. Thus, in our example,

$$[z^n]S(z) = 2^n, \quad \text{if } n \geq 1.$$

In addition to recording the size of objects, it is often useful to keep track of additional parameters in a *multivariate* generating function. For example, if $\omega : \mathcal{A} \rightarrow \mathbb{N}_0$ is a parameter of elements of combinatorial class \mathcal{A} , then the (ordinary) *bivariate* generating function for \mathcal{A} , in which w marks the parameter ω , is:

$$A(z, w) = \sum_{\alpha \in \mathcal{A}} z^{|\alpha|} w^{\omega(\alpha)} = \sum_{n \geq 0} \sum_{k \geq 0} a_{n,k} z^n w^k,$$

where $a_{n,k}$ is the number of elements $\alpha \in \mathcal{A}$ of size n for which $\omega(\alpha) = k$. This can be generalised for multiple parameters. Observe that $A(z, 1)$ is the ordinary univariate generating function.

For example, if we use t to mark the number of tails in a sequence of coin tosses, the corresponding bivariate generating function is

$$S(z, t) = \sum_{n=1}^{\infty} \sum_{k=0}^n \binom{n}{k} z^n t^k = \frac{(1+t)z}{1-(1+t)z}.$$

Why do we choose to use generating functions? Their extraordinary utility is beautifully elucidated in the consummate *magnum opus* of Flajolet & Sedgewick, *Analytic Combinatorics* [77]. Three aspects are pre-eminent in this work:

Firstly, it is possible to translate *directly* from a structural definition of a combinatorial class to equations that define the generating function for the class. It may then be possible to solve these functional equations to yield an explicit form for the generating function.

Secondly, generating functions enable us to answer questions concerning the *growth* of combinatorial classes. Given a generating function, $F(z) = \sum a_n z^n$, the asymptotic behaviour of the sequence (a_n) can be determined by treating $F(z)$ analytically as a complex function. The singularities of F provide full information on the asymptotic behaviour of its coefficients. Typically, for large n , $[z^n]F(z)$ behaves like $\gamma^n \theta(n)$, for some $\gamma > 0$ and subexponential function θ . The location of the singularities of F dictates the exponential growth rate, γ , of its coefficients. The nature of the singularities of F then determines the subexponential factor, $\theta(n)$.

Thirdly, it is possible to extract from multivariate generating functions precise information concerning the limiting distribution of parameter values. Thus, if ω is some parameter of elements of a combinatorial class, the asymptotic probability distribution $\mathbb{P}_n[\omega(\alpha) \leq k]$ for large objects of size n can be established.

A further benefit of possessing a generating function for a combinatorial class is that the nature of the generating function immediately reveals something of the structure of




Figure 1.1: The hierarchy of families of generating functions

objects in the class. Specifically, the extent to which a class is “well-behaved” depends on whether its generating function is rational, algebraic, *D*-finite, or not *D*-finite.

A *rational* function is the ratio of two polynomials, such as $S(z)$ above. Rational functions are the generating functions of deterministic finite state automata or, equivalently, of *regular* languages (see [77, Section I.4]). Objects in a class with a rational generating function therefore tend to exhibit a structure similar to the linear structure of words in a regular language. A function $F(z) = \sum f_n z^n$ is rational if and only if its coefficients satisfy a linear recurrence relation with constant coefficients.

A larger family than the rational functions is the family of *algebraic* functions. A function $F(z)$ is algebraic if it can be defined as the root of a polynomial equation. That is, there exists a bivariate polynomial $P(z, y)$ such that $P(z, F(z)) = 0$. Algebraic functions are the generating functions of unambiguous *context-free* languages (Chomsky & Schützenberger [60]). Objects in a class with an algebraic generating function tend to exhibit a branching structure similar to that of trees (see Bousquet-Mélou [53]). There is no simple characterization of the coefficients of algebraic generating functions. However, the coefficients of an algebraic function can be expressed in closed form as a finite linear combination of multinomial coefficients, which can be determined directly from the defining polynomial $P(z, y)$. This result is due to Flajolet & Soria; a proof can be found in [76, Theorem 8.10]; see also the presentation by Banderier & Drmota [31, Theorem 1], and [77, Note VII.34, p.495].

A more general family than that of rational or algebraic functions is the family of *D*-finite (“differentiably finite”) functions, also known as *holonomic* functions. A function $F(z)$ is *D*-finite if it satisfies a differential equation with coefficients that are polynomials in z . A power series is *D*-finite if and only if its coefficients satisfy a linear recurrence relation with polynomial coefficients. A sequence of numbers satisfying such a recurrence is said to be *P*-recursive (“polynomially recursive”).

Beyond D-finite power series lie those that are not D-finite. Given that there are only countably many D-finite functions, almost all of the uncountably many combinatorial classes are enumerated by non-D-finite generating functions. Such functions are considerably less amenable to analysis. One way of determining whether a generating function is D-finite or not is to make use of the fact that a D-finite function has only finitely many singularities. For more on ways of determining D-finiteness, see the papers by Guttman [93] and Flajolet, Gerhold & Salvy [75].

The symbolic method

There is a natural direct correspondence between the structure of combinatorial classes and their generating functions, as reflected in Table 1.1.

<i>Structure</i>	<i>Construction</i>	<i>Generating function</i>	<i>Condition</i>
Atom	$\mathcal{Z} = \{\bullet\}$	$Z(z) = z$	
Disjoint union	$\mathcal{A} = \mathcal{B} + \mathcal{C}$	$A(z) = B(z) + C(z)$	
Cartesian product	$\mathcal{A} = \mathcal{B} \times \mathcal{C}$	$A(z) = B(z)C(z)$	
Sequence	$\mathcal{A} = \text{SEQ}[\mathcal{B}]$	$A(z) = \frac{1}{1 - B(z)}$	$B(0) = 0$
Non-empty sequence	$\mathcal{A} = \text{SEQ}^+[\mathcal{B}]$	$A(z) = \frac{B(z)}{1 - B(z)}$	$B(0) = 0$
Pointing	$\mathcal{A} = \Theta\mathcal{B}$	$A(z) = z\partial_z B(z)$	
Marking	$\mathcal{A} = \mathfrak{u}\mathcal{B}$	$A(z) = \mathfrak{u}B(z)$	

Table 1.1: The correspondence between structure and generating functions

We always use \mathcal{Z} to denote the *atomic class* consisting of a single element of size 1.

We use $\mathcal{B} + \mathcal{C}$ or $\mathcal{B} \uplus \mathcal{C}$ to denote the *disjoint union* of classes \mathcal{B} and \mathcal{C} .

The *Cartesian product* of two classes contains all ordered pairs of elements of the classes. For example, suppose elements of class \mathcal{A} consist of ordered pairs of objects, $\alpha = (\beta, \gamma)$, where β and γ drawn from classes \mathcal{B} and \mathcal{C} respectively, and the size of α is given by $|\alpha| = |\beta| + |\gamma|$. Then the generating function for \mathcal{A} is given by

$$A(z) = \sum_{\alpha \in \mathcal{B} \times \mathcal{C}} z^{|\alpha|} = \sum_{\beta \in \mathcal{B}} \sum_{\gamma \in \mathcal{C}} z^{|\beta| + |\gamma|} = \sum_{\beta \in \mathcal{B}} z^{|\beta|} \sum_{\gamma \in \mathcal{C}} z^{|\gamma|} = B(z)C(z),$$

where $B(z)$ and $C(z)$ are the generating functions for \mathcal{B} and \mathcal{C} respectively.

We use $\text{SEQ}[\mathcal{B}]$ to represent the class of (possibly empty) sequences of elements of \mathcal{B} ,

and $\text{SEQ}^+[\mathcal{B}]$ to represent the class of non-empty sequences of elements of \mathcal{B} . The size of such a sequence is the sum of the sizes of its components. For example, if \mathcal{A} consists of non-empty sequences of elements of \mathcal{B} , then

$$\mathcal{A} = \text{SEQ}^+[\mathcal{B}] = \mathcal{B} + \mathcal{B}^2 + \mathcal{B}^3 + \dots,$$

where we write \mathcal{B}^2 for $\mathcal{B} \times \mathcal{B}$, etc. Thus,

$$A(z) = B(z) + B(z)^2 + B(z)^3 + \dots = \frac{B(z)}{1 - B(z)}.$$

The *pointing* construction, denoted here by Θ , represents the idea of “pointing at a distinguished atom”. For example, if $\mathcal{A} = \Theta\mathcal{B}$, then the generating function of \mathcal{A} is

$$A(z) = \sum_{n \geq 0} a_n z^n = \sum_{n \geq 0} n b_n z^n = z \partial_z B(z),$$

where ∂_z is used to denote the differential operator $\frac{d}{dz}$.

Marking enables us to record combinatorial parameters. We use a lowercase letter for marking in structural equations, corresponding to the variable in the generating function that is used to mark the parameter. For example, the class of sequences of coin tosses, with t marking tails has the following structure:

$$\mathcal{S} = \text{SEQ}^+[\mathcal{Z} + t\mathcal{Z}],$$

where the first term in the disjoint union corresponds to a throw of heads, and the second to a (marked) throw of tails.

Functional equations

Typically, to determine the generating function of a combinatorial class, we define the structure recursively. The prototypical example is that of *rooted plane trees*, where the size is the number of vertices.

Rooted plane trees consist of a root vertex joined to a (possibly empty) sequence of subtrees. Thus the class, \mathcal{T} , satisfies the recursive structural equation

$$\mathcal{T} = \mathcal{Z} \times \text{SEQ}[\mathcal{T}].$$

So, by the correspondence between structure and generating functions, we know that the




Figure 1.2: A rooted plane tree

generating function for rooted plane trees, $T(z)$, satisfies the equation

$$T(z) = z \frac{1}{1 - T(z)}.$$

This equation has two roots, but one of them has negative coefficients and so can be rejected. Hence,

$$T(z) = \frac{1}{2}(1 - \sqrt{1 - 4z}).$$

Thus, extracting coefficients, $T(z) = z + z^2 + 2z^3 + 5z^4 + 14z^5 + 42z^6 + \dots$, rooted plane trees being enumerated by the Catalan numbers.

Often, in order to derive the univariate generating function for a combinatorial class, multivariate functions are used, involving additional “catalytic” variables that record certain parameters of the objects in the class. These additional variables make it possible to establish functional equations, which can sometimes be solved to yield the required generating function. Typically, when employing a multivariate generating function, it is common to treat it simply as a function of the relevant catalytic variable, writing, for example, $F(u)$ rather than $F(z, u)$.

Another common technique is the use of *linear operators* on generating functions, for which we use the symbol Ω . Let us illustrate this by briefly considering the class of *stepped parallelogram polyominoes*, which we denote \mathcal{P} .




Figure 1.3: A stepped parallelogram polyomino

A stepped parallelogram polyomino is constructed from one or more rows, each consisting of a contiguous sequence of cells. Except for the bottom row, the leftmost cell of a row must occur to the right of the leftmost cell of the previous row but not to the right of the rightmost cell of the previous row, and the rightmost cell of a row must occur to the right

of the rightmost cell of the previous row. See Figure 1.3 for an example.

For our illustration, we consider the size of a polyomino to be given by its *width*. Let $P(u) = P(z, u)$ be the bivariate generating function for stepped parallelogram polyominoes, in which z marks the width and u marks the number of cells in the top row. Thus a polyomino that has k cells in its top row contributes a term to $P(u)$ in which u has exponent k . The generating function for elements of \mathcal{P} consisting of a single row is

$$P_1(u) = \text{SEQ}^+[uZ] = \frac{zu}{1-zu},$$

where, by mild abuse of notation, we identify the structure and the generating function.

The addition of a new row on top of a row with length k is reflected by the linear operator Ω_R defined by

$$\Omega_R[u^k] = (u^{k-1} + u^{k-2} + \dots + u)\text{SEQ}^+[uZ] = \frac{zu}{(1-u)(1-zu)}(u - u^k),$$

in which $\text{SEQ}^+[uZ]$ corresponds to the “overhanging” cells to the right.

Hence, the bivariate generating function $P(u)$ for stepped parallelogram polyominoes is defined by the recursive functional equation $P(u) = P_1(u) + \Omega_R[P(u)]$. That is,

$$P(u) = \frac{zu}{1-zu} + \frac{zu}{(1-u)(1-zu)}(uP(1) - P(u)), \quad (2)$$

an equation that relates $P(u)$ to $P(1)$.

The kernel method

Equations such as (2) can sometimes be solved by making use of what is known as the *kernel method*. For examples of its use, see the paper by Banderier, Bousquet-Mélou, Denise, Flajolet, Gardy & Gouyou-Beauchamps [30] and the expository article by Prodinger [139]. We illustrate how the kernel method works by deriving the generating function for the class of stepped parallelogram polyominoes.

To start, we express $P(u)$ in terms of $P(1)$, by expanding and rearranging (2) to give

$$P(u) = \frac{zu(1-u+uP(1))}{1-u+zu^2}.$$

Equivalently, we have the equation

$$(1-u+zu^2)P(u) = zu(1-u+uP(1)).$$

Now, if we set u to be a root of the multiplier of $P(u)$ on the left, then we obtain a linear equation for $P(1)$. This is known as “cancelling the kernel” (the multiplier being the kernel). The appropriate root to use can be identified from the combinatorial requirement that the series expansion of $P(1)$ contains no negative exponents and has only non-negative coefficients.

In this case, the correct root is $u = (1 - \sqrt{1 - 4z})/2z$, which yields the univariate generating function for \mathcal{P} ,

$$P(1) = \frac{1}{2}(1 - \sqrt{1 - 4z}).$$

This turns out to be the same as that for rooted plane trees, since stepped parallelogram polyominoes, counted by width, are also enumerated by the Catalan numbers (see Stanley [153]).

Analytic combinatorics

It is a remarkable fact that the asymptotic behaviour of the coefficients of a generating function $F(z)$ can be determined by considering the analytic properties of F considered as a complex function, that is, as a mapping of the complex plane to itself. We briefly present here the key facts.

We are often interested in determining how the number of objects in a combinatorial class grows with size. A fundamental quantity of interest is the exponential *growth rate*. The growth rate of a class \mathcal{C} is defined to be the limit

$$\text{gr}(\mathcal{C}) = \lim_{n \rightarrow \infty} |\mathcal{C}_n|^{1/n},$$

if it exists. We define the *upper* and *lower* growth rates similarly:

$$\overline{\text{gr}}(\mathcal{C}) = \limsup_{n \rightarrow \infty} |\mathcal{C}_n|^{1/n} \quad \underline{\text{gr}}(\mathcal{C}) = \liminf_{n \rightarrow \infty} |\mathcal{C}_n|^{1/n}.$$

Observe that combinatorial classes whose enumeration differs only by a polynomial factor have the same (upper/lower) growth rate. This fact follows directly from the definition of the growth rate.

Fundamental to determining exponential growth rates is Pringsheim’s Theorem. This result is concerned with the location of the singularities of functions with non-negative coefficients. Such functions include all enumerative generating functions.

Lemma 1.1 (Pringsheim [77, Theorem IV.6]). *If the power series for $F(z)$ has non-negative coefficients and radius of convergence ρ , then $z = \rho \in \mathbb{R}^+$ is a singularity of $F(z)$.*

The least singularity on the positive real axis is known as the *dominant singularity*. It is all that is required to determine the (upper) growth rate of a combinatorial class, as the following lemma reveals.

Lemma 1.2 (Exponential Growth Formula [77, Theorem IV.7]). *If $F(z)$ is analytic at 0 with non-negative coefficients and ρ is its dominant singularity, then*

$$\limsup_{n \rightarrow \infty} ([z^n]F(z))^{1/n} = \rho^{-1}.$$

Thus, the upper growth rate of a combinatorial class is equal to the reciprocal of the dominant singularity of its generating function.

Analytic combinatorics can give us more information than this. The *nature* of the dominant singularity prescribes the subexponential factor. To state the relevant results, we use the notation $F(n) \sim G(n)$ to denote the fact that asymptotically $F(n)$ and $G(n)$ are “approximately equal”. Formally,

$$F(n) \sim G(n) \quad \text{if} \quad \lim_{n \rightarrow \infty} \frac{F(n)}{G(n)} = 1.$$

For meromorphic functions (i.e., functions whose singularities are poles), we have the following:

Lemma 1.3 (see [77, Theorems IV.10 and VI.1]). *When the dominant singularity ρ of $F(z)$ is a pole of order r , then*

$$[z^n]F(z) \sim c\rho^{-n}n^{r-1} \quad \text{where} \quad c = \frac{\rho^{-r}}{(r-1)!} \lim_{z \rightarrow \rho} (\rho - z)^r F(z).$$

More generally, when the dominant singularity is not a pole, the following result can be employed:

Lemma 1.4 (see [77, Theorem VI.1]).

$$[z^n](1 - z/\rho)^{-\alpha} \sim \frac{\rho^{-n}n^{\alpha-1}}{\Gamma(\alpha)}, \quad \alpha \neq 0, -1, -2, \dots,$$

where $\Gamma(k) = (k-1)!$ and $\Gamma(k + \frac{1}{2}) = \frac{(2k)!}{4^k k!} \sqrt{\pi}$, if $k \in \mathbb{N}$.

This concludes our brief exposition of combinatorial enumeration. As mentioned above, it is also possible to extract distributional information about parameter values from generating functions. We make use of this in Chapter 16. The relevant results are presented there.

1.2 Classes of permutations

The combinatorial classes that we study in this thesis are classes of permutations. We begin by presenting some standard definitions, before giving a brief historical overview of research addressing enumerative questions concerning permutation classes.

Permutations

We consider a permutation to be simply an arrangement of the numbers $1, 2, \dots, k$ for some positive k . We use $|\sigma|$ to denote the length of permutation σ . It can be helpful to consider permutations graphically. If $\sigma = \sigma_1, \dots, \sigma_k$ is a permutation, its *plot* consists of the the points (i, σ_i) in the Euclidean plane, for $i = 1, \dots, k$. Often a permutation is identified with its plot.




Figure 1.4: The plot of permutation **31567482** with a **1324** subpermutation marked

A permutation τ is said to be *contained* in, or to be a *subpermutation* of, another permutation σ if σ has a subsequence whose terms are order isomorphic to (i.e. have the same relative ordering as) τ . From the graphical perspective, σ contains τ if the plot of τ results from erasing some points from the plot of σ and then “shrinking” the axes appropriately. We write $\tau \leq \sigma$ if τ is a subpermutation of σ . For example, **31567482** contains **1324** because the subsequence **3648** (among others) is ordered in the same way as **1324** (see Figure 1.4). If σ does not contain τ , we say that σ *avoids* τ . For example, **31567482** avoids **1243** since it has no subsequence ordered in the same way as **1243**. In the context of containment and avoidance, a permutation is often called a *pattern*.




Figure 1.5: The three left-to-right maxima and four right-to-left minima, and the two left-to-right minima and five right-to-left maxima, of a permutation

Sometimes we want to refer to the extremal points in a permutation. A value in a permutation is called a *left-to-right maximum* if it is larger than all the values to its left. *Left-to-right minima*, *right-to-left maxima* and *right-to-left minima* are defined analogously. See Figure 1.5 for an illustration.




Figure 1.6: The direct sum $2413 \oplus 4231$, the skew sum $2413 \ominus 4231$, and the layered permutation $21 \oplus 1 \oplus 321 \oplus 21$

Given two permutations σ and τ with lengths k and ℓ respectively, their *direct sum* $\sigma \oplus \tau$ is the permutation of length $k + \ell$ consisting of σ followed by a shifted copy of τ :

$$(\sigma \oplus \tau)(i) = \begin{cases} \sigma(i) & \text{if } i \leq k, \\ k + \tau(i - k) & \text{if } k + 1 \leq i \leq k + \ell. \end{cases}$$

The *skew sum* $\sigma \ominus \tau$ is defined analogously. See Figure 1.6 for an illustration.

A permutation is called *sum indecomposable* or just *indecomposable* if it cannot be expressed as the direct sum of two shorter permutations. For brevity, we call an indecomposable permutation simply an *indecomposable*. A permutation is *skew indecomposable* if it cannot be expressed as the skew sum of two shorter permutations. Note that every permutation has a unique representation as the direct sum of a sequence of indecomposables (and also as the skew sum of a sequence of skew indecomposables). If each indecomposable in a permutation is a decreasing sequence, then we say that the permutation is *layered*. See Figure 1.6 for an example.




Figure 1.7: A permutation and its ordered inversion graph

Sometimes it helps to consider permutations in a slightly broader context. A permutation can be considered to be a particular type of *ordered graph*. An ordered graph is a graph with a linear order on its vertices (it is natural to number the vertices from 1 to n). To each permutation we associate an ordered graph. The (ordered) (inversion) *graph* of a

permutation σ of length n has vertex set $\{1, \dots, n\}$ with an edge between vertices i and j if $i < j$ and $\sigma(i) > \sigma(j)$. A pair of terms in a permutation with this property is called an *inversion*. Thus, the graph of a permutation contains one edge for each inversion in the permutation. Note that a permutation graph is transitively closed. See Figure 1.7 for an illustration. We use G_σ to denote the graph of σ . An *induced ordered subgraph* of an ordered graph G is an induced subgraph of G that inherits its vertex ordering. It is easy to see that $\tau \leq \sigma$ if and only if G_τ is an induced ordered subgraph of G_σ .

Permutation classes

The subpermutation relation is a partial order on the set of all permutations. A classical *permutation class*, sometimes called a *pattern class*, is a set of permutations closed downwards (a down-set) under this partial order. Thus, if σ is a member of a permutation class \mathcal{C} and τ is contained in σ , then it must be the case that τ is also a member of \mathcal{C} . From a graphical perspective, this means that erasing points from the plot of a permutation in \mathcal{C} always results in the plot of another permutation in \mathcal{C} when the axes are rescaled appropriately.

It is common in the study of classical permutation classes to reserve the word “class” for sets of permutations closed downwards under containment, and to use the mathematical synonyms “set”, “collection” and “family” for other combinatorial classes. We do not follow this convention rigorously in this thesis.

It is natural to define a permutation class “negatively” by stating the minimal set of permutations that it avoids. This minimal forbidden set of patterns is known as the *basis* of the class. The class with basis B is denoted $\text{Av}(B)$, and we use $\text{Av}_n(B)$ for the permutations of length n in $\text{Av}(B)$. As a trivial example, $\text{Av}(\mathbf{21})$ is the class of increasing permutations (i.e. the identity permutation of each length). As another simple example, the **123**-avoiders, $\text{Av}(\mathbf{123})$, consists of those permutations that can be partitioned into two decreasing subsequences. The basis of a permutation class is an antichain (a set of pairwise incomparable elements) under the containment order, and may be infinite. Classes for which the basis consists of a single permutation are called *principal* classes.

Suppose that \mathcal{S} is a set of indecomposables that is downward closed under the subpermutation order (so if $\sigma \in \mathcal{S}$, $\tau \leq \sigma$ and τ is indecomposable, then $\tau \in \mathcal{S}$). The *sum closure* of \mathcal{S} , denoted $\oplus\mathcal{S}$, is then the class of permutations of the form $\sigma_1 \oplus \sigma_2 \oplus \dots \oplus \sigma_r$, where each $\sigma_i \in \mathcal{S}$. It is simple to check that $\oplus\mathcal{S}$ is, indeed, a permutation class. Such a class is *sum-closed*: if $\sigma, \tau \in \oplus\mathcal{S}$ then $\sigma \oplus \tau \in \oplus\mathcal{S}$. Furthermore, every sum-closed class is the sum closure of its set of indecomposables. It can easily be seen that a class is sum-closed if and only if all its basis elements are indecomposable.

In the context of graphs, a set closed under taking induced subgraphs is known as a *hereditary class*. It is an easy exercise using mathematical induction to prove that the graph of a permutation is an ordered graph that avoids the two induced ordered subgraphs $\bullet\bullet\bullet$ and $\bullet\overset{\curvearrowright}{\bullet}\bullet$. Each permutation class is thus isomorphic to a hereditary class of ordered graphs, and results can be transferred between the two domains.

We are interested in the enumeration of permutation classes. One natural question is to determine whether two classes, \mathcal{C} and \mathcal{D} , are equinumerous, i.e. $|\mathcal{C}_n| = |\mathcal{D}_n|$ for every n . Two classes that are equinumerous are said to be *Wilf-equivalent* and the equivalence classes are called *Wilf classes*. From the graphical perspective, it is clear that classes related by symmetries of the square are Wilf-equivalent. Thus, for example, $\text{Av}(\mathbf{132})$, $\text{Av}(\mathbf{231})$, $\text{Av}(\mathbf{213})$ and $\text{Av}(\mathbf{312})$ are equinumerous. However, not all Wilf-equivalences are a result of these symmetries. Indeed $\text{Av}(\mathbf{123})$ and $\text{Av}(\mathbf{132})$ are Wilf-equivalent.

Classical permutation classes are not the only sets of permutations that are of interest. A specific focus is classes avoiding certain *barred patterns*. A barred pattern is specified by a permutation with some entries barred ($\mathbf{4\bar{3}12}$, for example). For a permutation σ to avoid a barred pattern $\hat{\pi}$, whose underlying permutation is π , every occurrence in σ of the permutation order isomorphic to the non-barred entries in $\hat{\pi}$ ($\mathbf{312}$ in the example) must feature as part of an occurrence of π . Note that the class of permutations avoiding a barred pattern is not normally closed downwards under the subpermutation order. A combinatorial class of permutations that is not closed downwards is called a *non-classical class*.

1.3 Historical overview

We now present a brief, and somewhat selective, overview of the development of research into enumerative aspects of permutation patterns, a subject which now has a voluminous bibliography. For a more detailed survey of the same area, see Vatter's chapter [161] in the forthcoming *Handbook of Enumerative Combinatorics*. The topic is presented as part of a broader picture in the books by Kitaev [109] and Bóna [49]. Other useful sources include the volume of papers from the Permutation Patterns conference held in 2007 [120] and Steingrímsson's survey article produced for the 2013 British Combinatorial Conference [154].

Stacks, queues and dequeues

The study of permutation classes can reasonably be said to have begun with Knuth's investigations in the 1960s into sorting using stacks and dequeues (double-ended queues),

published in the first volume of his encyclopedic monograph, *The Art of Computer Programming* [112]. Knuth observed that a permutation can be sorted by a stack if and only if it does not contain the pattern **231**, and showed that this class is counted by the Catalan numbers. He also proved that the class of permutations sorted by an *input-restricted deque* (i.e. a deque with the push operation restricted to one end) is $\text{Av}(4231, 3241)$ and enumerated the class with what was possibly the first use of the kernel method. Knuth nicely described his approach in terms of railway “switchyard networks”, and this perspective was taken up and developed in subsequent papers by Even & Itai [72], Tarjan [155], and Pratt [138], which investigated networks of stacks, queues and deques.




Figure 1.8: A figure of a switchyard network from Tarjan’s paper [155]

The investigation of stack sorting was continued in the work of West [164, 165]. He considered the class of permutations that could be sorted by passing twice through a stack while requiring the contents of the stack to remain ordered (as in the Tower of Hanoi puzzle). These permutations, which now tend to be referred to as the West-2-stack-sortable permutations, constitute the non-classical class $\text{Av}(2341, 3\bar{5}241)$. West conjectured that this class had the same enumeration as non-separable planar maps. This conjecture was first proved by Zeilberger [175], using a computer to solve a complicated functional equation. Subsequently, Dulucq, Gire & West [70] and Goulden & West [89] found bijective proofs. An alternative approach to sorting with two ordered stacks was subsequently investigated by Atkinson, Murphy & Ruškuc [24], who determined both the infinite basis of the class and its algebraic generating function.

Despite this activity, most problems related to Knuth’s switchyard networks have turned out to be very hard. Fundamental questions remain unanswered, including the enumeration of permutations sortable by two stacks in parallel, the enumeration of permutations sortable by two stacks in series, and the enumeration of permutations sortable by a general deque. Albert, Atkinson & Linton [8] calculated lower and upper bounds on the growth rates of each of these classes. More recently, Albert & Bousquet-Mélou [9], in a paper that is a *tour de force* of analytic combinatorics, gave a pair of functional equations that characterise the generating function of permutations that can be sorted with two parallel stacks. For more on sorting with stacks, queues and deques, see the surveys by Bóna [45] and Atkinson [22].

Conjectures

Much of the research into permutation patterns has been driven by certain conjectures. The first of these was the Stanley–Wilf conjecture that every permutation class \mathcal{C} (excluding the class of all permutations) has a finite upper growth rate $\overline{\text{gr}}(\mathcal{C})$, i.e. that for any permutation σ there exists a constant c_σ such that, for every n , $|\text{Av}_n(\sigma)| \leq c_\sigma^n$. Arratia [18] observed that the Stanley–Wilf conjecture implies that every sum-closed class \mathcal{C} , and hence every principal class, has a growth rate $\text{gr}(\mathcal{C})$. Alon & Friedgut [15] managed to prove a result very close to the conjecture: that for any permutation σ there exists a constant c_σ such that, for every n , $|\text{Av}_n(\sigma)| \leq c_\sigma^{n^\gamma(n)}$, where γ is a function that grows extremely slowly. Bóna [44, 46] established that the conjecture was true for layered patterns. Meanwhile, Klazar [110] showed that the Stanley–Wilf conjecture was implied by a conjecture of Füredi & Hajnal [79] concerning 0-1 matrices. Finally, Marcus & Tardos [128] gave a proof of the Füredi–Hajnal conjecture, thus confirming the Stanley–Wilf conjecture. Thus, every principal class has a growth rate. There are no known examples of permutation classes that do not have a growth rate and it is widely believed that $\text{gr}(\mathcal{C})$ exists for every permutation class \mathcal{C} (see the first conjecture in [161]).

A second conjecture is that of Noonan & Zeilberger [132] that every finitely based permutation class has a D-finite generating function. Clearly, this is not the case for every permutation class since there are uncountably many permutation classes with distinct generating functions, but only countably many D-finite generating functions. This conjecture remains open. However, it is now generally believed to be false, Zeilberger (see [71]) counter-conjecturing that the Noonan–Zeilberger conjecture is false, and, in particular, the sequence $\text{Av}_n(\mathbf{1324})$ is not P-recursive. Recent work of Conway & Guttmann [65] strongly suggests that $\text{Av}(\mathbf{1324})$ does indeed have a non-D-finite generating function.¹

How fast can a permutation class grow? The proof, by Marcus and Tardos, of the Stanley–Wilf conjecture means that every principal class has a growth rate. Marcus and Tardos’ proof yielded a doubly exponential upper bound on $\text{gr}(\text{Av}(\beta))$ in terms of the length of β . Cibulka [63] was able to reduce it to the order of $2^{k \log k}$. Arratia [18] conjectured that the upper bound was much lower and that, for every permutation β of length k , $\text{gr}(\text{Av}(\beta)) \leq (k-1)^2$. Bóna [47] then strengthened this, by conjecturing that equality holds if and only if β is layered. Arratia’s conjecture was subsequently refuted by Albert, Elder, Rechnitzer, Westcott & Zabrocki [4] who showed that $\text{gr}(\text{Av}(\mathbf{1324}))$ exceeded 9.47.

However, evidence suggested that the fastest growing principal classes were those of layered permutations, and Bóna conjectured (see [48]) that the most easily avoided permutation of length k was $1 \oplus 21 \oplus 21 \oplus \dots \oplus 21$ for odd k and $1 \oplus 21 \oplus 21 \oplus \dots \oplus 21 \oplus 1$ for even k .

¹Following submission of this thesis, at the AMS meeting in Washington, DC, in March 2015, Scott Garrabrant announced a proof that the conjecture is false, the result of joint work with Igor Pak [81].

Claesson, Jelínek, and Steingrímsson [64] then proved that for every layered permutation β of length k , the growth rate of $\text{Av}(\beta)$ is less than $4k^2$, and Bóna [50] showed that for his conjectured easiest-to-avoid permutations, the growth rates were at most $9k^2/4$. It thus came as somewhat of a shock when Fox [78] recently proved (by considering the problem in the context of 0-1 matrices) that for almost all permutations β of length k , $\text{gr}(\text{Av}(\beta))$ is, in fact, of the order of 2^k . Therefore, in general, layered permutations are very far from being the easiest to avoid.

Specific classes

Another major strand in permutation class research has concerned the enumeration of classes with a few small basis elements. An up-to-date table of results in this area is recorded on the Wikipedia page [169]. Knuth's investigation of $\text{Av}(\mathbf{231})$ was not the first study of a permutation class. Half a century previously, MacMahon [126] had shown that $\text{Av}(\mathbf{123})$ was counted by the Catalan numbers. Knuth's matching result for $\text{Av}(\mathbf{231})$ thus gave the first example of a Wilf-equivalence not induced by symmetry, and revealed that there was only a single Wilf class for permutations of length 3.

One important class that was considered soon after Knuth's work was the *Baxter permutations*, previously studied by Baxter [34] in connection with an investigation into the behaviour of fixed points of commuting continuous functions. The Baxter permutations constitute the non-classical class $\text{Av}(\mathbf{25\bar{3}14}, \mathbf{41\bar{3}52})$ (see Gire [83]). They were enumerated by Chung, Graham, Hoggatt & Kleiman [62], who introduced *generating trees* as an enumerative mechanism, a technique later to be used more widely. Another significant early enumerative result was the proof by Regev [141] that $\text{gr}(\text{Av}(\mathbf{12\dots k})) = (k-1)^2$ for every k . Gessel [82] later gave an explicit enumeration in terms of determinants.

The first in-depth study of classical permutation classes was by Simion & Schmidt [148] who enumerated classes avoiding two patterns of length 3 and established that there were three Wilf classes. Subsequent results made heavy use of generating trees: West [166] showed that $\text{Av}(\mathbf{3142}, \mathbf{2413})$ is enumerated by the Schröder numbers. West also [167] enumerated classes avoiding one pattern of length 3 and one of length 4, and, in collaboration with Chow [61], those avoiding one pattern of length 3 and an increasing or decreasing sequence. Bóna [42] enumerated $\text{Av}(\mathbf{1342})$ by establishing a bijection between the class and $\beta(0, 1)$ trees.

Various other techniques have been used. Zeilberger [176] introduced *enumeration schemes* for automatic enumeration. These were later employed by Kremer & Shiu [117] to count several classes that avoid pairs of length 4 patterns. Enumeration schemes were further developed by Vatter [157], and extended by Pudwell [140] for use with barred pattern classes, and by Baxter & Pudwell [33] to enumerate classes avoiding *vincular* patterns, a

type of non-classical pattern introduced by Babson & Steingrímsson [26]. An alternative method that has been successful in some contexts is the *insertion encoding* of permutations introduced by Albert, Linton & Ruškuc [13]. Albert, Elder, Rechnitzer, Westcott, & Zabrocki [4] made use of the insertion encoding to determine a lower bound on the growth rate of $\text{Av}(\mathbf{1324})$. This technique was also utilized by Vatter [160] to enumerate two classes avoiding two patterns of length 4.

Another important approach has been the use of *grid classes*, an introduction to which we postpone until Chapter 2. Atkinson [20] determined the generating function for *skew-merged* permutations, which can be partitioned into an increasing sequence and a decreasing sequence, a class that Stankova [151] had previously determined to have the basis $\{\mathbf{2143}, \mathbf{3412}\}$. Atkinson [21] also made use of grid classes to enumerate a number of classes, including $\text{Av}(\mathbf{132}, \mathbf{4321})$. Much more recently, grid classes have been used to enumerate various classes avoiding two patterns of length 4 by Albert, Atkinson & Brignall [2, 3], Pantone [134], and Albert, Atkinson & Vatter [11]. Finally, a paper of Albert & Brignall [12] utilizes grid classes to enumerate $\text{Av}(\mathbf{4231}, \mathbf{35142}, \mathbf{42513}, \mathbf{351624})$, a class which arises in connection with algebraic geometry, specifically the categorization of Schubert varieties.

Parallel to the enumeration of specific classes went work on determining the Wilf classes. The first result of this sort was by West [164], who showed that, for any permutation σ , $\text{Av}(\mathbf{12} \oplus \sigma)$ and $\text{Av}(\mathbf{21} \oplus \sigma)$ are Wilf-equivalent. Hence, in particular, $\text{Av}(\mathbf{1234})$, $\text{Av}(\mathbf{1243})$ and $\text{Av}(\mathbf{2143})$ are in the same Wilf class. This result was subsequently generalised. Babson & West [27] proved the Wilf-equivalence of $\text{Av}(\mathbf{123} \oplus \sigma)$ and $\text{Av}(\mathbf{321} \oplus \sigma)$. Then, Backelin, West & Xin [28] demonstrated that $\text{Av}(\mathbf{12} \dots \mathbf{k} \oplus \sigma)$ and $\text{Av}(\mathbf{k} \dots \mathbf{21} \oplus \sigma)$ were in the same Wilf class for every k . It was also established by Stankova & West [150] that $\text{Av}(\mathbf{231} \oplus \sigma)$ was Wilf-equivalent to $\text{Av}(\mathbf{312} \oplus \sigma)$. In addition, Stankova [151] showed the Wilf-equivalence of $\text{Av}(\mathbf{1342})$ and $\text{Av}(\mathbf{2413})$. As well as these results on principal classes, papers by Bóna [43], Kremer [115, 116], and Le [118] together accomplished the Wilf classification of classes avoiding pairs of patterns of length four.

There is one obvious gap in this record of the enumeration of classes with small bases: $\text{Av}(\mathbf{1324})$. This class has been the *bête noire* of permutation class enumeration. Very little concrete progress has been made on it. We consider the $\mathbf{1324}$ -avoiders in Chapter 16, and postpone a historical overview until there.

Another important strand of research which we have ignored here is the question of determining the structure of the subset of the real line consisting of the growth rates of a permutation classes. We investigate this subject in Chapter 8, where we present the background to our work in this area.

1.4 Outline and list of main results

The rest of this thesis is divided into three parts. In Part **I**, we consider the enumeration of monotone grid classes of permutations, a family of permutation classes defined in terms of the permitted shape for plots of permutations in a class. The most important result in this part is an explicit formula for the growth rate of every permutation grid class. Part **II** is much shorter and concerns the structure of the set of growth rates of permutation classes. We establish a new upper bound on the value above which every real number is the growth rate of some permutation class. Finally, in Part **III**, we introduce a new technique that can be used for the enumeration of permutation classes, based on a graph associated with each permutation, which we call its Hasse graph. As well as using this method to determine the generating function for some previously unenumerated classes, we conclude the thesis by making use of our approach to provide an improved lower bound on the growth rate of $\text{Av}(\mathbf{1324})$.

The chapters in this thesis are of very unequal length, each one addressing a specific enumerative question. The work in Chapter 5 has been published (see [39]), as has that in Chapter 7 (see [37]). The research in Chapters 14 and 15 has been accepted for publication (see [41]), as has that in Chapter 16 (see [40]). The work in Chapter 8 has been submitted for publication (see [38]). Here, for reference, is a list of the primary results in this thesis:

Part I: Grid Classes

- Exact enumeration of skinny grid classes (Theorem 3.1).
- Exact enumeration of acyclic classes of gridded permutations (Theorem 4.3).
- Exact enumeration of unicyclic classes of gridded permutations (Theorem 4.5).
- The generating function of any unicyclic class of gridded permutations is algebraic (Theorem 4.6).
- The generating function of any class of gridded permutations is D-finite (Theorem 4.9).
- The growth rate of the family of balanced tours on a connected graph is the same as that of the family of all tours of even length on the graph (Theorem 5.8).
- The growth rate of a monotone grid class of permutations is equal to the square of the spectral radius of its row-column graph (Theorem 5.14).
- The growth rate of each monotone grid class is an algebraic integer (Corollary 5.15).
- A monotone grid class whose row-column graph is a cycle has growth rate 4 (Corollary 5.17).

- If the growth rate of a monotone grid class is less than 4, it is equal to $4 \cos^2\left(\frac{\pi}{k}\right)$ for some $k \geq 3$ (Corollary 5.19).
- For every $\gamma \geq 2 + \sqrt{5}$ there is a monotone grid class with growth rate arbitrarily close to γ (Corollary 5.24).
- Almost all large permutations in a monotone grid class have the same shape (Proposition 6.1).
- A technique for determining the limit shape of a permutation in a monotone grid class (Proposition 6.3).
- The growth rate of a geometric grid class is equal to the square of the largest root of the matching polynomial of the row-column graph of the double refinement of its gridding matrix (Theorem 7.2).
- The set of growth rates of geometric grid classes consists of the squares of the spectral radii of trees (Corollary 7.16).
- If $G(M)$ is connected, then $\text{gr}(\text{Geom}(M)) < \text{gr}(\text{Grid}(M))$ if and only if $G(M)$ contains a cycle (Corollary 7.25).
- A specification of the effect of the subdivision of an edge on the largest root of the matching polynomial of a graph (Lemma 7.27).

Part II: The Set of Growth Rates

- Let $\theta_B \approx 2.35526$ be the unique real root of $x^7 - 2x^6 - x^4 - x^3 - 2x^2 - 2x - 1$. The set of growth rates of permutation classes includes an infinite sequence of intervals whose infimum is θ_B (Theorem 8.1).
- Let $\lambda_B \approx 2.35698$ be the unique positive root of $x^8 - 2x^7 - x^5 - x^4 - 2x^3 - 2x^2 - x - 1$. Every value at least λ_B is the growth rate of a permutation class (Theorem 8.2).
- A specification of when the set of representations of numbers in non-integer bases, where each digit is drawn from a different set, is an interval (Lemma 8.3).

Part III: Hasse Graphs

- A new derivation of the algebraic generating function of $\text{Av}(1324, 2314)$, first enumerated by Kremer [115, 116] (Theorem 10.1).
- A functional equation for the generating function of $\text{Av}(1324, 1432)$ (Theorem 11.1).
- A new derivation of the algebraic generating function of forest-like permutations, $\text{Av}(1324, 21\bar{3}54)$, first enumerated by Bousquet-Mélou & Butler [55] (Theorem 12.1).
- A functional equation for the generating function of plane permutations, $\text{Av}(21\bar{3}54)$ (Theorem 13.1).

-
- The algebraic generating function of $\text{Av}(\mathbf{1234}, \mathbf{2341})$ (Theorem 14.1).
 - The algebraic generating function of $\text{Av}(\mathbf{1243}, \mathbf{2314})$ (Theorem 15.1).
 - The growth rate of $\text{Av}(\mathbf{1324})$ exceeds 9.81 (Theorem 16.1).
 - The number of occurrences of a fixed pattern in a Łukasiewicz path of length n exhibits a concentrated Gaussian limit distribution (Theorem 16.2).

Part I

GRID CLASSES

Chapter 2

Introducing grid classes

2.1 Grid classes and griddings

One approach to investigating permutation classes that has proven particularly fruitful involves the use of certain classes that are defined constructively, rather than in terms of their basis. The monotone *grid class* $\text{Grid}(M)$ ¹ is defined by a matrix M , all of whose entries are in $\{0, 1, -1\}$. This *gridding matrix* specifies the permitted shape for plots of permutations in the class. Each entry of M corresponds to a *cell* in a *gridding* of a permutation. If the entry is 1, any points in the cell must form an increasing sequence; if the entry is -1 , any points in the cell must form a decreasing sequence; if the entry is 0, the cell must be empty.




Figure 2.1: The seven griddings of permutation 31567482 in $\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$

For greater clarity, we denote grid classes by *cell diagrams* rather than by their matrices; for example, $\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} = \text{Grid}\begin{pmatrix} 1 & -1 & 0 \\ 0 & -1 & 1 \end{pmatrix}$. Sometimes, with a slight abuse of notation, we use a cell diagram to denote the gridding matrix itself. We say that a grid class $\text{Grid}(M)$ has *size* m if its matrix M has m non-zero entries. Note that a permutation may have multiple possible griddings in a grid class. See Figure 2.1 for an example.

When defining grid classes, to match the way we view permutations graphically, we index matrices from the lower left corner, with the order of the indices reversed from

¹Huczynska & Vatter [105] were the first to use the term “grid class” and the notation $\text{Grid}(M)$.

the normal convention. For example, a matrix with dimensions $t \times u$ has t columns and u rows, and $M_{2,1}$ is the entry in the second column from the left and in the bottom row of M .

If M is a $0/\pm 1$ matrix with t columns and u rows, then an M -gridding of a permutation σ of length k is a pair of integer sequences $0 = c_0 \leq c_1 \leq \dots \leq c_t = k$ (the *column dividers*) and $0 = r_0 \leq r_1 \leq \dots \leq r_u = k$ (the *row dividers*) such that for all $i \in \{0, \dots, t\}$ and $j \in \{0, \dots, u\}$, the subsequence of σ with indices in $(c_{i-1}, c_i]$ and values in $(r_{j-1}, r_j]$ is increasing if $M_{i,j} = 1$, decreasing if $M_{i,j} = -1$, and empty if $M_{i,j} = 0$. For example, in the leftmost gridding in Figure 2.1, $c_1 = 5$ and $r_1 = 2$.

The *grid class* $\text{Grid}(M)$ is then defined to be the set of all permutations that have an M -gridding. The *griddings* of a permutation in $\text{Grid}(M)$ are its M -griddings. We use $\text{Grid}_k(M)$ to denote the permutations in $\text{Grid}(M)$ of length k .

Sometimes we need to consider a permutation along with a specific gridding. In this case, we refer to a permutation together with an M -gridding as an M -gridded permutation. We use $\text{Grid}^\#(M)$ to denote the class of all M -gridded permutations, every permutation in $\text{Grid}(M)$ being present once with each of its griddings. We use $\text{Grid}_k^\#(M)$ for the set of M -gridded permutations of length k . An important observation is the following:

Lemma 2.1 (Vatter [159, Proposition 2.1]). *The upper/lower growth rate of a monotone grid class $\text{Grid}(M)$ is equal to the upper/lower growth rate of the corresponding class of gridded permutations $\text{Grid}^\#(M)$.*

Proof. Suppose that M has dimensions $t \times u$. Every permutation in $\text{Grid}(M)$ has at least one gridding in $\text{Grid}^\#(M)$, but no permutation in $\text{Grid}(M)$ of length k can have more than $P(k) = \binom{k+t-1}{t-1} \binom{k+u-1}{u-1}$ griddings in $\text{Grid}^\#(M)$ because $P(k)$ is the number of possible choices for the row and column dividers. Since $P(k)$ is a polynomial in k , the result follows from the definition of the upper/lower growth rate. \square

Row-column graphs

To each grid class, we associate a bipartite graph, which we call its *row-column graph*². If M has t rows and u columns, the row-column graph, $G(M)$, of $\text{Grid}(M)$ is the graph with vertices $r_1, \dots, r_t, c_1, \dots, c_u$ and an edge between r_i and c_j if and only if $M_{i,j} \neq 0$ (see Figure 2.2 for an example). Note that any bipartite graph is the row-column graph of some grid class, and that the size (number of edges) of the row-column graph is the same as the size (number of non-zero cells) of the grid class.

²Vatter [159] was the first to use the term “row-column graph”.




Figure 2.2: A grid class and its row-column graph

The row-column graph of a grid class captures a great deal of important structural information about the class, so it is common to apportion properties of the row-column graph directly to the grid class itself, for example speaking of a connected, acyclic or unicyclic grid class rather than of a grid class whose row-column graph is connected, acyclic or unicyclic. We follow this convention.

Of particular note, Murphy & Vatter [131] have shown that a grid class is *partially well-ordered* (contains no infinite antichains) if and only if its row-column graph has no cycles. A simpler proof of this result was subsequently given by Vatter & Waton [162]. Furthermore, Albert, Atkinson, Bouvel, Ruškuc & Vatter [10] proved a result that implies that if a grid class has an acyclic row-column graph then the class has a rational generating function.

It is generally believed that every grid class has a finite basis (see [105, Conjecture 2.3]), but this has so far resisted proof. Atkinson [21] proved that *skinny* grid classes (whose matrices have dimensions $m \times 1$ for some m) have a finite basis. Waton [163] established the same for $\begin{smallmatrix} \nearrow & \searrow \\ \nearrow & \searrow \end{smallmatrix}$, a result which has since been extended by Albert, Atkinson & Brignall to all 2×2 grid classes (in an unpublished note [7]). More recently, Albert, Atkinson, Bouvel, Ruškuc & Vatter [10] have proved that every grid class with an acyclic row-column graph is finitely based.

The concept of a grid class of permutations has been generalised, permitting arbitrary permutation classes in each cell (see Vatter [159, Section 2]). We only consider *monotone* grid classes here. Both monotone and generalised grid classes have played a key role in investigations of the set of permutation class growth rates (see [105, 159]). We explore this topic in Part II below. An interactive *Mathematica* demonstration of monotone grid classes is available online [36].

2.2 Outline of Part I

In the next five chapters, we explore the enumeration of grid classes from various angles.

Chapter 3 concerns *skinny* grid classes, whose matrices have dimensions $m \times 1$ for some m . A permutation in such a class consists of the juxtaposition of ascending and descending sequences. We present an effective procedure for calculating the generating function

for any skinny grid class.

In chapter 4, we investigate the enumeration of classes of gridded permutations. We exhibit an effective method for determining the generating function for any acyclic or unicyclic class of gridded permutations. In the process, we prove that unicyclic classes have algebraic generating functions. We also prove that the generating function of every class of gridded permutations is D-finite.

In Chapter 5, we turn away from exact enumeration and focus on determining the exponential growth rate of grid classes. We prove that the growth rate of a monotone grid class is given by the square of the spectral radius of its row-column graph. This is the first general result concerning the exact growth rates of a family of permutation classes. Our proof depends on relating classes of gridded permutations to certain families of tours on graphs, and in the process we establish a new result concerning these tours that is of independent interest. As a consequence of our theorem, we deduce a number of facts relating to the set of growth rates of grid classes. This work has been published in [39].

Chapter 6 concerns the shape of a “typical” large permutation in a monotone grid class. We show that almost all large permutations in a grid class do indeed look the same, and explain how to determine the limit shape.

The final chapter in this part of the thesis, Chapter 7, concerns *geometric* grid classes, a family of permutation classes closely related to monotone grid classes. We investigate the growth rates of these classes, and prove a result which relates the growth rate of a geometric grid class to the largest root of the *matching polynomial* of a graph. In the process we establish a new result concerning the effect of edge subdivision on the largest root of the matching polynomial. We also deduce a number of consequences including providing a characterisation of the growth rates of geometric grid classes in terms of the spectral radii of trees. This work was published in [37].

Chapter 3

Skinny grid classes

We say that the permutation grid class $\text{Grid}(V)$ is *skinny* if V is a ± 1 row vector. Thus, a permutation in a skinny grid class consists of the juxtaposition of ascending and descending sequences. For an illustration, see Figure 3.1. Skinny grid classes were previously studied by Atkinson, Murphy & Ruškuc [23] (under the name “monotone segment sets”). They proved that every skinny grid class can be described by a regular language and so has a rational generating function. In this chapter we present a way of determining the generating function for any skinny grid class $\text{Grid}(V)$.




Figure 3.1: A gridded permutation in skinny grid class $\boxed{\nearrow\searrow\searrow\searrow}$

To state our result, we need some definitions. Suppose $V = (v_1, \dots, v_k)$ is a ± 1 vector of length k . Then we define V^+ to be the vector (v_1, \dots, v_k, v_k) of length $k + 1$, created from V by repeating its last entry, and define V^- to be $(v_1, \dots, v_k, -v_k)$, created from V^+ by negating its last entry. For example, $\boxed{\nearrow\searrow}^+ = \boxed{\nearrow\searrow\searrow}$ and $\boxed{\nearrow\searrow}^- = \boxed{\nearrow\searrow\swarrow}$. Also, given a vector $V = (v_1, \dots, v_k)$ and some $j \leq k$, we use $V[j]$ to denote (v_1, \dots, v_j) , the prefix of V of length j . For example, $\boxed{\nearrow\searrow\searrow\searrow}[2] = \boxed{\nearrow\searrow}$.

We prove that the generating function for a skinny grid class can be computed as follows:

Theorem 3.1. *If V is a ± 1 vector of length k , then the ordinary generating function for skinny grid class $\text{Grid}(V)$ is given by*

$$G_V(z) = \frac{1}{z} \sum_{j=1}^k H_{V_{[j]}}(z, z)$$

where, for any ± 1 vector V , $H_V(x, y)$ is defined iteratively as follows:

$$H_{(1)}(x, y) = H_{(-1)}(x, y) = \frac{xy}{1-x},$$

$$H_{V^+}(x, y) = \frac{xy}{xy+x-y} \left(H_V(x, y) - H_V(y, y) - H_V\left(x, \frac{x}{1-x}\right) + H_V\left(\frac{x}{1-x}, \frac{x}{1-x}\right) \right),$$

$$H_{V^-}(x, y) = \frac{xy}{xy+x-y} \left(H_V\left(\frac{x}{1-x}, x\right) - H_V(y, x) \right).$$

One way of gridding a permutation in a skinny grid class is to process the permutation from left to right, placing points in cells as far to the left as possible while honouring the cell constraints. A new cell is used only when the next point “changes direction”. We call this process *greedy gridding*, and also call the resulting gridding of a permutation a *greedy gridding*. See Figure 3.2 for an illustration.




Figure 3.2: The greedy gridding of a permutation in skinny grid class $\boxed{\diagdown \diagup \diagdown \diagup \diagdown \diagup}$

Greedy gridding is guaranteed to produce a valid gridding for any permutation in a skinny grid class because placing a point in a cell as far to the left as possible can never make it harder to place subsequent points. Since there is exactly one greedy gridding of each permutation in a skinny grid class, we can enumerate skinny grid classes by counting greedy gridded permutations.

Given a permutation $\pi \in \text{Grid}(V)$, it may be the case that π has points in every cell of V when greedy gridded. We call such permutations *tightly gridded* and use $\text{Grid}^*(V)$ to denote the set of tightly gridded permutations in $\text{Grid}(V)$. For example, the permutation in Figure 3.2 is not in $\text{Grid}^*(\boxed{\diagdown \diagup \diagdown \diagup \diagdown \diagup})$, but is a member of $\text{Grid}^*(\boxed{\diagdown \diagdown \diagdown \diagdown})$. Clearly,

a skinny grid class is the disjoint union of the tightly gridded permutations in each of its prefixes:

$$\text{Grid}(V) = \bigsqcup_{1 \leq j \leq \text{len}(V)} \text{Grid}^*(V \llbracket j \rrbracket).$$

For example, $\text{Grid}(\begin{array}{|c|c|c|} \hline \diagdown & \diagup & \diagdown \\ \hline \end{array}) = \text{Grid}^*(\begin{array}{|c|} \hline \diagdown \\ \hline \end{array}) \uplus \text{Grid}^*(\begin{array}{|c|c|} \hline \diagdown & \diagup \\ \hline \end{array}) \uplus \text{Grid}^*(\begin{array}{|c|c|c|} \hline \diagdown & \diagup & \diagdown \\ \hline \end{array})$.

To enumerate skinny grid classes, we only need to enumerate tightly gridded permutations. This we do by using a bivariate generating function in which we parameterise twice by the position of the last point in the plot of the permutation. We use x to mark the position of the last point counting from the *bottom* and y to mark its position counting from the *top*:

$$G_V^*(x, y) = \sum_{\pi \in \text{Grid}^*(V)} x^{\pi(|\pi|)} y^{|\pi|+1-\pi(|\pi|)} = \sum_{r, s \geq 1} g_{r, s}^* x^r y^s,$$

where $g_{r, s}^*$ is the number of permutations π in $\text{Grid}^*(V)$ of length $r + s - 1$ such that $\pi(|\pi|) = r$. For example, the permutation **314652** would be represented by a contribution of $x^2 y^5$.

Clearly, $\text{Grid}^*(V)$ and $\text{Grid}^*(-V)$ are Wilf-equivalent. It helps to restrict our attention to the case in which the last cell of V is increasing. Thus we define

$$H_V(x, y) = \begin{cases} G_V^*(x, y), & \text{if the last cell of } V \text{ is increasing,} \\ G_V^*(y, x), & \text{if the last cell of } V \text{ is decreasing.} \end{cases} \quad (1)$$

Note that $\frac{1}{z} H_V(z, z)$ is the univariate generating function for $\text{Grid}^*(V)$, and hence the generating function for $\text{Grid}(V)$ is given by

$$G_V(z) = \frac{1}{z} \sum_{j=1}^k H_{V \llbracket j \rrbracket}(z, z)$$

as required.

We now need to determine functional equations for $H_V(x, y)$. To begin with, as the base case, we have $H_{(1)}(x, y) = H_{(-1)}(x, y) = \frac{xy}{1-x}$ since there is a single increasing permutation of each length in $\text{Grid}^*(1)$.

We now investigate the effect on the generating function of adding points when greedy gridding.

Adding a new point

Let us first consider how a greedy gridded permutation $\pi \in \text{Grid}^*(V)$ can be extended by adding a single point to its right. We assume, without loss of generality, that the last cell of V is increasing. If the new point is added above the last point of π , then the extended permutation is also in $\text{Grid}^*(V)$. On the other hand, if the new point is added below the last point of π , then it must be placed in a new cell and the extended permutation is in $\text{Grid}^*(V^+)$ or $\text{Grid}^*(V^-)$. See Figure 3.3 for an illustration.




Figure 3.3: Possibilities for adding a single point to a permutation in $\text{Grid}^*(\square/\square)$

For each of these three cases (remaining in $\text{Grid}^*(V)$, expanding to $\text{Grid}^*(V^+)$, and expanding to $\text{Grid}^*(V^-)$) we define a linear operator (Ω_0 , Ω_+ and Ω_- , respectively) acting on $H_V(x, y)$, that reflects in each case the effect of adding a single new point.

For a specific permutation, represented by $x^r y^s$, we have

$$\begin{aligned} \text{(i)} \quad \Omega_0[x^r y^s] &= \sum_{i=1}^s x^{r+i} y^{s+1-i} = \frac{xy}{x-y} (x^{r+s} - x^r y^s) \\ \text{(ii)} \quad \Omega_+[x^r y^s] &= \sum_{i=1}^r x^{r+1-i} y^{s+i} = \frac{xy}{x-y} (x^r y^s - y^{r+s}) \\ \text{(iii)} \quad \Omega_-[x^r y^s] &= \sum_{i=1}^r x^{s+i} y^{r+1-i} = \frac{xy}{x-y} (x^{r+s} - x^s y^r). \end{aligned}$$

Thus, the three operators are defined by

$$\begin{aligned} \Omega_0[f(x, y)] &= xy \frac{f(x, x) - f(x, y)}{x - y}, \\ \Omega_+[f(x, y)] &= xy \frac{f(x, y) - f(y, y)}{x - y}, \\ \Omega_-[f(x, y)] &= xy \frac{f(x, x) - f(y, x)}{x - y}. \end{aligned}$$

Note that $\Omega_-[f(x, y)] = \Omega_+[f(y, x)]$ as expected from (1).

Adding a new cell

Let us now consider how a permutation in $\text{Grid}^*(V^+)$ can be built by extending a permutation $\pi \in \text{Grid}^*(V)$. First, we need to add a point below the last point of π , and then add zero or more increasing points above this first new point. See Figure 3.4 for an illustration.




Figure 3.4: Building a permutation in $\text{Grid}^*(\boxtimes\boxtimes\boxtimes)$ from one in $\text{Grid}^*(\boxtimes\boxtimes)$

Thus, in this case, adding a new cell is reflected by a single application of Ω_+ followed by zero or more applications of Ω_0 . Analogously, building a permutation in $\text{Grid}^*(V^-)$ from one in $\text{Grid}^*(V)$ is reflected by a single application of Ω_- followed by zero or more applications of Ω_0 .

Let Ω_* be the linear operator that reflects the action of adding zero or more increasing points (i.e. it is equivalent to zero or more applications of Ω_0). Then Ω_* satisfies the following functional equation:

$$\Omega_*[f(x, y)] = f(x, y) + \Omega_0[\Omega_*[f(x, y)]].$$

This can be solved for Ω_* by using the kernel method (see Section 1.1). Expanding for Ω_0 and rearranging gives

$$(xy + x - y)\Omega_*[f(x, y)] = (x - y)f(x, y) + xy\Omega_*[f(x, x)]. \quad (2)$$

Cancelling the kernel, $(xy + x - y)$, by setting $y = \frac{x}{1-x}$, then yields

$$\Omega_*[f(x, x)] = f(x, \frac{x}{1-x}),$$

which we can use to substitute for $\Omega_*[f(x, x)]$ in (2), which we then solve for Ω_* :

$$\Omega_*[f(x, y)] = \frac{xyf(x, \frac{x}{1-x}) + (x - y)f(x, y)}{xy + x - y}.$$

We now have all we need to express the relationship between $H_{V^+}(x, y)$, $H_{V^-}(x, y)$ and

\square	$\frac{z}{1-z}$
$\square \diagdown$	$\frac{z}{1-2z}$
$\square \diagup$	$\frac{z-2z^2+2z^3}{(1-z)^2(1-2z)}$
$\square \diagdown \square$	$\frac{z-3z^2+3z^3}{(1-z)^2(1-3z)}$
$\square \diagup \square$	$\frac{z-6z^2+13z^3-9z^4}{(1-z)(1-2z)^2(1-3z)}$
$\square \diagdown \square \diagup$	$\frac{z-8z^2+26z^3-39z^4+30z^5-12z^6}{(1-z)^3(1-2z)^2(1-3z)}$

Table 3.1: Generating functions for small skinny grid classes

$H_V(x, y)$:

$$H_{V^+}(x, y) = \Omega_\star[\Omega_+[H_V(x, y)]]$$

$$H_{V^-}(x, y) = \Omega_\star[\Omega_-[H_V(x, y)]].$$

Expansion using the definitions of the linear operators then completes the proof of Theorem 3.1.

The resulting generating functions for a few small skinny grid classes are listed in Table 3.1.

Based on an inspection of the results for vectors whose entries are all ones, we conclude our considerations of skinny grid classes by proposing the following conjecture concerning permutations with no more than $k-1$ descents:

Conjecture 3.2. *The ordinary generating function for the $k \times 1$ skinny grid class $\text{Grid}(1, 1, \dots, 1)$ is given by*

$$G_{\underbrace{(1, 1, \dots, 1)}_k}(z) = -1 + \sum_{r=1}^k \frac{1}{1-rz} \left(\frac{rz}{rz-1} \right)^{k-r}.$$

Beyond skinny grid classes

Individual non-skinny monotone grid classes have been enumerated using various *ad hoc* methods. Nonetheless, finding a *general* procedure for the exact enumeration of monotone grid classes remains an open problem. The primary challenge is that, whereas each permutation in a skinny grid class has a unique greedy gridding, there is no apparent

way to choose a canonical gridding for permutations in an arbitrary grid class. Furthermore, in general, the generating function for such a class may not be rational or even algebraic.

There is, however, one family of permutation classes, closely related to grid classes, that *has* recently been enumerated. In [100], Homberger & Vatter present a way of exactly enumerating any *polynomial* permutation class. A class \mathcal{C} is said to be polynomial if, for all sufficiently large n , $|\mathcal{C}_n|$ is given by a polynomial. Equivalently, \mathcal{C} is polynomial if $\text{gr}(\mathcal{C}) \in \{0, 1\}$.

Each polynomial class can be expressed as the finite union of what Homberger & Vatter call *peg permutation grid classes*. Such a class can be defined by a matrix whose entries are drawn from $\{0, 1, -1, \bullet\}$ and in which each row and each column contains exactly one non-zero entry. As with monotone grid classes, the entries in the matrix specify the permitted pattern of points in the corresponding cell. If the entry is \bullet , the cell must contain only a single point or remain empty. An example of a peg permutation grid class is

$$\text{Av}(321, 2134) = \begin{array}{|c|c|c|c|c|} \hline & & & \bullet & \\ \hline & / & & & \\ \hline & & & & \bullet \\ \hline & / & & & / \\ \hline & & & / & \\ \hline & & & & / \\ \hline & / & & & \\ \hline \end{array},$$

an identity first proved by Atkinson in [21].

Homberger & Vatter present an effective method for enumerating any peg permutation grid class. It may be possible to extend the techniques they use so as to enable the enumeration of some non-polynomial non-skinny monotone grid classes.

In the next chapter, we turn to an investigation of how we can enumerate classes of M -gridded permutations, and explore how the enumeration of gridded permutations may be of use in enumerating monotone grid classes.

Chapter 4

Gridded permutations

Success in finding general methods for the exact enumeration of permutation grid classes is currently limited to skinny and polynomial classes. However, it is possible to exactly enumerate a somewhat broader family of classes of *gridded permutations*.




Figure 4.1: A $\begin{array}{|c|c|c|} \hline \text{X} & \text{X} & \text{X} \\ \hline \text{X} & \text{X} & \text{X} \\ \hline \end{array}$ -gridded permutation

In this chapter we present a procedure for determining the generating function for any class of gridded permutations whose row-column graph has no more than one cycle in any connected component. We show that, if $G(M)$ is acyclic, then the generating function of $\text{Grid}^\#(M)$ is rational. This actually follows from the proof by Albert, Atkinson, Bouvel, Ruškuc & Vatter [10] that acyclic grid classes have rational generating functions. Moreover, we prove that the generating function of a unicyclic class of gridded permutations is always algebraic. We also prove that, for any M , the generating function of M -gridded permutations is D-finite. Finally, we explore how the enumeration of gridded permutations can help when trying to enumerate grid classes.

It is possible to give an explicit expression for the number of gridded permutations of length k in any specified grid class. We do this by summing over the number of configurations with a specified number of points in each cell.

To express the result, we make use of *multinomial coefficients*, with their normal combinatorial interpretation, for which we use the standard notation

$$\binom{n}{k_1, k_2, \dots, k_r} = \frac{n!}{k_1! k_2! \dots k_r!}, \quad \text{where } \sum_{i=1}^r k_i = n,$$

to denote the number of ways of distributing n distinguishable objects between r (distinguishable) bins, such that bin i contains exactly k_i objects ($1 \leq i \leq r$).

Lemma 4.1. *If M has dimensions $r \times s$, then the number of gridded permutations of length k in $\text{Grid}^\#(M)$ is given by*

$$|\text{Grid}_k^\#(M)| = \sum \prod_{i=1}^r \binom{k_{i,1} + k_{i,2} + \dots + k_{i,s}}{k_{i,1}, k_{i,2}, \dots, k_{i,s}} \prod_{j=1}^s \binom{k_{1,j} + k_{2,j} + \dots + k_{r,j}}{k_{1,j}, k_{2,j}, \dots, k_{r,j}},$$

where the sum is over all combinations of non-negative $k_{i,j}$ ($1 \leq i \leq r$, $1 \leq j \leq s$) such that $\sum_{i,j} k_{i,j} = k$ and $k_{i,j} = 0$ if $M_{i,j} = 0$.

Proof. An M -gridded permutation consists of a number of points in each of the cells that correspond to a non-zero entry of M . For every permutation, the relative ordering of points (increasing or decreasing) within a particular cell is fixed by the value of the corresponding matrix entry. However, the relative interleaving between points in distinct cells in the same row or column can be chosen arbitrarily and independently for each row and column.

Each term in the sum is thus the number of gridded permutations in which there are $k_{i,j}$ points in the cell corresponding to $M_{i,j}$, the terms in the first product representing the number of ways of interleaving points in each column, and terms in the second product representing the number of ways of interleaving points in each row. The result follows by summing over configurations with a total of k points and no points in cells that correspond to a zero entry of M . \square

As a consequence of this result, it is clear that $|\text{Grid}_k^\#(M)| = |\text{Grid}_k^\#(M')|$ whenever M' results from permuting the rows and/or columns of M . Indeed, it is not hard to see that the enumeration of a class of gridded permutations depends only on its row-column graph, a fact that we prove formally later (Corollary 5.11).

We can “translate” Lemma 4.1 to give us an expression for the generating function of any class of gridded permutations. Note that, when enumerating gridded permutations, it turns out to be simpler if we include the zero-length permutation. However, when enumerating grid classes, we chose to exclude the zero-length permutation.

Central to this and following results is the *diagonalisation* of generating functions. Given a bivariate power series $f(x, y) = \sum a_{r,s} x^r y^s$, the *diagonal*, $\Delta(f)$, of f is the univariate series defined by $\Delta(f)(z) = \sum a_{n,n} z^n$. Equivalently, if $q(x, y)$ consists of the sum of just those terms of $f(x, y)$ for which the exponent of x is the same as that of y , then we have $\Delta(f)(z) = q(\sqrt{z}, \sqrt{z})$. Alternatively, $\Delta(f)(z) = [w^0] f(w\sqrt{z}, \sqrt{z}/w)$.

Lemma 4.2. *The generating function for $\text{Grid}^\#(M)$ is given by*

$$G_M^\#(z) = \left[\prod_{\substack{i,j \\ M_{i,j} \neq 0}} x_{i,j} \right]^0 \prod_i \left(1 - \sqrt{z} \sum_{\substack{j \\ M_{i,j} \neq 0}} x_{i,j} \right)^{-1} \prod_j \left(1 - \sqrt{z} \sum_{\substack{i \\ M_{i,j} \neq 0}} x_{i,j}^{-1} \right)^{-1},$$

in which there is one variable $x_{i,j}$ for each non-zero cell of M , and we extract only the constant terms as far as the $x_{i,j}$ are concerned.

Proof. Suppose that, for each non-zero entry $M_{i,j}$, we let the variable $u_{i,j}$ mark the number of points in the cell corresponding to $M_{i,j}$. Then the multivariate generating function for the ways in which the points in column i can be interleaved is

$$\frac{1}{1 - (u_{i,j_1} + u_{i,j_2} + \dots + u_{i,j_t})} = \left(1 - \sum_{\substack{j \\ M_{i,j} \neq 0}} u_{i,j} \right)^{-1},$$

where j_1, j_2, \dots, j_t are the values of j for which $M_{i,j} \neq 0$.

Now, let us *also* use $v_{i,j}$ to mark the number of points in the cell corresponding to non-zero entry $M_{i,j}$. If we use the $v_{i,j}$, rather than the $u_{i,j}$, when considering the interleaving of points in the same row, then the multivariate generating function for M -gridded permutations is given by those terms in

$$\prod_i \left(1 - \sum_{\substack{j \\ M_{i,j} \neq 0}} u_{i,j} \right)^{-1} \prod_j \left(1 - \sum_{\substack{i \\ M_{i,j} \neq 0}} v_{i,j} \right)^{-1} \quad (1)$$

for which the exponent of $u_{i,j}$ is the same as that of $v_{i,j}$ for all i and j .

The result follows by replacing each $u_{i,j}$ with $x_{i,j}\sqrt{z}$ and replacing each $v_{i,j}$ with $\sqrt{z}/x_{i,j}$, so that z marks the length of the permutation (which is the total number of points). Using \sqrt{z} in this way ensures that half of each point is counted when considering the interleaving of points in a column, and another half of each point is counted when considering the interleaving of points in a row. After this substitution, the terms required are those for which the exponent of each of $x_{i,j}$ is zero. \square

In general, coefficient extraction is hard, so Lemma 4.2 does not give us an effective proce-

ture for determining explicit expressions for the generating functions of arbitrary classes of gridded permutations. Nevertheless, it can be used to yield closed forms for the generating functions when the row-column graphs of the classes are acyclic or unicyclic.

4.1 Acyclic classes

We begin with acyclic classes. First, we describe a process for adding cells to a grid class that can be used repeatedly to construct any class whose row-column graph is a forest.



Figure 4.2: Extending an acyclic grid class by inserting three new columns, each with a single non-zero entry in the second row, and the corresponding extension of the row-column graph

In the context of graphs, it is clear that any forest can be constructed by starting with a 1-regular graph (a number of disconnected edges) and then repeatedly attaching some pendant edges to a leaf vertex of the current graph. The analogous method for grid classes consists of starting with a class in which each row and each column contains exactly one non-zero entry, and then repeatedly applying the following extension process: Choose a row or column that contains a single non-zero entry and form a new class by inserting additional columns or rows respectively, each containing a single non-zero entry in the chosen row or column. This grid class method corresponds exactly to the graph method applied to row-column graphs. See Figure 4.2 for an illustration of this extension process.

The following theorem enables us to use this extension process to enumerate any class of M -gridded permutations when $G(M)$ is a forest. For clarity, its statement only covers the case in which new columns are added to the class.

Theorem 4.3. *Suppose $G(M)$ is acyclic and that M has a single non-zero entry in some row r . Let M^+ be a matrix formed from M by inserting k additional columns, each containing a single non-zero entry in row r .*

If $G_M^\#(z, x)$ is the bivariate generating function for $\text{Grid}^\#(M)$, where z marks length and x marks the number of points in the cell corresponding to the non-zero entry of M in row r , then there exist polynomials P and Q such that

$$G_M^\#(z, x) = \frac{1}{1 + z(P(z) + xQ(z))}.$$

Moreover, if $G_{M^+}^\#(z, x_1, \dots, x_k)$ is the multivariate generating function for $\text{Grid}^\#(M^+)$ where x_1, \dots, x_k mark the number of points in each of the k new cells in row r , then

$$G_{M^+}^\#(z, x_1, \dots, x_k) = \frac{1}{1 + z(P(z) + Q(z) - (x_1 + \dots + x_k)(1 + zP(z)))}.$$

In proving this theorem, we make use of the following standard diagonalisation result concerning the extraction of coefficients from generating functions:

Lemma 4.4 (Stanley [152, Section 6.3]; see Furstenberg [80]). *If $g(x) = g(z, x)$ is a formal Laurent series, then the constant term $[x^0]g(x)$ is given by the sum of the residues of $x^{-1}g(x)$ at those poles α of $g(x)$ for which $\lim_{z \rightarrow 0} \alpha(z) = 0$.*

Proof of Theorem 4.3. For the base case, if each row and each column of M contains exactly one non-zero entry, then

$$G_M^\#(z, x) = \frac{1}{(1 - zx)(1 - z)^{m-1}},$$

where m is the number of non-zero entries in M . This can certainly be expressed in the form specified in the statement of the theorem.

We now consider the extension process. The argument is analogous to that used in the proof of Lemma 4.2. The multivariate generating function for the ways in which the points in row r of $\text{Grid}^\#(M^+)$ can be interleaved is

$$R(z, x) = R(z, x, x_1, \dots, x_k) = \frac{1}{1 - z(x + x_1 + \dots + x_k)}.$$

Therefore,

$$G_{M^+}^\#(z, x_1, \dots, x_k) = [x^0]G_M^\#(z, \frac{x}{\sqrt{z}})R(z, \frac{1}{x\sqrt{z}}) = [x^0]\frac{G_M^\#(z, x/\sqrt{z})}{1 - x^{-1}\sqrt{z} - zs},$$

where $s = x_1 + \dots + x_k$.

If we assume that

$$G_M^\#(z, x) = \frac{1}{1 + z(P(z) + xQ(z))},$$

then

$$G_{M^+}^\#(z, x_1, \dots, x_k) = [x^0]\frac{x}{(1 + zP(z) + x\sqrt{z}Q(z))(x(1 - zs) - \sqrt{z})}.$$

We now apply Lemma 4.4. The expression at the right has two poles. The root of the first factor of the denominator diverges at $z = 0$. The other pole, $\alpha = \sqrt{z}/(1 - zs)$, satisfies the necessary criterion.

Since α is a simple pole,

$$G_{M^+}^{\#}(z, x_1, \dots, x_k) = \frac{1}{(1 + zP(z) + \alpha\sqrt{z}Q(z))(1 - zs)},$$

which simplifies to yield the desired result.

Since every acyclic class can be constructed by multiple applications of the extension process, and the generating function resulting from the extension process can also be expressed in the required form, the result holds. \square

4.2 Unicyclic classes

We call a graph unicyclic if it is not acyclic and no connected component contains more than one cycle. A permutation class is unicyclic if its row-column graph is unicyclic.



Figure 4.3: Constructing a unicyclic grid class by identifying two cells of an acyclic class, and the corresponding operation on the row-column graph

In the context of graphs, a connected unicyclic graph can be constructed from a tree by identifying two pendant edges and their endvertices, oriented such that the two leaf vertices are not identified. We construct a unicyclic grid class in an analogous way. We begin with an acyclic class whose matrix M contains either $\begin{bmatrix} \diagup \\ \diagdown \end{bmatrix}$ or $\begin{bmatrix} \diagdown \\ \diagup \end{bmatrix}$ as a submatrix M' , such that one of the two non-zero cells in M' is the only non-zero cell in its column in M , and the other non-zero cell in M' is the only non-zero cell in its row in M . A unicyclic class is then constructed by combining the two columns of M that contain M' keeping all the non-zero entries, and similarly combining the two rows of M that contain M' . This corresponds to identifying the (pendant) edges of $G(M)$ that correspond to the two non-zero entries of M' . See Figure 4.3 for an illustration of this operation.

The following theorem enables us to use this operation to enumerate any unicyclic class of gridded permutations.

Theorem 4.5. *Suppose $G(M)$ is acyclic and that e_1 and e_2 are two pendant edges of $G(M)$ that correspond to the non-zero entries of a $\begin{bmatrix} \diagup \\ \diagdown \end{bmatrix}$ or $\begin{bmatrix} \diagdown \\ \diagup \end{bmatrix}$ submatrix of M . Let M° be the unicyclic matrix that results from combining the columns and rows of M containing the cells corresponding to e_1 and e_2 .*

If $G_M^{\#}(z, u, v)$ is the trivariate generating function for $\text{Grid}^{\#}(M)$, where z marks length, and

u and v mark the number of points in the cells corresponding to e_1 and e_2 , then there exist polynomials R, S, T and U such that

$$G_M^\#(z, u, v) = \frac{1}{1 + z(R(z) + uS(z) + vT(z) + uvzU(z))}.$$

Moreover, the generating function of $\text{Grid}^\#(M^\circ)$ is then given by

$$G_{M^\circ}^\#(z) = \frac{1}{\sqrt{(1 + zR(z) + zU(z))^2 - 4zS(z)T(z)}}.$$

We have as an immediate consequence:

Theorem 4.6. *The generating function of any unicyclic class of gridded permutations is algebraic.*

Proof of Theorem 4.5. The first part follows from the fact that the required form is preserved by the extension process of Theorem 4.3.

By construction, $G_{M^\circ}^\#(z)$ is given by those terms of $G_M^\#(z, u, v)$ for which u and v have the same exponent. Thus, if we assume that $G_M^\#(z, u, v)$ has the form specified in the statement of the theorem,

$$G_{M^\circ}^\#(z) = [w^0] G_M^\#(z, \frac{w}{\sqrt{z}}, \frac{1}{w\sqrt{z}}) = [w^0] \frac{w}{w^2 \sqrt{z} S(z) + w(1 + zR(z) + zU(z)) + \sqrt{z} T(z)}.$$

We now apply Lemma 4.4. The expression at the right has two poles. One diverges at $z = 0$. The other,

$$\alpha = \frac{-1 - zR(z) - zU(z) + \sqrt{(1 + zR(z) + zU(z))^2 - 4zS(z)T(z)}}{2\sqrt{z}S(z)},$$

satisfies the necessary criterion. Algebraic manipulation then yields the required result. \square

4.3 Exploitation

Theorems 4.3 and 4.5 can be used to help in enumerating acyclic and unicyclic monotone grid classes. One possible approach to the enumeration of a permutation grid class is to partition it into parts in such a way that each permutation in the class has a unique gridding in exactly one of the parts. The parts are then essentially sets of gridded permutations.

Acyclic classes

For example, it is not difficult to confirm that the acyclic grid class $\text{Grid}\left(\begin{array}{|c|} \hline \diagup \diagdown \\ \hline \end{array}\right)$ can be partitioned as follows:

$$\begin{array}{|c|} \hline \diagup \diagdown \\ \hline \end{array} = \begin{array}{|c|} \hline \bullet \\ \hline \end{array} \uplus \begin{array}{|c|} \hline \bullet \diagup \diagdown \\ \hline \end{array} \uplus \begin{array}{|c|} \hline \bullet \diagup \diagdown \\ \hline \end{array} \uplus \begin{array}{|c|} \hline \bullet \diagup \diagdown \\ \hline \end{array}.$$

We call diagrams like these *entanglement diagrams*. They are similar to the peg permutation grid class diagrams described in the previous chapter (see page 35). However, in an entanglement diagram, disks (\bullet) are placed on the intersections of grid lines, rather than in the cells, and they specify that any permutation in the set denoted by the diagram *must* have a (single) point at that location. Entanglement diagrams thus do not represent permutation classes closed under containment. With a slight abuse of terminology, we also use the term *entanglement diagram* to refer to the set of permutations represented by such a diagram.

Let us call a partition of a grid class into disjoint entanglement diagrams such that each permutation has a unique gridding a *proper partition* of the class. If we have a proper partition of a grid class, then the enumeration of the grid class can be achieved simply by counting gridded permutations in each of the diagrams. So, in our example, we have

$$G_{\begin{array}{|c|} \hline \diagup \diagdown \\ \hline \end{array}}(z) = zG_{\begin{array}{|c|} \hline \bullet \\ \hline \end{array}}^{\#}(z) + z^3G_{\begin{array}{|c|} \hline \bullet \diagup \diagdown \\ \hline \end{array}}^{\#}(z) + z^3G_{\begin{array}{|c|} \hline \bullet \diagup \diagdown \\ \hline \end{array}}^{\#}(z) + z^4G_{\begin{array}{|c|} \hline \bullet \diagup \diagdown \\ \hline \end{array}}^{\#}(z),$$

which, after numerous applications of Theorem 4.3, yields

$$G_{\begin{array}{|c|} \hline \diagup \diagdown \\ \hline \end{array}}(z) = \frac{1 - 5z + 8z^2 - 3z^3}{(1 - z)(1 - 2z)(1 - 3z + z^2)}.$$

It is elementary to determine a general technique for the proper partitioning of *skinny* grid classes into entanglement diagrams. This yields an alternative approach to their enumeration to that in the previous chapter. We leave the details as an exercise for the reader. However, attempts at discovering an effective procedure for proper partitioning applicable to all acyclic grid classes have so far been unsuccessful. However, it does seem reasonable to assume that every acyclic grid class can be enumerated in this way:

Conjecture 4.7. *Every acyclic monotone grid class can be properly partitioned into a finite number of acyclic entanglement diagrams.*

It is to be hoped that some effective procedure can be discovered for the partitioning so that it can be used to enumerate all unicyclic grid classes.

Unicyclic classes

When we consider unicyclic classes, the situation seems to be somewhat less straightforward. For example, it would appear that the “double chevron” grid class $\text{Grid}(\begin{smallmatrix} \diagup & \diagdown \\ \diagdown & \diagup \end{smallmatrix})$ can be properly partitioned as follows:

$$\begin{smallmatrix} \diagup & \diagdown \\ \diagdown & \diagup \end{smallmatrix} = \begin{smallmatrix} \diagup \\ \diagdown \end{smallmatrix} \uplus \begin{smallmatrix} \bullet & & & \\ \diagup & & & \\ & \bullet & & \\ \diagdown & & & \\ & & \bullet & \\ & & \diagup & \\ & & & \bullet \\ & & & \diagdown \end{smallmatrix} \uplus \begin{smallmatrix} \bullet & & & \\ \diagup & & & \\ & \bullet & & \\ \diagdown & & & \\ & & \bullet & \\ & & \diagup & \\ & & & \bullet \\ & & & \diagdown \end{smallmatrix}.$$

To achieve a proper partition of a unicyclic class, we seem to need to generalise our concept of an entanglement diagram, allowing some lines to cross cell boundaries.¹ It is not too hard to work out how to enumerate the gridded permutations in such generalised diagrams.

Assuming our partition of $\begin{smallmatrix} \diagup & \diagdown \\ \diagdown & \diagup \end{smallmatrix}$ is correct, it is then possible to deduce that

$$f_{\begin{smallmatrix} \diagup & \diagdown \\ \diagdown & \diagup \end{smallmatrix}}(z) = \frac{1}{\sqrt{1-4z}} - \frac{1-4z+5z^2}{(1-2z)(1-3z)}.$$

In the light of this, we tentatively make the following conjecture, which would hold if every unicyclic monotone grid class could be properly partitioned into a finite number of acyclic and unicyclic generalised entanglement diagrams:

Conjecture 4.8. *The generating function of any unicyclic monotone grid class is algebraic.*

Multicyclic classes

What can we say about classes with components having more than one cycle?

From (1) in the proof of Lemma 4.2, we know that the generating function of any class of gridded permutations results from the repeated diagonalisation of a rational function. The diagonal of a bivariate rational power series is always algebraic, a result first proved by Furstenberg [80]. Indeed the converse is also true, every algebraic function comes about by the diagonalisation of some rational function. However, diagonalisation does not preserve algebraicity. Nevertheless, Lipshitz [121] proved that *D-finite* power series

¹Such diagrams have row-column *hypergraphs*.

are closed under taking diagonals. (See Section 1.1 for the definition of a D-finite power series.) As a consequence, we have the following general result concerning gridded permutations:

Theorem 4.9. *The generating function of any class of gridded permutations is D-finite.*

As a consequence, to conclude, we dare to propose the following conjecture, which would hold if every monotone grid class could be properly partitioned into a finite number of suitably generalised entanglement diagrams, each of whose enumeration was D-finite:

Conjecture 4.10. *The generating function of any monotone grid class is D-finite.*

Chapter 5

Growth rates of grid classes

In this chapter and the next, we turn away from exact enumeration to the (slightly) easier question of determining the asymptotic growth rate of monotone grid classes of permutations. We prove that the exponential growth rate of $\text{Grid}(M)$ is equal to the square of the spectral radius of its row-column graph $G(M)$. Consequently, we utilize spectral graph theoretic results to characterise all slowly growing grid classes and to show that for every $\gamma \geq 2 + \sqrt{5}$ there is a grid class with growth rate arbitrarily close to γ .

To prove our main result, we establish bounds on the size of certain families of tours on graphs. In the process, we prove that the family of tours of even length on a connected graph grows at the same rate as the family of “balanced” tours on the graph (in which the number of times an edge is traversed in one direction is the same as the number of times it is traversed in the other direction).

5.1 Introduction

Our focus in this chapter is on the growth rates of grid classes. We prove the following theorem:

Theorem 5.14. *The growth rate of a monotone grid class of permutations exists and is equal to the square of the spectral radius of its row-column graph.*

The bulk of the work required to prove this theorem is concerned with carefully counting certain families of tours on graphs, in order to give bounds on their sizes. In particular, we consider “balanced” tours, in which the number of times an edge is traversed in one direction is the same as the number of times it is traversed in the other direction. As a consequence, we prove the following new result concerning tours on graphs:

Theorem 5.8. *The growth rate of the family of balanced tours on a connected graph is the same as that of the family of all tours of even length on the graph.*

As a consequence of Theorem 5.14, by using the machinery of spectral graph theory, we are able to deduce a variety of supplementary results. We give a characterisation of grid classes whose growth rates are no greater than $\frac{9}{2}$ (in a similar fashion to Vatter’s characterisation of “small” permutation classes in [159]). We also fully characterise all *accumulation points* of grid class growth rates, the least of which occurs at 4. Other results include:

Corollary 5.15. *The growth rate of every monotone grid class is an algebraic integer.*

Corollary 5.17. *A monotone grid class whose row-column graph is a cycle has growth rate 4.*

Corollary 5.19. *If the growth rate of a monotone grid class is less than 4, it is equal to $4 \cos^2\left(\frac{\pi}{k}\right)$ for some integer $k \geq 3$.*

Corollary 5.24. *For every $\gamma \geq 2 + \sqrt{5}$ there is a monotone grid class with growth rate arbitrarily close to γ .*

The remainder of this chapter is structured as follows: In Section 5.2, we introduce the particular families of tours on graphs that we study and present our results concerning these tours, culminating in the proof of Theorem 5.8. This is followed, in Section 5.3, by the application of these results to prove our grid class growth rate result, Theorem 5.14. To conclude, in Section 5.4, we present a number of consequences of Theorem 5.14 that follow from known spectral graph theoretic results.

5.2 Tours on graphs

In this section, we investigate families of tours on graphs, parameterised by the number of times each edge is traversed. We determine a lower bound on the size of families of “balanced” tours and an upper bound on families of arbitrary tours. Applying the upper bound to tours of even length gives us an expression compatible with the lower bound. Combining this with the fact that any balanced tour has even length enables us to prove Theorem 5.8 which reveals that even-length tours and balanced tours grow at the same rate. These bounds are subsequently used in Section 5.3 to relate tours on graphs to permutation grid classes.

To establish the lower and upper bounds, we first enumerate tours on trees. We then present a way of associating tours on an arbitrary connected graph G with tours on a related “partial covering” tree, which we employ to determine bounds for families of tours on arbitrary graphs. Let us begin by introducing the tours that we consider below.

5.2.1 Notation and definitions

A *walk*, of length k , on a graph is a non-empty alternating sequence of vertices and edges $v_0, e_1, v_1, e_2, v_2, \dots, e_k, v_k$ in which the endvertices of e_i are v_{i-1} and v_i . Neither the edges nor the vertices need be distinct. We say that such a walk *traverses* edges $\{e_1, \dots, e_k\}$ and *visits* vertices $\{v_1, \dots, v_{k-1}\}$. A *tour* (or *closed walk*) is a walk which starts and ends at the same vertex (i.e. $v_k = v_0$). Our interest is restricted to tours.

Below, when considering a graph with m edges, we denote its edges e_1, e_2, \dots, e_m . In any particular context, we can choose the ordering of the edges so as to simplify our presentation. We denote the edges incident to a given vertex v by $e_1^v, e_2^v, \dots, e_{d(v)}^v$, where $d(v)$ is the *degree* of v (number of edges incident to v). Again, we are free to choose the order of the edges incident to a vertex so as to clarify our arguments.

Families of tours

Our interest is in families of tours that are parameterised by the number of times each edge is traversed. Given non-negative integers h_1, h_2, \dots, h_m and some vertex u of a graph G , we use

$$\mathcal{W}_G((h_i); u) = \mathcal{W}_G(h_1, h_2, \dots, h_m; u)$$

to denote the family of tours on G which start and end at u and traverse each edge e_i exactly h_i times. (We use \mathcal{W} rather than \mathcal{T} for families of tours to avoid confusion when considering tours on trees.)

We use $h_1^v, h_2^v, \dots, h_{d(v)}^v$ for the number of traversals of edges incident to a vertex v in $\mathcal{W}_G((h_i); u)$. So, if v and w are the endvertices of e_i , h_i has two aliases h_j^v and $h_{j'}^w$ for some j and j' .

We use $W_G((h_i); u) = |\mathcal{W}_G((h_i); u)|$ to denote the number of these tours.

Note that for some values of h_1, \dots, h_m , the family $\mathcal{W}_G((h_i); u)$ is empty. In particular, if $E^+ = \{e_i \in E(G) : h_i > 0\}$ is the set of edges traversed by tours in the family, and $G^+ = G[E^+]$ is the subgraph of G induced by these edges, then if G^+ is disconnected or does not contain u , we have $\mathcal{W}_G((h_i); u) = \emptyset$. A family of tours may also be empty for ‘‘parity’’ reasons; for example, if T is a tree, then $\mathcal{W}_T((h_i); u) = \emptyset$ if any of the h_i are odd. Our counting arguments must remain valid for these empty families.

Of particular interest to us are tours in which the number of times an edge is traversed in one direction is the same as the number of times it is traversed in the other direction. We call such tours *balanced*.

Given non-negative integers k_1, k_2, \dots, k_m and some vertex u of a graph G , we use

$$\mathcal{W}_G^{\mathbf{B}}((k_i); u) = \mathcal{W}_G^{\mathbf{B}}(k_1, k_2, \dots, k_m; u)$$

to denote the family of balanced tours on G which start and end at u , and traverse each edge e_i exactly k_i times *in each direction*. Note that we parameterise balanced tours by *half* the number of traversals of each edge.

We use $k_1^v, k_2^v, \dots, k_{d(v)}^v$ for the number of traversals in either direction of edges incident to a vertex v in $\mathcal{W}_G^{\mathbf{B}}((k_i); u)$. So, if v and w are the endvertices of e_i , k_i has two aliases k_j^v and $k_{j'}^w$ for some j and j'

We use $W_G^{\mathbf{B}}((k_i); u) = |\mathcal{W}_G^{\mathbf{B}}((k_i); u)|$ to denote the number of these balanced tours.

As with $\mathcal{W}_G((h_i); u)$, $\mathcal{W}_G^{\mathbf{B}}((k_i); u)$ may be empty. Observe also that, since any tour on a forest is balanced, $\mathcal{W}_F((2k_i); u) = \mathcal{W}_F^{\mathbf{B}}((k_i); u)$ for any forest F and $u \in V(F)$. Moreover, for any graph G , we have $W_G^{\mathbf{B}}((k_i); u) \leq W_G((2k_i); u)$, with equality if and only if the component of G^+ containing u , if present, is acyclic, where G^+ is the subgraph of G induced by the edges that are actually traversed by tours in the family.

Visits and excursions

We use $\Psi(G, v)$ to denote the number of visits to v of any tour on G in some family (specified by the context). In practice, this notation is unambiguous because we only consider one family of tours on a particular graph at a time. Observe that any tour in $\mathcal{W}_G((h_i); u)$ visits vertex $v \neq u$ exactly $\frac{1}{2}(h_1^v + h_2^v + \dots + h_{d(v)}^v)$ times, and that for balanced tours in $\mathcal{W}_G^{\mathbf{B}}((k_i); u)$ we have $\Psi(G, v) = k_1^v + k_2^v + \dots + k_{d(v)}^v$.

If $\Psi(G, v)$ is positive, then separating the visits to v are $\Psi(G, v) - 1$ "subtours" starting and ending at v ; we refer to these subtours as *excursions* from v .

Multinomial coefficients

In our calculations, we make considerable use of multinomial coefficients as described on page 38. We make repeated use of the fact that a multinomial coefficient can be decomposed into a product of binomial coefficients: $\binom{n}{k_1, \dots, k_r} = \binom{k_1}{k_1} \binom{k_1+k_2}{k_2} \dots \binom{k_1+\dots+k_r}{k_r}$. We consider a multinomial coefficient that has one or more negative terms to be *zero*. This guarantees that the monotonicity condition $\binom{n}{k_1, \dots, k_r} \leq \binom{n+1}{k_1+1, \dots, k_r}$ holds for all possible sets of values.

5.2.2 Tours on trees

We begin by establishing bounds on the size of families of tours on *trees*. As we noted above, all such tours are balanced. We start with star graphs, giving an exact enumeration of any family:

Lemma 5.1. *If S_m is the star graph $K_{1,m}$ with central vertex u , then*

$$W_{S_m}^{\mathbf{B}}((k_i); u) = \binom{k_1 + k_2 + \dots + k_m}{k_1, k_2, \dots, k_m} = \binom{\Psi(S_m, u)}{k_1^u, k_2^u, \dots, k_{d(u)}^u}.$$

Proof. $W_{S_m}^{\mathbf{B}}((k_i); u)$ consists of all possible interleavings of k_i excursions from u out-and-back along each e_i . □

It is possible to extend our exact enumeration to those families of balanced tours on trees in which every internal (non-leaf) vertex is visited at least once. These families are never empty.

Lemma 5.2. *If T is a tree, $u \in V(T)$ and, for each $v \neq u$, e_1^v is the edge incident to v that is on the unique path between u and v , and if k_1^v is positive for all internal vertices v of T , then*

$$W_T^{\mathbf{B}}((k_i); u) = \binom{\Psi(T, u)}{k_1^u, k_2^u, \dots, k_{d(u)}^u} \prod_{v \neq u} \binom{\Psi(T, v) - 1}{k_1^v - 1, k_2^v, \dots, k_{d(v)}^v}.$$

Proof. We use induction on the number of internal vertices. Note that the multinomial coefficient for a leaf vertex simply contributes a factor of 1 to the product. Lemma 5.1 provides the base case.

Given a tree T with m edges e_1, \dots, e_m , and a leaf v of T , let T' be the tree “grown” from T by attaching r new pendant edges e_{m+1}, \dots, e_{m+r} to v .

If k_1^v is positive, since v is a leaf, each tour in $W_T^{\mathbf{B}}(k_1, \dots, k_m; u)$ visits v exactly k_1^v times, with $k_1^v - 1$ excursions from v along e_1^v separating these visits. Any such tour can be extended to a tour in $W_{T'}^{\mathbf{B}}(k_1, \dots, k_{m+r}; u)$ by arbitrarily interleaving k_{m+i} new excursions out-and-back along each new pendant edge e_{m+i} ($i = 1, \dots, r$) with the existing $k_1^v - 1$ excursions from v along e_1^v . □

This exact enumeration can be used to generate the following general bounds on the number of tours on trees:

Corollary 5.3. *If T is a tree, then for any vertex $u \in V(T)$, $W_T^B((k_i); u)$ satisfies the following bounds:*

$$\prod_{v \in V(T)} \binom{\Psi(T, v) - d(v)}{k_1^v - 1, k_2^v - 1, \dots, k_{d(v)}^v - 1} \leq W_T^B((k_i); u) \leq \prod_{v \in V(T)} \binom{\Psi(T, v)}{k_1^v, k_2^v, \dots, k_{d(v)}^v}.$$

Proof. If all the k_i are positive, then this follows directly from Lemma 5.2.

If one or more of the k_i is zero, then the lower bound is trivially true, because one of the multinomial coefficients is zero. The upper bound also holds trivially if there are no tours in the family. Otherwise, let T^+ be the subtree of T induced by the vertices actually visited by tours in $W_T^B((k_i); u)$. Then $W_T^B((k_i); u) = W_{T^+}^B((k_i); u)$. But we know that

$$W_{T^+}^B((k_i); u) \leq \prod_{v \in V(T^+)} \binom{\Psi(T^+, v)}{k_1^v, k_2^v, \dots, k_{d(v)}^v} = \prod_{v \in V(T)} \binom{\Psi(T, v)}{k_1^v, k_2^v, \dots, k_{d(v)}^v}$$

as a result of Lemma 5.2 and the fact that $k_i^v = 0$ for all edges e_i^v incident to unvisited vertices $v \in V(T) \setminus V(T^+)$. \square

5.2.3 Treeification

In order to establish the lower and upper bounds for tours on arbitrary connected graphs, we relate tours on a connected graph G to (balanced) tours on a related tree which we call a *treeification* of G . The process of treeification consists of repeatedly breaking cycles until the resulting graph is acyclic. This creates a sequence of graphs $G = G_0, G_1, \dots, G_t = T$ where T is a tree. We call this sequence a *treeification sequence*.

Formally, we define a treeification of a connected graph to be the result of the following (nondeterministic) process that transforms a connected graph into a tree with the same number of edges.




Figure 5.1: Splitting vertex v

To treeify a connected graph $G = G_0$, we first give an (arbitrary) order to its vertices. Then we apply the following vertex-splitting operation in turn to each G_j to create G_{j+1} ($j = 0, 1, \dots$), until no cycles remain:

1. Let v be the first vertex (in the ordering) that occurs in some cycle C of G_j .
2. Split vertex v by doing the following (see Figure 5.1):
 - (a) Delete an edge xv from $E(C)$ (there are two choices for vertex x).
 - (b) Add a new vertex v' (to the end of the vertex ordering).
 - (c) Add the pendant edge xv' (making v' a leaf).




Figure 5.2: A treeification sequence; numbers show the first few vertices in the ordering

Note that if a vertex v is split multiple times when treeifying a graph G , these splits occur contiguously (because of the ordering placed on the vertices of G). Thus, if v is split r times, there is a contiguous subsequence $G_j, G_{j+1}, \dots, G_{j+r}$ of the treeification sequence that corresponds to the splitting of v . See Figure 5.2 for an example of a treeification sequence.

There is a natural way to establish a relationship between tours on different graphs in a treeification sequence G_0, \dots, G_t . The treeification process induces *graph homomorphisms* (edge preserving maps) between the graphs in such a sequence. For all $i < j$, there is a *surjective* homomorphism from G_j onto G_i . This homomorphism is also *locally injective* since it maps neighbourhoods of G_j injectively into neighbourhoods of G_i . A locally injective map such as this is also known as a *partial cover*. In particular, for each $j < t$, there is a partial cover of G_{j+1} onto G_j that maps the new pendant edge xv' to the edge xv that it replaces. These homomorphisms impart a natural correspondence between families of tours on different graphs in the treeification sequence, which we employ later to determine our bounds.

Although the concept of treeification is a very natural one, these particular “partial covering trees” do not appear to have been studied before; their only previous use seems to be by Yarkony, Fowlkes and Ihler to address a problem in computer vision [174]. For a general introduction to graph homomorphisms, see the monograph by Hell & Nešetřil [97]. For more on partial maps and other locally constrained graph homomorphisms, see the survey article by Fiala & Kratochvíl [74].

If we have a treeification sequence $G = G_0, \dots, G_t = T$ for a connected graph G , we can use a three-stage process to establish a lower or upper bound for a family of tours on G . In

the first stage (“splitting once”), we relate the number of tours in the family on G_j ($j < t$) to the number of tours in a related family on G_{j+1} . In the second stage (“fully splitting one vertex”), for a vertex v , we consider the subsequence G_j, \dots, G_{j+r} that corresponds to the splitting of v and, iterating the inequality from the first stage, relate the number of tours on G_j to the number of tours on G_{j+r} . Finally (“fully splitting all vertices”), iterating the inequality from the second stage, we relate the number of tours on $G = G_0$ to the number of tours on $G_t = T$, and employ the bounds on tours on T from Corollary 5.3 to determine the bound for the family of tours on G .

In Subsection 5.2.4, we use this three-stage process to produce a lower bound on the value of $W_G^B((k_i); u)$. Then, in Subsection 5.2.5, we use the same three-stage process to establish an upper bound on $W_G((h_i); u)$.

5.2.4 The lower bound

Our lower bound is on the number of *balanced* tours. We only consider the families in which *every* edge is traversed at least once in each direction. On a connected graph, these families are never empty.

Lemma 5.4. *If G is a connected graph with m edges and k_1, \dots, k_m are all positive, then for any vertex $u \in V(G)$,*

$$W_G^B(k_1, k_2, \dots, k_m; u) \geq \prod_{v \in V(G)} \binom{k_1^v + k_2^v + \dots + k_{d(v)}^v - d(v)}{k_1^v - 1, k_2^v - 1, \dots, k_{d(v)}^v - 1}.$$

This lower bound does not hold in general for a *disconnected* graph since there are no tours possible if there is any positive k_i in a component not containing u .

Proof. Let T be some treeification of G with treeification sequence $G = G_0, \dots, G_t = T$ in which vertex u is never split. (This is possible by positioning u last in the ordering on the vertices.)

By exhibiting a surjection from $W_G^B((k_i); u)$ onto $W_T^B((k_i); u)$ that is consistent with the homomorphism from T onto G induced by the treeification process, we determine an inequality relating the number of tours in the two families.

I. Splitting once

Our first stage is to associate a number of tours on G_j , in $W_{G_j}^B((k_i); u)$, to each tour on G_{j+1} , in $W_{G_{j+1}}^B((k_i); u)$.

To simplify the notation, let $H_0 = G_j$ and $H = G_{j+1}$ for some $j < t$. Let v be the vertex of H_0 that is split in H , and let v' be the leaf vertex in H added when splitting v . Let e_1 be the (only) edge incident to v' in H ; we also use e_1 to refer to the corresponding edge (incident to v) in H_0 (see Figure 5.3).




Figure 5.3: Tours on H_0 corresponding to a tour on H

Any tour in $\mathcal{W}_H^{\mathbf{B}}((k_i); \mathbf{u})$ visits vertex v exactly $\Psi(H, v)$ times and visits vertex v' (along e_1) k_1 times. The corresponding tour on H_0 visits v exactly $\Psi(H_0, v) = \Psi(H, v) + k_1$ times. Of these visits there are k_1 which arrive along e_1 and then depart along e_1 .

Since $\Psi(H_0, v)$ is positive, separating the visits are $\Psi(H_0, v) - 1$ excursions from v . Depending on whether the final visit to v departs along e_1 or not, either $k_1 - 1$ or k_1 of these excursions begin with a traversal of e_1 ; these are interleaved with the other $\Psi(H, v)$ or $\Psi(H, v) - 1$ excursions which begin with a traversal of some other edge.

Changing the interleaving of these two sets of excursions (without altering their internal ordering) produces at least

$$\min \left[\binom{\Psi(H_0, v) - 1}{k_1 - 1}, \binom{\Psi(H_0, v) - 1}{k_1} \right] \geq \binom{\Psi(H_0, v) - 2}{k_1 - 1}$$

distinct tours in $\mathcal{W}_{H_0}^{\mathbf{B}}((k_i); \mathbf{u})$.

Note that there is only one interleaving of the sets of excursions that corresponds to a valid tour in $\mathcal{W}_H^{\mathbf{B}}((k_i); \mathbf{u})$: the one in which the excursions beginning with a traversal of e_1 away from v are arranged so they occur immediately following a traversal of e_1 towards v .

Hence we can deduce that

$$\mathcal{W}_{H_0}^{\mathbf{B}}((k_i); \mathbf{u}) \geq \binom{\Psi(H_0, v) - 2}{k_1 - 1} \mathcal{W}_H^{\mathbf{B}}((k_i); \mathbf{u}). \quad (1)$$

II. Fully splitting one vertex

For a given vertex v , let H_0, H_1, \dots, H_r be the subsequence of graphs that corresponds to the splitting of v . In our second stage, we relate the number of tours on H_0 to the number of tours on H_r .

Note that $\Psi(H_0, v) = \Psi(G, v)$ and $\Psi(H_r, v) = \Psi(T, v)$ since the splitting of other vertices cannot affect the number of visits to v .

Let e_1, \dots, e_r be the new pendant edges in H_r , and hence also in T , added when v is split, and let e_1, \dots, e_r also denote the corresponding edges in G . Then $\Psi(H_{i-1}, v) = \Psi(H_i, v) + k_i$ for $1 \leq i \leq r$, and thus $\Psi(H_{i-1}, v) = \Psi(T, v) + k_i + k_{i+1} + \dots + k_r$, and in particular $\Psi(G, v) = \Psi(T, v) + k_1 + \dots + k_r$.

Hence, by iterating inequality (1),

$$\begin{aligned}
W_{H_0}^{\mathbf{B}}((k_i); \mathbf{u}) &\geq \prod_{i=1}^r \binom{\Psi(H_{i-1}, v) - 2}{k_i - 1} W_{H_r}^{\mathbf{B}}((k_i); \mathbf{u}) \\
&= \prod_{i=1}^r \binom{\Psi(T, v) + (\sum_{j=i}^r k_j) - 2}{k_i - 1} W_{H_r}^{\mathbf{B}}((k_i); \mathbf{u}) \\
&\geq \prod_{i=1}^r \binom{\Psi(T, v) + (\sum_{j=i}^r (k_j - 1)) - 1}{k_i - 1} W_{H_r}^{\mathbf{B}}((k_i); \mathbf{u}) \\
&= \binom{\Psi(G, v) - (r + 1)}{\Psi(T, v) - 1, k_1 - 1, k_2 - 1, \dots, k_r - 1} W_{H_r}^{\mathbf{B}}((k_i); \mathbf{u}). \quad (2)
\end{aligned}$$

III. Fully splitting all vertices

Finally, our third stage is to relate the number of tours on G to the number of tours on T and then apply the tree bounds to establish the required lower bound.

For each $v \in V(G)$, let $r(v)$ be the number of times v is split. Note that $r(v)$ is less than the degree of v in G since $d_G(v) = d_T(v) + r(v)$.

Thus, with a suitable indexing of the edges around each vertex, if we iterate inequality (2) and combine with the lower bound on $W_T^{\mathbf{B}}((k_i); \mathbf{u})$ from Corollary 5.3, we get

$$\begin{aligned}
W_G^{\mathbf{B}}((k_i); \mathbf{u}) &\geq \prod_{v \in V(G)} \binom{\Psi(G, v) - (r(v) + 1)}{\Psi(T, v) - 1, k_1^v - 1, \dots, k_{r(v)}^v - 1} W_T^{\mathbf{B}}((k_i); \mathbf{u}) \\
&\geq \prod_{v \in V(G)} \binom{\Psi(G, v) - (r(v) + 1)}{\Psi(T, v) - 1, k_1^v - 1, \dots, k_{r(v)}^v - 1} \binom{\Psi(T, v) - d_T(v)}{k_{r(v)+1}^v - 1, \dots, k_{d_G(v)}^v - 1} \\
&\geq \prod_{v \in V(G)} \binom{\Psi(G, v) - d_G(v)}{k_1^v - 1, k_2^v - 1, \dots, k_{d_G(v)}^v - 1}
\end{aligned}$$

concluding the proof of Lemma 5.4. \square

5.2.5 The upper bound

Our upper bound applies to arbitrary families of tours $\mathcal{W}_G((h_i); u)$, without any restriction on the values of the h_i . Subsequently, we apply this result to families of tours of even length.

Lemma 5.5. *If G is a connected graph with m edges and u is any vertex of G , then*

$$\mathcal{W}_G(h_1, h_2, \dots, h_m; u) \leq (h + 2m)^m \prod_{v \in V(G)} \binom{k_1^v + k_2^v + \dots + k_{d(v)}^v}{k_1^v, k_2^v, \dots, k_{d(v)}^v}$$

for some $k_i \in [\frac{1}{2}h_i, \frac{1}{2}h_i + m]$ ($1 \leq i \leq m$), where $h = h_1 + \dots + h_m$ is the length of the tours in the family and $k_1^v, k_2^v, \dots, k_{d(v)}^v$ are the k_i corresponding to edges incident to v .

Proof. Let T be some treeification of G with treeification sequence $G = G_0, \dots, G_t = T$ in which vertex u is never split. (This is possible by positioning u last in the ordering on the vertices.)

We relate the number of (arbitrary) tours in $\mathcal{W}_G((h_i); u)$ to the number of (balanced) tours in $\mathcal{W}_T^B((k_i); u)$, for some k_i not much greater than $\frac{1}{2}h_i$. This is achieved by exhibiting a surjection from $\mathcal{W}_T^B((k_i); u)$ onto $\mathcal{W}_G((h_i); u)$ that is consistent with the homomorphism from T onto G induced by the treeification process.

The proof is broken down into the same three stages as for the proof of the lower bound. Initially, we restrict ourselves to the case in which all the $\Psi(G, v)$ are positive. The case of unvisited vertices is addressed in an additional stage at the end.

I. Splitting once

Our first stage is to associate to each tour on G_j a number of tours on G_{j+1} . However, unlike in the proof of the lower bound, the relationship is not between classes with the same parameterisation. Rather, we relate tours in $\mathcal{W}_{G_j}((h_i); u)$ to slightly longer tours in $\mathcal{W}_{G_{j+1}}((h'_i); u)$, for some h'_i such that, for each i , $h_i \leq h'_i \leq h_i + 2$.

As we did for the lower bound, let $H_0 = G_j$ and $H = G_{j+1}$ for some $j < t$. Let v be the vertex of H_0 that is split in H , and let v' be the leaf vertex in H added when splitting v .

Again, let e_1 be the (only) edge incident to v' in H ; we also use e_1 to refer to the corresponding edge (incident to v) in H_0 .

Let C be some cycle in H_0 containing e_1 , and let e_2 be the other edge on C that is incident to v (in both H_0 and H).



Figure 5.4: Tours on H corresponding to a tour on H_0 ; $k_1 = \lfloor \frac{1}{2}h_1 \rfloor + 1$

Given a tour on H_0 , we want to modify it so that the result is a valid tour on H . For a tour on H_0 to be valid on H , each traversal of e_1 towards v must be immediately followed by a traversal of e_1 from v . See Figure 5.4.

To achieve this, we make three kinds of changes to excursions from v :

1. Reverse the direction of some of the excursions.
2. Add one or two additional excursions (around C).
3. Modify the interleaving of excursions.

To manage the details, given a tour on H_0 , we consider the $\Psi(H_0, v) - 1$ excursions from v to be partitioned into subsets as follows:

- *-*: a_0 excursions that don't traverse e_1 at all
- 1-*: a_1 excursions that begin but don't end with a traversal of e_1
- *-1: a_2 excursions that end but don't begin with a traversal of e_1
- 1-1: a_3 excursions that both begin and end with traversals of e_1

We also refer to 1-* and 1-1 excursions as *1-initial*, and *-* and *-1 excursions as **-initial*.

We refer to the edge traversed in arriving for the first visit to v as the *arrival edge* and to the edge traversed in departing from the last visit to v as the *departure edge*. We call their traversals the *arrival* and the *departure* respectively. To account for these, we define a_1^+ to be $a_1 + 1$ if the departure edge is e_1 and to be a_1 otherwise, and define a_2^+ to be $a_2 + 1$ if the arrival edge is e_1 and to be a_2 otherwise.

So, to transform a tour on H_0 into one on H , we perform the following three steps:

1. If $a_2^+ > a_1 + 1$, reverse the direction of the last $\lfloor \frac{1}{2}(a_2^+ - a_1) \rfloor$ of the *-1 excursions (making them 1-*).
On the other hand, if $a_2^+ < a_1$, reverse the direction of the last $\lfloor \frac{1}{2}(a_1 - a_2^+) \rfloor$ of the 1-* excursions (making them *-1).
Update the values of a_1 and a_2 to reflect these reversals; we now have $a_2^+ = a_1$ or $a_2^+ = a_1 + 1$.

2. If $a_1^+ + a_2^+$ is even (h_1 is even) or $a_1^+ = a_1$ (the departure edge *isn't* e_1), add a new 1-* excursion consisting of a tour around the cycle C (returning to v along e_2); this should be added following all the existing excursions.

Also, if $a_1^+ + a_2^+$ is even (h_1 is even) or $a_1^+ = a_1 + 1$ (the departure edge *is* e_1), add a new *-1 excursion consisting of a tour around the cycle C (departing from v along e_2); this should be added following all the existing excursions.

Update the values of a_1 and a_2 to reflect the presence of the new excursion(s); we now have $a_2^+ = a_1^+$.

3. Change the interleaving of the 1-initial excursions with the *-initial excursions so that each visit to v along e_1 returns immediately along e_1 . This is always possible (see below) and there is only one way of doing it. We now have a valid tour on H .

	-1	3-2	3-1	1-1	3-1	2-1	2-1	3-2	1-3	3-1	2-2	1-1		2-
Step 1	-1	3-2	3-1	1-1	3-1	2-1	<i>1-2</i>	3-2	1-3	1-3	2-2	1-1		2-
Step 2	-1	3-2	3-1	1-1	3-1	2-1	1-2	3-2	1-3	1-3	2-2	1-1	1-2	2-
Step 3	-1	1-1	1-2	3-2	3-1	1-3	3-1	1-3	2-1	1-1	1-2	3-2	2-2	2-

	-3	1-1	1-3	1-3	3-1	1-2	3-3	1-2	2-1	1-1	3-1	2-3	1-1			1-
Step 1	-3	1-1	1-3	1-3	3-1	1-2	3-3	<i>2-1</i>	2-1	1-1	3-1	2-3	1-1			1-
Step 2	-3	1-1	1-3	1-3	3-1	1-2	3-3	2-1	2-1	1-1	3-1	2-3	1-1	1-2	2-1	1-
Step 3	-3	3-1	1-1	1-3	3-3	2-1	1-3	2-1	1-2	3-1	1-1	1-1	1-2	2-3	2-1	1-

Figure 5.5: Two examples of transforming tours by modifying excursions

Figure 5.5 shows two examples of this process. The two-digit entries in the table represent the initial and final edges traversed by excursions from v ; the single-digit entries give the arrival and departure edges; e_3 is an additional edge incident to v . 1-initial excursions (whose interleaving with the *-initial excursions is modified by Step 3) are shown in bold. In Step 1, excursions which are reversed are shown in italics.

Validation of Step 3

If we consider the 1-initial excursions and the *-initial excursions as two separate lists, with the *-initial excursions (together with the arrival and departure) as “fixed”, then we can insert 1-initial excursions into the list of *-initial excursions as follows:

Following each *-1 excursion (and the arrival if it is along e_1), place the next unused 1-* excursion together with any unused 1-1 excursions that precede it.

This procedure is successful, and ensures that each visit to v along e_1 returns immediately along e_1 as long as the number of traversals of e_1 towards v equals the number of traversals of e_1 away from v , unless either

- the departure edge is *not* e_1 and the last 1-initial excursion is a 1-1 excursion (the minimal example being -2 1-1 2-, using the notation of Figure 5.5), or
- the departure edge *is* e_1 and the last *-initial excursion is a ** excursion (the minimal example being -1 2-2 1-).

The rules controlling the addition of new final *-1 and 1-* excursions in Step 2 guarantee both that the number of traversals of e_1 towards v is the same as the number of traversals of e_1 away from v , and also that neither of the two exceptional cases occur. Thus Step 3 is always valid.

Counting

Step 2 can add at most two additional excursions from v (around C), so given a tour in $\mathcal{W}_{H_0}((h_i); u)$, this process produces a tour in $\mathcal{W}_H(2k_1, h'_2, \dots, h'_m; u)$ where $k_1 = \lfloor \frac{1}{2}h_1 \rfloor + 1$, and for each i , $h_i \leq h'_i \leq h_i + 2$.

After completing Step 1, there are $a_1 + a_2 + 1$ ways in which it could be undone (reverse no more than a_1 1-* excursions, reverse no more than a_2 *-1 excursions, or do nothing). Since $h_1 = a_1 + a_2 + 2a_3$, this does not exceed $h_1 + 1$.

Also, after Step 3, there are either k_1 or $k_1 - 1$ excursions that begin with a traversal of e_1 that could, prior to the step, have been arbitrarily interleaved with those that don't.

Thus we see that there are no more than

$$(h_1 + 1) \max \left[\binom{\Psi(H, v) + k_1 - 1}{k_1}, \binom{\Psi(H, v) + k_1 - 1}{k_1 - 1} \right] \leq 2k_1 \binom{\Psi(H, v) + k_1}{k_1}$$

distinct tours in $\mathcal{W}_{H_0}((h_i); u)$ that generate any specific tour in $\mathcal{W}_H(2k_1, h'_2, \dots, h'_m; u)$.

Hence,

$$\mathcal{W}_{H_0}((h_i); u) \leq 2k_1 \binom{\Psi(H, v) + k_1}{k_1} \mathcal{W}_H(2k_1, h'_2, \dots, h'_m; u). \quad (3)$$

Note also that either $\Psi(H_0, v) = \Psi(H, v) + k_1 - 2$ or $\Psi(H_0, v) = \Psi(H, v) + k_1 - 1$ (depending on whether h_1 is even or odd), and so

$$\Psi(H_0, v) < \Psi(H, v) + k_1. \quad (4)$$

Furthermore, $\Psi(H, v)$ is positive, since the additional excursion(s) ensure that h'_2 is positive.

II. Fully splitting one vertex

For a given vertex v , let H_0, H_1, \dots, H_r be the subsequence of graphs that corresponds to the splitting of v . In the second stage of our proof, we relate the number of tours on H_0 to the number of tours on H_r .

Note again that $\Psi(H_0, v) = \Psi(G, v)$ and $\Psi(H_r, v) = \Psi(T, v)$ since the splitting of other vertices cannot affect the number of visits to v .

We assume that $\Psi(G, v)$ is positive, and hence that $\Psi(H_0, v), \dots, \Psi(H_r, v) = \Psi(T, v)$ are all positive too.

Let e_1, \dots, e_r be the new pendant edges in T added when v is split, and let e_1, \dots, e_r also denote the corresponding edges in G . Then, by (4), for some k_1, \dots, k_r such that $\frac{1}{2}h_i \leq k_i \leq \frac{1}{2}h_i + i$, we have $\Psi(H_{i-1}, v) < \Psi(H_i, v) + k_i$, and thus

$$\Psi(H_{i-1}, v) < \Psi(T, v) + k_i + \dots + k_r.$$

Hence, by iterating inequality (3), if $h'_i = 2k_i$ for $1 \leq i \leq r$, then for some h'_{r+1}, \dots, h'_m such that $h_i \leq h'_i \leq h_i + 2r$,

$$\begin{aligned} W_{H_0}((h_i); u) &\leq 2^r \left(\prod_{i=1}^r k_i \binom{\Psi(H_i, v) + k_i}{k_i} \right) W_{H_r}((h'_i); u) \\ &< 2^r \left(\prod_{i=1}^r k_i \binom{\Psi(T, v) + \sum_{j=i}^r k_j}{k_i} \right) W_{H_r}((h'_i); u) \\ &= 2^r \left(\prod_{i=1}^r k_i \right) \binom{\Psi(T, v) + \sum_{i=1}^r k_i}{\Psi(T, v), k_1, \dots, k_r} W_{H_r}((h'_i); u). \end{aligned} \quad (5)$$

III. Fully splitting all vertices

In the third stage of the proof, we relate the number of tours on G to the number of tours on T and then apply the tree bounds to establish the required upper bound for the case in which all the $\Psi(G, v)$ are positive.

For each $v \in V(G)$, let $r(v)$ be the number of times v is split. Also, let $h = h_1 + \dots + h_m$ be the length of the tours in $\mathcal{W}_G((h_i); u)$.

Thus, with a suitable indexing of the edges around each vertex, if we iterate inequality (5) and combine with the upper bound on $W_T^B((k_i); u)$ from Corollary 5.3, we get, for some

k_1, \dots, k_m such that $\frac{1}{2}h_i \leq k_i \leq \frac{1}{2}h_i + m$,

$$\begin{aligned} W_G((h_i); \mathbf{u}) &\leq 2^m \left(\prod_{i=1}^m k_i \right) \prod_{v \in V(G)} \binom{\Psi(T, v) + \sum_{i=1}^{r(v)} k_i^v}{\Psi(T, v), k_1^v, \dots, k_{r(v)}^v} W_T^{\mathbf{B}}((k_i); \mathbf{u}) \\ &\leq (h + 2m)^m \prod_{v \in V(G)} \binom{\Psi(T, v) + \sum_{i=1}^{r(v)} k_i^v}{\Psi(T, v), k_1^v, \dots, k_{r(v)}^v} \binom{\Psi(T, v)}{k_{r(v)+1}^v, \dots, k_{d_G(v)}^v} \\ &= (h + 2m)^m \prod_{v \in V(G)} \binom{k_1^v + \dots + k_{d(v)}^v}{k_1^v, \dots, k_{d(v)}^v}, \end{aligned}$$

using the fact that for each i , we have $k_i \leq \frac{1}{2}h + m$.

IV. Unvisited vertices

Thus we have the desired result for the case in which all the $\Psi(G, v)$ are positive. To complete the proof, we consider families of tours in which some of the vertices are not visited.

If not all the $\Psi(G, v)$ are positive, then let G^+ be the subgraph of G induced by the vertices actually visited by tours in $\mathcal{W}_G((h_i); \mathbf{u})$. Then $W_G((h_i); \mathbf{u}) = W_{G^+}((h_i); \mathbf{u})$. But we know that

$$\begin{aligned} W_{G^+}((h_i); \mathbf{u}) &\leq (h + 2m)^m \prod_{v \in V(G^+)} \binom{k_1^v + \dots + k_{d(v)}^v}{k_1^v, \dots, k_{d(v)}^v} \\ &\leq (h + 2m)^m \prod_{v \in V(G)} \binom{k_1^v + \dots + k_{d(v)}^v}{k_1^v, \dots, k_{d(v)}^v} \end{aligned}$$

because the inclusion of the unvisited vertices in $V(G) \setminus V(G^+)$ cannot decrease the value of the product. So the bound holds for any family $\mathcal{W}_G((h_i); \mathbf{u})$.

This concludes the proof of Lemma 5.5. □

5.2.6 Tours of even length

In this subsection, we consider the family of all tours of *even length* on a graph and prove that it grows at the same rate as the more restricted family of all balanced tours.

To do this, we make use of the fact that the growth rate of a collection of objects does not change if we make “small” changes to what we are counting:

Observation 5.6. *If \mathcal{S} is a collection of objects, containing S_k objects of each size k , that has a finite growth rate, then for any positive polynomial P and fixed non-negative integers d_1, d_2 with $d_1 \leq d_2$,*

$$\lim_{k \rightarrow \infty} \left(P(k) \sum_{j=k+d_1}^{k+d_2} S_j \right)^{1/k} = \lim_{k \rightarrow \infty} S_k^{1/k} = \text{gr}(\mathcal{S}).$$

This follows directly from the definition of the growth rate and is a generalisation of an observation made in Section 1.1. We also use this observation when we consider the relationship between permutation grid classes and families of tours on graphs in the next section.

We can employ our upper bound for $W_G((h_i); u)$ to give us an upper bound for tours of a specific even length. We use $W_G(h; u)$ to denote the number of tours of length h starting and ending at vertex u .

Lemma 5.7. *If G is a connected graph with m edges and u is any vertex of G , then the number of tours of length $2k$ on G starting and ending at vertex u is bounded above as follows:*

$$W_G(2k; u) \leq (m+1)^m (2k+2m)^m \sum_{j=k}^{k+m^2} \sum_{k_1+\dots+k_m=j} \prod_{v \in V(G)} \binom{k_1^v + \dots + k_{d(v)}^v}{k_1^v, \dots, k_{d(v)}^v}.$$

Proof. From Lemma 5.5, for any vertex u of a graph G with m edges, we know that

$$\begin{aligned} W_G(2k; u) &= \sum_{h_1+\dots+h_m=2k} W_G((h_i); u) \\ &\leq (2k+2m)^m \sum_{h_1+\dots+h_m=2k} \prod_{v \in V(G)} \binom{k_1^v + k_2^v + \dots + k_{d(v)}^v}{k_1^v, k_2^v, \dots, k_{d(v)}^v} \end{aligned}$$

where each k_i is dependent on the sequence (h_i) with $\frac{1}{2}h_i \leq k_i \leq \frac{1}{2}h_i + m$.

There are no more than $(m+1)^m$ different values of the h_i that give rise to any specific set of k_i , and we have $k \leq k_1 + \dots + k_m \leq k + m^2$, so

$$W_G(2k; u) \leq (m+1)^m (2k+2m)^m \sum_{j=k}^{k+m^2} \sum_{k_1+\dots+k_m=j} \prod_{v \in V(G)} \binom{k_1^v + \dots + k_{d(v)}^v}{k_1^v, \dots, k_{d(v)}^v}. \quad \square$$

Now, drawing together our upper and lower bounds enables us to deduce that the family of balanced tours on a graph G grows at the same rate as the family of all tours of even length on G . We use W_G^B for the family of all balanced tours on G and W_G^E for the family of all tours of even length on G , where, in both cases, we consider the size of a tour to be *half* its length.

Theorem 5.8. *The growth rate of the family of balanced tours ($\mathcal{W}_G^{\mathbf{B}}$) on a connected graph is the same as growth rate of the family of all tours of even length ($\mathcal{W}_G^{\mathbf{E}}$) on the graph.*

Proof. From Lemma 5.4, we know that

$$\prod_{v \in V(G)} \binom{k_1^v + \dots + k_{d(v)}^v}{k_1^v, \dots, k_{d(v)}^v} \leq W_G^{\mathbf{B}}(k_1+1, \dots, k_m+1; u).$$

Substitution in the inequality in the statement of Lemma 5.7 then yields the following relationship between families of even-length and balanced tours:

$$W_G(2k; u) \leq (m+1)^m (2k+2m)^m \sum_{j=k+m}^{k+m+m^2} W_G^{\mathbf{B}}(j; u)$$

where $W_G^{\mathbf{B}}(j; u)$ is the number of balanced tours of length $2j$ on G starting and ending at u . Combining this with Observation 5.6 and the fact that $W_G^{\mathbf{B}}(k; u) \leq W_G(2k; u)$ produces the result $\text{gr}(\mathcal{W}_G^{\mathbf{B}}) = \text{gr}(\mathcal{W}_G^{\mathbf{E}})$. \square

Finally, before moving on to the relationship with permutation grid classes, we determine the value of the growth rate of the family of even-length tours $\mathcal{W}_G^{\mathbf{E}}$. This requires only elementary algebraic graph theory. We recall here the relevant concepts. The *adjacency matrix*, $A = A(G)$ of a graph G has rows and columns indexed by the vertices of G , with $A_{i,j} = 1$ or $A_{i,j} = 0$ according to whether vertices i and j are adjacent (joined by an edge) or not. The *spectral radius* $\rho(G)$ of a graph G is the largest eigenvalue (which is real and positive) of its adjacency matrix.

Lemma 5.9. *The growth rate of $\mathcal{W}_G^{\mathbf{E}}$ exists and is equal to the square of the spectral radius of G .*

Proof. If G has n vertices, then

$$W_G(2k) = \sum_{u \in V(G)} W_G(2k; u) = \text{tr}(A(G)^{2k}) = \sum_{i=1}^n \lambda_i^{2k},$$

where the λ_i are the (real) eigenvalues of $A(G)$, the adjacency matrix of G , since the diagonal entries of $A(G)^{2k}$ count the number of tours of length $2k$ starting at each vertex. Thus,

$$\text{gr}(\mathcal{W}_G^{\mathbf{E}}) = \lim_{k \rightarrow \infty} \left(\sum_{i=1}^n \lambda_i^{2k} \right)^{1/k}$$

Now the spectral radius is given by $\rho = \rho(G) = \max_{1 \leq i \leq n} \lambda_i$, so we can conclude that

$$\rho^2 = \lim_{k \rightarrow \infty} (\rho^{2k})^{1/k} \leq \lim_{k \rightarrow \infty} \left(\sum_{i=1}^n \lambda_i^{2k} \right)^{1/k} \leq \lim_{k \rightarrow \infty} ((n\rho)^{2k})^{1/k} = \rho^2.$$

Thus, $\text{gr}(\mathcal{W}_G^E) = \rho(G)^2$. □

5.3 Grid classes

In this section, we prove our main theorem, that the growth rate of a monotone grid class of permutations is equal to the square of the spectral radius of its row-column graph.

The proof is as follows: First, we present an explicit expression for the number of *gridded permutations* of a given length. Then, we use this to show that the class of gridded permutations grows at the same rate as the family of tours of even length on its row-column graph. Finally, we utilize the fact that the growth rate of a grid class is the same as the growth rate of the corresponding class of gridded permutations.

5.3.1 Counting gridded permutations

As we saw in the previous chapter, it is possible to give an explicit expression for the number of gridded permutations of length k in any specified grid class. This is essentially the same as Lemma 4.1, but using different notation. We repeat the proof for completeness. Observe the similarity to the formulae for numbers of tours.

Lemma 5.10. *If $G = G(M)$ is the row-column graph of $\text{Grid}(M)$, and G has m edges e_1, \dots, e_m , then the number of gridded permutations of length k in $\text{Grid}^\#(M)$ is given by*

$$|\text{Grid}_k^\#(M)| = \sum_{k_1 + \dots + k_m = k} \prod_{v \in V(G)} \binom{k_1^v + k_2^v + \dots + k_{d(v)}^v}{k_1^v, k_2^v, \dots, k_{d(v)}^v}$$

where $k_1^v, k_2^v, \dots, k_{d(v)}^v$ are the k_i corresponding to edges incident to v in G .

Proof. A gridded permutation in $\text{Grid}^\#(M)$ consists of a number of points in each of the cells that correspond to a non-zero entry of M . For every permutation, the relative ordering of points (increasing or decreasing) within a particular cell is fixed by the value of the corresponding matrix entry. However, the relative interleaving between points in distinct cells in the same row or column can be chosen arbitrarily and independently for each row and column.

Now, each vertex in G corresponds to either a row or a column in M , with an incident edge for each non-zero entry in that row or column. Thus, the number of gridded permutations with k_i points in the cell corresponding to edge e_i for each i is given by the following product of multinomial coefficients:

$$\prod_{v \in V(G)} \binom{k_1^v + k_2^v + \dots + k_{d(v)}^v}{k_1^v, k_2^v, \dots, k_{d(v)}^v}.$$

The result follows by summing over values of k_i that sum to k . \square

As an immediate consequence, we have the fact that the enumeration of a class of gridded permutations depends only on its row-column graph:

Corollary 5.11. *If $G(M) = G(M')$, then $\text{Grid}_k^\#(M) = \text{Grid}_k^\#(M')$ for all k .*

5.3.2 Gridded permutations and tours

We now use Lemmas 5.4 and 5.7 to relate the number of gridded permutations of length k in $\text{Grid}^\#(M)$ to the number of tours of length $2k$ on $G(M)$. We restrict ourselves to permutation classes with connected row-column graphs.

Lemma 5.12. *If $G(M)$ is connected, the growth rate of $\text{Grid}^\#(M)$ exists and is equal to the growth rate of $W_{G(M)}^E$.*

Proof. If matrix M has m non-zero entries (and thus $G(M)$ has m edges), then for any vertex u of $G(M)$, combining Lemmas 5.10 and 5.4, gives us

$$\begin{aligned} |\text{Grid}_k^\#(M)| &\leq \sum_{k_1 + \dots + k_m = k} W_{G(M)}^B(k_1 + 1, k_2 + 1, \dots, k_m + 1; u) \\ &\leq \sum_{k_1 + \dots + k_m = k} W_{G(M)}(2k_1 + 2, 2k_2 + 2, \dots, 2k_m + 2; u) \\ &\leq W_{G(M)}(2k + 2m). \end{aligned} \tag{6}$$

On the other hand, from Lemma 5.7, for any vertex u of a graph G with m edges,

$$W_G(2k; u) \leq (m + 1)^m (2k + 2m)^m \sum_{j=k}^{k+m^2} \sum_{k_1 + \dots + k_m = j} \prod_{v \in V(G)} \binom{k_1^v + \dots + k_{d(v)}^v}{k_1^v, \dots, k_{d(v)}^v}.$$

Let $W_G(h)$ be the number of tours of length h on G (starting at any vertex).

Now $W_G(\mathbf{h}) = \sum_{u \in V(G)} W_G(\mathbf{h}; u)$, so, using Lemma 5.10, if $G(M)$ has n vertices and m edges, we have

$$W_{G(M)}(2k) \leq n(m+1)^m (2k+2m)^m \sum_{j=k}^{k+m^2} |\text{Grid}_j^\#(M)|.$$

The multiplier on the right side of this inequality is a polynomial in k . Hence, using inequality (6) and Observation 5.6, we can conclude that

$$\text{gr}(\text{Grid}^\#(M)) = \text{gr}(W_{G(M)}^E)$$

if $G(M)$ is connected. □

5.3.3 Counting permutations

We nearly have the result we want. The final link is Lemma 2.1 which tells us that, as far as growth rates are concerned, classes of gridded permutations are indistinguishable from their grid classes. Thus, by Corollary 5.11:

Corollary 5.13. *Monotone grid classes with the same row-column graph have the same growth rate.*

5.3.4 The growth rate of grid classes

We now have all we need for the proof of our main theorem.

Theorem 5.14. *The growth rate of a monotone grid class of permutations exists and is equal to the square of the spectral radius of its row-column graph.*

Proof. For connected grid classes, the result follows immediately from Lemmas 5.9, 5.12 and 2.1. A little more work is required to handle the disconnected case.

If $G(M)$ is disconnected, then the growth rate of $\text{Grid}(M)$ is the maximum of the growth rates of the grid classes corresponding to the connected components of $G(M)$ (see Proposition 2.10 in Vatter [159]).

Similarly, the spectrum of a disconnected graph is the union (with multiplicities) of the spectra of the graph's connected components (see Theorem 2.1.1 in Cvetković, Rowlinson & Simić [68]). Thus the spectral radius of a disconnected graph is the maximum of the spectral radii of its components.

Combining these facts with Lemmas 5.9, 5.12 and 2.1 yields

$$\text{gr}(\text{Grid}(M)) = \rho(G(M))^2$$

as required. \square

5.4 Implications

As a consequence of Theorem 5.14, results concerning the spectral radius of graphs can be translated into facts about the growth rates of permutation grid classes. So we now present a number of corollaries that follow from spectral graph theoretic considerations. The two recent monographs by Cvetković, Rowlinson & Simić [68] and Brouwer & Haemers [58] provide a valuable overview of spectral graph theory, so, where appropriate, we cite the relevant sections of these (along with the original reference for a result).

As a result of Corollary 5.13, changing the sign of non-zero entries in matrix M has no effect on the growth rate of $\text{Grid}(M)$. For this reason, when considering particular collections of grid classes below, we choose to represent them by *grid diagrams* in which non-zero matrix entries are represented by a \boxtimes . As with grid classes, we freely apportion properties of a row-column graph to corresponding grid diagrams.




Figure 5.6: Some unicyclic grid diagrams that have the same row-column graph

Since transposing a matrix or permuting its rows and columns does not change the row-column graph of its grid class, there may be a number of distinct grid diagrams corresponding to a specific row-column graph (see Figure 5.6 for an example).

In many cases, we illustrate a result by showing a row-column graph and a corresponding grid diagram. We display just one of the possible grid diagrams corresponding to the row-column graph.

Our first result is the following elementary observation, which specifies a limitation on which numbers can be grid class growth rates. This is a consequence of the fact that the spectral radius of a graph is a root of the characteristic polynomial of an integer matrix.

Corollary 5.15. *The growth rate of a monotone grid class is an algebraic integer (the root of a monic polynomial).*

5.4.1 Slowly growing grid classes

Using results concerning graphs with small spectral radius, we can characterise grid classes with growth rates no greater than $\frac{9}{2}$. This is similar to Vatter’s characterisation of “small” permutation classes (with growth rate less than $\kappa \approx 2.20557$) in [159].

First, we recall that the growth rate of a disconnected grid class is the maximum of the growth rates of its components (see the proof of Theorem 5.14), so we only need to consider connected grid classes.




Figure 5.7: A cycle, an H graph and the three other Smith graphs, with corresponding grid diagrams

The connected graphs with spectral radius 2 are known as the *Smith* graphs. These are precisely the *cycle* graphs, the H graphs (paths with two pendant edges attached to both endvertices, including the star graph $K_{1,4}$), and the three other graphs shown in Figure 5.7.




Figure 5.8: A path, a Y graph and the three other connected proper subgraphs of Smith graphs, with corresponding grid diagrams

Similarly, the connected proper subgraphs of the Smith graphs are precisely the *path* graphs, the Y graphs (paths with two pendant edges attached to one endvertex) and the three other graphs in Figure 5.8. For details, see Smith [149] and Lemmens & Seidel [119]; also see [68, Theorem 3.11.1] and [58, Theorem 3.1.3].

With these, we can characterise all grid classes with growth rate no greater than 4:

Corollary 5.16. *If the growth rate of a connected monotone grid class equals 4, then its row-column graph is a Smith graph. If the growth rate of a connected monotone grid class is less than 4, then its row-column graph is a connected proper subgraph of a Smith graph.*

In particular, we have the following:

Corollary 5.17. *A monotone grid class of any size whose row-column graph is a cycle or an H graph has growth rate 4.*

In Appendix A of [159], Vatter considers *staircase* grid classes, whose row-column graphs are paths (see the leftmost grid diagram in Figure 5.8). The spectral radius of a path graph has long been known (Lovász & Pelikán [122]; also see [68, Theorem 8.1.17] and [58, 1.4.4]), from which we can conclude:

Corollary 5.18. *A monotone grid class of size m (having m non-zero cells) whose row-column graph is a path has growth rate $4 \cos^2\left(\frac{\pi}{m+2}\right)$. This is minimal for any connected grid class of size m .*

A Y graph of size m has spectral radius $2 \cos\left(\frac{\pi}{2m}\right)$, and the spectral radii of the three other graphs at the right of Figure 5.8 are $2 \cos\left(\frac{\pi}{12}\right)$, $2 \cos\left(\frac{\pi}{18}\right)$, and $2 \cos\left(\frac{\pi}{30}\right)$, from left to right (see [58, 3.1.1]). Thus we have the following characterisation of growth rates less than 4:

Corollary 5.19. *If the growth rate of a monotone grid class is less than 4, it is equal to $4 \cos^2\left(\frac{\pi}{k}\right)$ for some integer $k \geq 3$.*

The only grid class growth rates no greater than 3 are 1, 2, $\frac{1}{2}(3 + \sqrt{5}) \approx 2.618$, and 3.




Figure 5.9: E and F graphs

In order to characterise grid classes with growth rates slightly greater than 4, let an E graph be a tree consisting of three paths having one endvertex in common, and an F graph be a tree consisting of a path with a pendant edge attached to each of two distinct internal vertices (see Figure 5.9). Then, results of Brouwer & Neumaier [57] and Cvetković, Doob & Gutman [67] imply the following (also see [68, Theorem 3.11.2]):

Corollary 5.20. *If a connected monotone grid class has growth rate between 4 and $2 + \sqrt{5}$, then its row-column graph is an E or F graph.*

Thus, since $\sqrt{2 + \sqrt{5}}$ cannot be an eigenvalue of any graph (see [68, p. 93]), we can deduce the following:

Corollary 5.21. *If a monotone grid class properly contains a cycle then its growth rate exceeds $2 + \sqrt{5}$.*

More recently, Woo & Neumaier [173] have investigated the structure of graphs with spectral radius no greater than $\frac{3}{2}\sqrt{2}$ (also see [68, Theorem 3.11.3]). As a consequence, we have the following:

Corollary 5.22. *If the growth rate of a connected monotone grid class is no greater than $\frac{9}{2}$, then its row-column graph is one of the following:*

- (a) *a tree of maximum degree 3 such that all vertices of degree 3 lie on a path,*
- (b) *a unicyclic graph of maximum degree 3 such that all vertices of degree 3 lie on the cycle, or*
- (c) *a tree consisting of a path with three pendant edges attached to one endvertex.*

5.4.2 Accumulation points of grid class growth rates

Using graph theoretic results of Hoffman and Shearer, it is possible to characterise *all* accumulation points of grid class growth rates.

As we have seen, the growth rates of grid classes whose row-column graphs are paths and Y graphs grow to 4 from below; 4 is the least accumulation point of growth rates. The following characterises all accumulation points below $2 + \sqrt{5}$ (see Hoffman [98]):

Corollary 5.23. *For $k = 1, 2, \dots$, let β_k be the positive root of*

$$P_k(x) = x^{k+1} - (1 + x + x^2 + \dots + x^{k-1})$$

and let $\gamma_k = 2 + \beta_k + \beta_k^{-1}$. Then $4 = \gamma_1 < \gamma_2 < \dots$ are all the accumulation points of growth rates of monotone grid classes smaller than $2 + \sqrt{5}$.

The approximate values of the first eight accumulation points are: 4, 4.07960, 4.14790, 4.18598, 4.20703, 4.21893, 4.22582, 4.22988.

At $2 + \sqrt{5}$, things change dramatically; from this value upwards grid class growth rates are dense (see Shearer [145]):

Corollary 5.24. *Every $\gamma \geq 2 + \sqrt{5}$ is an accumulation point of growth rates of monotone grid classes.*




Figure 5.10: A grid diagram whose growth rate differs from 2π by less than 10^{-7} , and its caterpillar row-column graph

Thus, for every $\gamma \geq 2 + \sqrt{5} \approx 4.236068$, there is a grid class with growth rate arbitrarily close to γ . Indeed, for $\gamma > 2 + \sqrt{5}$, Shearer's proof provides an iterative process for generating a sequence of grid classes, each with a row-column graph that is a *caterpillar* (a tree such that all vertices of degree 2 or more lie on a path), with growth rates converging to γ from below. An example is shown in Figure 5.10.

5.4.3 Increasing the size of a grid class

We now consider the effect on the growth rate of making small changes to a grid class.

Firstly, growth rates of connected grid classes satisfy a *strict* monotonicity condition (see [68, Proposition 1.3.10]):

Corollary 5.25. *Adding a non-zero cell to a connected monotone grid class while preserving connectivity increases its growth rate.*

On the other hand, particularly surprising is the fact that grid classes with *longer internal paths* or cycles have *lower* growth rates.

An edge e of G is said to lie on an *endpath* of G if $G - e$ is disconnected and one of its components is a (possibly trivial) path. An edge that does not lie on an endpath is said to be *internal*. Note that a graph has an internal edge if and only if it contains either a cycle or non-star H graph.




Figure 5.11: Three unicyclic grid diagrams, of increasing size but decreasing growth rate from left to right, and their row-column graphs

An early result of Hoffman & Smith [99] shows that the subdivision of an internal edge *reduces* the spectral radius (also see [58, Proposition 3.1.4] and [68, Theorem 8.1.12]). Hence, we can deduce the following unexpected consequence for grid classes:

Corollary 5.26. *If $\text{Grid}(M)$ is connected, and $G(M')$ is obtained from $G(M)$ by subdividing an internal edge, then $\text{gr}(\text{Grid}(M')) < \text{gr}(\text{Grid}(M))$ unless $G(M)$ is a cycle or an H graph.*

For an example, see Figure 5.11.

5.4.4 Grid classes with extremal growth rates

Finally, we briefly consider grid classes with maximal or minimal growth rates for their size.



Figure 5.12: A skinny grid diagram and its row-column star graph

The row-column graph of a skinny grid class is a *star* (see Figure 5.12). Stars have maximal spectral radius among trees (see [68, Theorem 8.1.17]). This yields:

Corollary 5.27. *Among all connected acyclic monotone grid classes of size m , the skinny grid classes have the largest growth rate (equal to m).*

We have already seen (Corollary 5.18) that the connected grid classes with *least* growth rates are those whose row-column graph is a path. For unicyclic grid classes, we have the following (see [68, Theorem 8.1.18]):

Corollary 5.28. *Among all connected unicyclic monotone grid classes of size m , those whose row-column graph is a single cycle of length m have the smallest growth rate (equal to 4).*

There are many additional results known concerning graphs with extremal values for their spectral radii, especially for graphs with a small number of cycles. For an example, see the two papers by Simić [146, 147] on the largest eigenvalues of unicyclic and bicyclic graphs. Results like these can be translated into further facts concerning the growth rates of grid classes.

Chapter 6

Grid class limit shapes

In this chapter, we briefly investigate the shape of a “typical” large permutation in a monotone grid class. We show that almost all large permutations in a grid class do, in fact, have the same shape, and present a way of discovering this shape.

As we saw above (Lemma 4.1), if M has dimensions $r \times s$, then the number of gridded permutations of length n in $\text{Grid}^\#(M)$ is given by a sum of products of multinomial coefficients, one for each column and one for each row:

$$|\text{Grid}_n^\#(M)| = \sum \prod_{i=1}^r \binom{a_{i,1} + a_{i,2} + \dots + a_{i,s}}{a_{i,1}, a_{i,2}, \dots, a_{i,s}} \prod_{j=1}^s \binom{a_{1,j} + a_{2,j} + \dots + a_{r,j}}{a_{1,j}, a_{2,j}, \dots, a_{r,j}},$$

where the sum is over all combinations of non-negative $a_{i,j}$ such that $\sum_{i,j} a_{i,j} = n$ and $a_{i,j} = 0$ if $M_{i,j} = 0$. Each $a_{i,j}$ represents the number of points in the (i, j) cell.

To determine grid class limit shapes, we focus on the *proportion* $\alpha_{i,j} = a_{i,j}/n$ of points in each cell. (Vatter [159, Appendix A.4] takes a similar approach to yield a continued fraction expression for the growth rate of monotone grid classes whose row-column graph is a path.)

Given a $0/\pm 1$ matrix $M = (m_{i,j})$ with dimensions $r \times s$, let us use $\mathbf{A}(M)$ to denote the set of real $r \times s$ matrices $A = (\alpha_{i,j})$ with nonnegative entries, such that $\alpha_{i,j} = 0$ if $m_{i,j} = 0$ and $\sum_{i,j} \alpha_{i,j} = 1$. For example, $\begin{pmatrix} 0.2 & 0.7 \\ 0.1 & 0 \end{pmatrix} \in \mathbf{A} \begin{pmatrix} -1 & 1 \\ -1 & 0 \end{pmatrix}$.

Given a matrix $A = (\alpha_{i,j}) \in \mathbf{A}(M)$, let $\text{Grid}^\#(A)$ be the set of gridded permutations in $\text{Grid}^\#(M)$ for which the proportion of points in each cell is given by the values of the $\alpha_{i,j}$. Thus,

$$\text{Grid}^\#(M) = \bigsqcup_{A \in \mathbf{A}(M)} \text{Grid}^\#(A).$$

Note that, if $A = (\alpha_{i,j})$, then $\text{Grid}_n^\#(A)$ is nonempty only if, for each i, j , the number $\alpha_{i,j}n$ is an integer.

We claim that, asymptotically, the proportion of points in each cell is the same for almost all large gridded permutations in $\text{Grid}^\#(M)$. To state this formally, we require the following definition: For a real matrix A , let $\mathbf{A}_\varepsilon(A)$ be the set of matrices that are ε -close to A :

$$\mathbf{A}_\varepsilon(A) = \{B : \|B - A\|_{\max} \leq \varepsilon\},$$

where $\|(a_{i,j})\|_{\max} = \max_{i,j} |a_{i,j}|$ is the ‘‘entrywise’’ max norm.

Proposition 6.1. *For any connected class of gridded permutations, $\text{Grid}^\#(M)$, there is a unique $A_M \in \mathbf{A}(M)$, such that for all $\varepsilon > 0$,*

$$|\text{Grid}_n^\#(M)| \sim \sum_{B \in \mathbf{A}_\varepsilon(A_M)} |\text{Grid}_n^\#(B)|$$

Rather than giving a formal proof, we just demonstrate the result by considering the simplest nontrivial grid class, $\mathcal{C} = \text{Grid}^\#(\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix})$.

If αn is an integer, then the number of gridded permutations in \mathcal{C} of length n that have αn points in the left cell and $(1 - \alpha)n$ points in the right cell is $\binom{n}{\alpha n}$. Applying Stirling’s approximation, asymptotically we have

$$|\text{Grid}_n^\#(\alpha, 1 - \alpha)| \sim \frac{1}{\sqrt{2\pi\alpha(1 - \alpha)n}} \left(\alpha^{-\alpha} (1 - \alpha)^{-(1 - \alpha)} \right)^n = c_\alpha n^{-1/2} \gamma_\alpha^n.$$

It can easily be confirmed using elementary calculus that the exponential term, γ_α , is unimodal and takes a maximum value (of 2) when $\alpha = \frac{1}{2}$. Let $A_M = (\frac{1}{2}, \frac{1}{2})$.

Partitioning \mathcal{C} between, firstly, those gridded permutations whose distribution of points is ε -close to A_M and, secondly, those whose distribution of points is *not* ε -close to A_M yields

$$|\mathcal{C}_n| = \sum_{B \in \mathbf{A}_\varepsilon(\frac{1}{2}, \frac{1}{2})} |\text{Grid}_n^\#(B)| + \sum_{B \in \mathbf{A}(1,1) \setminus \mathbf{A}_\varepsilon(\frac{1}{2}, \frac{1}{2})} |\text{Grid}_n^\#(B)|. \quad (1)$$

Now, there is always some $\alpha \in [\frac{1}{2} - \frac{1}{2n}, \frac{1}{2}]$ for which $\text{Grid}_n^\#(\alpha, 1 - \alpha)$ is nonempty. Hence, for large enough n , the first summand in (1) is at least $c_{\alpha_1} n^{-1/2} \gamma_{\alpha_1}^n$, where $\alpha_1 = \frac{1}{2} - \frac{\varepsilon}{2}$. On the other hand, the second summand is no more than $n c_{\alpha_2} n^{-1/2} \gamma_{\alpha_2}^n$, where $\alpha_2 = \frac{1}{2} - \varepsilon$. So, since $\gamma_{\alpha_1} > \gamma_{\alpha_2}$, the set of gridded permutations whose asymptotic distribution of points is ε -close to A_M grows exponentially faster than the set of those whose distribution of points is not ε -close to A_M . Thus, asymptotically, almost all of the gridded permutations in \mathcal{C} have a distribution of points ε -close to A_M .

We claim that this is true of any class of gridded permutations as long as its row-column graph is connected.

Now, we know (see Lemma 2.1) that the maximum number of griddings of permutations in a monotone grid class grows only polynomially with length. Hence, asymptotically almost all permutations in $\text{Grid}(M)$ have griddings whose distribution is ε -close to A_M . Moreover, it can be shown that, of all the ways of interleaving points in cells in the same row or column, asymptotically almost all of them distribute the points evenly. Determining A_M is thus sufficient to give us the *limit shape* of permutations in $\text{Grid}(M)$, showing what a typical large gridded permutation in the class looks like.




Figure 6.1: Limit shapes of monotone grid classes whose row-column graphs are paths

Let us now consider how to ascertain the limiting distribution of points. When applied to a multinomial coefficient, Stirling's approximation gives the following asymptotic form:

$$\binom{\tau n}{\beta_1 n, \beta_2 n, \dots, \beta_k n} \sim \sqrt{\frac{\tau}{(2\pi)^{k-1} \beta_1 \beta_2 \dots \beta_k}} n^{-(k-1)/2} \left(\tau^\tau \beta_1^{-\beta_1} \beta_2^{-\beta_2} \dots \beta_k^{-\beta_k} \right)^n,$$

where $\tau = \beta_1 + \beta_2 + \dots + \beta_k$ and each β_i is positive.

Thus if $A = (\alpha_{i,j})$ has dimensions $r \times s$ and all the $\alpha_{i,j}$ are rational,

$$\overline{\text{gr}}(\text{Grid}^\#(A)) = \prod_{i=1}^r \frac{\kappa_i^{\kappa_i}}{\prod_j \alpha_{i,j}^{\alpha_{i,j}}} \times \prod_{j=1}^s \frac{\rho_j^{\rho_j}}{\prod_i \alpha_{i,j}^{\alpha_{i,j}}},$$

where $\kappa_i = \sum_j \alpha_{i,j}$ is the proportion of points in column i and $\rho_j = \sum_i \alpha_{i,j}$ is the proportion of points in row j , and we take sums and products over the nonzero $\alpha_{i,j}$ only.

Hence, by an argument analogous to that used above in analysing the asymptotic structure of $\text{Grid}^\#(\text{grid})$, and the fact that $\text{gr}(\text{Grid}(M)) = \text{gr}(\text{Grid}^\#(M))$ (see Lemma 2.1), we have the following consequence of Proposition 6.1:

Lemma 6.2. *The growth rate of a connected monotone grid class, $\text{Grid}(\mathcal{M})$, is given by*

$$\text{gr}(\text{Grid}(\mathcal{M})) = \max_{(\alpha_{i,j}) \in \mathcal{A}(\mathcal{M})} \prod_{i=1}^r \frac{\kappa_i^{\kappa_i}}{\prod_j \alpha_{i,j}^{\alpha_{i,j}}} \times \prod_{j=1}^s \frac{\rho_j^{\rho_j}}{\prod_i \alpha_{i,j}^{\alpha_{i,j}}},$$

where $\kappa_i = \sum_j \alpha_{i,j}$ and $\rho_j = \sum_i \alpha_{i,j}$.

Moreover, the maximum is achieved uniquely by the matrix $A_{\mathcal{M}}$, specified in Proposition 6.1, which indicates the limiting distribution of points in each cell.

To determine the limiting distribution, we can make use of the following result:

Proposition 6.3. *Let $A = (\alpha_{i,j})$ be the distribution of points in a gridding of a typical large permutation in a connected monotone grid class. Then there is a constant C such that, for every nonzero $\alpha_{i,j}$, we have $\alpha_{i,j}^2 / \kappa_i \rho_j = C$, where $\kappa_i = \sum_j \alpha_{i,j}$ and $\rho_j = \sum_i \alpha_{i,j}$.*

Proof. We use the method of Lagrange multipliers. Let $f(A)$ represent the expression from Lemma 6.2 that we need to maximise. We introduce the auxiliary function

$$\Lambda(A, \lambda) = \log f(A) - \lambda \left(1 - \sum_{i,j} \alpha_{i,j}\right).$$

The limiting distribution is then given by the solution to the set of equations

$$\frac{\partial \Lambda(A, \lambda)}{\partial \alpha_{i,j}} = 0,$$

one for each nonzero $\alpha_{i,j}$, together with the original constraint $\sum_{i,j} \alpha_{i,j} = 1$.

Now, for each $\alpha_{i,j}$,

$$\frac{\partial \Lambda(A, \lambda)}{\partial \alpha_{i,j}} = \lambda - 2 \log \alpha_{i,j} + \log \kappa_i + \log \rho_j.$$

Rearrangement and exponentiation then yields the fact that, for every $\alpha_{i,j}$, the expression $\alpha_{i,j}^2 / \kappa_i \rho_j$ has the same value (e^λ). \square

Proposition 6.3 can be used to generate a set of polynomial equations for $A_{\mathcal{M}}$ which can, in theory, be solved. By applying Lemma 6.2, this gives a method for determining the growth rate of a grid class. In practice, this is rather harder than finding the growth rate by evaluating the spectral radius of the row-column graph. However it is useful if we desire to know the limit shape.

Let's work through an example: $\text{Grid}^\# \left(\begin{array}{|c|c|c|} \hline \mathbb{N} & \mathbb{N} & \mathbb{N} \\ \hline \mathbb{N} & & \\ \hline \mathbb{N} & & \\ \hline \end{array} \right)$.

For this class, the equations for the entries in $A_M = \begin{pmatrix} \alpha_{1,2} & \alpha_{2,2} & \alpha_{3,2} \\ \alpha_{1,1} & 0 & 0 \end{pmatrix}$ are

$$\frac{\alpha_{1,1}}{\alpha_{1,1} + \alpha_{1,2}} = \frac{\alpha_{1,2}^2}{(\alpha_{1,1} + \alpha_{1,2})(\alpha_{1,2} + \alpha_{2,2} + \alpha_{3,2})} = \frac{\alpha_{2,2}}{\alpha_{1,2} + \alpha_{2,2} + \alpha_{3,2}} = \frac{\alpha_{3,2}}{\alpha_{1,2} + \alpha_{2,2} + \alpha_{3,2}},$$


together with

$$\alpha_{1,1} + \alpha_{1,2} + \alpha_{2,2} + \alpha_{3,2} = 1.$$

These equations have the unique positive solution

$$\alpha_{1,1} = \frac{1}{4}(2 - \sqrt{2}), \quad \alpha_{1,2} = \frac{1}{2\sqrt{2}}, \quad \alpha_{2,2} = \alpha_{3,2} = \frac{1}{4}.$$

Thus, we discover that, in $\text{Grid}^\# \left(\begin{array}{|c|c|c|} \hline \text{N} & \text{N} & \text{N} \\ \hline \text{N} & & \\ \hline \text{N} & & \\ \hline \end{array} \right)$, almost all large permutations look like this:



The limit shapes of some other grid classes are illustrated in Figure 6.1. We encounter similar pictures when considering the structure of permutations in the class $\text{Av}(1324)$ in the final chapter of this thesis.

One interpretation of the limit shape of a monotone grid class is as the measure-theoretic limit of convergent sequences of permutations. We conclude this chapter with a very brief sketch of this concept.

One of the most exciting recent developments in graph theory is the idea of the limit of a convergent sequence of finite graphs, introduced by Lovász & Szegedy [125, 123], and known as a *graphon*. A graphon is a Lebesgue measurable function W from $[0, 1]^2$ to $[0, 1]$ that is symmetric in its two arguments, i.e., $W(x, y) = W(y, x)$ for every $x, y \in [0, 1]$.

A similar analytic limit object, called a *permuton* [84], has also been defined recently for sequences of permutations by Hoppen, Kohayakawa, Sampaio and their collaborators [32, 101, 102, 103, 104]. Formally, a permuton is a probability measure μ on the σ -algebra of Borel sets of the unit square $[0, 1]^2$ such that μ has uniform marginals, i.e.,

$$\mu([\alpha, \beta] \times [0, 1]) = \mu([0, 1] \times [\alpha, \beta]) = \beta - \alpha$$

for every $0 \leq \alpha \leq \beta \leq 1$.

One way of understanding the idea of the limit of a sequence of permutations is to associate a permuton μ_σ with each permutation σ . If $|\sigma| = n$, then for each $i = 1, \dots, n$,

$$\mu_\sigma([[(i-1)/n, i/n] \times [[(\sigma(i)-1)/n, \sigma(i)/n]]) = 1.$$




Figure 6.2: Permutons associated with permutations in \boxtimes

Over regions that don't intersect these n squares, μ_σ is zero. See Figure 6.2 for some examples. A sequence (σ_k) of permutations of increasing length is then said to converge to a permuton μ if, for every measurable region $R \subseteq [0, 1]^2$, we have $\mu_{\sigma_k}(R) \rightarrow \mu(R)$ as k tends to infinity.

In the graph-theoretical context, Janson [106, 95] has initiated a study of the graphons that occur as limits of sequences of graphs in a graph class. He discovered an intriguing relationship between the growth rate of such a class and the “entropy” of its graphons. Sadly, this theorem does not transfer naturally to the world of permutations. However, an investigation of the permutons that occur as limits of sequences of permutations in a permutation class may yield interesting results. One initial elementary observation is that the *support* of any such permuton must have zero measure.

Chapter 7

Geometric grid classes

7.1 Introduction

Closely related to monotone grid classes are *geometric* grid classes, as investigated by Albert, Atkinson, Bouvel, Ruškuc & Vatter [10] and Albert, Ruškuc & Vatter [14]. The geometric grid class $\text{Geom}(M)$ is a subclass of $\text{Grid}(M)$, permutations in $\text{Geom}(M)$ satisfying an additional “geometric” constraint.




Figure 7.1: At left: The standard figure for $\text{Geom}\left(\begin{smallmatrix} 1 & 0 & -1 \\ & 1 & -1 \\ & & 1 \end{smallmatrix}\right)$, showing two plots of the permutation 1527634 with distinct gridings. At right: Its row-column graph; positive edges are shown as solid lines, negative edges are dashed.

Like a monotone grid class, a geometric grid class is specified by a $0/\pm 1$ matrix which represents the shape of plots of permutations in the class. As before, to match the Cartesian coordinate system, we index these matrices from the lower left, by column and then by row. If M is such a matrix, then we say that the *standard figure* of M , denoted Λ_M , is the subset of \mathbb{R}^2 consisting of the union of oblique open line segments $L_{i,j}$ with slope $M_{i,j}$ for each i, j for which $M_{i,j}$ is nonzero, where $L_{i,j}$ extends from $(i-1, j-1)$ to (i, j) if $M_{i,j} = 1$, and from $(i-1, j)$ to $(i, j-1)$ if $M_{i,j} = -1$.

The geometric grid class $\text{Geom}(M)$ is then defined to be the set of those permutations $\sigma_1 \sigma_2 \dots \sigma_n$ that can be plotted as a subset of the standard figure, i.e. for which there exists a sequence of points $(x_1, y_1), \dots, (x_n, y_n) \in \Lambda_M$ such that $x_1 < x_2 < \dots < x_n$ and the sequence y_1, \dots, y_n is order-isomorphic to $\sigma_1, \dots, \sigma_n$. See Figure 7.1 for an example.

We define the row-column graph $G(M)$ of geometric grid class $\text{Geom}(M)$ in the same way as for a monotone grid class (see Section 2.1), except that we label each edge $r_i c_j$ with the value of $M_{i,j}$. Edges labelled $+1$ are called *positive*; edges labelled -1 are called *negative*. See Figure 7.1 for an illustration.

Clearly $\text{Geom}(M)$ is a subset of $\text{Grid}(M)$. Indeed, $\text{Grid}(M)$ consists of those permutations that can be plotted as a subset of some figure consisting of the union of *any monotonic curves* $\Gamma_{i,j}$ with the same endpoints as the $L_{i,j}$ in Λ_M . This permits greater flexibility in the positioning of points in the cells, so one would expect that for some matrices, $\text{Grid}(M)$ would contain permutations that cannot be plotted on Λ_M . This is the case. For example, it can easily be shown that **2413** and **3142** lie in $\text{Grid}(\boxtimes)$ but not in $\text{Geom}(\boxtimes)$. In fact, the geometric grid class $\text{Geom}(M)$ and the monotone grid class $\text{Grid}(M)$ are identical if and only if $G(M)$ is acyclic [10, Theorem 3.2].




Figure 7.2: At left: The standard figure of $\begin{pmatrix} 1 & 0 & -1 \\ 1 & -1 & 1 \end{pmatrix}^{\times 2}$, with a consistent orientation marked. At right: Its row-column graph.

To state our result, we need one final definition related to geometric grid classes. If M is a $0/\pm 1$ matrix of dimensions $t \times u$, we define the *double refinement* $M^{\times 2}$ of M to be the $0/\pm 1$ matrix of dimensions $2t \times 2u$ obtained from M by replacing each 0 with $\begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$, each 1 with $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, and each -1 with $\begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$. See Figure 7.2 for an example. Note that the standard figure of $M^{\times 2}$ is essentially a scaled copy of the standard figure of M , so we have:

Observation 7.1. $\text{Geom}(M^{\times 2}) = \text{Geom}(M)$ for any $0/\pm 1$ matrix M .

We demonstrate a connection between the growth rate of $\text{Geom}(M)$ and the *matching polynomial* of the graph $G(M^{\times 2})$, the row-column graph of the double refinement of M . A k -*matching* of a graph is a set of k edges, no pair of which have a vertex in common. For example, the negative (dashed) edges in the graph in Figure 7.2 constitute a 4-matching. If, for each k , $m_k(G)$ denotes the number of distinct k -matchings of a graph G with n vertices, then the *matching polynomial* $\mu_G(z)$ of G is defined to be

$$\mu_G(z) = \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k m_k(G) z^{n-2k}. \quad (1)$$

Observe that the exponents of the variable z enumerate *defects* in k -matchings: the number of vertices which are *not* endvertices of an edge in such a matching. If n is even, $\mu_G(z)$ is an even function; if n is odd, $\mu_G(z)$ is an odd function.

With the relevant definitions complete, we can now state our theorem:

Theorem 7.2. *The growth rate of geometric grid class $\text{Geom}(M)$ exists and is equal to the square of the largest root of the matching polynomial $\mu_{G(M^{\times 2})}(z)$, where $G(M^{\times 2})$ is the row-column graph of the double refinement of M .*

In the next section, we prove this theorem by utilizing the link between geometric grid classes and trace monoids, and their connection to rook numbers and the matching polynomial. Then, in Section 7.3 we investigate a number of implications of this result by utilizing properties of the matching polynomial, especially the fact that the moments of $\mu_G(z)$ enumerate certain closed walks on G . Firstly, we characterise the growth rates of geometric grid classes in terms of the spectral radii of trees. Then, we explore the influence of cycle parity on growth rates and relate the growth rates of geometric grid classes to those of monotone grid classes. Finally, we consider the effect of subdividing edges in the row-column graph, proving some new results regarding how edge subdivision affects the largest root of the matching polynomial.

7.2 Proof of growth rate theorem

In order to prove our result, we make use of the connection between geometric grid classes and *trace monoids*. This relationship was first used by Vatter & Waton [162] to establish certain structural properties of grid classes, and was developed further in [10] from where we use a number of results. To begin with, we need to consider griddings of permutations.

We define M -griddings and M -gridded permutations as for monotone grid classes (see Section 2.1). A permutation may have multiple distinct griddings in a given geometric grid class; see Figure 7.1 for an example. We use $\text{Geom}^\#(M)$ to denote the set of all M -gridded permutations in $\text{Geom}(M)$. As is the case for monotone grid classes (Lemma 2.1), the growth rate of $\text{Geom}(M)$ is equal to the growth rate of the corresponding class of M -gridded permutations $\text{Geom}^\#(M)$. So we can restrict our considerations to M -gridded permutations.

To determine the growth rate of $\text{Geom}^\#(M)$, we relate M -gridded permutations to words in a trace monoid. To achieve this, one additional concept is required, that of a *consistent orientation* of a standard figure. If $\Lambda_M = \bigcup \{L_{i,j} : M_{i,j} \neq 0\}$ is the standard figure of

a $0/\pm 1$ matrix M , then a *consistent orientation* of Λ_M consists of an orientation of each oblique line $L_{i,j}$ such that in each column either all the lines are oriented leftwards or all are oriented rightwards, and in each row either all the lines are oriented downwards or all are oriented upwards.¹ See Figures 7.2 and 7.3 for examples.

It is not always possible to consistently orient a standard figure. The ability to do so depends on the *cycles* in the row-column graph. We say that the *parity* of a cycle in $G(M)$ is the product of the labels of its edges, a *positive cycle* is one which has parity $+1$, and a *negative cycle* is one with parity -1 . The following result relates cycle parity to consistent orientations:

Lemma 7.3 (see Vatter & Waton [162, Proposition 2.1]). *The standard figure Λ_M has a consistent orientation if and only if its row-column graph $G(M)$ contains no negative cycles.*

For example, $G(\begin{smallmatrix} \diagdown & \diagup \\ \diagup & \diagdown \end{smallmatrix})$ contains a negative cycle so its standard figure has no consistent orientation (see Figure 7.1), whereas $G(\begin{smallmatrix} \diagdown & \diagdown \\ \diagup & \diagup \end{smallmatrix})$ has no negative cycles so its standard figure has a consistent orientation (see Figure 7.3).

On the other hand, we can always consistently orient the standard figure of the double refinement of a matrix by orienting each oblique line towards the centre of its 2×2 block (as in Figure 7.2). So we have the following:

Lemma 7.4 (see [10, Proposition 4.1]). *If M is any $0/\pm 1$ matrix, then $\Lambda_{M \times 2}$ has a consistent orientation.*

Thus, by Lemma 7.3, the row-column graph of the double refinement of a matrix never contains a negative cycle. Figure 7.2 shows a consistent orientation of the standard figure of the double refinement of a matrix whose standard figure (shown in Figure 7.1) doesn't itself have a consistent orientation.




Figure 7.3: The plots of permutation 1527634 in $\text{Geom}\left(\begin{smallmatrix} -1 & 0 & -1 \\ 1 & -1 & 1 \end{smallmatrix}\right)$ associated with the words $a_{32}a_{32}a_{11}a_{12}a_{21}a_{31}a_{32}$ and $a_{11}a_{32}a_{21}a_{32}a_{31}a_{12}a_{32}$. Both plots correspond to the same gridding.

¹For ease of exposition, we use the concept of a consistent orientation rather than the approach used previously involving partial multiplication matrices; results from [10] follow *mutatis mutandis*.

We are now in a position to describe the association between words and M -gridded permutations. If M is a $0/\pm 1$ matrix, then we let $\Sigma_M = \{a_{ij} : M_{ij} \neq 0\}$ be an alphabet of symbols, one for each nonzero cell in M . If we have a consistent orientation for Λ_M , then we can associate to each finite word $w_1 \dots w_n$ over Σ_M a specific plot of a permutation in $\text{Geom}(M)$ as follows: If $w_k = a_{ij}$, include the point at distance $k\sqrt{2}/(n+1)$ along line segment $L_{i,j}$ according to its orientation. See Figure 7.3 for two examples. Clearly, this induces a mapping from the set of all finite words over Σ_M to $\text{Geom}^\#(M)$. In fact, it can readily be shown that this map is surjective, every M -gridded permutation corresponding to some word over Σ_M ([10, Proposition 5.3]).

As can be seen in Figure 7.3, distinct words may be mapped to the same gridded permutation. This occurs because the order in which two consecutive points are included is immaterial if they occur in cells that are neither in the same column nor in the same row. From the perspective of the words, adjacent symbols corresponding to such cells may be interchanged without changing the gridded permutation. This corresponds to a structure known as a trace monoid.

If we have a consistent orientation for standard figure Λ_M , then we define the *trace monoid of M* , which we denote by $\mathbb{M}(M)$, to be the set of equivalence classes of words over Σ_M in which a_{ij} and a_{kl} commute (i.e. $a_{ij}a_{kl} = a_{kl}a_{ij}$) whenever $i \neq k$ and $j \neq l$. It is then relatively straightforward to show equivalence between gridded permutations and elements of the trace monoid:

Lemma 7.5 (see [10, Proposition 7.1]). *If the standard figure Λ_M has a consistent orientation, then gridded n -permutations in $\text{Geom}^\#(M)$ are in bijection with equivalence classes of words of length n in $\mathbb{M}(M)$.*

Hence, by combining Lemmas 7.4 and 7.5 with Observation 7.1, we know that the growth rate of $\text{Geom}(M)$ is equal to the growth rate of $\mathbb{M}(M^{\times 2})$ if it exists. All that remains is to determine the growth rate of the trace monoid of a matrix.

Trace monoids were first studied by Cartier & Foata [59]. Using extended Möbius inversion, they determined the general form of the generating function, as follows:

Lemma 7.6 ([59]; see also Flajolet & Sedgewick [77, Note V.10]). *The ordinary generating function for $\mathbb{M}(M)$ is given by*

$$f_M(z) = \frac{1}{\sum_{k \geq 0} (-1)^k r_k(M) z^k}$$

where $r_k(M)$ is the number of k -subsets of Σ_M whose elements commute pairwise.

Since symbols in $\mathbb{M}(M)$ commute if and only if they correspond to cells that are neither in the same column nor in the same row, it is easy to see that $r_k(M)$ is the number of

distinct ways of placing k chess rooks on the nonzero entries of M in such a way that no two rooks attack each other by being in the same column or row. The numbers $r_k(M)$ are known as the *rook numbers* for M (see Riordan [143]). Moreover, a matching in the row-column graph $G(M)$ also corresponds to a set of cells no pair of which share a column or row. So the rook numbers for M are the same as the numbers of matchings in $G(M)$:

Observation 7.7. For all $k \geq 0$, $r_k(M) = m_k(G(M))$.

Now, by elementary analytic combinatorics (see Section 1.1), we know that the growth rate of $\mathbb{M}(M)$ is given by the reciprocal of the root of the denominator of $f_M(z)$ that has least magnitude (Lemma 1.2). The fact that this polynomial has a *unique* root of smallest modulus was proved by Goldwurm & Santini in [88]. It is real and positive by Pringsheim's Theorem (Lemma 1.1).

But the *reciprocal* of the *smallest* root of a polynomial is the same as the *largest* root of the *reciprocal* polynomial (obtained by reversing the order of the coefficients). Hence, if M has dimensions $t \times u$ and $n = t + u$, then the growth rate of $\mathbb{M}(M)$ is the largest (positive real) root of the polynomial

$$g_M(z) = \frac{1}{z^{\lfloor n/2 \rfloor} f_M\left(\frac{1}{z}\right)} = \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k r_k(M) z^{\lfloor n/2 \rfloor - k}. \quad (2)$$

Here, $g_M(z)$ is the reciprocal polynomial of $(f_M(z))^{-1}$ multiplied by some nonnegative power of z , since $r_k(M) = 0$ for all $k > \lfloor n/2 \rfloor$. Note also that n is the number of vertices in $G(M)$.

If we now compare the definition of $g_M(z)$ in (2) with that of the matching polynomial $\mu_G(z)$ in (1) and use Observation 7.7, then we see that:

$$g_M(z^2) = \begin{cases} \mu_{G(M)}(z), & \text{if } n \text{ is even;} \\ z^{-1} \mu_{G(M)}(z), & \text{if } n \text{ is odd.} \end{cases}$$

Hence, the largest root of $g_M(z)$ is the square of the largest root of $\mu_{G(M)}(z)$.

We now have all we need to prove Theorem 7.2: The growth rate of $\text{Geom}(M)$ is equal to the growth rate of $\mathbb{M}(M^{\times 2})$ which equals the square of the largest root of $\mu_{G(M^{\times 2})}(z)$.

In the above argument, we only employ the double refinement $M^{\times 2}$ to ensure that a consistent orientation is possible. By Lemma 7.3, we know that if $G(M)$ is free of negative cycles then Λ_M can be consistently oriented. Thus, we have the following special case of Theorem 7.2:

Corollary 7.8. *If $G(M)$ contains no negative cycles, then the growth rate of $\text{Geom}(M)$ is equal to the square of the largest root of $\mu_{G(M)}(z)$.*

7.3 Consequences

In this final section, we investigate some of the implications of Theorem 7.2. By considering properties of the matching polynomial, we characterise the growth rates of geometric grid classes in terms of the spectral radii of trees, prove a monotonicity result, and explore the influence of cycle parity on growth rates. We then compare the growth rates of geometric grid classes with those of monotone grid classes. Finally, we consider the effect of subdividing edges in the row-column graph.

Let's begin by introducing some notation. We use $G + H$ to denote the graph composed of two disjoint subgraphs G and H . The graph resulting from deleting the vertex v (and all edges incident to v) from a graph G is denoted $G - v$. Generalising this, if H is a subgraph of G , then $G - H$ is the graph obtained by deleting the vertices of H from G . In contrast, we use $G \setminus e$ to denote the graph that results from deleting the edge e from G . The number of connected components of G is represented by $\text{comp}(G)$. The characteristic polynomial of a graph G is denoted $\Phi_G(z)$. We use $\rho(G)$ to denote the spectral radius of G , the largest root of $\Phi_G(z)$. Finally, we use $\lambda(G)$ for the largest root of the matching polynomial $\mu_G(z)$.

The matching polynomial was independently discovered a number of times, beginning with Heilmann & Lieb [96] when investigating monomer-dimer systems in statistical physics. It was first studied from a combinatorial perspective by Farrell [73] and Gutman [92]. The theory was then further developed by Godsil & Gutman [87] and Godsil [85]. An introduction can be found in the books by Godsil [86] and Lovász & Plummer [124].

The facts concerning the matching polynomial that we use are covered by three lemmas. As a consequence of the first, we only need to consider connected graphs:

Lemma 7.9 (Farrell [73]; Gutman [92]). *The matching polynomial of a graph is the product of the matching polynomials of its connected components.*

Thus, in particular:

Corollary 7.10. *For any graphs G and H , we have $\lambda(G + H) = \max(\lambda(G), \lambda(H))$.*

The second lemma relates the matching polynomial to the characteristic polynomial.

Lemma 7.11 (Godsil & Gutman [87]). *If \mathcal{C}_G consists of all nontrivial subgraphs of G which are unions of vertex-disjoint cycles (i.e., all subgraphs of G which are regular of degree 2), then*

$$\mu_G(z) = \Phi_G(z) + \sum_{C \in \mathcal{C}_G} 2^{\text{comp}(C)} \Phi_{G-C}(z),$$

where $\Phi_{G-C}(z) = 1$ if $C = G$.

As an immediate consequence, we have the following:

Corollary 7.12 (Sachs [144]; Mowshowitz [130]; Lovász & Pelikán [122]). *The matching polynomial of a graph is identical to its characteristic polynomial if and only if the graph is acyclic.*

In particular, their largest roots are identical:

Corollary 7.13. *If G is a forest, then $\lambda(G) = \rho(G)$.*

Thus, using Corollaries 7.8 and 7.10, we have the following alternative characterisation for the growth rates of acyclic geometric grid classes:

Corollary 7.14. *If $G(M)$ is a forest, then $\text{gr}(\text{Geom}(M)) = \rho(G(M))^2$.*

This result is to be expected, since we know that $\text{Geom}(M) = \text{Grid}(M)$ if $G(M)$ is acyclic, and have shown (Theorem 5.14) that $\text{gr}(\text{Grid}(M)) = \rho(G(M))^2$.

The last, and most important, of the three lemmas allows us to determine the largest root of the matching polynomial of a graph from the spectral radius of a related tree. It is a consequence of the fact, determined by Godsil in [85], that the moments (sums of the powers of the roots) of $\mu_G(z)$ enumerate certain closed walks on G , which he calls *tree-like*.² This is analogous to the fact that the moments of $\Phi_G(z)$ count *all* closed walks on G . On a tree, all closed walks are tree-like.




Figure 7.4: Expanding G at u along uv ; H is the component of $G - u$ that contains v

²The author is grateful to Brendan McKay, whose helpful response to a question on *MathOverflow* [129] alerted him to the significance of this result.

Lemma 7.15 (Godsil [85]; see also [86] and [124]). *Let G be a graph and let u and v be adjacent vertices in a cycle of G . Let H be the component of $G - u$ that contains v . Now let K be the graph constructed by taking a copy of $G \setminus uv$ and a copy of H and joining the occurrence of u in the copy of $G \setminus uv$ to the occurrence of v in the copy of H (see Figure 7.4). Then $\lambda(G) = \lambda(K)$.*

The process that is described in Lemma 7.15 we call “*expanding G at u along uv* ”. Each such expansion of a graph G produces a graph with fewer cycles than G . Repeated application of this process thus eventually results in a forest F such that $\lambda(F) = \lambda(G)$. We shall say that F results from *fully expanding* G . Hence, by Corollaries 7.10 and 7.13, the largest root of the matching polynomial of a graph equals the spectral radius of some tree: for any graph G , there is a tree T such that $\lambda(G) = \rho(T)$.

It is readily observed that every tree is the row-column graph of some geometric grid class. Thus we have the following characterisation of geometric grid class growth rates.

Corollary 7.16. *The set of growth rates of geometric grid classes consists of the squares of the spectral radii of trees.*

The spectral radii of connected graphs satisfy the following strict monotonicity condition:

Lemma 7.17 ([68, Proposition 1.3.10]). *If G is connected and H is a proper subgraph of G , then we have $\rho(H) < \rho(G)$.*

Lemma 7.15 enables us to prove the analogous fact for the largest roots of matching polynomials, from which we can deduce a monotonicity result for geometric grid classes:

Corollary 7.18. *If G is connected and H is a proper subgraph of G , then $\lambda(H) < \lambda(G)$.*

Proof. Suppose we fully expand H (at vertices u_1, \dots, u_k , say), then the result is a forest F such that $\lambda(H) = \rho(F)$. Now suppose that we repeatedly expand G analogously at u_1, \dots, u_k , and then continue to fully expand the resulting graph. The outcome is a tree T (since G is connected) such that F is a proper subgraph of T and $\lambda(G) = \rho(T)$. The result follows from Lemma 7.17. □

Adding a non-zero cell to a $0/\pm 1$ matrix M adds an edge to $G(M)$. Thus, geometric grid classes satisfy the following monotonicity condition:

Corollary 7.19. *If $G(M)$ is connected and M' results from adding a non-zero cell to M in such a way that $G(M')$ is also connected, then $\text{gr}(\text{Geom}(M')) > \text{gr}(\text{Geom}(M))$.*

7.3.1 Cycle parity

The growth rate of a geometric grid class depends on the parity of its cycles. Consider the case of $G(M)$ being a cycle graph C_n . If $G(M)$ is a negative cycle, then $G(M^{\times 2}) = C_{2n}$. Now, by Lemma 7.15, we have $\lambda(C_n) = \rho(P_{2n-1})$, where P_n is the path graph on n vertices. The spectral radius of a path graph on n vertices is $2 \cos \frac{\pi}{n+1}$. So,

$$\text{gr}(\text{Geom}(M)) = \begin{cases} 4 \cos^2 \frac{\pi}{2n}, & \text{if } G(M) \text{ is a positive cycle;} \\ 4 \cos^2 \frac{\pi}{4n}, & \text{if } G(M) \text{ is a negative cycle.} \end{cases} \quad (3)$$

Thus the geometric grid class whose row-column graph is a negative cycle has a greater growth rate than the class whose row-column graph is a positive cycle. As another example,

$$\text{gr}(\text{Geom}\left(\begin{pmatrix} 1 & 0 & -1 \\ 1 & -1 & 1 \end{pmatrix}\right)) = 3 + \sqrt{2} \approx 4.41421, \quad (4)$$

whereas

$$\text{gr}(\text{Geom}\left(\begin{pmatrix} -1 & 0 & -1 \\ 1 & -1 & 1 \end{pmatrix}\right)) = 4. \quad (5)$$

The former, containing a negative cycle, has a greater growth rate than the latter, whose cycle is positive. This is typical; we prove the following result:

Corollary 7.20. *If $G(M)$ is connected and contains no negative cycles, and M_1 results from changing the sign of a single entry of M that is in a cycle (thus making one or more cycles in $G(M_1)$ negative), then $\text{gr}(\text{Geom}(M_1)) > \text{gr}(\text{Geom}(M))$.*

In order to do this, we need to consider the structure of $G(M^{\times 2})$. The graph $G(M^{\times 2})$ can be constructed from $G(M)$ as follows: If $G(M)$ has vertex set $\{v_1, \dots, v_n\}$, then we let $G(M^{\times 2})$ have vertices v_1, \dots, v_n and v'_1, \dots, v'_n . If $v_i v_j$ is a positive edge in $G(M)$, then in $G(M^{\times 2})$ we add an edge between v_i and v_j and also between v'_i and v'_j . On the other hand, if $v_i v_j$ is a negative edge in $G(M)$, then in $G(M^{\times 2})$ we join v_i to v'_j and v'_i to v_j . The correctness of this construction follows directly from the definitions of double refinement and of the row-column graph of a matrix. For an illustration, compare the graph in Figure 7.2 against that in Figure 7.1.

Note that if v_1, \dots, v_k is a positive k -cycle in $G(M)$, then $G(M^{\times 2})$ contains two vertex-disjoint positive k -cycles, the union of whose vertices is $\{v_1, \dots, v_k, v'_1, \dots, v'_k\}$. In contrast, if v_1, \dots, v_ℓ is a negative ℓ -cycle in $G(M)$, then $G(M^{\times 2})$ contains a (positive) 2ℓ -cycle on the vertices $\{v_1, \dots, v_\ell, v'_1, \dots, v'_\ell\}$ in which v_i is opposite v'_i (i.e. v'_i is at distance ℓ from v_i around the cycle) for each i , $1 \leq i \leq \ell$. We make the following additional observations:

Observation 7.21. *If $G(M)$ has no negative cycles, then $G(M^{\times 2}) = G(M) + G(M)$.*

Observation 7.22. *If $G(M)$ is connected and has a negative cycle, then $G(M^{\times 2})$ is connected.*

We now have all we require to prove our cycle parity result.




Figure 7.5: Graphs used in the proof of Corollary 7.20

Proof of Corollary 7.20. Let $G = G(M)$ and $G_1 = G(M_1^{\times 2})$, and let uv be the edge in G corresponding to the entry in M that is negated to create M_1 . Since G contains no negative cycles, by Observation 7.21, $G(M^{\times 2}) = G + G$. Thus, since G is connected, it has the form at the left of Figure 7.5, in which H is the component of $G - u$ containing v . Moreover, we have $\text{gr}(\text{Geom}(M)) = \lambda(G)$. (This also follows from Corollary 7.8.) Now, if we expand G at u along uv , by Lemma 7.15, $\lambda(G) = \lambda(K)$, where K is the graph in the centre of Figure 7.5.

On the other hand, G_1 is obtained from $G(M^{\times 2})$ by removing the edges uv and $u'v'$, and adding uv' and $u'v$, as shown at the right of Figure 7.5. It is readily observed that K is a proper subgraph of G_1 (see the shaded box in Figure 7.5), and hence, by Corollary 7.18, $\lambda(K) < \lambda(G_1)$. Since $\text{gr}(\text{Geom}(M_1)) = \lambda(G_1)$, the result follows. \square

Thus, making the first negative cycle increases the growth rate. We suspect, in fact, that the following stronger statement is also true:

Conjecture 7.23. *If $G(M)$ is connected and M_1 results from negating a single entry of M that is in one or more positive cycles but in no negative cycle, then*

$$\text{gr}(\text{Geom}(M_1)) > \text{gr}(\text{Geom}(M)).$$

To prove this more general result seems to require some new ideas. If $G(M)$ already contains a negative cycle, then $G(M^{\times 2})$ is connected, and, when this is the case, there appears to be no obvious way to generate a subgraph of $G(M_1^{\times 2})$ by expanding $G(M^{\times 2})$.

7.3.2 Monotone grid classes

Typically, the growth rate of a monotone grid class is greater than that of the corresponding geometric grid class. For example, if $G(M)$ is a cycle then $\text{gr}(\text{Grid}(M)) = 4$, whereas

from (3) we know that $\text{gr}(\text{Geom}(M)) < 4$. As another example,

$$\text{gr}(\text{Grid}\begin{pmatrix} 1 & 0 & -1 \\ 1 & -1 & 1 \end{pmatrix}) = \text{gr}(\text{Grid}\begin{pmatrix} -1 & 0 & -1 \\ 1 & -1 & 1 \end{pmatrix}) = \frac{1}{2}(5 + \sqrt{17}) \approx 4.56155,$$

which should be compared with (4) and (5).

The fact that the growth rate of the monotone grid class is strictly greater is a consequence of the fact that, if G is connected and not acyclic, then $\lambda(G)$ and $\rho(G)$ are distinct:

Lemma 7.24 (Godsil & Gutman [87]). *If G is connected and contains a cycle, then $\lambda(G) < \rho(G)$.*

Proof. By Lemma 7.17, if C is a nonempty subgraph of G , then $\rho(G - C) < \rho(G)$. So we have $\Phi_{G-C}(z) > 0$ for all $z \geq \rho(G)$. Moreover, $\Phi_G(z) \geq 0$ for $z \geq \rho(G)$. So, since G contains a cycle, from Lemma 7.11 we can deduce that $\mu_G(z) > 0$ if $z \geq \rho(G)$, and thus $\lambda(G) < \rho(G)$. \square

Note that, analogously to Observation 7.1, $\text{Grid}(M^{\times 2}) = \text{Grid}(M)$. Hence it must be the case that $\rho(G(M^{\times 2})) = \rho(G(M))$, the growth rate of a monotone grid class thus being independent of the parity of its cycles. As a consequence, from Lemma 7.24 we can deduce that in the non-acyclic case there is a strict inequality between the growth rate of a geometric grid class and the growth rate of the corresponding monotone grid class:


Corollary 7.25. *If $G(M)$ is connected, then $\text{gr}(\text{Geom}(M)) < \text{gr}(\text{Grid}(M))$ if and only if $G(M)$ contains a cycle.*

7.3.3 Subdivision of edges

One particularly surprising result in Chapter 5 concerning the growth rates of monotone grid classes is the fact that classes whose row-column graphs have longer internal paths or cycles exhibit *lower* growth rates (Corollary 5.26). Recall that an edge e of a graph G is said to lie on an *endpath* of G if $G \setminus e$ is disconnected and one of its components is a (possibly trivial) path. An edge that does not lie on an endpath is said to be *internal*. The following result of Hoffman & Smith states that the subdivision of an edge increases or decreases the spectral radius of the graph depending on whether the edge lies on an endpath or is internal:

Lemma 7.26 (Hoffman & Smith [99]). *Let G be a connected graph and G' be obtained from G by subdividing an edge e . If e lies on an endpath, then $\rho(G') > \rho(G)$. Otherwise (if e is an*

internal edge), $\rho(G') \leq \rho(G)$, with equality if and only if G is a cycle or has the following form (which we call an “H graph”):



Thus for monotone grid classes, if $G(M)$ is connected, and $G(M')$ is obtained from $G(M)$ by the subdivision of one or more internal edges, then $\text{gr}(\text{Grid}(M')) \leq \text{gr}(\text{Grid}(M))$.

As we see shortly, the situation is not as simple for geometric grid classes. The effect of edge subdivision on the largest root of the matching polynomial does not seem to have been addressed previously. In fact, the subdivision of an edge that is in a cycle may cause $\lambda(G)$ to increase or decrease, or may leave it unchanged. See Figures 7.8–7.10 for illustrations of the three cases. We investigate this further below. However, if the edge being subdivided is not on a cycle in G , then the behaviour of $\lambda(G)$ mirrors that of $\rho(G)$, as we now demonstrate:

Lemma 7.27. *Let G be a connected graph and G' be obtained from G by subdividing an edge e . If e lies on an endpath, then $\lambda(G') > \lambda(G)$. However, if e is an internal edge and not on a cycle, then $\lambda(G') \leq \lambda(G)$, with equality if and only if G is an H graph.*



Figure 7.6: Graphs used in the proof of Lemma 7.27

Proof. If e lies on an endpath, then G is a proper subgraph of G' and so the result follows from Corollary 7.18. On the other hand, if e is internal and G is acyclic, the conclusion is a consequence of Corollary 7.13 and Lemma 7.26. Thus, we need only consider the situation in which e is internal and G contains a cycle. We proceed by induction on the number of cycles in G , acyclic graphs constituting the base case. Let uv be an edge in a cycle of G such that u is not an endvertex of e . Now, let K be the result of expanding G at u along uv , and let K' , analogously, be the result of expanding G' at u along uv .

We consider the effect of the expansion of G upon e and the effect of the expansion of G' upon the two edges resulting from the subdivision of e . If e is in the component of $G - u$ containing v , then e is duplicated in K , both copies of e remaining internal (see Figure 7.6). Moreover, K' results from subdividing both copies of e in K . Conversely, if e is in a component of $G - u$ not containing v , then e is not duplicated in K (and remains

internal). In this case, K' results from subdividing e in K . In either case, K' is the result of subdividing internal edges of K (a graph with fewer cycles than G), and so the result follows from the induction hypothesis. \square

Now, the subdivision of an edge of a row-column graph that is not on a cycle has no effect on the parity of the cycles. Hence, we have the following conclusion for the growth rates of geometric grid classes:

Corollary 7.28. *If $G(M)$ is connected, and $G(M')$ is obtained from $G(M)$ by the subdivision of one or more internal edges not on a cycle, then $\text{gr}(\text{Geom}(M')) \leq \text{gr}(\text{Geom}(M))$, with equality if and only if $G(M)$ is an H graph.*




Figure 7.7: Graphs used in Lemma 7.29

Let us now investigate the effect of subdividing an edge e that lies on a cycle. We restrict our attention to graphs in which there is a vertex u such that the two endvertices of e are in distinct components of $(G \setminus e) - u$. See the graph at the left of Figure 7.7 for an illustration. We leave the consideration of multiply-connected graphs that fail to satisfy this condition for future study.

Lemma 7.29. *Let G be a connected graph and $e = x_1x_2$ an edge on a cycle C of G . Let u be a vertex on C , and let H_1 and H_2 be the distinct components of $(G \setminus e) - u$ that contain x_1 and x_2 respectively. Finally, let G' be the graph obtained from G by subdividing e .*

- (a) *If, for $i \in \{1, 2\}$, H_i is a (possibly trivial) path of which x_i is an endvertex, then $\lambda(G') > \lambda(G)$.*
- (b) *If, for $i \in \{1, 2\}$, H_i is not a path or is a path of which x_i is not an endvertex, then $\lambda(G') < \lambda(G)$.*

Proof. Let K be the result of repeatedly expanding G at u along every edge joining u to H_1 . K has the form shown at the right of Figure 7.7. Also let K' be the result of repeatedly expanding G' (G with edge e subdivided) in an analogous way at u . Clearly K' is the same as the graph that results from subdividing the copies of e in K .




Figure 7.8: Standard figures and row-column graphs of geometric grid classes whose growth rates increase from left to right




Figure 7.9: Standard figures and row-column graphs of geometric grid classes whose growth rates are all the same (equal to 5)




Figure 7.10: Standard figures and row-column graphs of geometric grid classes whose growth rates decrease from left to right

Now, for part (a), since H_1 is a path with an end at x_1 , and also H_2 is a path with an end at x_2 , we see that K' is the result of subdividing edges of K that are on endpaths. Hence, by the first part of Lemma 7.27, we have $\lambda(G') > \lambda(G)$ as required.

For part (b), since H_1 is not a path with an end at x_1 , and nor is H_2 a path with an end at x_2 , we see that K' is the result of subdividing internal edges of K . Since K is not an H graph, by Lemma 7.27, we have $\lambda(G') < \lambda(G)$ as required. \square

If the conditions for parts (a) and (b) of this lemma both fail to be satisfied (i.e. H_1 is a suitable path and H_2 isn't, or *vice versa*), then the proof fails. This is due to the fact that expansion leads to at least one copy of e in K being internal and to another copy of e in K being on an endpath. Subdivision of the former decreases $\lambda(G)$ whereas subdivision of the latter causes it to increase. Sometimes, as in Figure 7.9, these effects balance exactly; on other occasions one or the other dominates. We leave a detailed analysis of such cases for later study.

To conclude, we state the consequent result for the growth rates of geometric grid classes. To simplify its statement and avoid having to concern ourselves directly with cycle parities, we define $G^\times(M)$ to be $G(M)$ when $G(M)$ has no negative cycles and $G^\times(M)$ to be $G(M^{\times 2})$ otherwise.

Corollary 7.30. *Suppose $G^\times(M)$ is connected.*

- (a) *If $G^\times(M')$ is obtained from $G^\times(M)$ by subdividing one or more edges that satisfy the conditions of part (a) of Lemma 7.29, then $\text{gr}(\text{Geom}(M')) > \text{gr}(\text{Geom}(M))$.*
- (b) *If $G^\times(M')$ is obtained from $G^\times(M)$ by subdividing one or more edges that satisfy the conditions of part (b) of Lemma 7.29, then $\text{gr}(\text{Geom}(M')) < \text{gr}(\text{Geom}(M))$.*

Figure 7.8 provides an illustration of part (a) and Figure 7.10 an illustration of part (b).

Part II

THE SET OF GROWTH RATES

Chapter 8

Intervals of growth rates

8.1 Introduction

What does the set of possible permutation class growth rates look like? The last few years have seen substantial progress on answering this question, with particular focus on significant phase transition values. Kaiser & Klazar [108] characterised all growth rates up to 2, showing that the golden ratio $\varphi \approx 1.61803$ (the unique positive root of $x^2 - x - 1$) is the least growth rate greater than 1. Indeed, they prove a stronger statement, known as the *Fibonacci dichotomy*, that if $|\mathcal{C}_n|$ is ever less than the n th Fibonacci number, then $|\mathcal{C}_n|$ is eventually polynomial. Kaiser & Klazar also determine that 2 is the least limit point in the set of growth rates. Huczynska & Vatter [105] later gave a shorter proof of the Fibonacci dichotomy, by characterising polynomial permutation classes in terms of monotone grid classes (see the brief discussion above on page 35).

Klazar [111] considered the least growth rate admitting uncountably many permutation classes (which he denoted κ) and proved that κ is at least 2 and is no greater than approximately 2.33529 (a root of a quintic). Vatter [159] determined the exact value of κ to be the unique real root of $x^3 - 2x^2 - 1$ (approximately 2.20557) and completed the characterisation of all growth rates up to κ , proving that there are uncountably many permutation classes with growth rate κ and only countably many with growth rate less than κ . Central to Vatter's work is the use of (generalised) grid classes. The phase transition at κ also has enumerative ramifications: Albert, Ruškuc & Vatter [14] have shown that every permutation class with growth rate less than κ has a rational generating function.

Balogh, Bollobás & Morris [29] extended Kaiser & Klazar's work to the more general setting of ordered graphs. They conjectured that the set of growth rates of hereditary classes of ordered graphs contains no limit points from above, and also that all such growth rates are integers or algebraic irrationals. These conjectures were disproved by

Albert & Linton [5] who exhibited an uncountable *perfect set* (a closed set every member of which is a limit point) of growth rates of permutation classes and in turn conjectured that the set of growth rates includes some interval (λ, ∞) .

Their conjecture was established by Vatter [158] who proved that there are permutation classes of every growth rate at least $\lambda_A \approx 2.48187$ (the unique real root of $x^5 - 2x^4 - 2x^2 - 2x - 1$). In addition, he conjectured that this value was optimal, the set of growth rates below λ_A being nowhere dense. By generalising Vatter's constructions, we now refute his conjecture by proving the following two results:

Theorem 8.1. *Let $\theta_B \approx 2.35526$ be the unique real root of $x^7 - 2x^6 - x^4 - x^3 - 2x^2 - 2x - 1$. For any $\varepsilon > 0$, there exist δ_1 and δ_2 with $0 < \delta_1 < \delta_2 < \varepsilon$ such that every value in the interval $[\theta_B + \delta_1, \theta_B + \delta_2]$ is the growth rate of a permutation class.*

Theorem 8.2. *Let $\lambda_B \approx 2.35698$ be the unique positive root of $x^8 - 2x^7 - x^5 - x^4 - 2x^3 - 2x^2 - x - 1$. Every value at least λ_B is the growth rate of a permutation class.*

The proofs of these theorems are based upon an analysis of expansions of real numbers in non-integer bases, the study of which was initiated by Rényi in the 1950s. In particular, we prove two generalisations of a result of Pedicini concerning expansions in which the digits are drawn from sets of allowed values. These results have been submitted for publication (see [38]).

In order to complete the characterisation of all permutation class growth rates, further investigation is required of the values between κ and λ_B . This interval contains the first limit point from above (which we denote ξ) and the first perfect set (whose infimum we denote η), as well as the first interval (whose infimum we denote θ). In [158], Vatter showed that $\xi_A \approx 2.30522$ (the unique real root of $x^5 - 2x^4 - x^2 - x - 1$) is the infimum of a perfect set in the set of growth rates and hence a limit point from above, and conjectured that this is the least such limit point (in which case $\eta = \xi$). If Γ denotes the set of all permutation class growth rates, then the current state of our knowledge can be summarised by the following table:

κ	$\inf\{\gamma : \Gamma \cap (1, \gamma) > \aleph_0\}$	$\kappa \approx 2.20557$
ξ	$\inf\{\gamma : \gamma \text{ is a limit point from above in } \Gamma\}$	$\xi \leq \xi_A \approx 2.30522$
η	$\inf\{\gamma : \gamma \in \mathcal{P}, \mathcal{P} \subset \Gamma, \mathcal{P} \text{ is a perfect set}\}$	$\eta \leq \xi_A \approx 2.30522$
θ	$\inf\{\gamma : (\gamma, \delta) \subset \Gamma\}$	$\theta \leq \theta_B \approx 2.35526$
λ	$\inf\{\gamma : (\gamma, \infty) \subset \Gamma\}$	$\lambda \leq \lambda_B \approx 2.35698$

Further research is required to determine whether the upper bounds on ξ , η , θ and λ are, in fact, equalities. Currently, the only proven lower bound we have for any of these is κ .

Clearly, there are analogous phase transition values in the set of growth rates of hereditary classes of ordered graphs, and the permutation class upper bounds also bound the corresponding ordered graph values from above. The techniques used here could readily be applied directly to ordered graphs to yield smaller upper bounds to the ordered graph phase transitions. We leave such a study for the future.

In the next section, we investigate expansions of real numbers in non-integer bases. In [136], Pedicini considered the case in which the digits are drawn from some allowable digit set and determined the conditions on the base under which every real number in an interval can be expressed by such an expansion. We prove two generalisations of Pedicini's result, firstly permitting each digit to be drawn from a different allowable digit set, and secondly allowing the use of what we call *generalised digits*.

In Section 8.3, we consider the growth rates of sum-closed permutation classes (whose members are constructed from a sequence of indecomposable permutations). Building on our results in the previous section concerning expansions of real numbers in non-integer bases, we determine sufficient conditions for the set of growth rates of a family of sum-closed permutation classes to include an interval of the real line. These conditions relate to the enumeration of the indecomposable permutations in the classes.

In the final section, we prove our two theorems by constructing families of permutation classes whose indecomposable permutations are enumerated by sequences satisfying the required conditions. The key permutations in our constructions are formed from permutations known as *increasing oscillations* by *inflating* their ends. Our constructions are similar to, but considerably more general than, those used by Vatter to prove that the set of permutation class growth rates includes the interval $[\lambda_A, \infty)$.

8.2 Non-integer bases and generalised digits

Suppose we want a way of representing any value in an interval of the real line. As is well known, infinite sequences of small non-negative integers suffice. Given some integer $\beta > 1$ and any number $x \in [0, 1]$ there is some sequence (a_n) where each $a_n \in \{0, 1, \dots, \beta - 1\}$ such that

$$x = \sum_{n=1}^{\infty} a_n \beta^{-n}.$$

This is simply the familiar "base β " expansion of x , normally written $x = 0.a_1a_2a_3\dots$, the a_n being " β -ary digits". For simplicity, we use $(a_n)_\beta$ to denote the sum above. For

example, if $a_n = 1$ for each $n \geq 1$ then $(a_n)_3 = \frac{1}{2}$.

In a seminal paper, Rényi [142] generalized these expansions to arbitrary real bases $\beta > 1$, and since then many papers have been devoted to the various connections with other branches of mathematics, such as measure theory and dynamical systems. For further information, see the recent survey of Komornik [114]. Of particular relevance to us, Pedicini [136] explored the further generalisation in which the digits are drawn from some allowable digit set $A = \{a_0, \dots, a_m\}$, with $a_0 < \dots < a_m$. He proved that every point in the interval $[\frac{a_0}{\beta-1}, \frac{a_m}{\beta-1}]$ has an expansion of the above form with $a_n \in A$ for all $n \geq 1$ if and only if $\max_{1 \leq j \leq m} (a_j - a_{j-1}) \leq \frac{a_m - a_0}{\beta-1}$.

We establish two additional generalisations of Pedicini's result. Firstly, we permit each digit to be drawn from a different allowable digit set. Specifically, for each n we allow the n^{th} digit a_n to be drawn from some finite set A_n of permitted values. For instance, we could have $A_n = \{1, 4\}$ for odd n and $A_n = \{1, 3, 5, 7, 9\}$ for even n . We return to this example below.

Let $(A_n)_\beta = \{(a_n)_\beta : a_n \in A_n\}$ be the set of values for which there is an expansion in base β . The following lemma, generalising the result of Pedicini, gives sufficient conditions on β for $(A_n)_\beta$ to be an interval of the real line.

Lemma 8.3. *Given a sequence (A_n) of non-empty finite sets of non-negative real numbers, let ℓ_n and u_n denote the least and greatest ('lower' and 'upper') elements of A_n respectively, and let $\Delta_n = u_n - \ell_n$. Also, let δ_n be the largest gap between consecutive elements of A_n (with $\delta_n = 0$ if $|A_n| = 1$).*

If (u_n) is bounded, $\beta > 1$, and for each n , β satisfies the inequality

$$\delta_n \leq \sum_{i=1}^{\infty} \Delta_{n+i} \beta^{-i},$$

then $(A_n)_\beta = [(\ell_n)_\beta, (u_n)_\beta]$.

Proof. If we let $A'_n = \{a - \ell_n : a \in A_n\}$ for each n , then the set of values that can be represented by digits from the A_n is the same as the set of values representable using digits from the A'_n shifted by $(\ell_n)_\beta$:

$$(A_n)_\beta = \{(a_n)_\beta : a_n \in A_n\} = \{(a'_n)_\beta + (\ell_n)_\beta : a'_n \in A'_n\}.$$

So we need only consider cases in which each ℓ_n is zero, whence we have $\Delta_n = u_n$ for all n .

Given some x in the specified interval, we choose each digit maximally such that no partial sum exceeds x , i.e. for each n , we select a_n to be the greatest element of A_n such

that

$$S(n) = \sum_{i=1}^n a_i \beta^{-i} \leq x.$$

It follows that $(a_n)_\beta = \lim_{n \rightarrow \infty} S(n) \leq x$. We claim that $(a_n)_\beta = x$.

If, for some n , $a_n < u_n$ then there is some element of A_n greater than a_n but no larger than $a_n + \delta_n$, so by our choice of the a_n , $S(n) + \delta_n \beta^{-n} > x$. Hence, $|x - S(n)| < \delta_n \beta^{-n}$. Thus, if $a_n < u_n$ for infinitely many values of n , $(a_n)_\beta = \lim_{n \rightarrow \infty} S(n) = x$, since the δ_n are bounded and $\beta > 1$.

Suppose now that $a_n = u_n$ for all but finitely many n . If, in fact, $a_n = u_n$ for all n , then we have $x = (u_n)_\beta = (a_n)_\beta$ as required. Indeed, this is the only possibility given the way we have chosen the a_n .¹ Let us assume that N is the greatest index for which $a_N < u_N$, so that $S(N) + \delta_N \beta^{-N} > x$.

Now, by the constraints on β and the fact that $\Delta_i = u_i$ for all i , we have

$$\delta_N \beta^{-N} \leq \sum_{i=N+1}^{\infty} u_i \beta^{-i},$$

and thus

$$(a_n)_\beta = S(N) + \sum_{i=N+1}^{\infty} u_i \beta^{-i} > x,$$

which produces a contradiction. □

We refer to the conditions relating β to the values of the δ_n and Δ_n as the *gap inequalities*. These reoccur in subsequent lemmas.

Let us return to our example. If $A_n = \{1, 4\}$ for odd n and $A_n = \{1, 3, 5, 7, 9\}$ for even n , then to produce an interval of values it is sufficient for β to satisfy the two gap inequalities:

$$3 \leq 8\beta^{-1} + 3\beta^{-2} + 8\beta^{-3} + 3\beta^{-4} + \dots = \frac{8\beta+3}{\beta^2-1}, \quad \text{so } \beta \leq \frac{1}{3}(4 + \sqrt{34}) \approx 3.27698,$$

$$2 \leq 3\beta^{-1} + 8\beta^{-2} + 3\beta^{-3} + 8\beta^{-4} + \dots = \frac{3\beta+8}{\beta^2-1}, \quad \text{so } \beta \leq \frac{1}{4}(3 + \sqrt{89}) \approx 3.10850.$$

Thus, if $1 < \beta \leq \frac{1}{4}(3 + \sqrt{89})$ then $(A_n)_\beta$ is an interval. For instance, with $\beta = 3$, we have $(A_n)_3 = [\frac{1}{8}, \frac{9}{4}]$.

We now generalise our concept of a digit by permitting members of the A_n to have the

¹For example, our approach would select 0.50000... rather than 0.49999... as the decimal representation of $\frac{1}{2}$, but 0.99999... as the only possible representation for 1.

form $c_0.c_1 \dots c_k$ for non-negative integers c_i , where this is to be interpreted to mean

$$c_0.c_1 \dots c_k = \sum_{i=0}^k c_i \beta^{-i}$$

as would be expected. We call $c_0.c_1 \dots c_k$ a *generalised digit* or a *generalised β -digit* of length $k + 1$ with *subdigits* c_0, c_1, \dots, c_k . Note that the value of a generalised β -digit depends on the value of β .

As an example of sets of generalised digits, consider $A_n = \{1.1, 1.11, 1.12, 1.2, 1.21, 1.22\}$ for odd n and $A_n = \{0\}$ for even n . This is similar to what we use below to create intervals of permutation class growth rates. We return to this example below.

As long as the sets of generalised digits are suitably bounded, the same gap constraints on β suffice as before to ensure that $(A_n)_\beta$ is an interval. Thus we have the following generalisation of Lemma 8.3.

Lemma 8.4. *Given a sequence (A_n) of non-empty finite sets of generalised β -digits, for a fixed β let ℓ_n and u_n denote the least and greatest elements of A_n respectively, and let $\Delta_n = u_n - \ell_n$. Also, let δ_n be the largest gap between consecutive elements of A_n (with $\delta_n = 0$ if $|A_n| = 1$).*

If both the length and the subdigits of all the generalised digits in the A_n are bounded, $\beta > 1$, and for each n , β satisfies the inequality

$$\delta_n \leq \sum_{i=1}^{\infty} \Delta_{n+i} \beta^{-i},$$

then $(A_n)_\beta = [(\ell_n)_\beta, (u_n)_\beta]$.

We omit the proof since it is essentially the same as that for Lemma 8.3.

Let us analyse our generalised digit example. If $A_n = \{1.1, 1.11, 1.12, 1.2, 1.21, 1.22\}$ for odd n and $A_n = \{0\}$ for even n , then for odd n we have $\Delta_n = 0.12 = \beta^{-1} + 2\beta^{-2}$. However, the determination of δ_n depends on whether 1.2 exceeds 1.12 or not. If it does (which is the case when β exceeds 2), then δ_n is the greater of β^{-2} and $\beta^{-1} - 2\beta^{-2}$. If $\beta < 2$, then δ_n is the greatest of β^{-2} , $\beta^{-1} - \beta^{-2}$ and $-\beta^{-1} + 2\beta^{-2}$. Careful solving of all the resulting inequalities reveals that if $1 < \beta \leq \frac{1}{2}(1 + \sqrt{13}) \approx 2.30278$, then $(A_n)_\beta$ is an interval.²

Let us now begin to relate this to the construction of permutation classes.

²When the A_n contain generalised digits, the solution set need not consist of a single interval. For example, if $A_n = \{0.5, 0.501, 0.502, 0.51, 0.511, 0.512, 0.52, 0.521, 1.3\}$ for odd n and $A_n = \{0\}$ for even n , then $(A_n)_\beta$ is an interval if $1 < \beta \leq 2.732\dots$ and also if $2.879\dots \leq \beta \leq 2.923\dots$

8.3 Sum-closed permutation classes

Recall, from Section 1.2, the definition of the direct sum, $\sigma \oplus \tau$, of two permutations and the concept of an indecomposable, and also that of a sum-closed class, and of the sum closure $\bigoplus \mathcal{S}$ of a downward closed set of indecomposables. All the permutation classes that we construct to prove our results are sum-closed, and we define them by specifying the (downward closed) sets of indecomposables from which they are composed. The remainder of this section is devoted to study the growth rates of sum-closed classes, and our main results, Lemmas 8.7 and 8.8, provide conditions on when the growth rates of a family of sum-closed permutation classes form an interval. In Section 8.4, we construct families of sum-closed classes satisfying these conditions, thereby proving Theorems 8.1 and 8.2.

We say that a set of permutations \mathcal{S} is *positive* and *bounded* (by a constant c) if $1 \leq |\mathcal{S}_n| \leq c$ for all n . As long as its set of indecomposables is positive and bounded, the growth rate of a sum-closed permutation class is easy to determine. Variants of the following lemma can be found in the work of Albert & Linton [5] and of Vatter [158].

Lemma 8.5. *If \mathcal{S} is a downward closed set of indecomposables that is positive and bounded, then $\text{gr}(\bigoplus \mathcal{S})$ is the unique γ such that $\gamma > 1$ and $S(\gamma^{-1}) = \sum_{n=1}^{\infty} |\mathcal{S}_n| \gamma^{-n} = 1$.*

This follows from the following standard analytic combinatorial result.

Lemma 8.6 ([77, Theorem V.1]). *Suppose $\mathcal{F} = \text{SEQ}[\mathcal{G}]$ and the corresponding generating functions, F and G , are analytic at 0, with $G(0) = 0$ and*

$$\lim_{x \rightarrow \rho^-} G(x) > 1,$$

where ρ is the radius of convergence of G . If, in addition, G is aperiodic, i.e. there does not exist an integer $d \geq 2$ such that $G(z) = H(z^d)$ for some H analytic at 0, then $\text{gr}(\mathcal{F}) = \sigma^{-1}$, where σ is the only root in $(0, \rho)$ of $G(x) = 1$.

Proof of Lemma 8.5. The permutation class $\bigoplus \mathcal{S}$ is in bijection with $\text{SEQ}[\mathcal{S}]$. Since $|\mathcal{S}_n|$ is positive and bounded, $S(z)$ is aperiodic, its radius of convergence is 1, and $\lim_{x \rightarrow 1^-} S(x) = +\infty$. Thus the requirements of Proposition 8.6 are satisfied, from which the desired result follows immediately. \square

Building on our earlier results concerning expansions of real numbers in non-integer bases, we are now in a position to describe conditions under which the set of growth rates of a family of sum-closed permutation classes includes an interval of the real line.

To state the lemma, we introduce a notational convenience. If (a_n) is a bounded sequence of positive integers that enumerates a downward closed set of indecomposables \mathcal{S} (i.e. $|\mathcal{S}_n| = a_n$ for all n), let us use $\text{gr}(\bigoplus(a_n))$ to denote the growth rate of the sum-closed permutation class $\bigoplus\mathcal{S}$.

Lemma 8.7. *Given a sequence (A_n) of non-empty finite sets of positive integers, let ℓ_n and u_n denote the least and greatest elements of A_n respectively, and let $\Delta_n = u_n - \ell_n$. Also, let δ_n be the largest gap between consecutive elements of A_n (with $\delta_n = 0$ if $|A_n| = 1$).*

If (u_n) is bounded, then for every real number γ such that $\text{gr}(\bigoplus(\ell_n)) \leq \gamma \leq \text{gr}(\bigoplus(u_n))$ and

$$\delta_n \leq \sum_{i=1}^{\infty} \Delta_{n+i} \gamma^{-i}$$

for each n , there is some sequence (a_n) with each $a_n \in A_n$ such that $\gamma = \text{gr}(\bigoplus(a_n))$.

Proof. Let $\gamma_1 = \text{gr}(\bigoplus(\ell_n))$ and $\gamma_2 = \text{gr}(\bigoplus(u_n))$, so $\gamma_1 \leq \gamma \leq \gamma_2$. Now, by Lemma 8.5,

$$(\ell_n)_\gamma = \sum_{n=1}^{\infty} \ell_n \gamma^{-n} \leq \sum_{n=1}^{\infty} \ell_n \gamma_1^{-n} = 1 = \sum_{n=1}^{\infty} u_n \gamma_2^{-n} \leq \sum_{n=1}^{\infty} u_n \gamma^{-n} = (u_n)_\gamma.$$

Hence, $(\ell_n)_\gamma \leq 1 \leq (u_n)_\gamma$ and so, since γ satisfies the gap inequalities, by Lemma 8.3 there is some sequence (a_n) with each $a_n \in A_n$ such that $(a_n)_\gamma = 1$ and thus, by Lemma 8.5, $\gamma = \text{gr}(\bigoplus(a_n))$. \square

This result generalises Proposition 2.4 in Vatter [158], which treats only the case in which all the A_n are identical and consist of an interval of integers $\{\ell, \ell + 1, \dots, u\}$.

As a consequence of this lemma, given a suitable sequence (A_n) of sets of integers, if we could construct a family of sum-closed permutation classes such that every sequence (a_n) with each $a_n \in A_n$ enumerates the indecomposables of some member of the family, then the set of growth rates of the classes in the family would include an interval of the real line. Returning to our original example in which $A_n = \{1, 4\}$ for odd n and $A_n = \{1, 3, 5, 7, 9\}$ for even n , suppose it were possible to construct sum-closed classes whose indecomposables were enumerated by each sequence $\{(a_n) : a_n \in A_n\}$. Then the set of growth rates would include the interval from $\text{gr}(\bigoplus(1, 1, \dots)) = 2$ to the lesser of $\text{gr}(\bigoplus(4, 9, 4, 9, \dots)) = 2 + \sqrt{14}$ and $\frac{1}{4}(3 + \sqrt{89})$, the greatest value permitted by the gap inequalities. (Clearly, the fact that there are only countably many permutation classes with growth rate less than $\kappa \approx 2.20557$ means that such a construction is, in fact, impossible.)

We now broaden this result to handle the case in which the A_n contain generalised digits. In the integer case, a term a_n corresponds to a set consisting of a_n indecomposables of length n . In the generalised digit case, $a_n = c_0.c_1 \dots c_k$ corresponds to a set consisting of

c_0 indecomposables of length n , c_1 indecomposables of length $n + 1$, and so on, up to c_k indecomposables of length $n + k$.

Before stating the generalised digit version of our lemma, we generalise our terminology and notation. Let us say that a sequence of generalised digits (a_n) enumerates a set \mathcal{S} of permutations if \mathcal{S} is the union $\uplus \mathcal{S}^{(n)}$ of disjoint sets $\mathcal{S}^{(n)}$, such that, for each n , if we have $a_n = c_0.c_1 \dots c_k$ then $\mathcal{S}^{(n)}$ consists of exactly $a_n + c_i$ permutations of length $n + i$ for $i = 0, \dots, k$. If (a_n) is a sequence of generalised digits that enumerates a downward closed set of indecomposables \mathcal{S} , let us use $\text{gr}(\oplus(a_n))$ to denote the growth rate of the sum-closed permutation class $\oplus\mathcal{S}$.

In the proofs of our theorems, we make use of the following elementary fact. Given a sequence of generalised digits that enumerates a set \mathcal{S} , there is a unique sequence of integers that also enumerates \mathcal{S} . For example, if $a_n = 1.12$ for odd n and $a_n = 3.1$ for even n , then (a_n) and the sequence $(t_n) = (1, 4, 4, 4, \dots)$ are equivalent in this way. We write such equivalences $(a_n) \equiv (t_n)$.

Lemma 8.8. *Given a sequence (A_n) of non-empty finite sets of generalised γ -digits, for a fixed γ let ℓ_n and u_n denote the least and greatest elements of A_n respectively, and let $\Delta_n = u_n - \ell_n$. Also, let δ_n be the largest gap between consecutive elements of A_n (with $\delta_n = 0$ if $|A_n| = 1$).*

If the length and subdigits of all the generalised γ -digits in the A_n are bounded, then for every real number γ such that $\text{gr}(\oplus(\ell_n)) \leq \gamma \leq \text{gr}(\oplus(u_n))$ and

$$\delta_n \leq \sum_{i=1}^{\infty} \Delta_{n+i} \gamma^{-i}$$

for each n , there is some sequence (a_n) with each $a_n \in A_n$ such that $\gamma = \text{gr}(\oplus(a_n))$.

We omit the proof since it is essentially the same as that for Lemma 8.7, but making use of Lemma 8.4 rather than Lemma 8.3.

Let us revisit our generalised digit example in which $A_n = \{1.1, 1.11, 1.12, 1.2, 1.21, 1.22\}$ for odd n and $A_n = \{0\}$ for even n . Suppose it were possible to construct sum-closed classes whose indecomposables were enumerated by each sequence $\{(a_n) : a_n \in A_n\}$. Then the set of growth rates of the classes would include the interval extending from $\text{gr}(\oplus(1, 1, \dots)) = 2$ to the lesser of $\text{gr}(\oplus(1, 2, 3, 2, 3, \dots)) \approx 2.51155$ (the real root of a cubic) and $\frac{1}{2}(1 + \sqrt{13})$, the maximum value permitted by the gap inequalities. (As with our previous example, such a construction is, in fact, impossible.)

We also need to use the following elementary analytic result, also found in Vatter [158].

Lemma 8.9 ([158, Proposition 2.3]). *Given $\varepsilon > 0$ and $c > 0$, there is a positive integer m , dependent only on ε and c , such that if (r_n) and (s_n) are two sequences of positive integers*

bounded by c , and $r_n = s_n$ for all $n \leq m$, then $\text{gr}(\oplus(r_n))$ and $\text{gr}(\oplus(s_n))$ differ by no more than ε .

Proof. Let $\gamma_1 = \text{gr}(\oplus(r_n))$ and $\gamma_2 = \text{gr}(\oplus(s_n))$ and suppose that $\gamma_1 < \gamma_2$. Note that $\gamma_1 \geq 2$ because (r_n) is positive, and $\gamma_2 \leq c + 1$ because (s_n) is bounded by c .

From Lemma 8.5, we have $\sum_{n=1}^{\infty} r_n \gamma_1^{-n} = \sum_{n=1}^{\infty} s_n \gamma_2^{-n}$. Hence, since $r_n = s_n$ for $n \leq m$,

$$\frac{\gamma_2 - \gamma_1}{(c + 1)^2} \leq \gamma_1^{-1} - \gamma_2^{-1} \leq \sum_{n \leq m} r_n (\gamma_1^{-n} - \gamma_2^{-n}) = \sum_{n > m} s_n \gamma_2^{-n} - \sum_{n > m} r_n \gamma_1^{-n} \leq \frac{c}{2^m}.$$

Thus, $m = \left\lceil \log_2 \frac{c(c+1)^2}{\varepsilon} \right\rceil$ suffices. □

8.4 Constructions that yield intervals of growth rates

To prove our theorems, we construct families of permutation classes whose indecomposables are enumerated by sequences satisfying the requirements of Lemma 8.8. In doing this, it helps to take a graphical perspective. Recall from Section 1.2 that the (ordered) graph, G_σ , of a permutation σ of length n has vertex set $\{1, \dots, n\}$ with an edge between vertices i and j if $i < j$ and $\sigma(i) > \sigma(j)$. Recall also that $\tau \leq \sigma$ if and only if G_τ is an induced ordered subgraph of G_σ . Observe too that σ is indecomposable if and only if G_σ is *connected*. It tends to be advantageous to think of indecomposables as those permutations whose graphs are connected.




Figure 8.1: Primary and secondary oscillations

All our constructions are based on permutations whose graphs are *paths*. Such permutations are called *increasing oscillations*. We distinguish two cases, the *primary oscillation* of length n , which we denote ω_n , whose *least* entry corresponds to an end of the path, and the *secondary oscillation* of length n , denoted $\bar{\omega}_n$ whose *first* entry corresponds to an end of the path. See Figure 8.1 for an illustration. Trivially, we have $\omega_1 = \bar{\omega}_1 = 1$ and $\omega_2 = \bar{\omega}_2 = 21$. The (lower and upper) *ends* of an increasing oscillation are the entries corresponding to the (lower and upper) ends of its path graph.




Figure 8.2: The set of permutations $R_7^{4,3}$

The key permutations in our constructions are formed by *inflating* (i.e. replacing) each end of a primary oscillation of odd length with an increasing permutation. For $n \geq 4$ and $r, s \geq 1$, let $\omega_n^{r,s}$ denote the permutation of length $n - 2 + r + s$ formed by replacing the lower and upper ends of ω_n with the increasing permutations of length r and s respectively. The graph of $\omega_n^{r,s}$ thus consists of a path on $n - 2$ vertices with r pendant edges attached to its lower end and s pendant edges attached to its upper end. See Figure 8.2 for some examples. We also occasionally make use of $\bar{\omega}_n^{r,s}$, which we define analogously.

The building blocks for our constructions are the following sets of indecomposables. Given a primary oscillation with both ends inflated, i.e. some $\omega_n^{r,s}$ with $r, s \geq 2$, we define $R_n^{r,s}$ to be the set of primary oscillations with both ends inflated that are subpermutations of $\omega_n^{r,s}$:

$$R_n^{r,s} = \{\omega_n^{u,v} : 2 \leq u \leq r, 2 \leq v \leq s\}.$$

We only make use of $R_n^{r,s}$ for odd n . See Figure 8.2 for an example.

Let us investigate the properties of these sets. Firstly, if $n \neq m$, then no element of $R_n^{r,s}$ is a subpermutation of an element of $R_m^{r,s}$. This is a consequence of the following elementary observation, which follows directly from consideration of the (ordered) graphs of the permutations.

Observation 8.10. *If $n, m \geq 4$, $n \neq m$ and $r, s, u, v \geq 2$, then $\omega_n^{r,s}$ and $\omega_m^{u,v}$ are incomparable under the subpermutation order.*

Thus for fixed r and s , and varying n , the $R_n^{r,s}$ form a collection of sets of permutations, no member of one being a subpermutation of a member of another. This should be compared and contrasted with the concept of an *antichain*, which is a set of permutations, none of which is a subpermutation of another.

We build permutation classes by specifying that, for some fixed r and s , their indecomposables must include some subset of $R_n^{r,s}$ for each (odd) n . Because, if $n \neq m$, any element of $R_n^{r,s}$ is incomparable with any element of $R_m^{r,s}$, the choice of subset to include can be made independently for each n . This provides the flexibility we need to construct

families of classes whose growth rates include an interval. (Vatter’s simpler constructions in [158], which are also based on inflating the ends of increasing oscillations, rely on being able to choose any subset of an infinite antichain to include in the indecomposables of a class.)

Clearly, for each n , the subset of $R_n^{r,s}$ included in the indecomposables must be downward closed. It has been observed heuristically that subsets containing $\omega_n^{3,2}$ are better for generating intervals of growth rates. In light of this, we define $\mathcal{F}_n^{r,s}$ to be the family of downward closed subsets of $R_n^{r,s}$ that contain $\omega_n^{3,2}$ (and hence also contain $\omega_n^{2,2}$).

For example, $\mathcal{F}_n^{4,3}$ consists of the seven downsets whose sets of maximal elements are

$$\{\omega_n^{3,2}\}, \{\omega_n^{3,2}, \omega_n^{2,3}\}, \{\omega_n^{3,3}\}, \{\omega_n^{4,2}\}, \{\omega_n^{4,2}, \omega_n^{2,3}\}, \{\omega_n^{4,2}, \omega_n^{3,3}\}, \text{ and } \{\omega_n^{4,3}\}.$$

To each downset S in $\mathcal{F}_n^{r,s}$, we can associate a generalised digit $c_0.c_1 \dots c_k$ such that c_i records the number of elements in S of length $n + 2 + i$ for each $i \geq 0$. For example, the generalised digits associated with the elements of $\mathcal{F}_n^{4,3}$ are (in the same order as above)

$$1.1, 1.2, 1.21, 1.11, 1.21, 1.22 \text{ and } 1.221.$$

We also need to take into account the indecomposables that are subpermutations of $\omega_n^{r,s}$ but are *not* elements of $R_n^{r,s}$. For odd n , these are of the following types:

- Primary oscillations ω_m and secondary oscillations $\bar{\omega}_m$.
- Permutations whose graphs are stars $K_{1,u}$ and whose first (and greatest) entry corresponds to the internal vertex; we use ψ_u to denote these star permutations.
- Increasing oscillations with just one end inflated: $\omega_m^{u,1}$, $\omega_m^{1,v}$ and $\bar{\omega}_m^{1,v}$.




Figure 8.3: The elements of $Q^{4,3}$ of size 6

Given $r, s \geq 2$, let $Q^{r,s}$ be the (infinite) set of permutations that are subpermutations of $\omega_n^{r,s}$ for some odd $n \geq 5$, but are not elements of $R_n^{r,s}$. See Figure 8.3 for an illustration.

We now have the building blocks we need. Given $r \geq 3$, $s \geq 2$, and odd $k \geq 5$, we define $\Phi_{r,s,k}$ to be the family of those permutation classes whose indecomposables are the union of $Q^{r,s}$ together with an element of $\mathcal{F}_n^{r,s}$ for each odd $n \geq k$:

$$\Phi_{r,s,k} = \left\{ \bigoplus (Q^{r,s} \cup S_k \cup S_{k+2} \cup S_{k+4} \cup \dots) : S_n \in \mathcal{F}_n^{r,s}, n = k, k+2, \dots \right\}.$$

The sequence of families $(\Phi_{5,3,k})_{k=5,7,\dots}$ is what we need to prove our first theorem, which we restate below. The set of growth rates of permutation classes in each family $\Phi_{5,3,k}$ consists of an interval, I_k , such that the sequence (I_k) of these intervals approaches θ_B from above.

We make use of one additional notational convenience. A sequence of integers (a_n) whose entries have the same value for all $n \geq k$ is denoted $(a_1, a_2, \dots, a_{k-1}, \overline{a_k})$. For example, $(1, \overline{4}) = (1, 4, 4, 4, \dots)$. We also use $(0^k, a_{k+1}, \dots)$ for a sequence whose first k terms are zero. For example, $(0^3, 1) = (0, 0, 0, 1)$.

Theorem 8.1. *Let $\theta_B \approx 2.35526$ be the unique real root of $x^7 - 2x^6 - x^4 - x^3 - 2x^2 - 2x - 1$. For any $\varepsilon > 0$, there exist δ_1 and δ_2 with $0 < \delta_1 < \delta_2 < \varepsilon$ such that every value in the interval $[\theta_B + \delta_1, \theta_B + \delta_2]$ is the growth rate of a permutation class.*

Proof. Let (q_n) be the sequence that enumerates $Q^{5,3}$. It can readily be checked that we have $(q_n) = (1, 1, 2, 3, 5, 7, \overline{8})$, and that $\text{gr}(\oplus(q_n)) = \theta_B$.

Let $k \geq 5$ be odd. For each odd $n \geq k + 2$, let

$$F_n = \{1.1, 1.11, 1.111, 1.2, 1.21, 1.211, 1.22, 1.221, 1.222, 1.2221\}$$

be the set of generalised digits associated with the sets of indecomposables in $\mathcal{F}_{n-2}^{5,3}$. Otherwise (if n is even or $n \leq k$), let $F_n = \{0\}$. Now, for each n , let $A_n = \{q_n + f_n : f_n \in F_n\}$.

So, by construction, for every permutation class $\oplus S \in \Phi_{5,3,k}$ there is a corresponding sequence (a_n) , with each $a_n \in A_n$, that enumerates S .

Let $\ell_n = q_n + 1.1$ for odd $n \geq k + 2$ and $\ell_n = q_n$ otherwise. Similarly, let $u_n = q_n + 1.2221$ for odd $n \geq k + 2$ and $u_n = q_n$ otherwise. We have the following equivalences:

$$\begin{aligned} (\ell_n) &\equiv (1, 1, 2, 3, 5, 7, \overbrace{8, 8, \dots, 8}^{k-5}, \overline{9}) \\ (u_n) &\equiv (1, 1, 2, 3, 5, 7, \overbrace{8, 8, \dots, 8}^{k-5}, \overline{9, 10, 11, 12}) \end{aligned}$$

We now apply Lemma 8.8. It can be checked that the gap inequalities necessitate only that the growth rate does not exceed $\gamma_{\max} \approx 2.470979$, the unique positive root of the quartic $x^4 - 2x^3 - x^2 - 1$. This is independent of the value of k , and is greater than $\text{gr}(\oplus(u_n))$ for all odd $k \geq 5$ since $\text{gr}(\oplus(u_n)) \approx 2.362008$ if $k = 5$.

So, for each k we have an interval of growth rates: If γ is such that it is the case that $\text{gr}(\oplus(\ell_n)) \leq \gamma \leq \text{gr}(\oplus(u_n))$, then there is some permutation class in $\Phi_{5,3,k}$ whose growth rate is γ .

Moreover, by Lemma 8.9,

$$\lim_{k \rightarrow \infty} \text{gr}(\oplus(\ell_n)) = \lim_{k \rightarrow \infty} \text{gr}(\oplus(u_n)) = \text{gr}(\oplus(q_n)) = \theta_B,$$

so these intervals can be found arbitrarily close to θ_B . \square

For our second theorem, we need to add extra sets of indecomposables to our constructions. As before, we start with $r \geq 3$, $s \geq 2$, and odd $k \geq 5$. A suitable collection, \mathcal{H} , of extra sets of indecomposables satisfies the following two conditions:

- Each set in \mathcal{H} is disjoint from $Q^{r,s}$ and also disjoint from each $R_n^{r,s}$ for odd $n \geq k$.
- For each set $H \in \mathcal{H}$, the union $H \cup Q^{r,s}$ is a downward closed set of indecomposables.

Given these conditions, we define $\Phi_{r,s,k,\mathcal{H}}$ to be the family of those permutation classes whose indecomposables are the union of $Q^{r,s}$ together with an element of \mathcal{H} and an element of $\mathcal{F}_n^{r,s}$ for each odd $n \geq k$:

$$\Phi_{r,s,k,\mathcal{H}} = \left\{ \oplus(Q^{r,s} \cup H \cup S_k \cup S_{k+2} \cup S_{k+4} \cup \dots) : H \in \mathcal{H} \text{ and } S_n \in \mathcal{F}_n^{r,s}, n = k, k+2, \dots \right\}.$$

We define our extra sets of indecomposables by specifying an upper set of *maximal* indecomposables, U , and a lower set of *required* indecomposables, L . If S is a set of indecomposables, let $\downarrow S$ denote the downset consisting of those indecomposables that are subpermutations of elements of S . Then, given $r, s \geq 2$ and suitable sets of indecomposables U and L , we use $U \downarrow_{r,s} L$ to denote the collection of those downward closed subsets of $\downarrow U \setminus Q^{r,s}$ that include the set of required indecomposables L . For instance, we have $\mathcal{F}_n^{r,s} = \{\omega_n^{r,s}\} \downarrow_{r,s} \{\omega_n^{3,2}\}$.

Let's consider as an example the family $\Phi_{5,3,5,\mathcal{H}}$ where $\mathcal{H} = \{\omega_5^{7,1}\} \downarrow_{5,3} \{\psi_7\}$; see π_0 and μ_1 in Figure 8.4 below. The sets in \mathcal{H} consist of indecomposables that are subpermutations of $\omega_5^{7,1}$ but are not in $Q^{5,3}$. There are six such indecomposables: ψ_7 , ψ_8 , $\omega_4^{6,1}$, $\omega_4^{7,1}$, $\omega_5^{6,1}$, and $\omega_5^{7,1}$. The collection \mathcal{H} consists of the nine downward closed subsets of these six that contain ψ_7 .

The set of indecomposables in a permutation class in our family consists of $Q^{5,3}$ together with an element of $\mathcal{F}_n^{5,3}$ for each odd $n \geq 5$ and an extra set from \mathcal{H} . As in the proof of Theorem 8.1, $Q^{5,3}$ contributes $(q_n) = (1, 1, 2, 3, 5, 7, \bar{8})$ to the enumeration of the indecomposables, and for each odd $n \geq 5$, there are ten distinct generalised digits associated with sets of indecomposables in $\mathcal{F}_n^{5,3}$, ranging between 1.1 and 1.2221. Let F_n consist of this set of generalised digits for odd $n \geq 7$ and otherwise contain only 0. The extra sets in \mathcal{H} have seven distinct enumerations. These can be represented by the set of generalised

digits $H_8 = \{1, 1.1, 1.11, 1.2, 1.21, 1.22, 1.221\}$; let $H_n = \{0\}$ for $n \neq 8$. Now, for each n , let $A_n = \{q_n + f_n + h_n : f_n \in F_n, h_n \in H_n\}$.

So, by construction, for every permutation class $\oplus \mathcal{S} \in \Phi_{5,3,5,\mathcal{H}}$ there is a corresponding sequence (a_n) , with each $a_n \in A_n$, that enumerates \mathcal{S} . The minimal enumeration sequence is $(\ell_n) \equiv (1, 1, 2, 3, 5, 7, 9, 10, \bar{9})$ for which the growth rate is $\text{gr}(\oplus(\ell_n)) \approx 2.36028$. Similarly, the maximal enumeration sequence is $(u_n) \equiv (1, 1, 2, 3, 5, 7, 9, 11, 13, 14, 13, \bar{12})$ for which the growth rate is $\text{gr}(\oplus(u_n)) \approx 2.36420$. The gap inequalities only necessitate that the growth rate not exceed $\gamma_{\max} \approx 2.47098$. Thus the growth rates of permutation classes in our example family $\Phi_{5,3,5,\mathcal{H}}$ form an interval.

The proof of our second theorem follows similar lines to this example.

Theorem 8.2. *Let $\lambda_B \approx 2.35698$ be the unique positive root of $x^8 - 2x^7 - x^5 - x^4 - 2x^3 - 2x^2 - x - 1$. Every value at least λ_B is the growth rate of a permutation class.*

Proof. In [158], Vatter has shown that there are permutation classes of every growth rate at least $\lambda_A \approx 2.48187$ (the unique real root of $x^5 - 2x^4 - 2x^2 - 2x - 1$). It thus suffices to exhibit permutation classes whose growth rates cover the interval $[\lambda_B, \lambda_A]$. With the π_i and μ_j as in Figures 8.4–8.6 below, we claim that the permutation classes in the following five families meet our needs:

- Family A:** $\Phi_{5,3,7,\mathcal{A}}$ where $\mathcal{A} = \{\pi_1\} \downarrow_{5,3} \{\mu_1\}$
- Family B:** $\Phi_{5,3,5,\mathcal{B}}$ where $\mathcal{B} = \{\pi_2\} \downarrow_{5,3} \emptyset$
- Family C:** $\Phi_{9,8,5,\mathcal{C}}$ where $\mathcal{C} = \{\pi_3\} \downarrow_{9,8} \{\mu_2\}$
- Family D:** $\Phi_{5,3,5,\mathcal{D}}$ where $\mathcal{D} = \{\pi_4, \pi_5\} \downarrow_{5,3} \{\mu_3\}$
- Family E:** $\Phi_{5,5,5,\mathcal{E}}$ where $\mathcal{E} = \{\pi_6, \pi_7, \pi_8\} \downarrow_{5,5} \{\mu_3, \mu_6\} \cup \{\pi_6, \pi_7, \pi_8\} \downarrow_{5,5} \{\mu_2, \mu_4, \mu_5\}$




Figure 8.4: Permutations used to define Families A and B


We briefly outline the calculations concerning each of these families. The details can be fleshed out in the same way as the example above.

Family A: $\Phi_{5,3,7,\mathcal{A}}$

- $Q^{5,3}$ is enumerated by $(1, 1, 2, 3, 5, 7, \bar{8})$.
- For each odd $n \geq 7$, there are 10 distinct generalised digits associated with sets of indecomposables in $\mathcal{F}_n^{5,3}$, ranging between 1.1 and 1.2221.
- There are 47 distinct enumerations of sets of indecomposables in $\mathcal{A} = \{\pi_1\} \downarrow_{5,3} \{\mu_1\}$, ranging between $(0^7, 1)$ and $(0^7, 1, 2, 3, 4, 4, 3, 2, 1)$.
- The indecomposables in the smallest permutation class in Family A are enumerated by $(\ell_n) \equiv (1, 1, 2, 3, 5, 7, 8, \bar{9})$.
- The indecomposables in the largest permutation class in Family A are enumerated by $(u_n) \equiv (1, 1, 2, 3, 5, 7, 8, 9, 11, 13, 15, 16, 15, 14, 13, \bar{12})$.
- The gap inequalities require the growth rate not to exceed $\gamma_{\max} \approx 2.470979$.
- $\text{gr}(\oplus(\ell_n)) = \lambda_B \approx 2.356983$; $\text{gr}(\oplus(u_n)) \approx 2.359320$.

Family B: $\Phi_{5,3,5,\mathcal{B}}$

- $Q^{5,3}$ is enumerated by $(1, 1, 2, 3, 5, 7, \bar{8})$.
- For each odd $n \geq 5$, there are 10 distinct generalised digits associated with sets of indecomposables in $\mathcal{F}_n^{5,3}$, ranging between 1.1 and 1.2221.
- There are 29 distinct enumerations of sets of indecomposables in $\mathcal{B} = \{\pi_2\} \downarrow_{5,3} \emptyset$, ranging between (0) and $(0^5, 1, 2, 3, 3, 1)$.
- The indecomposables in the smallest permutation class in Family B are enumerated by $(\ell_n) \equiv (1, 1, 2, 3, 5, 7, \bar{9})$.
- The indecomposables in the largest permutation class in Family B are enumerated by $(u_n) \equiv (1, 1, 2, 3, 5, 8, 11, 13, 14, 13, \bar{12})$.
- The gap inequalities require the growth rate not to exceed $\gamma_{\max} \approx 2.470979$.
- $\text{gr}(\oplus(\ell_n)) \approx 2.359304$; $\text{gr}(\oplus(u_n)) \approx 2.375872$.

**Figure 8.5:** Permutations used to define Families C, D and E**Family C:** $\Phi_{9,8,5,\mathcal{C}}$

- $Q^{9,8}$ is enumerated by $(1, 1, 2, 3, 5, 7, 9, 11, 13, 15, \bar{17})$.

- For each odd $n \geq 5$, there are 574 distinct generalised digits associated with sets of indecomposables in $\mathcal{F}_n^{9,8}$, ranging between 1.1 and 1.2345677654321.
- There are 19 distinct enumerations of sets of indecomposables in $\mathcal{C} = \{\pi_3\} \downarrow_{9,8} \{\mu_2\}$, ranging between $(0^5, 1)$ and $(0^5, 1, 3, 4, 3, 1)$.
- The indecomposables in the smallest permutation class in Family C are enumerated by $(\ell_n) \equiv (1, 1, 2, 3, 5, 8, 10, 12, 14, 16, \overline{18})$.
- The indecomposables in the largest permutation class in Family C are enumerated by $(u_n) \equiv (1, 1, 2, 3, 5, 8, 13, 17, 20, 22, 26, 29, 33, 36, 39, 41, 43, 44, \overline{45})$.
- The gap inequalities require the growth rate not to exceed $\gamma_{\max} \approx 2.786389$.
- $\text{gr}(\oplus(\ell_n)) \approx 2.373983$; $\text{gr}(\oplus(u_n)) \approx 2.389043$.

Family D: $\Phi_{5,3,5,\mathcal{D}}$

- $Q^{5,3}$ is enumerated by $(1, 1, 2, 3, 5, 7, \overline{8})$.
- For each odd $n \geq 5$, there are 10 distinct generalised digits associated with sets of indecomposables in $\mathcal{F}_n^{5,3}$, ranging between 1.1 and 1.2221.
- There are 37 distinct enumerations of sets of indecomposables in $\mathcal{D} = \{\pi_4, \pi_5\} \downarrow_{5,3} \{\mu_3\}$, ranging between $(0^3, 1)$ and $(0^3, 1, 2, 2, 2, 1)$.
- The indecomposables in the smallest permutation class in Family D are enumerated by $(\ell_n) \equiv (1, 1, 2, 4, 5, 7, \overline{9})$.
- The indecomposables in the largest permutation class in Family D are enumerated by $(u_n) \equiv (1, 1, 2, 4, 7, 9, 11, \overline{12})$.
- The gap inequalities require the growth rate not to exceed $\gamma_{\max} \approx 2.470979$.
- $\text{gr}(\oplus(\ell_n)) \approx 2.389038$; $\text{gr}(\oplus(u_n)) \approx 2.430059$.




Figure 8.6: Permutations used to define Family E

Family E: $\Phi_{5,5,5,\mathcal{E}}$

- $Q^{5,5}$ is enumerated by $(1, 1, 2, 3, 5, 7, 9, \overline{10})$.
- For each odd $n \geq 5$, there are 26 distinct generalised digits associated with sets of indecomposables in $\mathcal{F}_n^{5,5}$, ranging between 1.1 and 1.234321.

- There are 61 distinct enumerations of sets of indecomposables in $\mathcal{E} = \{\pi_6, \pi_7, \pi_8\} \downarrow_{5,5} \{\mu_3, \mu_6\} \cup \{\pi_6, \pi_7, \pi_8\} \downarrow_{5,5} \{\mu_2, \mu_4, \mu_5\}$, ranging between $(0^3, 1, 2, 1)$ and $(0^3, 2, 3, 5, 4, 1)$.
- The indecomposables in the smallest permutation class in Family E are enumerated by $(\ell_n) \equiv (1, 1, 2, 4, 7, 8, 10, \overline{11})$.
- The indecomposables in the largest permutation class in Family E are enumerated by $(u_n) \equiv (1, 1, 2, 5, 8, 12, 14, 13, 14, 16, 17, \overline{18})$.
- The gap inequalities require the growth rate to be at least $\gamma_{\min} \approx 2.363728$, but not to exceed $\gamma_{\max} \approx 2.489043$.
- $\text{gr}(\oplus(\ell_n)) \approx 2.422247$; $\text{gr}(\oplus(u_n)) \approx 2.485938 > \lambda_A$.

Here is a summary:

	<i>Enumeration of smallest and largest sets of indecomposables</i>	<i>Interval covered</i>
A	$(1, 1, 2, 3, 5, 7, 8, \overline{9})$ $(1, 1, 2, 3, 5, 7, 8, 9, 11, 13, 15, 16, 15, 14, 13, \overline{12})$	2.356983 – 2.359320
B	$(1, 1, 2, 3, 5, 7, \overline{9})$ $(1, 1, 2, 3, 5, 8, 11, 13, 14, 13, \overline{12})$	2.359304 – 2.375872
C	$(1, 1, 2, 3, 5, 8, 10, 12, 14, 16, \overline{18})$ $(1, 1, 2, 3, 5, 8, 13, 17, 20, 22, 26, 29, 33, 36, 39, 41, 43, 44, \overline{45})$	2.373983 – 2.389043
D	$(1, 1, 2, 4, 5, 7, \overline{9})$ $(1, 1, 2, 4, 7, 9, 11, \overline{12})$	2.389038 – 2.430059
E	$(1, 1, 2, 4, 7, 8, 10, \overline{11})$ $(1, 1, 2, 5, 8, 12, 14, 13, 14, 16, 17, \overline{18})$	2.422247 – 2.485938

Thus we have five intervals of growth rates that cover $[\lambda_B, \lambda_A]$. □

Part III

HASSE GRAPHS

Chapter 9

Introducing Hasse graphs

We now turn to the investigation of the structure of various classes of permutations, by considering how the *Hasse graphs* of permutations in these classes can be built from a sequence of rooted source graphs. For most of the classes, the analysis leads to functional equations for their generating functions. In some cases, we are able to use the kernel method to solve these equations and determine the corresponding algebraic generating function. Our approach is similar to that of “adding a slice”, used previously to enumerate various classes of polyominoes and other combinatorial structures.

9.1 Building Hasse graphs

Corresponding to each permutation σ , we define an ordered plane graph H_σ , which we call its *Hasse graph*. To create the Hasse graph for a permutation $\sigma = \sigma_1 \dots \sigma_n$, let vertex i be the point (i, σ_i) in the Euclidean plane. Now, for each pair i, j such that $i < j$, add an edge between vertices i and j , if and only if $\sigma(i) < \sigma(j)$ and there is no vertex k such that $i < k < j$ and $\sigma(i) < \sigma(k) < \sigma(j)$. Note that the edges of H_σ correspond to the edges of the Hasse diagram of the sub-poset, P_σ , of \mathbb{N}^2 consisting of the points (i, σ_i) . See Figure 9.1 for an illustration. Observe, also, that the Hasse graph of a permutation σ is the transitive reduction of the reversal of the ordered (inversion) graph of the reversal of σ . In practice, we tend not to distinguish between a permutation and its Hasse graph. The minimal elements of the poset P_σ are the left-to-right minima of the permutation σ . Similarly, maximal elements of P_σ are the right-to-left maxima of σ .

Hasse graphs of permutations were previously considered by Bousquet-Mélou & Butler [55], who determined the algebraic generating function of the class of *forest-like* permutations whose Hasse graphs are acyclic. We revisit forest-like permutations in Chapter 12 below.




Figure 9.1: The Hasse graph of a permutation, showing its four source graphs, rooted at the circled vertices

Given a permutation σ , we partition the vertices of H_σ by spanning it with a sequence of graphs, which we call the *source graphs* of σ . There is one source graph for each left-to-right minimum of σ . Suppose u_1, \dots, u_m are the vertices of H_σ corresponding to the left-to-right minima of σ , listed from left to right. Then the k th source graph G_k is the graph induced by u_k and those vertices of H_σ lying above and to the right of u_k that are not in G_1, \dots, G_{k-1} . We refer to u_k as the *root* of source graph G_k . See Figure 9.1 for an illustration. In the figures, edges of source graphs are shown as thick lines.

The structure of the source graphs of permutations in a specific permutation class is constrained by the need to avoid the patterns in the basis of the class. If the source graphs for a class are acyclic, we refer to them as *source trees*.

The *bottom subgraph* of a Hasse graph is the graph induced by its lowest vertex (the least entry in the permutation) and all the vertices lying above and to its right. Observe that the bottom subgraph may contain vertices from more than one source graph. For example, the bottom subgraph of the Hasse graph in Figure 9.1 contains vertices from three source graphs. Bottom subgraphs of permutations in a specific permutation class satisfy the same structural restrictions as do the source graphs. We refer to an acyclic bottom subgraph as a *bottom subtree*.

We build the Hasse graph of a permutation by starting with a single source graph and then repeatedly adding another source graph to the lower right, rooted at a new left-to-right minimum. The technique is similar to that of “adding a slice”, which has been used successfully to enumerate constrained compositions and other classes of polyominoes, a topic of interest in statistical mechanics (see, for example, Bousquet-Mélou’s review paper [52], the books of van Rensburg [156] and Guttmann [94], and [77, Examples III.22 and V.20]). When a source graph is added, its vertices are interleaved horizontally with the non-root vertices of the bottom subgraph of the graph built from the previous source graphs. Typically, the positioning of the vertices of the new source graph relative to those

of the bottom subgraph is constrained by the need to avoid forbidden patterns.

When a new source graph is added, additional edges may be added, each joining a vertex in the source graph to a vertex in the bottom subgraph. We call these edges *tendrils* and say that the vertices at the upper ends of the tendrils are *tethered* by the tendrils. Tendrils are represented by thin edges in Figure 9.1 and in other figures below. Clearly, the positioning of the vertices of the source graph relative to those of the bottom subgraph determines which tendrils are added, and *vice versa*. Rather than analysing the possible positions for the vertices of the source graph, it is often easier to consider the permissible choices for the tendrils instead.

9.2 Avoiding 1324

A number of the classes that we analyse in the following chapters are subclasses of $\text{Av}(1324)$. What can we say about the Hasse graphs of permutations in such classes?




Figure 9.2: The graph of a permutation avoiding 1324, showing its construction from five source trees

If a permutation avoids 1324, then its Hasse graph does not have the diamond graph $H_{1324} = \nearrow$ as a minor. In particular, each source graph of H_σ is a *tree*, and so is its bottom subgraph. We call the uppermost (leftmost) branch (path towards the upper right from the root) of a source tree its *trunk*. The topmost vertex of the trunk is its *tip*. When adding a new source tree to the lower right of a Hasse graph H , to avoid creating a 1324, any vertices of the source tree not in the trunk must be positioned to the right of all vertices of H . Thus the lower ends of any new tendrils created are on the trunk of the new source tree. Moreover, when a new source tree is added, no vertex in (the bottom subtree of) H is tethered by more than one tendril. See Figure 9.2 for an illustration.

9.3 Outline of Part III

In the remaining chapters of this thesis, we study seven different permutation classes. In four cases, analysis of the Hasse graph structure enables us to ascertain the algebraic generating function for the class. Two of these results are new. In two others, we are only able to establish functional equations. In the final case, the **1324**-avoiders, we achieve neither of these, but are able to determine a new lower bound for the growth rate of the class.

In Chapters 10–12, we consider three classes that avoid **1324**. In Chapter 10, we look at $\text{Av}(\mathbf{1324}, \mathbf{2314})$, which is counted by the (large) Schröder numbers, and present a new derivation of its algebraic generating function. This class was first enumerated by Kremer [115, 116].

In Chapter 11, we consider $\text{Av}(\mathbf{1324}, \mathbf{1432})$. The source graphs in this class have a very simple form. We derive a (very complicated) functional equation for the generating function of this class, which, unfortunately, we are unable to solve.

In Chapter 12, we use our methods to give a simplified derivation of the algebraic generating function for *forest-like* permutations, originally enumerated by Bousquet-Mélou & Butler [55]. This is the class whose Hasse graphs are acyclic. It is not a classical permutation class: the Hasse graph of **2143** is a cycle, so **2143** is not forest-like, but **2143** is contained in **21354**, the Hasse graph of which is a tree. The class of forest-like permutations can be defined as $\text{Av}(\mathbf{1324}, \mathbf{21\bar{3}54})$, where **21 $\bar{3}$ 54** is a barred pattern (see Section 1.2).

In Chapter 13, we investigate the class of *plane permutations*, those permutations whose Hasse graph is plane (i.e. non-crossing). This is the barred pattern class $\text{Av}(\mathbf{21\bar{3}54})$. We establish a nice functional equation for the generating function of this class, which we are not able to solve.

In Chapters 14 and 15, we enumerate two more permutation classes that each avoid two patterns of length 4. We establish the generating function for $\text{Av}(\mathbf{1234}, \mathbf{2341})$ in Chapter 14. In doing so, the kernel method is used six times to solve the relevant functional equations. In Chapter 15, we derive functional equations for $\text{Av}(\mathbf{1243}, \mathbf{2314})$, and solve them to determine its algebraic generating function. This requires an unusual simultaneous double application of the kernel method. These results have been accepted for publication (see [41]).

Finally, in Chapter 16, we use our Hasse graph techniques in an investigation of the structure of the notorious class $\text{Av}(\mathbf{1324})$. As a result, we establish a new lower bound of 9.81 for its growth rate. As a consequence of our examination of the asymptotic distribution of certain substructures in the Hasse graphs of **1324**-avoiders, we are able to prove that,

in the limit, patterns in Łukasiewicz paths exhibit a concentrated Gaussian distribution. This work has been accepted for publication (see [\[40\]](#)).

Chapter 10

Avoiding 1324 and 2314

We begin our investigation of Hasse graphs of permutations by studying $\text{Av}(1324, 2314)$. This class was first enumerated by Kremer [115, 116] using generating trees. It is counted by A006318 in OEIS [133]. In this chapter, we determine its generating function by analysing its Hasse graphs, which have a relatively simple structure.

Let us use \mathcal{R} to denote $\text{Av}(1324, 2314)$. Since permutations in \mathcal{R} avoid 1324, the source graphs and bottom subgraphs are trees. The avoidance of 2314 places no further restrictions on the structure of source trees because 2314 is not contained in any permutation whose Hasse graph is a tree rooted at the lower left.

In order to avoid 2314, a permutation can only contain occurrences of 213 for which there is no point to the lower left of the 2. Thus the 2 in any 213 must be a left-to-right minimum, the root of a source tree. In any 213, the edge joining the 1 to the 3 must be a tendril. Hence, when adding a new source tree, the upper end of each tendril must be a child of the root of the bottom subtree. Specifically, if the root of the bottom subtree has k children, v_1, \dots, v_k , from left to right, then the upper ends of the tendrils are v_i, \dots, v_k for some $i \in \{1, \dots, k\}$. See Figure 10.1 for an illustration. Clearly, the key parameter to keep account of is the root degree of the bottom subtree.

Let us imagine that source trees come with tendrils already attached to some of their trunk vertices. We use \mathcal{S} to denote the class of extended source trees. As we have noted, the set of tethered vertices in the bottom subtree can be determined from the number of tendrils. (Indeed, if there are m tendrils, they are the rightmost m children of the root of the bottom subtree.) Thus, an element of \mathcal{S} specifies both a specific source tree and also the positioning of its vertices relative to those of the bottom subtree.

Let us determine the structure of elements of \mathcal{S} . We use z to mark the number of vertices in the source tree. We do not count the (tethered) vertices at the upper ends of the tendrils.




Figure 10.1: The Hasse graph of a permutation in class \mathcal{R} , showing its construction from four source trees

In addition, we use y to mark the root degree (including counting any tendrils attached to the root). After a new source tree has been attached, this value is the root degree of the bottom subtree in the resulting permutation. Finally, we use x to mark the total number of tendrils. For example, the last source tree in Figure 10.1 would be represented by $z^3y^3x^3$.

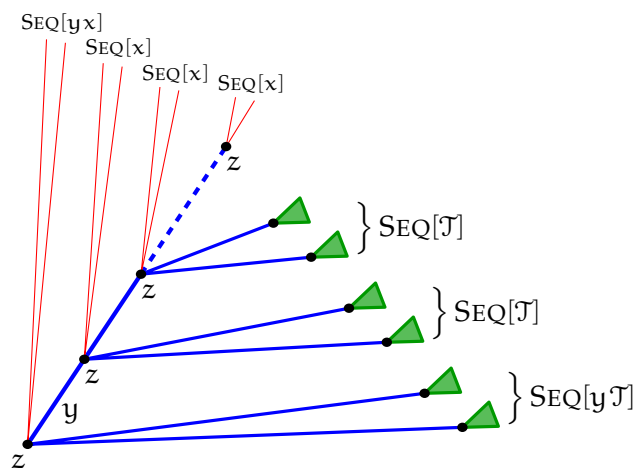


Figure 10.2: The structure of extended source trees for class \mathcal{R}

The class, \mathcal{S} , of source trees with tendrils attached satisfies the following structural equation:

$$\mathcal{S} = z\text{SEQ}[yx] + (z\text{SEQ}[yx]y\text{SEQ}[y\mathcal{T}] \times \text{SEQ}[z\text{SEQ}[x]\text{SEQ}[\mathcal{T}]] \times z\text{SEQ}[x]), \quad (1)$$

where \mathcal{T} represents plane trees.

The first summand corresponds to single vertex source trees, and the second to source trees with more than one vertex (and hence more than one trunk vertex). As illustrated in Figure 10.2, the first term in the parenthesized product corresponds to the root vertex and structures attached to it, the second term corresponds to intermediate vertices in the

trunk and structures attached to them, and the third term corresponds to the tip of the trunk and tendrils attached to it.

To account for child vertices of the root of the bottom subtree which are not tethered (vertices v_1, \dots, v_{i-1} above), we can consider these vertices as being the upper ends of “virtual” tendrils which are not attached at the lower end to any vertex of the source tree. Let \mathcal{S}^* denotes the extension of class \mathcal{S} to include these virtual tendrils. Thus, an element of \mathcal{S}^* consists of a source tree with tendrils attached to its trunk along with zero or more virtual tendrils. We have $\mathcal{S}^* = \text{SEQ}[x]\mathcal{S}$. Expanding (1) then gives us the following trivariate generating function for \mathcal{S}^* :

$$S^*(z, y, x) = \frac{z(x(1-x)(1-y) - z(1-y + y(1-yx)(t(z) - x)))}{(1-y + zy^2)(1-x)(1-yx)(x - x^2 - z)},$$

where $t(z) = \frac{1}{2}(1 - \sqrt{1 - 4z})$ is the generating function for \mathcal{T} , the class of rooted plane trees.

Let $R(y) = R(z, y)$ be the bivariate generating function for \mathcal{R} , in which y marks the root degree of the bottom subtree. Since a vertex in the bottom subtree can only be tethered by one tendril, adding a new source tree corresponds to the linear operator on $R(y)$ that maps y^d to $[x^d]S^*(z, y, x)$.

The result now follows by algebraic manipulation and the kernel method. It can be checked (using a computer algebra system, or otherwise) that

$$[x^d]S^*(z, y, x) = \frac{z}{1-y + zy^2} (1 + y(t(z) - 1 - y^d + z^{-d}t(z)^d)).$$

The class of initial source trees is defined by $\mathcal{Z} \times \text{SEQ}[y\mathcal{T}]$, so \mathcal{R} satisfies the functional equation

$$R(y) = \frac{z}{1-yt(z)} + \frac{z}{1-y + zy^2} (R(1) + y(t(z) - 1 - R(y) + R(t(z)/z))).$$

This can be solved for $R(1)$ by using the kernel method. Rearranging yields

$$(1 - y + zy + zy^2)R(y) = z(1 + R(1) + y(R(t(z)/z) - 1 - \sqrt{1 - 4z})). \quad (2)$$

The kernel can be cancelled by setting $y = (1 - z - \sqrt{z^2 - 6z + 1})/2z$, which results in an equation that can be solved to give $R(t(z)/z)$ in terms of $R(1)$. Finally, substituting for $R(t(z)/z)$ in (2), setting y to 1 and solving for $R(1)$ yields the following:

Theorem 10.1. *The class of permutations avoiding **1324** and **2314** has the algebraic generating function*

$$\frac{1}{2} \left(1 - z - \sqrt{1 - 6z + z^2} \right).$$

The corresponding numbers are known as the (large) Schröder numbers. These numbers also count the number of paths from the southwest corner $(0,0)$ of an $n \times n$ grid to the northeast corner (n,n) , using only single steps east, northeast, or north, that do not rise above the SW-NE diagonal. The following question may be worth consideration:

Question 10.2. *Is there a nice bijection between these lattice paths and the Hasse graphs of permutations in $\text{Av}(\mathbf{1324}, \mathbf{2314})$?*

Chapter 11

Avoiding 1324 and 1432

In this chapter, we investigate $\text{Av}(1324, 1432)$. This is one of the classes, listed in [169], with a basis consisting of two permutations of length 4, that remains to be enumerated. The source graphs in this class have a very simple form. Despite this, the functional equation we derive for its generating function is rather complicated, and we are unable to solve it.

Let us use \mathcal{S} to denote $\text{Av}(1324, 1432)$. As in the previous chapter, since \mathcal{S} avoids 1324, source graphs and bottom subgraphs are trees. The further requirement to avoid 1432 severely restricts the structure of both source trees and bottom subtrees: they must either be paths or else have a single fork. Let's call these *spindly trees*.




Figure 11.1: Spindly trees: a path, a forked tree, and a path with a tendrils attached

Let \mathcal{S} denote the class of spindly trees. Let us use

- w to mark vertices in path trees (except the root) and in the *right* (lower) branch of forked trees,
- v to mark vertices in the *left* (upper) branch of forked trees, and
- u to mark vertices in the *stem* of forked trees: the forked vertex and vertices below the fork (except the root).

There are thus two cases, source trees that are paths, the class of which we denote \mathcal{S}_P , and forked source trees, for which we use \mathcal{S}_F . These can be defined by the following two structural equations.

Firstly,

$$\mathcal{S}_P = \mathcal{Z} \times \text{SEQ}[w\mathcal{Z}].$$

Secondly,

$$\mathcal{S}_F = \mathcal{Z} \times \text{SEQ}[u\mathcal{Z}] \times \text{SEQ}^+[v\mathcal{Z}] \times \text{SEQ}^+[w\mathcal{Z}].$$

The corresponding generating functions are thus

$$S_P(w) = \frac{z}{1 - zw}$$

and

$$S_F(u, v, w) = \frac{z^3 vw}{(1 - zu)(1 - zv)(1 - zw)}.$$

The sum of these give the following generating function for \mathcal{S} :

$$S(u, v, w) = \frac{z(1 - zu(1 - zv) - zv(1 - zu))}{(1 - zu)(1 - zv)(1 - zw)}.$$

It is also helpful to consider the class consisting of path source trees with a tendril attached to any one of the vertices. We denote this class $\widehat{\mathcal{S}}_P$. We use w to mark the vertices on the path above the tendril and u to mark the rest except the root. Thus,

$$\widehat{\mathcal{S}}_P = \mathcal{Z} \times \text{SEQ}[u\mathcal{Z}] \times \text{SEQ}[w\mathcal{Z}],$$

for which the generating function is

$$\widehat{S}_P(u, w) = \frac{z}{(1 - zu)(1 - zw)}.$$

Let $G_C(u, v, w) = G_C(z, u, v, w)$ be the multivariate generating function for connected permutations in \mathcal{G} , in which u, v and w mark the non-root vertices in the spindly bottom subtree as described above.

When adding a new spindly source tree, no more than two tendrils can be added. We count *connected* (skew-indecomposable) permutations. The addition of a new spindly source tree can be broken down into four cases. We denote the corresponding linear operators on the generating function by $\Omega_1, \Omega_2, \Omega_3$ and Ω_4 .

Cases I, II and III are only applicable when the source tree is a path. These are illustrated in Figure 11.2.




Figure 11.2: Cases I, II and III: Three methods for adding a path

Case I

In case I, a path tree is inserted below a path bottom subtree or the right branch of a forked bottom subtree so that the resultant bottom subtree is a path, a tendrill being attached to the tip of the path. Thus,

$$\Omega_1[u^r v^s w^t] = \sum_{i=1}^t S_P(w) w^i = \frac{zw(1-w^t)}{(1-w)(1-zw)},$$

where i counts the tethered vertex in the right branch of the bottom subtree, starting from the tip. Hence,

$$\Omega_1[f(u, v, w)] = \frac{zw}{(1-w)(1-zw)} (f(1, 1, 1) - f(1, 1, w)).$$

Case II

In case II, a path tree is inserted to the left of the fork in a forked bottom subtree, a tendrill being attached to the tip of the path. The resultant bottom subtree is forked. Thus,

$$\Omega_2[u^r v^s w^t] = \sum_{i=1}^r S_P(u) u^i v^s w^t = \frac{zu(1-u^r)v^s w^t}{(1-u)(1-zu)},$$

where i counts the tethered vertex in the stem of the bottom subtree, starting from the forked vertex. Hence,

$$\Omega_2[f(u, v, w)] = \frac{zu}{(1-u)(1-zu)} (f(1, v, w) - f(u, v, w)).$$

Case III

In case III, a path tree is inserted below the left branch of a forked bottom subtree. The leftmost vertex in the right branch is always tethered to the tip of the path tree. The resultant bottom subtree is forked. Thus,

$$\Omega_3[u^r v^s w^t] = \sum_{i=1}^s \widehat{S}_P(u, w) v^i w^t = \frac{zv(1-v^s)w^t}{(1-zu)(1-v)(1-zw)},$$

where i counts the tethered vertex in the left branch of the bottom subtree, starting from the tip. Hence,

$$\Omega_3[f(u, v, w)] = \frac{zv}{(1-zu)(1-v)(1-zw)} (f(1, 1, w) - f(1, v, w)).$$




Figure 11.3: Case IV: The addition of a path and of a forked tree

Case IV

Case IV is illustrated in Figure 11.3. In this case, a path or forked tree is inserted below a path bottom subtree or the right branch of a forked bottom subtree so that the resultant bottom subtree is forked. In this case, the tendrils may not be attached to the tip of a path source tree. Thus,

$$\Omega_4[u^r v^s w^t] = \sum_{i=1}^t (\widehat{S}_P(u, w) - S_P(u) + S_F(u, v, w)) v^i = \frac{z^2vw(1-v^t)}{(1-zu)(1-v)(1-zv)(1-zw)},$$

and hence

$$\Omega_4[f(u, v, w)] = \frac{z^2vw}{(1-zu)(1-v)(1-zv)(1-zw)} (f(1, 1, 1) - f(1, 1, v)).$$

Thus we have the following:

Theorem 11.1. *The generating function $G_C(u, v, w)$ for connected permutations in the class $\text{Av}(1324, 1432)$ satisfies the functional equation*

$$G_C(u, v, w) = S(u, v, w) + (\Omega_1 + \Omega_2 + \Omega_3 + \Omega_4)[G_C(u, v, w)].$$

This is a very complex algebraic identity relating $G_C(u, v, w)$ to $G_C(1, v, w)$, $G_C(1, 1, w)$, $G_C(1, 1, v)$ and $G_C(1, 1, 1)$.

The kernel method can be applied to this to yield $G_C(1, v, w)$ as a function of $G_C(1, 1, w)$, $G_C(1, 1, v)$ and $G_C(1, 1, 1)$, and then used again to give $G_C(1, 1, w)$ in terms of $G_C(1, 1, 1)$ and $G_C(1, 1, \beta(w))$ for some algebraic function $\beta(w)$. This does not appear to help us to get an explicit form for $G_C(1, 1, 1)$. Despite the simplicity of the structure of the source trees of \mathcal{G} , its exact enumeration remains an open question.

However, the recurrence in Theorem 11.1 can be iterated to extract the coefficients of z^n from $G_C(1, 1, 1)$ and hence enumerate $\mathcal{G} = \text{Av}(1324, 1432)$ for small n . Values up to $n = 40$ (determined in seven hours by *Mathematica* [172]) can be found at [A165542](#) in OEIS [133].

Postscript: Following submission of this thesis, at the AMS meeting in Washington, DC, in March 2015, Nathaniel Shar presented a much more efficient recurrence for class \mathcal{G} ; his work is currently unpublished.

Chapter 12

Forest-like permutations

In this chapter, we consider *forest-like* permutations, $\text{Av}(\mathbf{1324}, \mathbf{21\bar{3}54})$, permutations whose Hasse graph is a plane forest. These are $\mathbf{1324}$ -avoiders whose Hasse graph is plane. This class is counted by [A111053](#) in OEIS [133]. It was first enumerated by Bousquet-Mélou & Butler [55]. Their approach makes use of a rather complicated analysis of the Hasse graphs of permutations. We present here a somewhat simpler derivation of the algebraic generating function for this class.

Since permutations in the class avoid $\mathbf{1324}$, the “tethering rules” for $\mathbf{1324}$ -avoiders presented in Section 9.2 apply. Source and bottom subgraphs are trees. The planarity of the Hasse graphs implies that the only vertices in the bottom subtree that can be tethered are those in its lowermost (rightmost) branch. As a consequence, to avoid creating a $\mathbf{1324}$, at most one tendril can be used when adding a new source tree. We count *connected* (skew-indecomposable) permutations in the class, those constructed by using exactly one tendril when attaching each new source tree. Clearly, we must keep track of the length of the lowermost branch of the bottom subtree.

We begin by presenting structural equations for source trees. We use u to mark vertices in the uppermost branch (the trunk) and v to mark vertices in the lowermost branch, excluding the root vertex. There are two cases, source trees that are paths, the class of which we denote \mathcal{S}_P , and source trees with at least one fork, for which we use \mathcal{S}_F . These can be defined by the following two structural equations.

Firstly,

$$\mathcal{S}_P = uZ \times \text{SEQ}[uvZ],$$

the first term corresponding to the root and the second to the remaining vertices.

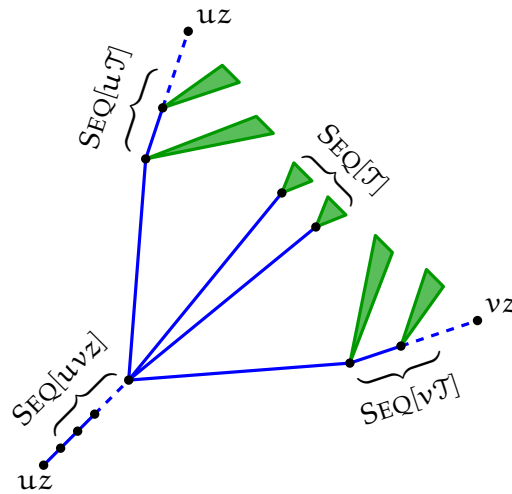


Figure 12.1: The structure of source trees for forest-like permutations

Secondly,

$$S_F = uz \times \text{SEQ}[uvz] \times (\text{SEQ}[u\mathcal{T}] \times uz) \times (\text{SEQ}[v\mathcal{T}] \times vz) \times \text{SEQ}[\mathcal{T}],$$

where \mathcal{T} represents plane trees.

This is illustrated in Figure 12.1. The first term in the product corresponds to the root. The second corresponds to the *stem* of the tree (vertices up to and including the first forked vertex). The third and fourth terms correspond to vertices above the stem lying on the uppermost and lowermost branches respectively, along with subtrees rooted at these vertices, the subterms uz and vz corresponding to the tips of the relevant branches. The final term corresponds to the remaining vertices, which lie on subtrees whose roots are children of the first forked vertex that are not on the uppermost or lowermost branch.

The corresponding trivariate generating functions are thus:

$$S_P(u, v) = \frac{zu}{1 - zuv},$$

$$S_F(u, v) = \frac{z^2 u^2 v t(z)}{(1 - zuv)(1 - ut(z))(1 - vt(z))},$$

where $t(z)$ is the generating function for plane trees. Let $S(u, v) = S_P(u, v) + S_F(u, v)$ be the generating function for all source trees.

We now designate one vertex in the trunk to which a tendril is to be attached. Once we have done this, we no longer need to keep track of the number of vertices in the trunk. We use the pointing construction (see Section 1.1) to derive the appropriate generating functions.

For source trees that are paths, we exclude the case in which the tendrils are attached to the tip (because we need to handle this case separately). If $\widehat{S}_P(v)$ is the corresponding bivariate generating function for paths with a tendrils attached to a non-tip vertex, then

$$\widehat{S}_P(v) = \partial_u S_P(u, v) \Big|_{u=1} - S_P(1, v) = \frac{z^2 v}{(1 - zv)^2}.$$

Analogously, if we use $\widehat{S}_F(v)$ for forked source trees with a tendrils attached to a trunk vertex, then

$$\widehat{S}_F(v) = \partial_u S_F(u, v) \Big|_{u=1} = \frac{z^2 vt(z)(2 - zv - t(z))}{(1 - zv)^2 (1 - t(z))^2 (1 - vt(z))}.$$

Let $\widehat{S}(v) = \widehat{S}_P(v) + \widehat{S}_F(v)$.

Let $L_C(v) = L_C(z, v)$ be the bivariate generating function for connected forest-like permutations, in which v marks the length (number of edges in) the lowermost branch of the bottom subtree.




Figure 12.2: The addition of a path and of a tree

There are two methods by which a new source tree can be added to a bottom subtree, as shown in Figure 12.2. In the first, a path source tree is added with a tendrils attached to its tip. In this case, the lowermost branch of the result may contain vertices from the bottom subtree. The action of adding a path source tree in this way is thus seen to be reflected by the linear operator Ω_1 , acting on $L_C(v)$, that satisfies

$$\Omega_1[v^k] = \sum_{i=1}^k S_P(1, v) v^i = \frac{zv(1 - v^k)}{(1 - zv)(1 - v)},$$

where i counts the tethered vertex in the lowermost branch of the bottom subtree, starting from the tip. Hence,

$$\Omega_1[f(v)] = \frac{zv(f(1) - f(v))}{(1 - zv)(1 - v)}.$$

In the second method, a source tree (which may be a path) is added in such a way that some of its vertices are positioned to the right of the bottom subtree. This action is reflected by the linear operator Ω_2 that satisfies

$$\Omega_2[v^k] = k\widehat{S}(v),$$

since there are k choices for the upper end of the tendril. Hence,

$$\Omega_2[f(v)] = \widehat{S}(v) f'(1).$$

The class of initial source trees is enumerated by $S(1, v)$. Hence, the generating function for connected forest-like permutations, $L_C(v)$, satisfies the functional equation

$$L_C(v) = S(1, v) + (\Omega_1 + \Omega_2)[L_C(v)],$$

which expands to give

$$L_C(v) = S(1, v) + \frac{zv(L_C(1) - L_C(v))}{(1-zv)(1-v)} + \widehat{S}(v)L'_C(1). \quad (1)$$

Taking the limit as v tends to 1, we get

$$L_C(1) = S(1, 1) + \frac{z}{(1-z)}L'_C(1) + \widehat{S}(1)L'_C(1). \quad (2)$$

If we eliminate $L'_C(1)$ from (1) and (2), solve for $L_C(v)$, expand and simplify, then

$$L_C(v) = z \frac{(1-v)(1-2zv - (1-zv)t(z)) - v(zv - (2-v)(1-t(z)))L_C(1)}{(1-zv - t(z))(1-v + zv^2)}.$$

The kernel method now applies. Cancelling the kernel, by letting $v = t(z)/z$, and solving for $L_C(1)$ yields

$$L_C(1) = \frac{z - 3z^2 + (1 - 5z)t(z)}{2 - 9z},$$

and so the generating function for all forest-like permutations is given by

$$\frac{L_C(1)}{1 - L_C(1)} = \frac{z^2 + z^3 + (1 - 5z)t(z)}{1 - 5z + 2z^2 - z^3}.$$

Hence, expanding $t(z)$, we have:

Theorem 12.1. *The class of forest-like permutations has the algebraic generating function*

$$\frac{2z^2(1+z) + (1-5z)(1-\sqrt{1-4z})}{2(1-5z+2z^2-z^3)}.$$

Chapter 13

Plane permutations

In [55], Bousquet-Mélou & Butler ask the question, “What is the number of plane permutations?”. Plane permutations are those whose Hasse graph is plane (i.e. non-crossing). See Figure 13.1 for an example. (The permutations in the illustration on page vii are also plane permutations.) Plane permutations are characterised by avoiding the barred pattern $21\bar{3}54$. In this chapter, we investigate this class of permutations.

We determine a functional equation for this class. Our first observation is that plane permutations of size n are in bijection with plane source graphs with $n + 1$ vertices, by the addition/removal of a (single) root vertex to the lower left. So we consider plane source graphs, and then discount the root vertex.




Figure 13.1: A plane source graph with its spanning tree built from eight branches; the tips of the branches are numbered in the order they are added

To build a plane source graph, we repeatedly add *branches* (maximal increasing paths) towards the lower right. The initial branch consists of the left-to-right maxima of the permutation. These branches build a rooted spanning tree. See Figure 13.1 for an illustration.

Specifically, suppose we are given a plane source graph G and have already constructed a partial spanning tree, T , of G from a sequence of branches. Let x be the maximal (most north-easterly) vertex in T for which there is some vertex in $G \setminus T$ above it (to its north-east); x is among the right-left minima of T . Then the next branch is joined to T at x and consists of those left-to-right maxima of $G \setminus T$ that are above x . See Figure 13.1 for an illustration of this process. We call x the *joining vertex*.




Figure 13.2: Adding a new branch; open vertices, both before and after adding the branch, are circled

To keep track of the possibilities, at each step, we *mark* each of the current right-to-left minima as either *open* or *closed*. Open vertices are those to which the next branch may be joined. In the initial branch, all vertices except its tip are open. After a branch is added, the open vertices are those to the lower left of its tip; the tip and any vertices to its upper right are closed. See Figure 13.2 for an example of adding a branch. In this example, prior to the new branch being added, there were 6 open vertices and 5 closed vertices. Afterwards, there are 8 open vertices and 4 closed vertices.

Note that vertices in a new branch are positioned to the right of the marked (open or closed) vertex immediately above the joining vertex. They may be arbitrarily interleaved horizontally with any further marked vertices above the joining vertex, occurring immediately above the joining vertex to avoid creating a **2143**.

Given this construction process, it is now straightforward to specify a functional equation for the generating function. Let us use u for open vertices and v for closed vertices. Then the process of adding a new branch to a partial spanning tree T corresponds to the linear operator Ω satisfying the equation

$$\Omega[u^r v^s] = \sum_{i=1}^r \sum_{j=1}^{s+i-1} u^{r+1-i} \text{SEQ}[zu]^j z v^{s+i-j}. \quad (1)$$

Here, i indexes (from the upper right) the choice of joining vertex from among the open vertices of T , and j indexes (from the lower left) the choice for the relative position of

the tip of the branch from among the marked vertices of T above the joining vertex. For example, in Figure 13.2, the choices are $i = 4$ and $j = 5$. Each of the j occurrences of $\text{SEQ}[zu]$ corresponds to a sequence of new *closed* vertices inserted to the right of a marked vertex of T . In Figure 13.2, these sequences have lengths 3, 0, 0, 1, 1 from left to right.

We can now take the expansion and linear extension of (1) and combine it with an expression that enumerates the initial branch to yield a functional equation for the generating function of plane permutations. After discounting the root vertex, this gives us the following result:

Theorem 13.1. *The generating function for plane permutations is given by $P(1, 1)$, where*

$$P(u, v) = \frac{zuv}{1 - zu} + \frac{zuv}{1 - v - zuv} \left(\frac{P(u, q) - P(q, q)}{u - q} - \frac{P(u, v) - P(v, v)}{u - v} \right),$$

in which $q = \frac{1}{1 - zu}$.

It is not known how to solve this equation. However, this recurrence can be iterated to extract the coefficients of z^n from $M(1, 1)$ and hence enumerate $\text{Av}(21\bar{3}54)$ for small n . Values up to $n = 37$ (determined in twelve hours by *Mathematica* [172]) can be found at [A117106](#) in OEIS [133].

Conjectures

Two intriguing conjectures have been made concerning the enumeration of the class of plane permutations. Firstly, in the OEIS entry, Mark van Hoeij posits a connection between plane permutations and the sequence of *Apéry numbers*

$$a_n = \sum_{k=0}^n \binom{n}{k}^2 \binom{n+k}{k}.$$

These are known as Apéry numbers because, along with other sequences, they were used in Apéry's celebrated proof of the irrationality of $\zeta(2)$ and $\zeta(3)$ [16, 17] (see [137]). Van Hoeij's conjecture is as follows:

Conjecture 13.2 (van Hoeij; [A117106](#) in OEIS [133]). *The number of plane permutations, p_n , of length $n \geq 2$ is given by*

$$p_n = \frac{24((5n^3 - 5n + 6)a_{n+1} - (5n^2 + 15n + 18)a_n)}{5(n-1)n^2(n+2)^2(n+3)^2(n+4)},$$

where $a_n = \sum_{k=0}^n \binom{n}{k}^2 \binom{n+k}{k}$.

We have confirmed that this conjecture is consistent with the values calculated using the recurrence in Theorem 13.1.

Apéry found that a_n satisfies the following recurrence:

$$a_n = \frac{1}{n^2} ((11n^2 - 11n + 3)a_{n-1} + (n-1)^2 a_{n-2}).$$

This can be combined with Conjecture 13.2, enabling us to conjecture the following:

Conjecture 13.3. *The number of plane permutations, p_n , of length $n \geq 3$ satisfies the recurrence*

$$p_n = \frac{(11n^2 + 11n - 6)p_{n-1} + (n-2)(n-3)p_{n-2}}{(n+3)(n+4)}.$$

If Conjecture 13.2 holds, then the growth rate of plane permutations, $\lim_{n \rightarrow \infty} \sqrt[n]{p_n}$, is the same as that for the Apéry numbers. From the definition of the a_n , it is clear that this is given by

$$\max_{0 \leq \lambda \leq 1} \lim_{n \rightarrow \infty} \sqrt[n]{\binom{n}{\lambda n}^2 \left(\frac{1+\lambda}{\lambda}\right)^n},$$

which can easily be computed using Stirling's approximation and elementary calculus:

Conjecture 13.4. *The growth rate of the class of plane permutations is $\frac{1}{2}(11+5\sqrt{5}) \approx 11.09017$.*

A rather different conjecture was made by Carla Savage during her plenary talk at the Permutation Patterns 2014 conference in Tennessee. It relates plane permutations to certain *inversion sequences*¹, which are sequences of nonnegative integers $e_1 e_2 \dots e_n$ such that $e_i < i$ for each i . These sequences are known as inversion sequences because there is a map, φ , which maps each permutation σ to an inversion sequence that counts the inversions of σ :

$$\varphi(\sigma) = e_1 \dots e_n, \quad \text{where } e_j = |\{i : i < j \text{ and } \sigma(i) > \sigma(j)\}|.$$

This map is a length-preserving bijection (see Knuth [113]).

In an analogous manner to permutations, we say that an inversion sequence $\eta = e_1 \dots e_n$ *contains* a pattern $\sigma = s_1 \dots s_k$ if there is a subsequence of η that is order isomorphic to σ , and say that η *avoids* σ if there is no such subsequence of η . In the context of inversion sequences, a pattern may be an arbitrary sequence of nonnegative integers. For example, **01031** contains two occurrences of **010**, but avoids **201**.

Savage's conjecture proposes a bijection between plane permutations and a particular class of inversion sequences:

¹Also known as Lehmer codes.

Conjecture 13.5 (Savage, 2014). *The number of plane permutations of length n is the same as the number of inversion sequences of length n avoiding **101** and **201**.*

We have confirmed that the classes are indeed equinumerous for $n \leq 36$.

The conjectures of van Hoeij and Savage appear to be unrelated, but there is in fact a possible connection. It turns out that the Apéry numbers enumerate sets of lattice points in a certain sequence of lattices, and inversion sequences can also be interpreted from the perspective of lattice point enumeration.

For each $d \geq 0$, the A_d lattice consists of the set of those points in \mathbb{Z}^{d+1} whose coordinate sum is zero. Given a lattice L , the *crystal ball* numbers, $G_L(r)$, for L , denote the number of points in L that are within r steps of the origin. Crystal ball numbers for the A_d lattices were first determined, independently, by Conway & Sloane [66] and Baake & Grimm [25]. The Apéry numbers a_n are the “diagonal” crystal ball numbers for the A_d lattices: for each n , they enumerate the number of points in A_n that are within n steps of the origin. In symbols,

$$a_n = G_{A_n}(n).$$

On the other hand, inversion sequences of length n can be viewed as lattice points in a (half-open) $1 \times 2 \times \dots \times n$ box in the cubic lattice \mathbb{Z}^n . Indeed, Savage and her co-authors make much use of this lattice point enumeration perspective in their work (for example, see [56]). Perhaps there is some way of leveraging these connections to lattices to prove both conjectures.

Chapter 14

Avoiding 1234 and 2341

In this chapter and the next, we enumerate two classes whose bases consist of two permutations of length 4. The first of these is $\text{Av}(\mathbf{1234}, \mathbf{2341})$. We use our Hasse graph approach to prove the following theorem:

Theorem 14.1. *The class of permutations avoiding $\mathbf{1234}$ and $\mathbf{2341}$ has the algebraic generating function*

$$\frac{2 - 10z + 9z^2 + 7z^3 - 4z^4 - (2 - 8z + 9z^2 - 3z^3)\sqrt{1 - 4z}}{(1 - 3z + z^2)((1 - 5z + 4z^2) + (1 - 3z)\sqrt{1 - 4z})}.$$

Its growth rate is equal to 4.

In doing so, we use the kernel method six times to solve the relevant functional equations.

Let us use \mathcal{F} to denote $\text{Av}(\mathbf{1234}, \mathbf{2341})$. The structure of class \mathcal{F} depends critically on the presence or absence of occurrences of the pattern $\mathbf{123}$. In light of this, to enumerate this class, we partition it into three sets \mathcal{A} , \mathcal{B} and \mathcal{C} as follows:

- $\mathcal{A} = \text{Av}(\mathbf{123})$.
- $\mathcal{B} = \text{Av}(\mathbf{1234}, \mathbf{2341}, \mathbf{13524}, \mathbf{14523}) \setminus \mathcal{A}$. Every permutation in \mathcal{B} contains at least one occurrence of a $\mathbf{123}$, but avoids $\mathbf{13524}$ and $\mathbf{14523}$.
- $\mathcal{C} = \text{Av}(\mathbf{1234}, \mathbf{2341}) \setminus (\mathcal{A} \cup \mathcal{B})$. Every permutation in \mathcal{C} contains a $\mathbf{13524}$ or a $\mathbf{14523}$.

We refer to a permutation in \mathcal{A} as an \mathcal{A} -permutation, and similarly for \mathcal{B} and \mathcal{C} .

The addition of a source graph to a \mathcal{C} -permutation can only yield another \mathcal{C} -permutation (since it can't cause the removal a $\mathbf{13524}$ or $\mathbf{14523}$ pattern). Similarly, the addition of a source graph to a \mathcal{B} -permutation can't result in an \mathcal{A} -permutation. Hence, we can enumerate \mathcal{A} without first considering \mathcal{B} or \mathcal{C} , and can enumerate \mathcal{B} before considering \mathcal{C} .




Figure 14.1: A permutation in class \mathcal{F} , spanned by four source graphs

Before investigating the structure of permutations in \mathcal{A} , \mathcal{B} and \mathcal{C} , let us briefly examine what a typical source graph in \mathcal{F} looks like. Firstly, the avoidance of **1234** means that the non-root vertices of any source graph form a **123**-avoider. Secondly, the avoidance of **2341** presents no additional restriction on the structure of a source graph, because the presence of a **2341** would imply the presence of a **123** in the non-root vertices. Thus a source graph in \mathcal{F} consists of a root together with a **123**-avoider formed from the non-root vertices.

The structure of set \mathcal{A}

We begin by looking at $\mathcal{A} = \text{Av}(\mathbf{123})$. As is very well known, this class is enumerated by the Catalan numbers. However, we need to keep track of the structure of the bottom subgraph. So we must determine the appropriate bivariate generating function.

Let \mathcal{A}_S denote the set of source graphs in set \mathcal{A} . Now, each member of \mathcal{A}_S is simply a *fan*, a root vertex connected to a (possibly empty) sequence of pendant edges. Bottom subgraphs are also fans. Thus source graphs and bottom subgraphs of \mathcal{A} are acyclic.

When enumerating \mathcal{A} , we use the variable u to mark the *number of leaves* (non-root vertices) in the bottom subtree. The generating function for \mathcal{A}_S is thus given by

$$\mathcal{A}_S(u) = z + z^2u + z^3u^2 + \dots = \frac{z}{1 - zu}.$$

We now consider the process of building an \mathcal{A} -permutation from a sequence of source trees. When a source tree is added to an \mathcal{A} -permutation, the root vertex of the source tree may be inserted to the left of zero or more of the leaves of the bottom subtree. See Figure 14.2 for an illustration. Note that, in this and other similar figures, the original bottom subgraph is displayed to the upper left, with the new source graph to the lower right.




Figure 14.2: Adding a source tree to the bottom subtree of an \mathcal{A} -permutation

The action of adding a source tree is thus seen to be reflected by the linear operator $\Omega_{\mathcal{A}}$ whose effect on u^k is given by

$$\Omega_{\mathcal{A}}[u^k] = A_{\mathcal{S}}(u)(1 + u + \dots + u^k) = A_{\mathcal{S}}(u) \frac{1 - u^{k+1}}{1 - u}.$$

Hence, the bivariate generating function $A(u)$ for \mathcal{A} is defined by the following recursive functional equation:

$$A(u) = A_{\mathcal{S}}(u) + A_{\mathcal{S}}(u) \frac{A(1) - uA(u)}{1 - u}.$$

This equation can be solved using the kernel method. To start, we express $A(u)$ in terms of $A(1)$, by expanding and rearranging to give

$$A(u) = \frac{z(1 - u + A(1))}{1 - u + zu^2}. \quad (1)$$

Now, cancelling the kernel by setting $u = (1 - \sqrt{1 - 4z})/2z$ yields the univariate generating function for \mathcal{A} ,

$$A(1) = \frac{1 - \sqrt{1 - 4z}}{2z} - 1.$$

This is the generating function for the Catalan numbers as expected.

Finally, by substituting for $A(1)$ back into (1) we get the following explicit algebraic expression for $A(u)$:

$$A(u) = \frac{1 - 2zu - \sqrt{1 - 4z}}{2(1 - u + zu^2)}.$$

The structure of set \mathcal{B}

We now consider set \mathcal{B} . Recall that sets \mathcal{B} and \mathcal{C} consist of those permutations in class \mathcal{F} that contain at least one occurrence of a **123**. We need to keep track of the position of the leftmost occurrence of a **3** in such a pattern. Given a permutation in \mathcal{B} or \mathcal{C} , let us call the vertex corresponding to the leftmost **3** in a **123** the *spike*. In the figures, the spike is marked with a star.

We now make a key observation. When adding a source graph to a permutation containing a **123**, no vertex of the source graph may be positioned to the right of the spike, or else a **2341** would be created. Hence, the spike in any permutation in classes \mathcal{B} or \mathcal{C} occurs in its bottom subgraph. When enumerating sets \mathcal{B} and \mathcal{C} , we use the variable u to mark *the number of vertices to the left of the spike* in its bottom subgraph.




Figure 14.3: A source graph in set \mathcal{B}

Let \mathcal{B}_S be the set of source graphs in set \mathcal{B} . Since \mathcal{B} -permutations contain a **123** but avoid **13524** and **14523**, the non-root vertices of a permutation in \mathcal{B}_S consist of two descending sequences, the second sequence beginning (with the spike) above the last vertex in the first sequence. See Figure 14.3 for an illustration. If we consider the non-root vertices in order from top to bottom, then it can be seen that \mathcal{B}_S is defined by the structural equation

$$\mathcal{B}_S = uZ \times \text{SEQ}[uZ] \times Z \times \text{SEQ}[uZ + Z] \times uZ \times \text{SEQ}[Z].$$

The first term on the right corresponds to the root and the remaining terms deal with the non-root vertices in order from top to bottom, vertices to the left of the spike being marked with u . The third term corresponds to the spike and the fifth represents the lowest point to the left of the spike (the rightmost **2** of a **123**). Hence, the generating function for \mathcal{B}_S is

$$B_S(u) = \frac{z^3 u^2}{(1-z)(1-zu)(1-z-zu)}.$$

We now study the process of building a \mathcal{B} -permutation from a sequence of source graphs. There are two cases. A permutation in \mathcal{B} may result either from the addition of a source graph to an \mathcal{A} -permutation, or else from adding a source graph to another \mathcal{B} -permutation. We address these two cases in turn.

One way to create a \mathcal{B} -permutation from an \mathcal{A} -permutation is to add a source graph from \mathcal{B}_S , positioning its root to the left of zero or more of the leaves of the bottom subtree of the \mathcal{A} -permutation and its non-root vertices to the right of the bottom subtree. In this case, the new permutation inherits its spike from the added source graph. This is illustrated in the left diagram in Figure 14.4. The generating function for this set of permutations is




Figure 14.4: Ways to create a \mathcal{B} -permutation by adding a source graph to the bottom subtree of an \mathcal{A} -permutation

thus given by

$$B_{AB1}(u) = B_S(u) \frac{A(1) - uA(u)}{1 - u}.$$

For simplicity, we choose not to present the expanded form of $B_{AB1}(u)$, or that of most subsequent expressions. They can all be represented in the form $(p + q\sqrt{1 - 4z})/r$ for appropriate polynomials p , q and r .

The other possibility for creating a \mathcal{B} -permutation from an \mathcal{A} -permutation involves the positioning of some non-root vertices of the source graph to the left of some of the leaves in the bottom subtree, making one of the vertices in the original bottom subtree the spike. The source graph may be drawn from either \mathcal{A}_S or \mathcal{B}_S , as illustrated in the centre and right diagrams in Figure 14.4.

In this situation, if the source graph has a spike, it must be positioned to the right of all leaves in the bottom subtree, or else a **1234** would be created. Furthermore, any source graph vertices placed to the left of leaves in the bottom subtree must occur at the same position in the bottom subtree, or else a **13524** would be created. This position may be chosen independently of where the root vertex is placed.

From these considerations, it can be determined that the resulting set of permutations has the generating function defined by

$$B_{AB2}(u) = \left(B_S(u) + \frac{z^2 u^2}{(1 - z)(1 - zu)} \right) \frac{1}{1 - u} \left(A'(1) - \frac{u}{1 - u} (A(1) - A(u)) \right),$$

where the presence of the derivative A' is a consequence of the independent choice of two positions in the bottom subtree.

Finally, we consider the addition of a source graph to a \mathcal{B} -permutation. As we noted in our key observation on page 148, no vertex of the source graph may be positioned to the right of the spike in the bottom subgraph. As a result, the new source graph may not contain a **123** or else a **1234** would be created, so the source graph must be a member of




Figure 14.5: Adding a source tree to the bottom subgraph of a \mathcal{B} -permutation

\mathcal{A}_S (a fan). Moreover, the leaves of the source tree must be positioned *immediately* to the left of the spike, or else a **1234** would be created. See Figure 14.5 for an illustration.

Note that, as a consequence of these restrictions, it is impossible for the addition of a source graph to a \mathcal{B} -permutation to create a **13524** or **14523**. So it is not possible to extend a \mathcal{B} -permutation so as to create a \mathcal{C} -permutation.

Thus the bivariate generating function $B(u)$ of set \mathcal{B} is defined by the following recursive functional equation:

$$B(u) = B_S(u) + B_{AB_1}(u) + B_{AB_2}(u) + \frac{zu}{1-zu} \frac{B(1) - B(u)}{1-u},$$

where the final term reflects the addition of a source tree to a \mathcal{B} -permutation.

This equation is amenable to the kernel method. It can be rearranged to express $B(u)$ in terms of $B(1)$. The kernel can then be cancelled by setting $u = (1 - \sqrt{1 - 4z})/2z$, which yields an expression for $B(1)$:

$$B(1) = \frac{-1 + 8z - 19z^2 + 12z^3 + (1 - 6z + 9z^2 - 2z^3)\sqrt{1 - 4z}}{2z^3(1 - 4z)}.$$

Substitution then results in an explicit algebraic expression for $B(u)$, which we refrain from presenting explicitly due to its size.




Figure 14.6: A source graph in set \mathcal{C}

The structure of set \mathcal{C}

We begin our enumeration of \mathcal{C} by counting its set of source graphs, which we denote \mathcal{C}_S . Rather than doing this directly, we enumerate all the source graphs that contain a **123** (i.e. those in either \mathcal{B}_S or \mathcal{C}_S) and then subtract those in \mathcal{B}_S . To begin, we consider how we might build an *arbitrary* source graph in class \mathcal{F} by adding vertices from left to right.

Suppose we have a partly formed source graph with at least one non-root vertex, whose rightmost vertex is v , and we want to add further vertices to its right. What are the options? If v is not the lowest of the non-root vertices, then any subsequent vertices must be placed lower than v . The only other restriction is that vertices must be positioned higher than the root. If we use y to mark the number of positions in which a vertex may be inserted, then the action of adding a new vertex can be seen to be reflected by the following linear operator:

$$\Omega_{\downarrow}[f(y)] = zy^2 \frac{f(1) - f(y)}{1 - y}.$$

We choose to denote this operator Ω_{\downarrow} because it corresponds to the action used in building a Łukasiewicz path.

Now let us consider source graphs whose rightmost vertex is a spike. These are in \mathcal{B}_S , so let's call this set \mathcal{B}_{S0} . As usual, we mark with u those vertices to the left of the spike. If, in addition, we mark with y those vertices not above the spike, then \mathcal{B}_{S0} is defined by the structural equation

$$\mathcal{B}_{S0} = uZ \times \text{SEQ}[uZ] \times \text{SEQ}^+[uyZ] \times yZ.$$

It is readily seen that y correctly marks the number of positions in which an additional vertex may be inserted to the right.

Let $\mathcal{D}_S = \mathcal{B}_S \cup \mathcal{C}_S$. Since every member of \mathcal{D}_S is built from an element of \mathcal{B}_{S0} by applying Ω_{\downarrow} zero or more times, it follows that the generating function for \mathcal{D}_S is defined by the recursive functional equation

$$D_S(y) = \frac{z^3 y^2 u^2}{(1 - zu)(1 - zyu)} + zy^2 \frac{D_S(1) - D_S(y)}{1 - y}.$$

This equation can be solved for $D_S(1)$ by the kernel method, using $y = (1 - \sqrt{1 - 4z})/2z$ to cancel the kernel. The generating function for \mathcal{C}_S is then defined by

$$C_S(u) = D_S(1) - B_S(u).$$

We now study the process of building a \mathcal{C} -permutation from a sequence of source graphs. As with set \mathcal{B} , there are two cases. A permutation in \mathcal{C} may result either from the addition

of a source graph to an \mathcal{A} -permutation, or else from adding a source graph to another \mathcal{C} -permutation. (As we observed above, it is not possible to create a \mathcal{C} -permutation by adding a source graph to a \mathcal{B} -permutation.) We address the two cases in turn.




Figure 14.7: Ways to create a \mathcal{C} -permutation by adding a source graph to the bottom subtree of an \mathcal{A} -permutation

One way to create a \mathcal{C} -permutation from an \mathcal{A} -permutation is to add a source graph from \mathcal{C}_S , positioning its root to the left of zero or more of the leaves of the bottom subtree of the \mathcal{A} -permutation and its non-root vertices to the right of the bottom subtree. This is illustrated in the left diagram in Figure 14.7. The generating function for this set of permutations is thus

$$C_{\mathcal{A}\mathcal{C}1}(u) = C_S(u) \frac{A(1) - uA(u)}{1 - u}.$$

The other method for creating a \mathcal{C} -permutation from an \mathcal{A} -permutation involves the positioning of some non-root vertices of the source graph to the left of some of the leaves in the bottom subtree. This is illustrated in the right diagram in Figure 14.7. In analysing this method, it is more convenient to look, more generally, at how an \mathcal{A} -permutation can be extended to yield a permutation containing a **123**, in either \mathcal{B} or \mathcal{C} . We can then subtract those members of \mathcal{B} that are enumerated by $B_{\mathcal{A}\mathcal{B}2}$.

We achieve the enumeration by adding vertices from left to right in four steps:

1. The first step adds the root.
2. The second step adds the first non-root vertex, which determines the position of the new spike, and also any other vertices positioned to the left of the spike.
3. The third step adds any additional vertices to the right of the spike but to the left of some other leaves in the bottom subtree. The addition of such vertices creates occurrences of **13524**.
4. Finally, the fourth step adds any vertices to the right of the bottom subtree.

Step 1: Permutations that result from the addition of the root vertex are enumerated by

$$D_1(u) = zu \frac{A(1) - A(u)}{1 - u}.$$

Step 2: In this step, we insert the descending sequence of vertices that creates the new spike. In the generating function for permutations resulting from this action, we introduce two additional catalytic variables that we require for steps 3 and 4. For use in step 3, v marks the number of source tree leaves to the right of the new spike. For step 4, we use y to mark valid positions for the insertion of subsequent vertices, as we did previously. The generating function is

$$D_2(v) = \frac{zy^2u^2}{1-zyu} \frac{D_1(v) - D_1(u)}{v-u}.$$

Step 3: The effect of adding additional vertices to the right of the spike but to the left of some other leaves in the bottom subtree is represented by the recursive functional equation

$$D_3(y, v) = D_2(v) + zyv \frac{D_3(y, 1) - D_3(y, v)}{1-v}.$$

Again, the kernel method can be used to solve this for $D_3(y, 1)$, the kernel being cancelled by setting $v = 1/(1 - zy)$.

Step 4: Finally, the addition of vertices to the right of the bottom subtree is reflected by the Łukasiewicz operator Ω_{\downarrow} , giving rise to the recursive functional equation

$$D_4(y) = D_3(y, 1) + zy^2 \frac{D_4(1) - D_4(y)}{1-y},$$

which can be solved for $D_4(1)$ by cancelling the kernel with $y = (1 - \sqrt{1 - 4z})/2z$.

The generating function for the set of permutations resulting from the second way of creating a \mathcal{C} -permutation from an \mathcal{A} -permutation is then defined by

$$C_{AC2}(u) = D_4(1) - B_{AB2}(u).$$

Our work is almost complete. We only have to consider how a source graph may be added to a \mathcal{C} -permutation. In fact, the situation is extremely constrained. First, as noted earlier, the source graph must be positioned to the left of the spike. Furthermore, the presence of a **13524** or **14523** means that the addition of a source graph with even a single non-root vertex would create a **1234**. So the only possibility is the addition of a trivial (single vertex) source tree. Thus the bivariate generating function $C(u)$ of set \mathcal{C} is defined by the following recursive functional equation:

$$C(u) = C_S(u) + C_{AC1}(u) + C_{AC2}(u) + zu \frac{C(1) - C(u)}{1-u}.$$

where the final term reflects the addition of a trivial source tree to a \mathcal{C} -permutation. This equation can be solved to yield the following expression for $C(1)$ by a sixth and final

application of the kernel method, cancelling the kernel by setting $u = 1/(1 - z)$:

$$\frac{-1 + 10z - 35z^2 + 52z^3 - 35z^4 + 12z^5 + (1 - 8z + 21z^2 - 22z^3 + 11z^4 - 2z^5)\sqrt{1 - 4z}}{2z^3(1 - 4z)(1 - 3z + z^2)}.$$

We now have all we need to prove Theorem 14.1 by obtaining an explicit expression for the generating function that enumerates class \mathcal{F} . Since \mathcal{F} is the disjoint union of \mathcal{A} , \mathcal{B} and \mathcal{C} , its generating function is given by $A(1) + B(1) + C(1)$. Thus, by appropriate expansion and simplification, the generating function for $\text{Av}(\mathbf{1234}, \mathbf{2341})$ can be shown to be equal to

$$\frac{2 - 10z + 9z^2 + 7z^3 - 4z^4 - (2 - 8z + 9z^2 - 3z^3)\sqrt{1 - 4z}}{(1 - 3z + z^2)((1 - 5z + 4z^2) + (1 - 3z)\sqrt{1 - 4z})}.$$

This has singularities at $z = \frac{1}{4}$, $z = \frac{1}{2}(3 - \sqrt{5})$ and $z = \frac{1}{2}(3 + \sqrt{5})$. Hence, the growth rate of $\text{Av}(\mathbf{1234}, \mathbf{2341})$ is equal to 4, the reciprocal of the least of these.

The first twelve terms of the sequence $|\mathcal{F}_n|$ are 1, 2, 6, 22, 89, 376, 1611, 6901, 29375, 123996, 518971, 2155145. More values can be found at [A165540](#) in OEIS [133].

Chapter 15

Avoiding 1243 and 2314

In this chapter, we analyse the structure of the Hasse graphs of permutations in the class $\text{Av}(1243, 2314)$ and prove the following theorem:

Theorem 15.1. *The class of permutations avoiding 1243 and 2314 has an algebraic generating function $F(z)$ which satisfies the cubic polynomial equation*

$$(z - 3z^2 + 2z^3) - (1 - 5z + 8z^2 - 5z^3)F(z) + (2z - 5z^2 + 4z^3)F(z)^2 + z^3F(z)^3 = 0.$$

Its growth rate is approximately 5.1955, the greatest real root of the quintic polynomial

$$2 - 41z + 101z^2 - 97z^3 + 36z^4 - 4z^5.$$

The proof requires an unusual simultaneous double application of the kernel method.




Figure 15.1: A permutation in class \mathcal{E} , spanned by three source graphs

Let us use \mathcal{E} to denote $\text{Av}(1243, 2314)$. What can we say about the structure of source graphs in \mathcal{E} ? Firstly, since $H_{1243} = \curvearrowright$ may not occur as a subgraph, only the root of a

source graph may fork towards the upper right. Secondly, each source graph in \mathcal{E} is *plane*. This is the case because every non-plane graph contains a $H_{2143} = \curvearrowright$, and, furthermore, any **2143** in a source graph occurs as part of a **13254** (where the **1** is the root of the source graph). But this is impossible in \mathcal{E} , since **13254** does not avoid **1243**.




Figure 15.2: A source graph for class \mathcal{E} , constructed from four *u*-trees

If we combine these two observations, we see that the non-root vertices of a source graph consist of a sequence of inverted subtrees whose roots are right-to-left maxima. The avoidance of $H_{2314} = \curvearrowleft$ places restrictions on the structure of the subtrees, so that they must consist of a path at the lower right, which we call the *trunk*, with pendant edges attached to its left. It is readily seen that these correspond to permutations in $\text{Av}(\mathbf{132}, \mathbf{231})$. We call trees of this form *u*-trees, short for *unbalanced* trees. See Figure 15.2 for an illustration of a source graph constructed from *u*-trees.

The class \mathcal{U} of *u*-trees satisfies the structural equation

$$\mathcal{U} = \mathcal{Z} \times \text{SEQ}[\mathcal{Z} \times \text{SEQ}[\mathcal{Z}]]$$

where the first term on the right represents the lowest leaf at the tip of the trunk and the second represents the remainder of the vertices in the trunk, each with a (possibly empty) sequence of pendant edges attached to the upper left. Hence the generating function for \mathcal{U} is

$$U(z) = \frac{z(1-z)}{1-2z}.$$

If we use u to mark the number of *u*-trees, the class \mathcal{S} of source graphs satisfies the structural equation

$$\mathcal{S} = \mathcal{Z} \times \text{SEQ}[u\mathcal{U}]$$

and thus has bivariate generating function

$$S(u) = S(z, u) = \frac{z(1-2z)}{1-(2+u)z+uz^2}.$$

Let us now examine how a permutation in \mathcal{E} can be built from a sequence of source graphs. Observe that, when a source graph is added, no vertex of the source graph can be positioned between two vertices of a u-tree in the bottom subgraph, because otherwise a **2314** would be created. In addition, there are strong constraints on when u-trees in the new source graph can be positioned to the left of a u-tree in the bottom subgraph.




Figure 15.3: The two methods for adding a source graph in class \mathcal{E} ; u-trees are shown schematically as filled triangles

These conditions result in there being two distinct ways in which a source graph may be added. These are illustrated in Figure 15.3. In the first method, the root of the source graph is positioned to the left of zero or more u-trees in the bottom subgraph and the u-trees in the source graph are positioned to the right of the bottom subgraph.

The second method is more subtle. It is only applicable if the rightmost u-tree of the bottom subgraph is a path. If that is the case, then an initial sequence of u-trees in the source graph can be positioned to the left of this path subtree, as long as each of them, except possibly the last, consists of a single vertex. If the rightmost u-tree of the bottom subgraph were not a path, then a **1243** would be created. Similarly, if a non-final u-tree consisted of more than one vertex, then a **2314** would be created.

In order to handle this second method, we need to keep track of those source graphs in which the rightmost u-tree is a path. Let \mathcal{S}_P be the class of such graphs. It satisfies the structural equation

$$\mathcal{S}_P = \mathcal{Z} \times \text{SEQ}[\mathbf{u}\mathcal{U}] \times \mathbf{u}\text{SEQ}^+[\mathcal{Z}],$$

where \mathbf{u} marks the number of u-trees as before. This class thus has bivariate generating function

$$\mathcal{S}_P(\mathbf{u}) = \mathcal{S}_P(z, \mathbf{u}) = \frac{\mathbf{u}z^2(1-2z)}{(1-z)(1-(2+\mathbf{u})z+\mathbf{u}z^2)}.$$

In order to distinguish between those situations when the second method of adding a source graph is applicable and those when it isn't, let us use \mathcal{P} to denote the set of those

permutations in \mathcal{E} whose Hasse graphs have bottom subgraphs in which the rightmost u -tree is a path.

We are interested in determining the two bivariate generating functions $E(u) = E(z, u)$ and $P(u) = P(z, u)$ for \mathcal{E} and \mathcal{P} respectively, where u marks the number of u -trees *in the bottom subgraph*. To do this, we establish four linear operators on these generating functions that reflect the different ways in which a source graph can be added.

The action of adding a source graph using the first method is readily seen to be reflected by the following linear operator:

$$\Omega_{EE}[f(u)] = S(u) \frac{f(1) - uf(u)}{1 - u}.$$

The first method creates a member of \mathcal{P} from an arbitrary element of \mathcal{E} whenever the source graph is in \mathcal{S}_P (i.e. its rightmost u -tree is a path). Thus the appropriate linear operator is

$$\Omega_{EP}[f(u)] = S_P(u) \frac{f(1) - uf(u)}{1 - u}.$$

Now let us determine the linear operators corresponding to the second method of adding a source graph.

The set, \mathcal{S}^* , of source graphs that can be added using the second method satisfies the structural equation

$$\mathcal{S}^* = \mathcal{Z} \times \text{SEQ}[\mathcal{Z}] \times u\mathcal{U} \times \text{SEQ}[u\mathcal{U}],$$

in which the third term on the right identifies the u -tree which is positioned immediately to the left of the rightmost (path) u -tree in the bottom subgraph. This specification thus counts multiple times those source graphs that can be added in more than one way due to the presence of a non-empty initial sequence of single-vertex u -trees. Note also that we don't mark the initial sequence of single-vertex u -trees with u . The generating function for \mathcal{S}^* is

$$S^*(u) = \frac{uz^2}{1 - (2 + u)z + uz^2}.$$

The action of adding a source graph using the second method is then seen to be reflected by the following linear operator:

$$\Omega_{PE}[f_P(u)] = S^*(u) \frac{f_P(1) - f_P(u)}{1 - u}.$$

Finally, let us consider when adding a source graph to an arbitrary member of \mathcal{P} creates another permutation in \mathcal{P} . The second method creates an element of \mathcal{P} if the source graph is in \mathcal{S}_P and its rightmost (path) u -tree is added to the right of the bottom subgraph. An element of \mathcal{P} is also created if the source graph has a single path u -tree or consists of a

single vertex (the root). Thus the set, S_P^* , of source graphs, counted with multiplicity, that can be added to create an element of \mathcal{P} satisfies the structural equation

$$S_P^* = Z \times \text{SEQ}[Z] \times \text{SEQ}^+[u\mathcal{U}] \times u\text{SEQ}^+[Z] + Z \times u\text{SEQ}[Z].$$

Its generating function is

$$S_P^*(u) = \frac{uz(1-2z)(1-uz)}{(1-z)(1-(2+u)z+uz^2)},$$

and the corresponding linear operator is

$$\Omega_{PP}[f_P(u)] = S_P^*(u) \frac{f_P(1) - f_P(u)}{1-u}.$$

We are now in a position to derive the generating function for \mathcal{E} and hence prove Theorem 15.1. From the analysis above, we know that the bivariate generating function $E(u) = E(z, u)$ of class \mathcal{E} is defined by the following pair of mutually recursive functional equations:

$$\begin{aligned} E(u) &= S(u) + \Omega_{EE}[E(u)] + \Omega_{PE}[P(u)] \\ P(u) &= S_P(u) + \Omega_{EP}[E(u)] + \Omega_{PP}[P(u)]. \end{aligned}$$

These can be expanded to give the following:

$$E(u) = z \frac{(1-2z)(1-u+E(1)-uE(u)) + uz(P(1)-P(u))}{(1-u)(1-(2+u)z+uz^2)}, \quad (1)$$

$$P(u) = uz(1-2z) \frac{z(1-u+E(1)-uE(u)) + (1-uz)(P(1)-P(u))}{(1-u)(1-z)(1-(2+u)z+uz^2)}. \quad (2)$$

An unusual simultaneous double application of the kernel method can then be used to yield the algebraic generating function for class \mathcal{E} as follows.

First, we eliminate $P(u)$ from (1) and (2), and express $E(u)$ in terms of $E(1)$ and $P(1)$ as a rational function. Cancelling the resulting kernel,

$$(1-3z+2z^2) - (2-7z+7z^2-z^3)u + (1-3z+3z^2)u^2 - (z-3z^2+3z^3)u^3,$$

with the appropriate root then gives us an equation relating $E(1)$ and $P(1)$.

Secondly, we eliminate $E(u)$ from (1) and (2), and express $P(u)$ in terms of $E(1)$ and $P(1)$. Cancelling the (same) kernel (using a different root) gives a second equation relating $E(1)$ and $P(1)$.

Finally, we eliminate $P(1)$ from these two equations to yield the following explicit expression for $E(1)$:

$$E(1) = \frac{(1 - \rho)((1 - z)(1 - 2z)(1 - \sigma + z\sigma^2(1 - \rho)) - \alpha\rho(1 - \sigma) + z^2(1 - z - z\sigma)\rho\sigma)}{(\alpha - z^2(1 - z)\rho)\rho(1 - \sigma) - (1 - 2z)((1 - z)(1 - \sigma + z\sigma^2) - z(1 - 2z)\rho\sigma^2)},$$

where

$$\alpha = 1 - 4z + 5z^2 - z^3,$$

and ρ and σ are the appropriate roots of the kernel:

$$\begin{aligned}\rho &= \frac{1}{3z} \left(1 - \frac{2\omega\eta}{\xi^{1/3}} - \frac{\omega^2\xi^{1/3}}{2\beta} \right) \\ \sigma &= \frac{1}{3z} \left(1 - \frac{2\eta}{\xi^{1/3}} - \frac{\xi^{1/3}}{2\beta} \right),\end{aligned}$$

with

$$\begin{aligned}\xi &= 12\sqrt{-3z^6(1 - 2z)\beta^3\theta} - 4\beta^2\zeta, \\ \omega &= \frac{1}{2}(-1 + i\sqrt{3}), \\ \beta &= 1 - 3z + 3z^2, \\ \eta &= 1 - 9z + 24z^2 - 21z^3 + 3z^4, \\ \theta &= 4 - 36z + 97z^2 - 101z^3 + 41z^4 - 2z^5, \\ \zeta &= 2 - 24z + 96z^2 - 144z^3 + 63z^4.\end{aligned}$$

Thus, using a computer algebra system to handle the details of the algebraic manipulation, it can be determined that the generating function $F(z) = E(1)$ for $\text{Av}(1243, 2314)$ has the minimal polynomial

$$(z - 3z^2 + 2z^3) - (1 - 5z + 8z^2 - 5z^3)F(z) + (2z - 5z^2 + 4z^3)F(z)^2 + z^3F(z)^3,$$

and, by determining the location of the singularities, that the growth rate of the class is approximately 5.1955, the greatest real root of the quintic polynomial

$$-z^5\theta(z^{-1}) = 2 - 41z + 101z^2 - 97z^3 + 36z^4 - 4z^5,$$

as required.

The first twelve terms of the sequence $|\mathcal{E}_n|$ are 1, 2, 6, 22, 88, 367, 1571, 6861, 30468, 137229, 625573, 2881230. More values can be found at [A165539](#) in OEIS [133].

Chapter 16

Avoiding 1324

16.1 Introduction

The class $\text{Av}(\mathbf{1324})$, of permutations avoiding the pattern $\mathbf{1324}$, is the only class avoiding a single pattern of length four that is yet to be enumerated exactly. There are three Wilf classes for single permutations of length four.

The classes $\text{Av}(\mathbf{1234})$, $\text{Av}(\mathbf{1243})$, $\text{Av}(\mathbf{1432})$ and $\text{Av}(\mathbf{2143})$ all have the same enumeration. This fact is a consequence of the work of Babson & West [27], which was later generalised by Backelin, West & Xin [28] (see Section 1.3). Gessel [82] derived an explicit form for the (D-finite) generating functions of the classes $\text{Av}(\mathbf{12} \dots \mathbf{k})$ in terms of determinants, for which Bousquet-Mélou [54] later gave an alternative derivation using the kernel method. For $k = 4$, the enumeration can be expressed explicitly:

$$\text{Av}_n(\mathbf{1234}) = \frac{1}{(n+1)^2(n+2)} \sum_{k=0}^n \binom{2k}{k} \binom{n+1}{k+1} \binom{n+2}{k+1}.$$

The growth rate of this class is 9, a specific case of the general result of Regev [141] that $\text{gr}(\text{Av}(\mathbf{12} \dots \mathbf{k})) = (k-1)^2$.

The permutation classes $\text{Av}(\mathbf{1342})$ and $\text{Av}(\mathbf{2413})$ constitute another Wilf equivalence class (see Stankova [151]). Bóna [42] determined the algebraic generating function of $\text{Av}(\mathbf{1342})$ to be

$$\frac{32z}{1 + 20z - 8z^2 - (1 - 8z)^{3/2}},$$

from which its growth rate of 8 can readily be obtained.

The remaining class, $\mathbf{1324}$ -avoiding permutations, remains unenumerated. The notoriety of this problem is renowned. Indeed, as reported in [71], at the Permutation Patterns

conference in Florida in 2005, Zeilberger made the unorthodox metaphysical conjecture that “not even God knows $|\text{Av}_{1000}(\mathbf{1324})|$ ”. Conway & Guttman [65] recently presented evidence strongly suggesting that the generating function for this class does not have an algebraic singularity, and is thus not D-finite. Even the growth rate of the $\mathbf{1324}$ -avoiders is currently unknown.

In this chapter, by considering certain large subsets of $\text{Av}(\mathbf{1324})$, which consist of permutations with a particularly regular structure, we prove that the growth rate of the class exceeds 9.81. This improves on a previous lower bound of 9.47. Central to our proof is an examination of the asymptotic distributions of certain substructures in the Hasse graphs of the permutations. In this context, we consider occurrences of patterns in Łukasiewicz paths and prove that in the limit they exhibit a concentrated Gaussian distribution.

In a recent paper, Conway & Guttman [65] calculate the number of permutations avoiding $\mathbf{1324}$ up to length 36, building on earlier work by Johansson & Nakamura [107]. They then analyse the sequence of values and give an estimate for the growth rate of $\text{Av}(\mathbf{1324})$ of 11.60 ± 0.01 . However, rigorous bounds still differ from this value quite markedly.

The last few years have seen a steady reduction in upper bounds on the growth rate, based on a colouring scheme of Claesson, Jelínek & Steingrímsson [64] which yields a value of 16. Bóna [51] has now reduced this to 13.73718 by employing a refined counting argument.

As far as lower bounds go, Albert, Elder, Rechnitzer, Westcott & Zabrocki [4] established that the growth rate is at least 9.47, by using the *insertion encoding* of $\mathbf{1324}$ -avoiders to construct a sequence of finite automata that accept subclasses of $\text{Av}(\mathbf{1324})$. The growth rate of a subclass is then determined from the transition matrix of the corresponding automaton. Our main result is an improvement on this lower bound:

Theorem 16.1. $\text{gr}(\text{Av}(\mathbf{1324})) > 9.81$.

As observed in Section 9.2, the subgraph of the Hasse graph H_σ induced by a left-to-right minimum of σ and the points to its upper right is a tree. By symmetry, this is also the case for the subgraph induced by a right-to-left maximum of σ and the points to its lower left.

What does a *typical* $\mathbf{1324}$ -avoider look like? Figures 16.1 and 16.2 contain illustrations of large $\mathbf{1324}$ -avoiders.¹ As is noted by Flajolet & Sedgewick ([77, p.169]), the fact that a single example can be used to illustrate the asymptotic structure of a large random combinatorial object can be attributed to concentration of distributions, of which we make much use below in determining our lower bound. Observe the cigar-shaped boundary

¹The data for Figure 16.2 was provided by Einar Steingrímsson from the investigations he describes in [154, Footnote 4].




Figure 16.1: The plot of a **1324**-avoider of length 187 and its Hasse graph

regions consisting of numerous small subtrees, and also the relative scarcity of points in the interior, which tend to be partitioned into a few paths connecting the two boundaries. Many questions concerning the shape of a typical large **1324**-avoider remain to be answered or even to be posed precisely. The recent investigations of Madras & Liu [127] and Atapour & Madras [19] provide a starting point.

We restrict our attention to **1324**-avoiders whose Hasse graphs are spanned by a disjoint sequence of trees, rooted at alternate boundaries. In our investigation of how these trees can interact, we consider the asymptotic distribution of certain substructures of the Hasse graphs. In doing so, we exploit the fact that plane trees are in bijection with Łukasiewicz paths. A Łukasiewicz path of length n is a sequence of integers y_0, \dots, y_n such that $y_0 = 0$, $y_i \geq 1$ for $i \geq 1$, and each step $s_i = y_i - y_{i-1} \leq 1$. Thus, at each step, a Łukasiewicz path may rise by at most one, but may fall by any amount as long as it doesn't drop to zero or below.

In particular, we investigate the distribution of *patterns* in Łukasiewicz paths. A pattern ω of length m in such a path is a sequence of steps $\omega_1, \dots, \omega_m$ that occur contiguously in the path (i.e. there is some $k \geq 0$ such that $\omega_j = s_{k+j}$ for $1 \leq j \leq m$), with the restriction that the *height* $\sum_{j=1}^i \omega_j$ after the i th step is positive for $1 \leq i \leq m$. Note that multiple occurrences of a given pattern may overlap in a Łukasiewicz path. See Figure 16.3 for an illustration.

Under very general conditions, substructures of recursively defined combinatorial classes can be shown to be distributed normally in the limit. By generalising the correlation polynomial of Guibas & Odlyzko, and combining it with an application of the kernel method, we prove that patterns in Łukasiewicz paths also satisfy the conditions necessary for




Figure 16.2: The plot of a 1324-avoider of length 1000 and part of its Hasse graph




Figure 16.3: The plot of a Łukasiewicz path that contains three occurrences of the pattern $1,0,1$, two of which overlap

asymptotic normality:

Theorem 16.2. *The number of occurrences of a fixed pattern in a Łukasiewicz path of length n exhibits a Gaussian limit distribution with mean and standard deviation asymptotically linear in n .*

In the next section, we introduce certain subsets of $\text{Av}(\mathbf{1324})$ for consideration, which consist of permutations having a particularly regular structure, and explore restrictions on their structure. We follow this in Section 16.3 by looking at a number of parameters that record the distribution of substructures in our permutations. Key to our result is the fact that these are asymptotically concentrated, and in this section we prove three of the four concentration results we need. Section 16.4 is reserved for the proof of Theorem 16.2, concerning the distribution of patterns in Łukasiewicz paths. To conclude, in Section 16.5, we use Theorem 16.2 to prove our final concentration result, and then pull everything together to calculate a lower bound for $\text{gr}(\text{Av}(\mathbf{1324}))$, thus proving Theorem 16.1.

16.2 Permutations with a regular structure

In this section, we present the structure and substructures of the permutations that we investigate. Let \mathcal{W} be the set of all permutations avoiding $\mathbf{1324}$ whose Hasse graphs are spanned by a sequence of trees rooted alternately at the lower left and the upper right. See Figure 16.4 for an example.

Trees rooted at a left-to-right minimum we colour red, and trees rooted at a right-to-left maximum we colour blue. We refer to these as *red trees* and *blue trees* respectively. As a mnemonic, note that Red trees grow towards the Right and bLue trees grow towards the Left.

Observe that the root of each non-initial blue tree is the uppermost point below the root of the previous red tree, and the root of each non-initial red tree is the leftmost point to the right of the root of the previous blue tree. Note that \mathcal{W} does not contain every $\mathbf{1324}$ -avoider. For example, $\mathbf{2143} \notin \mathcal{W}$.

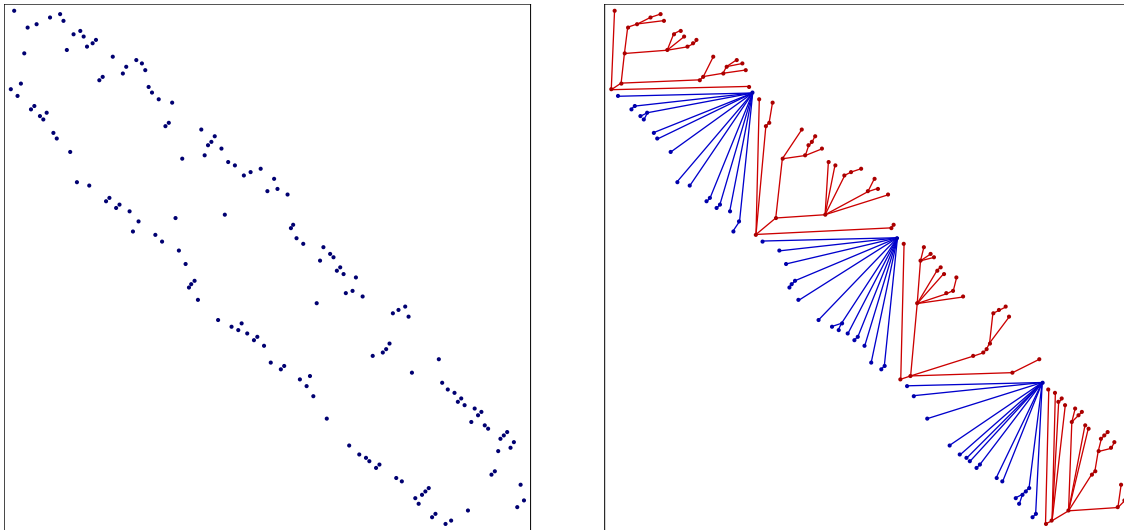


Figure 16.4: A permutation in $\mathcal{W}(3, 25, 19, 12)$ and the spanning of its Hasse graph by red and blue trees

We consider elements of \mathcal{W} with a particularly regular structure. Each red tree has the same number of vertices. Similarly, each blue tree has the same number of vertices. Moreover, every blue tree also has the same root degree. Specifically, for any positive t , k , ℓ and d , let $\mathcal{W}(t, k, \ell, d)$ be the set of those permutations in \mathcal{W} which satisfy the following four conditions:

1. Its Hasse graph is spanned by $t + 1$ red trees and t blue trees.
2. Each red tree has k vertices.
3. Each blue tree has ℓ vertices.
4. Each blue tree has root degree d .

See Figure 16.4 for an illustration of a permutation in $\mathcal{W}(3, 25, 19, 12)$.

To simplify our presentation, we use the term *blue subtree* to denote a principal subtree of a blue tree. (The principal subtrees of a rooted tree are the connected components resulting from deleting the root.) Thus each blue tree consists of a root vertex and a sequence of blue subtrees. We also refer to the roots of blue subtrees simply as *blue roots*.

Our goal is to determine a lower bound for the growth rate of the union of all the $\mathcal{W}(t, k, \ell, d)$. To achieve this, our focus is on sets in which the number and sizes of the trees grow together along with the root degree of blue trees. Specifically, we consider the parameterised sets

$$\mathcal{W}_{\lambda, \delta}(k) = \mathcal{W}(k, k, \lceil \lambda k \rceil, \lceil \delta \lambda k \rceil),$$

for some $\lambda > 0$ and $\delta \in (0, 1)$, consisting of $k + 1$ k -vertex red trees and k $\lceil \lambda k \rceil$ -vertex blue trees each having root degree $\lceil \delta \lambda k \rceil$. Thus, λ is the asymptotic ratio of the size of blue

trees to red trees, and δ is the limiting ratio of the root degree of each blue tree to its size. Note that, asymptotically, $1/\delta$ is the mean number of vertices in a blue subtree. Typically these subtrees are small.

Let $g(\lambda, \delta)$ denote the upper growth rate of $\bigcup_k \mathcal{W}_{\lambda, \delta}(k)$:

$$g(\lambda, \delta) = \lim_{k \rightarrow \infty} |\mathcal{W}_{\lambda, \delta}(k)|^{1/n(k, \lambda)},$$

where $n(k, \lambda) = k(k + \lceil \lambda k \rceil + 1)$ is the length of each permutation in $\mathcal{W}_{\lambda, \delta}(k)$. In order to prove Theorem 16.1, we show that there is some λ and δ for which $g(\lambda, \delta) > 9.81$.




Figure 16.5: Valid and invalid horizontal interleavings, and valid and invalid vertical interleavings; occurrences of **1324** are shown with thicker edges

$\mathcal{W}(t, k, \ell, d)$ consists precisely of those permutations that can be built by starting with a k -vertex red tree and repeating the following two steps exactly t times (see Figure 16.4):

1. Place an ℓ -vertex blue tree with root degree d below the previous red tree (with its root to the right of the red tree), horizontally interleaving its non-root vertices with the non-root vertices of the previous red tree in any way that avoids creating a **1324**.
2. Place a k -vertex red tree to the right of the previous blue tree (with its root below the blue tree), vertically interleaving its non-root vertices with the non-root vertices of the previous blue tree without creating a **1324**.

See Figure 16.5 for illustrations of valid interleavings of the non-root vertices of red and blue trees, and also of invalid interleavings containing occurrences of **1324**. The configurations that have to be avoided when interleaving are shown schematically in Figure 16.6.

We simply call a valid interleaving of the non-root vertices of a red tree with those of a blue tree an *interleaving* of the trees. Note that the choice of interleaving at each step is completely independent of the interleaving at any previous or subsequent step. The only requirement is that no **1324** is created by any of the interleavings.

The key to our result is thus an analysis of how vertices of red and blue trees may be interleaved without forming a **1324**. The remainder of this chapter consists of this analysis.



Figure 16.6: Possible causes of a **1324** when interleaving horizontally and vertically

In what follows, we work exclusively with interleavings of red and blue trees in elements of $\mathcal{W}_{\lambda,\delta}(k)$, for some given $\lambda > 0$ and $\delta \in (0, 1)$. Thus, we assume, without restatement, that a red tree has k vertices, and that a blue tree has $\ell = \lceil \lambda k \rceil$ vertices and is composed of $d = \lceil \delta \lambda k \rceil$ blue subtrees.

We now consider how to avoid creating a **1324**. Without loss of generality, we limit our discussion to the horizontal case.




Figure 16.7: An interleaving of red and blue trees in which no blue subtree (shown in a shaded rectangle) is split by a red vertex

One way to guarantee that no **1324** is created when interleaving trees is to ensure that no blue subtree is split by a red vertex, since the pattern **1324** is avoided in any interleaving in which no red vertex occurs between two blue vertices of the same blue subtree. See Figure 16.7 for an illustration.

Let $\mathcal{W}_{\lambda,\delta}^0(k)$ be the subset of $\mathcal{W}_{\lambda,\delta}(k)$ in which red vertices are interleaved with blue subtrees in this manner in each interleaving.

$\mathcal{W}_{\lambda,\delta}^0(k)$ is easy to enumerate since trees and interleavings can be chosen independently.

Indeed,

$$|\mathcal{W}_{\lambda,\delta}^0(k)| = R_k^{k+1} \times B_k^k \times P_k^{2k},$$

where R_k is the number of distinct red trees, B_k is the number of distinct blue trees and P_k is the number of distinct ways of interleaving red vertices with blue subtrees.

$R_k = \frac{1}{k} \binom{2k-2}{k-1}$, $B_k = \frac{d}{\ell-1} \binom{2\ell-3-d}{\ell-2}$ (see [77, Example III.8]), and $P_k = \binom{k-1+d}{d}$. Hence, by applying Stirling's approximation we obtain the following expression for the growth rate of $\mathcal{W}_{\lambda,\delta}^0(k)$:

$$g_0(\lambda, \delta) = \lim_{k \rightarrow \infty} |\mathcal{W}_{\lambda,\delta}^0(k)|^{1/n(k,\lambda)} = E(\lambda, \delta)^{1/(1+\lambda)}, \quad (1)$$

where

$$E(\lambda, \delta) = 4 \frac{(2-\delta)^{(2-\delta)\lambda} (1+\delta\lambda)^{2(1+\delta\lambda)}}{(1-\delta)^{(1-\delta)\lambda} (\delta\lambda)^{2\delta\lambda}}. \quad (2)$$

It is now elementary to determine the maximum value of this growth rate. For fixed λ , $E(\lambda, \delta)$ is maximal when δ has the value

$$\delta_\lambda = \frac{2\lambda - 1 + \sqrt{1 + 4\lambda + 8\lambda^2}}{2\lambda(2 + \lambda)}.$$

Thence, numerically maximising $g_0(\lambda, \delta)$ by setting $\lambda \approx 0.61840$ (with $\delta_\lambda \approx 0.86238$) yields a preliminary lower bound for $\text{gr}(\text{Av}(\mathbf{1324}))$ of 9.40399. It is rather a surprise that such a simple construction exhibits a growth rate as large as this.

From this analysis, we see that we have complete freedom in choosing the positions of the blue roots (roots of blue subtrees) relative to the vertices of the red tree. In the light of this, we divide the process of interleaving into two stages:

1. Freely interleave the blue roots with the red vertices.
2. Select positions for the non-root vertices of each blue subtree, while avoiding the creation of a **1324**.

We call the outcome of the first stage a *pre-interleaving*. A pre-interleaving is thus a sequence consisting of $k-1$ red vertices and $d = \lceil \delta\lambda k \rceil$ blue vertices (the blue roots); the non-root vertices of the blue subtrees play no role in a pre-interleaving.

Note that in the second stage, each blue subtree can be considered independently since no **1324** can contain vertices from more than one blue subtree. We now consider where the non-root vertices may be positioned.

Our first observation is as follows: Suppose v is the nearest red vertex to the right of the root u of some blue subtree T . Now let x be the parent of v in the red tree. Then no vertex of T can be positioned to the left of x , since otherwise a **1324** would be created in which xuv would be the **324**. Thus, vertices of T can only be interleaved with those red vertices




Figure 16.8: An interleaving of a blue subtree T with its two-component red forest

positioned between u and x . We call the graph induced by this set of red vertices (which may be empty) a *red forest*. See Figure 16.8 for an illustration.




Figure 16.9: Two interleavings of a blue subtree T with a red forest; the red fringes consist of the vertices in the shaded regions

Our second (elementary) observation is as follows: Suppose u is the root of some blue subtree T , and y is the next blue root to the left of u . Then all the non-root vertices of T must occur to the right of y (else T would not be a tree). Note that y may occur either to the left of x or to its right. See Figure 16.9 for illustrations of both of these situations.

These two observations provide two independent constraints on the set of red vertices with which the non-root vertices of a blue subtree may be interleaved, the first determined by the structure of the red tree and the second by the pre-interleaving. This set consists of those vertices of the red forest situated to the right of both x and y . These red vertices induce a subgraph of the red forest which we call its *red fringe*. In the examples in Figure 16.9, the red fringes consist of those vertices in the shaded regions. The key fact that motivates the rest of our analysis is that vertices of a blue subtree may only be interleaved with vertices of its red fringe.

The size of a red fringe depends on both the size of the corresponding red forest and also on the location of the next blue root to the left. Let us call the number of red vertices positioned between a blue root u and the next blue root to its left (y) the *gap size* of u ; the gap size may be zero. The number of vertices in the red fringe is thus the smaller of the

gap size and the number of vertices in the red forest.

If we combine this fact with results concerning the limiting distributions of blue subtrees and red fringes, then we can establish a lower bound for $g(\lambda, \delta)$. This is the focus of the next section.

16.3 Concentration of distributions

To determine our lower bound, we depend critically on the fact that the asymptotic distributions of substructures of permutations in $\mathcal{W}_{\lambda, \delta}(k)$ are *concentrated*. In this section we introduce certain parameters counting these substructures, show how their concentration enables us to bound $g(\lambda, \delta)$ from below, and prove three of the four concentration results we require.

It is frequently the case that distributions of parameters counting the proportion of particular substructures in combinatorial classes have a convergent mean and a variance that vanishes asymptotically. As a direct consequence of Chebyshev's inequality, such distributions have the following concentration property (see [77, Proposition III.3]):

Proposition 16.3. *If ξ_n is a sequence of random variables with means $\mu_n = \mathbb{E}[\xi_n]$ and variances $v_n = \mathbb{V}[\xi_n]$ satisfying the conditions*

$$\lim_{n \rightarrow \infty} \mu_n = \mu, \quad \lim_{n \rightarrow \infty} v_n = 0,$$

for some constant μ , then ξ_n is concentrated at μ in the sense that, for any $\varepsilon > 0$, given sufficiently large n ,

$$\mathbb{P}[|\xi_n - \mu| \leq \varepsilon] > 1 - \varepsilon.$$

In practice this often means that we can work on the assumption that the value of any such parameter is entirely concentrated at its limiting mean. This is the case for the parameters in which we are interested.

We also make use of the following result concerning *multiple* concentrated parameters.

Proposition 16.4. *If ξ_n and ξ'_n are two sequences of random variables on the same sample space concentrated at μ and μ' respectively, then they are jointly concentrated in the sense that, for any $\varepsilon > 0$, given sufficiently large n ,*

$$\mathbb{P}[|\xi_n - \mu| \leq \varepsilon \text{ and } |\xi'_n - \mu'| \leq \varepsilon] > 1 - \varepsilon.$$

Proof. For any $\eta > 0$ and sufficiently large n , the probability that ξ_n differs from μ by less than η exceeds $1 - \eta$, and similarly for ξ' with μ' . Hence the probability that both are simultaneously η -close to their asymptotic means is at least $1 - 2\eta$. Let $\eta = \varepsilon/2$. \square

We now introduce the parameters we need:

Blue subtrees β_k : For each plane tree \mathbf{T} , let $\beta_k(\mathbf{T})$ be the random variable that records the proportion of blue subtrees in a blue tree that are isomorphic to \mathbf{T} .

Gap sizes γ_k : For each $j \geq 0$, let $\gamma_k(j)$ be the random variable that records the proportion of blue roots in a pre-interleaving that have gap size j . Also, let $\gamma_k(> j)$ record the proportion of blue roots in a pre-interleaving whose gap size exceeds j .

Red forests ρ_k : For each plane forest \mathbf{F} , let $\rho_k(\mathbf{F})$ be the random variable that records the proportion of positions in a red tree whose red forest is isomorphic to \mathbf{F} . Also, let $\rho_k(\mathbf{F}^+)$ record the proportion of positions in a red tree whose red forest has at least $|\mathbf{F}|$ vertices, and for which the graph induced by the rightmost $|\mathbf{F}|$ vertices of the forest is isomorphic to \mathbf{F} .

Below, we prove that each of these parameters is concentrated, and calculate their asymptotic means. First we describe how the parameters are combined.

Our first combined parameter counts red fringes. Given the combination of a red tree and a pre-interleaving of its vertices with a sequence of blue roots, let $\varphi_k(\mathbf{F})$ be the random variable that records the proportion of blue roots whose red fringe is isomorphic to \mathbf{F} . Now, occurrences of blue roots with a given gap size j are spread almost uniformly across the positions in a red tree, non-uniformity only occurring for the j leftmost positions. This is also the case for the distribution of occurrences of blue roots whose gap size is at least j . Hence, by the definition of a red fringe at the end of Section 16.2, given any $\varepsilon > 0$, if k is large enough, $\varphi_k(\mathbf{F})$ differs from

$$\gamma_k(|\mathbf{F}|)\rho_k(\mathbf{F}^+) + \gamma_k(>|\mathbf{F}|)\rho_k(\mathbf{F}) \tag{3}$$

by less than ε .

Our second combined parameter concerns pairs consisting of a blue subtree and a red fringe. Given a red tree, a blue tree and a pre-interleaving of their red vertices and blue roots, let $\psi_k(\mathbf{T}, \mathbf{F})$ be the random variable that records the proportion of blue subtrees that are isomorphic to \mathbf{T} and have a red fringe that is isomorphic to \mathbf{F} . We call such a blue subtree a (\mathbf{T}, \mathbf{F}) -subtree. Given that occurrences of a given blue subtree are distributed uniformly across the blue roots, we have

$$\psi_k(\mathbf{T}, \mathbf{F}) = \beta_k(\mathbf{T})\varphi_k(\mathbf{F}). \tag{4}$$

Since, as we show below, β_k, γ_k and ρ_k are concentrated, it follows that ψ_k is also concentrated. Let $\mu(\mathbf{T}, \mathbf{F})$ denote the limiting mean of $\psi_k(\mathbf{T}, \mathbf{F})$ as k tends to infinity.

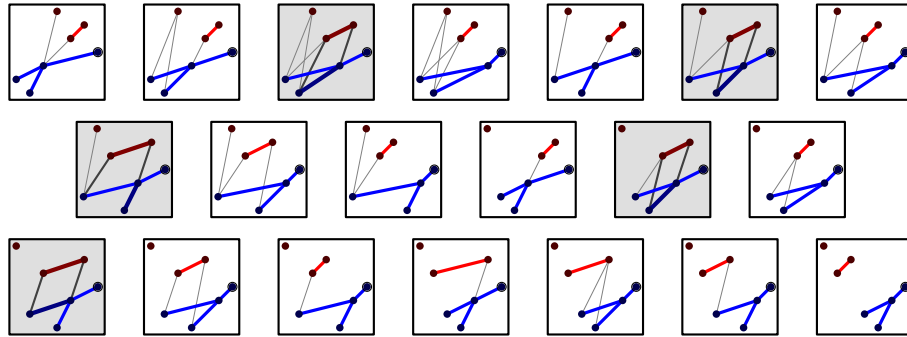


Figure 16.10: $Q(2134, 312) = 15$; the five shaded interleavings contain a 1324

Finally, given a blue subtree \mathbf{T} and a red fringe \mathbf{F} , let $Q(\mathbf{T}, \mathbf{F})$ denote the number of distinct ways of interleaving the non-root vertices of \mathbf{T} and the vertices of \mathbf{F} without creating a 1324. See Figure 16.10 for an example.

With all the relevant parameters defined, we are now in a position to present a lower bound on the value of $g(\lambda, \delta)$.

Proposition 16.5. *Let \mathcal{S} be any finite set of pairs (\mathbf{T}, \mathbf{F}) composed of a plane tree \mathbf{T} and a plane forest \mathbf{F} . Then*

$$g(\lambda, \delta) \geq E(\lambda, \delta)^{1/(1+\lambda)} \times \prod_{(\mathbf{T}, \mathbf{F}) \in \mathcal{S}} Q(\mathbf{T}, \mathbf{F})^{2\delta\lambda\mu(\mathbf{T}, \mathbf{F})/(1+\lambda)},$$

where $E(\lambda, \delta)$ is as defined in equation (2) on page 169.

Proof. Consider a red tree and a blue tree together with a pre-interleaving of their red vertices and blue roots. By Propositions 16.3 and 16.4, for any $\varepsilon > 0$, if k is large enough, then with probability exceeding $1 - \varepsilon$, it is the case that $|\psi_k(\mathbf{T}, \mathbf{F}) - \mu(\mathbf{T}, \mathbf{F})| \leq \varepsilon$ for every $(\mathbf{T}, \mathbf{F}) \in \mathcal{S}$.

So the proportion of pre-interleaved pairs of trees with at least $\lceil \delta\lambda k \rceil (\mu(\mathbf{T}, \mathbf{F}) - \varepsilon)$ occurrences of (\mathbf{T}, \mathbf{F}) -subtrees for every $(\mathbf{T}, \mathbf{F}) \in \mathcal{S}$ exceeds $1 - \varepsilon$.

Elements of $\mathcal{W}_{\lambda, \delta}(k)$ are constructed by independently choosing trees and interleavings. Thus, the size of $\mathcal{W}_{\lambda, \delta}(k)$ is bounded below by

$$|\mathcal{W}_{\lambda, \delta}(k)| \geq |\mathcal{W}_{\lambda, \delta}^0(k)| \times \left(\prod_{(\mathbf{T}, \mathbf{F}) \in \mathcal{S}} (1 - \varepsilon) Q(\mathbf{T}, \mathbf{F})^{\lceil \delta\lambda k \rceil (\mu(\mathbf{T}, \mathbf{F}) - \varepsilon)} \right)^{2k}.$$

Recall that

$$g(\lambda, \delta) = \lim_{k \rightarrow \infty} |\mathcal{W}_{\lambda, \delta}(k)|^{1/n(k, \lambda)},$$

where $n(k, \lambda) = k(k + \lceil \lambda k \rceil + 1)$ is the length of each permutation in $\mathcal{W}_{\lambda, \delta}(k)$. The desired result follows after expanding and taking the limit, making use of (1). \square

To determine the asymptotic mean and variance of our parameters, we utilise bivariate generating functions. The following standard result enables us to obtain the required moments directly as long as we can extract coefficients. We use $[z^n]f(z)$ to denote the coefficient of z^n in the series expansion of $f(z)$; we also use f_x for $\frac{\partial f}{\partial x}$ and f_{xx} for $\frac{\partial^2 f}{\partial x^2}$.

Proposition 16.6 ([77, Proposition III.2]). *Suppose $A(z, x)$ is the bivariate generating function for some combinatorial class, in which z marks size and x marks the value of a parameter ξ . Then the mean and variance of ξ for elements of size n are given by*

$$\mathbb{E}_n[\xi] = \frac{[z^n]A_x(z, 1)}{[z^n]A(z, 1)} \quad \text{and} \quad \mathbb{V}_n[\xi] = \frac{[z^n]A_{xx}(z, 1)}{[z^n]A(z, 1)} + \mathbb{E}_n(\xi) - \mathbb{E}_n(\xi)^2$$

respectively.

The proofs of our first three concentration results each follow a similar pattern: establish the generating function; extract the coefficients; apply Proposition 16.6; take limits using Stirling’s approximation; finally apply Proposition 16.3.

First, we consider blue subtrees. Recall that the random variable $\beta_k(\mathbf{T})$ records the proportion of principal subtrees in a $\lceil \lambda k \rceil$ -vertex plane tree with root degree $\lceil \delta \lambda k \rceil$ that are isomorphic to \mathbf{T} .

Proposition 16.7. *Let $i = |\mathbf{T}|$. $\beta_k(\mathbf{T})$ is concentrated at*

$$\mu_{\beta}(\mathbf{T}) = \frac{(1 - \delta)^{i-1}}{(2 - \delta)^{2i-1}}.$$

Proof. Let $T(z) = \frac{1}{2}(1 - \sqrt{1 - 4z})$ be the generating function for plane trees. Then the bivariate generating function for plane trees with root degree d , in which z marks vertices and u marks principal subtrees isomorphic to \mathbf{T} , is given by

$$B(z, u) = z(T(z) + (u - 1)z^i)^d.$$

Extracting coefficients yields

$$\begin{aligned} [z^\ell]B(z, 1) &= \frac{d}{\ell-1} \binom{2\ell-d-3}{\ell-2}, \\ [z^\ell]B_u(z, 1) &= \frac{d(d-1)}{\ell-i-1} \binom{2\ell-2i-d-2}{\ell-i-2}, \\ [z^\ell]B_{uu}(z, 1) &= \frac{d(d-1)(d-2)}{\ell-2i-1} \binom{2\ell-4i-d-1}{\ell-2i-2}. \end{aligned}$$

Hence, with $\ell = \lceil \lambda k \rceil$ and $d = \lceil \delta \lambda k \rceil$, applying Proposition 16.6 and taking limits gives

$$\lim_{k \rightarrow \infty} \mathbb{E}[\beta_k(\mathbf{T})] = \frac{(1 - \delta)^{i-1}}{(2 - \delta)^{2i-1}} \quad \text{and} \quad \lim_{k \rightarrow \infty} k \mathbb{V}[\beta_k(\mathbf{T})] = v_\beta(\mathbf{T}),$$

where $v_\beta(\mathbf{T})$ is some rational function in δ . So, $\beta_k(\mathbf{T})$ satisfies the conditions for Proposition 16.3 and is thus concentrated at $\mu_\beta(\mathbf{T})$ as required. \square

Secondly, we consider gap size. Recall that, given a pre-interleaving of the non-root vertices of a k -vertex red tree and $\lceil \delta \lambda k \rceil$ blue roots, the random variable $\gamma_k(j)$ records the proportion of blue roots that have gap size j . Similarly, $\gamma_k(> j)$ records the proportion that have gap size exceeding j .

Proposition 16.8. $\gamma_k(j)$ is concentrated at

$$\mu_\gamma(j) = \frac{\delta \lambda}{(1 + \delta \lambda)^{j+1}}.$$

Also, $\gamma_k(> j)$ is concentrated at

$$\mu_\gamma(> j) = \frac{1}{(1 + \delta \lambda)^{j+1}}.$$

Proof. The bivariate generating function for pre-interleavings containing d blue roots, in which z marks red vertices and v marks gaps of size j , is given by

$$G(z, v) = \frac{z}{1-z} \left(\frac{1}{1-z} + (v-1)z^j \right)^d.$$

Extracting coefficients yields

$$\begin{aligned} [z^k] G(z, 1) &= \binom{k+d-1}{d}, \\ [z^k] G_v(z, 1) &= d \binom{k-j+d-2}{d-1}, \\ [z^k] G_{vv}(z, 1) &= d(d-1) \binom{k-2j+d-3}{d-2}. \end{aligned}$$

Hence, with $d = \lceil \delta \lambda k \rceil$, applying Proposition 16.6 and taking limits gives

$$\lim_{k \rightarrow \infty} \mathbb{E}[\gamma_k(j)] = \frac{\delta \lambda}{(1 + \delta \lambda)^{j+1}} \quad \text{and} \quad \lim_{k \rightarrow \infty} k \mathbb{V}[\gamma_k(j)] = v_\gamma(j),$$

where $v_\gamma(j)$ is some rational function in δ and λ . So, $\gamma_k(j)$ satisfies the conditions for Proposition 16.3 and is thus concentrated at $\mu_\gamma(j)$ as required.

Also, since

$$\lim_{k \rightarrow \infty} \mathbb{E}[\gamma_k(> j)] = 1 - \sum_{i=0}^j \mu_\gamma(i) = \frac{1}{(1 + \delta \lambda)^{j+1}},$$

$\gamma_k(>j)$ is concentrated at $\mu_\gamma(>j)$ as required. \square

Thirdly, we consider red forests. Recall that the random variable $\rho_k(\mathbf{F})$ records the proportion of positions in a k -vertex red tree whose red forest is isomorphic to \mathbf{F} .

Proposition 16.9. *Let $m = |\mathbf{F}|$. $\rho_k(\mathbf{F})$ is concentrated at*

$$\mu_\rho(\mathbf{F}) = \frac{1}{2^{2m+1}}.$$

Proof. If \mathbf{F} has h components, then an occurrence of \mathbf{F} in a red tree comprises the leftmost h subtrees of some vertex x that has at least one additional child vertex to the right. See Figure 16.9 for an illustration. Hence, if \mathcal{R} is the class of red trees augmented by marking occurrences of \mathbf{F} with w , then \mathcal{R} satisfies the structural equation

$$\mathcal{R} = z(\text{SEQ}[\mathcal{R}] + (w-1)z^m \text{SEQ}^+[\mathcal{R}]).$$

So the corresponding bivariate generating function, $R(z, w)$, satisfies the functional equation

$$R(z, w) = \frac{z(1 + (w-1)z^m R(z, w))}{1 - R(z, w)},$$

and hence

$$R(z, w) = \frac{1}{2} \left(1 + (1-w)z^{m+1} - \sqrt{(1 + (1-w)z^{m+1})^2 - 4z} \right).$$

Extracting coefficients then yields

$$\begin{aligned} [z^k]R(z, 1) &= \frac{1}{k} \binom{2k-2}{k-1}, \\ [z^k]R_w(z, 1) &= \binom{2k-2m-3}{k-m-1}, \\ [z^k]R_{ww}(z, 1) &= (k-2m-2) \binom{2k-4m-4}{k-2m-2}. \end{aligned}$$

Hence, applying Proposition 16.6 and taking limits gives

$$\lim_{k \rightarrow \infty} \mathbb{E}[\rho_k(\mathbf{F})] = \frac{1}{2^{2m+1}} \quad \text{and} \quad \lim_{k \rightarrow \infty} k\mathbb{V}[\rho_k(\mathbf{F})] = \nu_\rho(\mathbf{F}),$$

where $\nu_\rho(\mathbf{F})$ depends only on $|\mathbf{F}|$. So, $\rho_k(\mathbf{F})$ satisfies the conditions for Proposition 16.3 and is thus concentrated at $\mu_\rho(\mathbf{F})$ as required. \square

Our fourth and final concentration result concerns red fringes. Recall that the random variable $\rho_k(\mathbf{F}^+)$ records the proportion of positions in a red tree whose red forest has at least $|\mathbf{F}|$ vertices, and for which the graph induced by the rightmost $|\mathbf{F}|$ vertices of the forest is isomorphic to \mathbf{F} .




Figure 16.11: A partial red tree and the corresponding Łukasiewicz path; the three marked red fringes correspond to the occurrences of the pattern 1, 0, 1

We would like to determine the bivariate generating function for red trees in which occurrences of the red fringe F are marked. This is considerably less straightforward than was the case for the other parameters. Primarily, this is because distinct occurrences of F may overlap. See the left of Figure 16.11 for an illustration. To achieve our goal, it is convenient to rephrase our problem in terms of Łukasiewicz paths.

Recall from Section 16.1 that a Łukasiewicz path of length n is a sequence of integers y_0, \dots, y_n such that $y_0 = 0$, $y_i \geq 1$ for $i \geq 1$, and each step $s_i = y_i - y_{i-1} \leq 1$. It is easy to see that Łukasiewicz paths are in bijection with red trees: visit the vertices of the tree from *right to left* and let the height of the path be equal to the number of components in the forest induced by the vertices visited so far. Thus, for each leaf vertex, the path contains an up-step, and for each internal vertex with r children, the path contains a $(1-r)$ -step. See Figure 16.11 for an illustration.

Recall also that a pattern ω of length m in a Łukasiewicz path is a sequence of contiguous steps $\omega_1, \dots, \omega_m$ in the path such that $\sum_{j=1}^i \omega_j > 0$ for $1 \leq i \leq m$. We do not consider sequences of steps for which the height drops to zero or below. Thus, a pattern in a Łukasiewicz path corresponds to an occurrence of a red fringe in a red tree. Again, see Figure 16.11, where this is illustrated.

16.4 Patterns in Łukasiewicz paths

The asymptotic distribution of patterns in *words* has been investigated before. For an exposition, see [77, Examples I.12, III.26 and IX.13]. The approach taken there makes use of the *correlation polynomial* of a pattern, introduced by Guibas & Odlyzko in [91] to analyse pattern-matching in strings, and also employs the *cluster method* of Goulden & Jackson [90]. We refine this approach for use with patterns in Łukasiewicz paths by utilising a generalisation of the correlation polynomial and combining it with an application of the

kernel method.

It is readily seen that the bivariate generating function, $L(z, y)$, for Łukasiewicz paths, in which z marks length and y marks height, satisfies the functional equation

$$L(z, y) = zy + \frac{zy}{1-y} (L(z, 1) - yL(z, y)). \quad (5)$$

Given a pattern $\omega = \omega_1, \dots, \omega_m$, let us use $h_i(\omega) = \sum_{j=1}^i \omega_j$ to denote the height after the i th step of ω , and let us call $h_m(\omega)$ the *final height* of ω .

The correlation polynomial of Guibas & Odlyzko is univariate. For our purposes, we define the *bivariate autocorrelation polynomial*, $\hat{a}_\omega(z, y)$, for a pattern $\omega = \omega_1, \dots, \omega_m$ in a Łukasiewicz path as follows:

$$\hat{a}_\omega(z, y) = \sum_{i=1}^{m-1} c_i z^i y^{h_i(\omega)},$$

where

$$c_i = \begin{cases} 1, & \text{if } \omega_{i+1}, \dots, \omega_m = \omega_1, \dots, \omega_{m-i}; \\ 0, & \text{otherwise.} \end{cases}$$

Thus, c_i records whether ω matches itself when shifted (left or right) by i , the variable z marks the shift, and y marks the height. For example, $\hat{a}_{1,1,0,1,1}(z, y) = z^3 y^2 + z^4 y^3$.

Given a fixed pattern ω of length m and final height h , we want to determine the trivariate generating function, $L_\omega(z, y, u)$, for Łukasiewicz paths, where u marks the number of occurrences of the pattern ω in a path. In order to achieve this, we first consider the class of Łukasiewicz paths augmented by distinguishing an arbitrary selection of occurrences of ω . Let $M_\omega(z, y, v)$ be the corresponding generating function, in which v marks distinguished occurrences of the pattern in a path. By the standard inclusion-exclusion principle (see [77, p.208]), we know that

$$L_\omega(z, y, u) = M_\omega(z, y, u - 1). \quad (6)$$

In order to construct a functional equation for M_ω , we consider subpaths each consisting of a maximal collection of overlapping distinguished occurrences of ω . These collections are called *clusters*. It is readily seen that the generating function for clusters is

$$C_\omega(z, y, v) = \frac{z^m y^h v}{1 - v \hat{a}_\omega(z, y)}, \quad (7)$$

where v is used to mark distinguished occurrences of ω in a cluster.

Furthermore, we have

$$M_\omega(z, y, v) = zy + \frac{zy}{1-y} (M_\omega(z, 1, v) - yM_\omega(z, y, v)) + M_\omega(z, y, v)C_\omega(z, y, v), \quad (8)$$

since a path grows either by adding an arbitrary step, as in (5), or else by adding a cluster.²

Combining equations (6), (7) and (8) and rearranging gives us the following functional equation for $L_\omega(z, y, u)$:

$$L_\omega(z, y, u) = \frac{zy(1 + (1-u)\widehat{a}_\omega(z, y))(1-y + L_\omega(z, 1, u))}{z^m y^h (1-y)(1-u) + (1-y + zy^2)(1 + (1-u)\widehat{a}_\omega(z, y))}.$$

This equation is susceptible to the kernel method, so $L_\omega(z, 1, u) = y_0(z, u) - 1$, where y_0 is the appropriate root for y of the denominator. Rearranging, we obtain the following polynomial functional equation for $L = L(z, u) = L_\omega(z, 1, u)$, the bivariate generating function for Łukasiewicz paths in which u marks occurrences of ω :

$$L = z(1 + L)^2 - (1-u) \left(z^m L(1 + L)^h + (L - z(1 + L)^2) \widehat{a}_\omega(z, 1 + L) \right). \quad (9)$$

The fact that L satisfies this equation enables us to demonstrate that patterns in Łukasiewicz paths are concentrated, and moreover are distributed normally in the limit. The following proposition gives very general conditions for this to be the case for some parameter.

Proposition 16.10 ([77, Proposition IX.17 with Theorem IX.12]; see also [69, Theorem 1]). *Let $F(z, u)$ be a bivariate function, analytic at $(0, 0)$ and with non-negative Taylor coefficients, and let ξ_n be the sequence of random variables with probability generating functions*

$$\frac{[z^n]F(z, u)}{[z^n]F(z, 1)}.$$

Assume that $F(z, u)$ is a solution for y of the equation

$$y = \Phi(z, u, y),$$

where Φ is a polynomial of degree at least two in y , $\Phi(z, 1, y)$ has non-negative Taylor coefficients and is analytic in some domain $|z| < R$ and $|y| < S$, $\Phi(0, 1, 0) = 0$, $\Phi_y(0, 1, 0) \neq 1$, $\Phi_{yy}(z, 1, y) \neq 0$, and there exist positive $z_0 < R$ and $y_0 < S$ satisfying the pair of equations

$$\Phi(z_0, 1, y_0) = y_0, \quad \Phi_y(z_0, 1, y_0) = 0.$$

Then, as long as its asymptotic variance is non-zero, ξ_n converges in law to a Gaussian distribu-

²This equation excludes distinguished occurrences of ω that begin with the first step of the path; this simplifies the algebra somewhat while having no effect on the asymptotics.

tion with mean and standard deviation asymptotically linear in n .

All that remains is to check that L satisfies the relevant requirements.

Theorem 16.2. *The number of occurrences of a fixed pattern in a Łukasiewicz path of length n exhibits a Gaussian limit distribution with mean and standard deviation asymptotically linear in n .*

Proof. From (9), it can easily be seen that $L(z, u)$ satisfies the requirements of Proposition 16.10, with $\Phi(z, 1, y) = z(1 + y)^2$, $z_0 = \frac{1}{4}$ and $y_0 = 1$. \square

16.5 Summing up

Since patterns in Łukasiewicz paths are in bijection with red fringes in red trees, $L(z, u)$ is also the bivariate generating function for red trees in which u marks occurrences of the red fringe \mathbf{F} corresponding to the pattern ω , with $m = |\mathbf{F}|$ and h the number of components of \mathbf{F} . Thus, we know that $\rho_k(\mathbf{F}^+)$ is concentrated. It remains for us to determine the limiting mean.

Proposition 16.11. *Let $m = |\mathbf{F}|$ and h be the number of components of \mathbf{F} . $\rho_k(\mathbf{F}^+)$ is concentrated at*

$$\mu_\rho(\mathbf{F}^+) = \frac{1}{2^{2m-h}}.$$

Proof. Let $F(z) = \frac{1}{2z}(1 - \sqrt{1 - 4z})$ be the generating function for plane forests.

Solving (9) with $u = 1$ gives $L(z, 1) = F(z) - 1$ (as expected).

Similarly, differentiating (9) with respect to u , setting $u = 1$, and solving the resulting equation gives

$$L_u(z, 1) = \frac{z^m F(z)^h (1 - (1 - 2z)F(z))}{1 - 4z}.$$

Then, extracting coefficients yields

$$\begin{aligned} [z^k]L(z, 1) &= \frac{1}{k+1} \binom{2k}{k}, \\ [z^k]L_u(z, 1) &= \binom{2k-2m+h}{k-m-1}. \end{aligned}$$

Hence, applying Proposition 16.6 and taking limits,

$$\lim_{k \rightarrow \infty} \mathbb{E}[\rho_k(\mathbf{F}^+)] = \frac{1}{2^{2m-h}}.$$

Concentration follows from Theorem 16.2. \square

We are finally in a position to compute a lower bound for the growth rate of the class of permutations avoiding **1324**, proving our main theorem.

Theorem 16.1. $\text{gr}(\text{Av}(\mathbf{1324})) > 9.81$.

Proof. We calculate the contribution to the growth rate from pairs consisting of a tree and a forest of bounded size. From Proposition 16.5, we know that, for each $N > 0$, the growth rate is at least

$$g_N(\lambda, \delta) = E(\lambda, \delta)^{1/(1+\lambda)} \times \prod_{|\mathbf{T}|+|\mathbf{F}|\leq N} Q(\mathbf{T}, \mathbf{F})^{2\delta\lambda\mu(\mathbf{T}, \mathbf{F})/(1+\lambda)},$$

where

$$\mu(\mathbf{T}, \mathbf{F}) = \mu_\beta(\mathbf{T}) (\mu_\gamma(|\mathbf{F}|)\mu_\rho(\mathbf{F}^+) + \mu_\gamma(>|\mathbf{F}|)\mu_\rho(\mathbf{F})),$$

as follows from (3) and (4) and Propositions 16.7, 16.8, 16.9 and 16.11.

Using *Mathematica* [172] to evaluate $Q(\mathbf{T}, \mathbf{F})$ and $\mu(\mathbf{T}, \mathbf{F})$ and then to apply numerical maximisation over values of λ and δ yields

$$g_{14}(\lambda, \delta) > 9.81056$$

with $\lambda \approx 0.69706$ and $\delta \approx 0.75887$. □

The determination of this value requires the processing of more than 1.6 million pairs consisting of a tree and a forest. Larger values of N would require more sophisticated programming techniques. However, increasing N is unlikely to lead to a significantly improved lower bound; although the rate of convergence at $N = 14$ is still quite slow, numerical analysis of the computational data suggests that $\lim_{N \rightarrow \infty} \max_{\lambda, \delta} g_N(\lambda, \delta)$ is probably not far from 9.82.

We conclude with the observation that in the construction that gives our bound, the mean number of vertices in a blue subtree, $1/\delta$, is less than 1.32. We noted earlier that the cigar-shaped boundary regions of a typical **1324**-avoider contain numerous small subtrees (although it is not immediately obvious how one should identify such a boundary tree). Is it the case that the mean size of these subtrees is asymptotically bounded? Perhaps, on the contrary, their average size grows unboundedly (but very slowly), and understanding how (and the rate at which) this occurs would lead to an improved lower bound. In the meantime, the following question might be somewhat easier to answer:

Question 16.12. *Asymptotically, what proportion of the points in a typical **1324**-avoider are left-to-right minima or right-to-left maxima?*

Bibliography

- [1] Martin Aigner and Günter M. Ziegler. *Proofs from THE BOOK*. Springer-Verlag, fifth edition, 2014. Cited on page [v](#).
- [2] M. H. Albert, M. D. Atkinson, and R. Brignall. The enumeration of permutations avoiding 2143 and 4231. *Pure Math. Appl.*, 22(2):87–98, 2011. Cited on page [18](#).
- [3] M. H. Albert, M. D. Atkinson, and Robert Brignall. The enumeration of three pattern classes using monotone grid classes. *Electron. J. Combin.*, 19(3): Paper 20, 34 pp, 2012. Cited on page [18](#).
- [4] M. H. Albert, M. Elder, A. Rechnitzer, P. Westcott, and M. Zabrocki. On the Stanley-Wilf limit of 4231-avoiding permutations and a conjecture of Arratia. *Adv. in Appl. Math.*, 36(2):96–105, 2006. Cited on pages [16](#), [18](#), and [162](#).
- [5] M. H. Albert and S. A. Linton. Growing at a perfect speed. *Combin. Probab. Comput.*, 18(3):301–308, 2009. Cited on pages [100](#) and [105](#).
- [6] Michael Albert. PermLab: Software for permutation patterns. <http://www.cs.otago.ac.nz/PermLab>, 2012. Cited on page [vi](#).
- [7] Michael Albert, Mike Atkinson, and Robert Brignall. Notes on the basis of monotone grid classes. Unpublished, 2010. Cited on page [27](#).
- [8] Michael Albert, Mike Atkinson, and Steve Linton. Permutations generated by stacks and dequeues. *Ann. Comb.*, 14(1):3–16, 2010. Cited on page [15](#).
- [9] Michael Albert and Mireille Bousquet-Mélou. Permutations sortable by two stacks in parallel and quarter plane walks. *European J. Combin.*, 43:131–164, 2015. Cited on page [15](#).
- [10] Michael H. Albert, M. D. Atkinson, Mathilde Bouvel, Nik Ruškuc, and Vincent Vatter. Geometric grid classes of permutations. *Trans. Amer. Math. Soc.*, 365(11):5859–5881, 2013. Cited on pages [27](#), [37](#), [81](#), [82](#), [83](#), [84](#), and [85](#).
- [11] Michael H. Albert, M. D. Atkinson, and Vincent Vatter. Inflations of geometric grid classes: three case studies. *Australasian J. Combin.*, 58(1):27–47, 2014. Cited on page [18](#).
- [12] Michael H. Albert and Robert Brignall. Enumerating indices of Schubert varieties defined by inclusions. *J. Combin. Theory Ser. A*, 123(1):154–168, 2014. Cited on page [18](#).
- [13] Michael H. Albert, Steve Linton, and Nik Ruškuc. The insertion encoding of permutations. *Electron. J. Combin.*, 12: Research paper 47, 31 pp, 2005. Cited on page [18](#).
- [14] Michael H. Albert, Nik Ruškuc, and Vincent Vatter. Inflations of geometric grid classes of permutations. *Israel J. Math.*, 205(1):73–108, 2015. Cited on pages [81](#) and [99](#).
- [15] Noga Alon and Ehud Friedgut. On the number of permutations avoiding a given pattern. *J. Combin. Theory Ser. A*, 89(1):133–140, 2000. Cited on page [16](#).
- [16] Roger Apéry. Irrationalité de $\zeta(2)$ et $\zeta(3)$. *Astérisque*, 61:11–13, 1979. Cited on page [141](#).

- [17] Roger Apéry. Interpolation de fractions continues et irrationalité de certaines constantes. *Bull. Sec. Sci. CTHS*, III:37–53, 1981. Cited on page 141.
- [18] Richard Arratia. On the Stanley-Wilf conjecture for the number of permutations avoiding a given pattern. *Electron. J. Combin.*, 6: Note 1, 4 pp, 1999. Cited on page 16.
- [19] Mahshid Atapour and Neal Madras. Large deviations and ratio limit theorems for pattern-avoiding permutations. *Combin. Probab. Comput.*, 23:161–200, 2014. Cited on page 163.
- [20] M. D. Atkinson. Permutations which are the union of an increasing and a decreasing subsequence. *Electron. J. Combin.*, 5: Research paper 6, 13 pp, 1998. Cited on page 18.
- [21] M. D. Atkinson. Restricted permutations. *Discrete Math.*, 195:27–38, 1999. Cited on pages 18, 27, and 35.
- [22] M. D. Atkinson. Permuting machines and permutation patterns. In *Permutation patterns*, volume 376 of *London Math. Soc. Lecture Note Ser.*, pages 67–88. Cambridge Univ. Press, Cambridge, 2010. Cited on page 15.
- [23] M. D. Atkinson, M. M. Murphy, and N. Ruškuc. Partially well-ordered closed sets of permutations. *Order*, 19(2):101–113, 2002. Cited on page 29.
- [24] M. D. Atkinson, M. M. Murphy, and N. Ruškuc. Sorting with two ordered stacks in series. *Theoret. Comput. Sci.*, 289(1):205–223, 2002. Cited on page 15.
- [25] M. Baake and U. Grimm. Coordination sequences for root lattices and related graphs. *Zeitschrift für Kristallographie*, 212(4):253–256, 1997. Cited on page 143.
- [26] Eric Babson and Einar Steingrímsson. Generalized permutation patterns and a classification of the Mahonian statistics. *Sém. Lothar. Combin.*, 44: Art. B44b, 18 pp., 2000. Cited on page 18.
- [27] 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. Cited on pages 18 and 161.
- [28] Jörgen Backelin, Julian West, and Guoce Xin. Wilf-equivalence for singleton classes. *Adv. in Appl. Math.*, 38(2):133–148, 2007. Cited on pages 18 and 161.
- [29] József Balogh, Béla Bollobás, and Robert Morris. Hereditary properties of ordered graphs. In *Topics in Discrete Mathematics*, volume 26 of *Algorithms Combin.*, pages 179–213. Springer, 2006. Cited on page 99.
- [30] Cyril Banderier, Mireille Bousquet-Mélou, Alain Denise, Philippe Flajolet, Danièle Gardy, and Dominique Gouyou-Beauchamps. Generating functions for generating trees. *Discrete Math.*, 246:29–55, 2002. Formal power series and algebraic combinatorics (Barcelona, 1999). Cited on page 8.
- [31] Cyril Banderier and Michael Drmota. Coefficients of algebraic functions: formulae and asymptotics. *DMTCS Proceedings*, pages 1065–1076, 2013. 25th International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2013). Cited on page 4.
- [32] Antônio J.O. Bastos, Carlos Hoppen, Yoshiharu Kohayakawa, and Rudini M. Sampaio. Every hereditary permutation property is testable. *Electronic Notes in Discrete Mathematics*, 38:123–128, 2011. Cited on page 79.
- [33] Andrew M. Baxter and Lara K. Pudwell. Enumeration schemes for vincular patterns. *Discrete Math.*, 312(10):1699–1712, 2012. Cited on page 17.
- [34] Glen Baxter. On fixed points of the composite of commuting functions. *Proc. Amer. Math. Soc.*, 15:851–855, 1964. Cited on page 17.
- [35] David Bevan. Sets of points determining only acute angles and some related colouring problems. *Electron. J. Combin.*, 13(1): Research paper 12, 24 pp, 2006. Cited on page v.

- [36] David Bevan. Permutation grid classes. <http://demonstrations.wolfram.com/PermutationGridClasses>, Wolfram Demonstrations Project, 2012. Cited on page 27.
- [37] David Bevan. Growth rates of geometric grid classes of permutations. *Electron. J. Combin.*, 21(4): Paper 4.51, 17 pp, 2014. Cited on pages 19 and 28.
- [38] David Bevan. Intervals of permutation class growth rates. *Preprint. arXiv:1410.3679*, 2014. Cited on pages 19 and 100.
- [39] David Bevan. Growth rates of permutation grid classes, tours on graphs, and the spectral radius. *Trans. Amer. Math. Soc.*, 367(8):5863–5889, 2015. Cited on pages 19 and 28.
- [40] David Bevan. Permutations avoiding 1324 and patterns in Łukasiewicz paths. *J. London Math. Soc.*, in press; [arXiv:1406.2890](https://arxiv.org/abs/1406.2890). Cited on pages 19 and 123.
- [41] David Bevan. The permutation classes $\text{Av}(1234, 2341)$ and $\text{Av}(1243, 2314)$. *Australasian J. Combin.*, in press; [arXiv:1407.0570](https://arxiv.org/abs/1407.0570). Cited on pages 19 and 122.
- [42] Miklós Bóna. Exact enumeration of 1342-avoiding permutations: a close link with labeled trees and planar maps. *J. Combin. Theory Ser. A*, 80(2):257–272, 1997. Cited on pages 17 and 161.
- [43] Miklós Bóna. The permutation classes equinumerous to the smooth class. *Electron. J. Combin.*, 5: Research paper 31, 12 pp., 1998. Cited on page 18.
- [44] Miklós Bóna. The solution of a conjecture of Stanley and Wilf for all layered patterns. *J. Combin. Theory Ser. A*, 85(1):96–104, 1999. Cited on page 16.
- [45] Miklós Bóna. A survey of stack-sorting disciplines. *Electron. J. Combin.*, 9(2): Article 1, 16 pp., 2003. Permutation patterns (Otago, 2003). Cited on page 15.
- [46] Miklós Bóna. A simple proof for the exponential upper bound for some tenacious patterns. *Adv. in Appl. Math.*, 33(1):192–198, 2004. Cited on page 16.
- [47] Miklós Bóna. The limit of a Stanley-Wilf sequence is not always rational, and layered patterns beat monotone patterns. *J. Combin. Theory Ser. A*, 110(2):223–235, 2005. Cited on page 16.
- [48] Miklós Bóna. New records in Stanley-Wilf limits. *European J. Combin.*, 28(1):75–85, 2007. Cited on page 16.
- [49] Miklós Bóna. *Combinatorics of Permutations*. Discrete Mathematics and its Applications. CRC Press, second edition, 2012. Cited on page 14.
- [50] Miklós Bóna. On the best upper bound for permutations avoiding a pattern of a given length. *Preprint. arXiv:1209.2404*, 2012. Cited on page 17.
- [51] Miklós Bóna. A new record for 1324-avoiding permutations. *European J. Math.*, 1(1):198–206, 2015. Cited on page 162.
- [52] Mireille Bousquet-Mélou. A method for the enumeration of various classes of column-convex polygons. *Discrete Math.*, 154:1–25, 1996. Cited on page 120.
- [53] Mireille Bousquet-Mélou. Rational and algebraic series in combinatorial enumeration. In *International Congress of Mathematicians. Vol. III*, pages 789–826. Eur. Math. Soc., Zürich, 2006. Cited on page 4.
- [54] Mireille Bousquet-Mélou. Counting permutations with no long monotone subsequence via generating trees and the kernel method. *J. Algebraic Combin.*, 33(4):571–608, 2011. Cited on page 161.
- [55] Mireille Bousquet-Mélou and Steve Butler. Forest-like permutations. *Ann. Comb.*, 11:335–354, 2007. Cited on pages 20, 119, 122, 135, and 139.
- [56] Katie L. Bright and Carla D. Savage. The geometry of lecture hall partitions and quadratic permutation statistics. In *22nd International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2010)*, pages 569–580. Assoc. Discrete Math. Theor. Comput. Sci., 2010. Cited on page 143.

- [57] A. E. Brouwer and A. Neumaier. The graphs with spectral radius between 2 and $\sqrt{2 + \sqrt{5}}$. *Linear Algebra Appl.*, 114/115:273–276, 1989. Cited on page 70.
- [58] Andries E. Brouwer and Willem H. Haemers. *Spectra of Graphs*. Universitext. Springer, 2012. Cited on pages 68, 69, 70, and 72.
- [59] P. Cartier and D. Foata. *Problèmes combinatoires de commutation et réarrangements*. Lecture Notes in Mathematics, No. 85. Springer-Verlag, 1969. Cited on page 85.
- [60] N. Chomsky and M. P. Schützenberger. The algebraic theory of context-free languages. In *Computer programming and formal systems*, pages 118–161. North-Holland, Amsterdam, 1963. Cited on page 4.
- [61] Timothy Chow and Julian West. Forbidden subsequences and Chebyshev polynomials. *Discrete Math.*, 204(1-3):119–128, 1999. Cited on page 17.
- [62] F. R. K. Chung, R. L. Graham, V. E. Hoggatt, Jr., and M. Kleiman. The number of Baxter permutations. *J. Combin. Theory Ser. A*, 24(3):382–394, 1978. Cited on page 17.
- [63] Josef Cibulka. On constants in the Füredi-Hajnal and the Stanley-Wilf conjecture. *J. Combin. Theory Ser. A*, 116(2):290–302, 2009. Cited on page 16.
- [64] Anders Claesson, Vít Jelínek, and Einar Steingrímsson. Upper bounds for the Stanley–Wilf limit of 1324 and other layered patterns. *J. Combin. Theory Ser. A*, 119(8):1680–1691, 2012. Cited on pages 17 and 162.
- [65] Andrew R. Conway and Anthony J. Guttmann. On 1324-avoiding permutations. *Adv. in Appl. Math.*, 64(0):50–69, 2015. Cited on pages 16 and 162.
- [66] J. H. Conway and N. J. A. Sloane. Low-dimensional lattices. VII. Coordination sequences. *Proc. Royal Soc. London Ser. A*, 453:2369–2389, 1997. Cited on page 143.
- [67] Dragoš Cvetković, Michael Doob, and Ivan Gutman. On graphs whose spectral radius does not exceed $(2 + \sqrt{5})^{1/2}$. *Ars Combin.*, 14:225–239, 1982. Cited on page 70.
- [68] Dragoš Cvetković, Peter Rowlinson, and Slobodan Simić. *An Introduction to the Theory of Graph Spectra*, volume 75 of *London Mathematical Society Student Texts*. Cambridge University Press, 2010. Cited on pages 67, 68, 69, 70, 72, 73, and 89.
- [69] Michael Drmota. Systems of functional equations. *Random Structures Algorithms*, 10(1-2):103–124, 1997. Cited on page 179.
- [70] S. Dulucq, S. Gire, and J. West. Permutations with forbidden subsequences and nonseparable planar maps. In *Proceedings of the 5th Conference on Formal Power Series and Algebraic Combinatorics (Florence, 1993)*, volume 153, pages 85–103, 1996. Cited on page 15.
- [71] Murray Elder and Vince Vatter. Problems and conjectures presented at the Third International Conference on Permutation Patterns, University of Florida, March 7-11, 2005. Preprint. [arXiv:math/0505504](https://arxiv.org/abs/math/0505504), 2005. Cited on pages 16 and 161.
- [72] S. Even and A. Itai. Queues, stacks, and graphs. In *Theory of machines and computations (Proc. Internat. Sympos., Technion, Haifa, 1971)*, pages 71–86. Academic Press, New York, 1971. Cited on page 15.
- [73] E. J. Farrell. An introduction to matching polynomials. *J. Combin. Theory Ser. B*, 27(1):75–86, 1979. Cited on page 87.
- [74] Jiří Fiala and Jan Kratochvíl. Locally constrained graph homomorphisms—structure, complexity, and applications. *Computer Science Review*, 2(2):97–111, 2008. Cited on page 53.
- [75] Philippe Flajolet, Stefan Gerhold, and Bruno Salvy. On the non-holonomic character of logarithms, powers, and the n th prime function. *Electron. J. Combin.*, 11(2): Article 2, 16 pp, 2005. Cited on page 5.
- [76] Philippe Flajolet and Robert Sedgewick. Analytic combinatorics: Functional equations, rational and algebraic functions. Rapport de recherche 4103, INRIA, January 2001. Cited on page 4.

- [77] Philippe Flajolet and Robert Sedgewick. *Analytic Combinatorics*. Cambridge University Press, 2009. Cited on pages [3](#), [4](#), [9](#), [10](#), [85](#), [105](#), [120](#), [162](#), [169](#), [171](#), [174](#), [177](#), [178](#), and [179](#).
- [78] J. Fox. Stanley–Wilf limits are typically exponential. *Preprint*. [arXiv:1310.8378](#), 2013. Cited on page [17](#).
- [79] Zoltán Füredi and Péter Hajnal. Davenport–Schinzel theory of matrices. *Discrete Math.*, 103(3):233–251, 1992. Cited on page [16](#).
- [80] Harry Furstenberg. Algebraic functions over finite fields. *J. Algebra*, 7:271–277, 1967. Cited on pages [41](#) and [45](#).
- [81] Scott Garrabrant and Igor Pak. Pattern avoidance is not P-recursive. *Preprint*. [arXiv:1505.0650](#), 2015. Cited on page [16](#).
- [82] Ira M. Gessel. Symmetric functions and P-recursiveness. *J. Combin. Theory Ser. A*, 53(2):257–285, 1990. Cited on pages [17](#) and [161](#).
- [83] S. Gire. *Arbres, permutations à motifs exclus et cartes planaires: quelques problèmes algorithmiques et combinatoires*. PhD thesis, Université Bordeaux I, 1993. Cited on page [17](#).
- [84] Roman Glebov, Andrzej Grzesik, Tereza Klimošová, and Daniel Král’. Finitely forcible graphons and permutons. *J. Combin. Theory Ser. B*, 110:112–135, 2015. Cited on page [79](#).
- [85] C. D. Godsil. Matchings and walks in graphs. *J. Graph Theory*, 5(3):285–297, 1981. Cited on pages [87](#), [88](#), and [89](#).
- [86] C. D. Godsil. *Algebraic combinatorics*. Chapman & Hall, 1993. Cited on pages [87](#) and [89](#).
- [87] C. D. Godsil and I. Gutman. On the theory of the matching polynomial. *J. Graph Theory*, 5(2):137–144, 1981. Cited on pages [87](#), [88](#), and [92](#).
- [88] Massimiliano Goldwurm and Massimo Santini. Clique polynomials have a unique root of smallest modulus. *Inform. Process. Lett.*, 75(3):127–132, 2000. Cited on page [86](#).
- [89] I. P. Goulden and J. West. Raney paths and a combinatorial relationship between rooted nonseparable planar maps and two-stack-sortable permutations. *J. Combin. Theory Ser. A*, 75(2):220–242, 1996. Cited on page [15](#).
- [90] Ian P. Goulden and David M. Jackson. *Combinatorial Enumeration*. Dover Publications Inc., 2004. Cited on page [177](#).
- [91] L. J. Guibas and A. M. Odlyzko. String overlaps, pattern matching, and nontransitive games. *J. Combin. Theory Ser. A*, 30(2):183–208, 1981. Cited on page [177](#).
- [92] Ivan Gutman. The acyclic polynomial of a graph. *Publ. Inst. Math. (Beograd) (N.S.)*, 22(36):63–69, 1977. Cited on page [87](#).
- [93] Anthony J. Guttmann. The analytic structure of lattice models — Why can’t we solve most models? *Pramana*, 64(6):829–846, 2005. Cited on page [5](#).
- [94] Anthony J. Guttmann, editor. *Polygons, Polyominoes and Polycubes*, volume 775 of *Lecture Notes in Physics*. Springer, 2009. Cited on page [120](#).
- [95] H. Hatami, S. Janson, and B. Szegedy. Graph properties, graph limits and entropy. *Preprint*. [arXiv:1312.5626](#), 2013. Cited on page [80](#).
- [96] Ole J. Heilmann and Elliott H. Lieb. Monomers and dimers. *Phys. Rev. Lett.*, 24:1412–1414, 1970. Cited on page [87](#).
- [97] Pavol Hell and Jaroslav Nešetřil. *Graphs and homomorphisms*, volume 28 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, 2004. Cited on page [53](#).
- [98] Alan J. Hoffman. On limit points of spectral radii of non-negative symmetric integral matrices. In *Graph Theory and Applications*, volume 303 of *Lecture Notes in Mathematics*, pages 165–172. Springer, 1972. Cited on page [71](#).

- [99] Alan J. Hoffman and John Howard Smith. On the spectral radii of topologically equivalent graphs. In *Recent Advances in Graph Theory (Proc. Second Czechoslovak Sympos., Prague, 1974)*, pages 273–281. Academia, 1975. Cited on pages [72](#) and [92](#).
- [100] Cheyne Homberger and Vincent Vatter. On the effective and automatic enumeration of polynomial permutation classes. *Preprint. arXiv:1308.4946*, 2013. Cited on page [35](#).
- [101] Carlos Hoppen, Yoshiharu Kohayakawa, Carlos Gustavo Moreira, Balázs Ráth, and Rudini Menezes Sampaio. Limits of permutation sequences. *J. Combin. Theory Ser. B*, 103(1):93–113, 2013. Cited on page [79](#).
- [102] Carlos Hoppen, Yoshiharu Kohayakawa, Carlos Gustavo Moreira, and Rudini Menezes Sampaio. Limits of permutation sequences through permutation regularity. *Preprint. arXiv:1106.1663*, 2011. Cited on page [79](#).
- [103] Carlos Hoppen, Yoshiharu Kohayakawa, Carlos Gustavo Moreira, and Rudini Menezes Sampaio. Testing permutation properties through subpermutations. *Theoret. Comput. Sci.*, 412(29):3555–3567, 2011. Cited on page [79](#).
- [104] Carlos Hoppen, Yoshiharu Kohayakawa, and Rudini M. Sampaio. A note on permutation regularity. *Discrete Appl. Math.*, 160(18):2716–2727, 2012. Cited on page [79](#).
- [105] Sophie Huczynska and Vincent Vatter. Grid classes and the Fibonacci dichotomy for restricted permutations. *Electron. J. Combin.*, 13(1): Research paper 54, 14 pp, 2006. Cited on pages [25](#), [27](#), and [99](#).
- [106] Svante Janson. Graph limits and hereditary properties. *Preprint. arXiv:1102.3571*, 2011. Cited on page [80](#).
- [107] Fredrik Johansson and Brian Nakamura. Using functional equations to enumerate 1324-avoiding permutations. *Adv. in Appl. Math.*, 56:20–34, 2014. Cited on page [162](#).
- [108] Tomáš Kaiser and Martin Klazar. On growth rates of closed permutation classes. *Electron. J. Combin.*, 9(2): Research paper 10, 20 pp, 2003. Cited on page [99](#).
- [109] Sergey Kitaev. *Patterns in Permutations and Words*. Springer, 2011. Cited on page [14](#).
- [110] Martin Klazar. The Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture. In *Formal Power Series and Algebraic Combinatorics (Moscow, 2000)*, pages 250–255. Springer, 2000. Cited on page [16](#).
- [111] Martin Klazar. On the least exponential growth admitting uncountably many closed permutation classes. *Theoret. Comput. Sci.*, 321(2-3):271–281, 2004. Cited on page [99](#).
- [112] Donald E. Knuth. *The Art of Computer Programming. Vol. 1: Fundamental Algorithms*. Addison-Wesley Publishing Co., 1969. Cited on page [15](#).
- [113] Donald E. Knuth. *The art of computer programming. Vol. 3: Sorting and searching*. Addison-Wesley Publishing Co., 1973. Cited on page [142](#).
- [114] Vilmos Komornik. Expansions in noninteger bases. *Integers*, 11B: A9, 30 pp, 2011. Cited on page [102](#).
- [115] Darla Kremer. Permutations with forbidden subsequences and a generalized Schröder number. *Discrete Math.*, 218:121–130, 2000. Cited on pages [18](#), [20](#), [122](#), and [125](#).
- [116] Darla Kremer. Postscript: “Permutations with forbidden subsequences and a generalized Schröder number” [Discrete Math. 218 (2000) 121–130]. *Discrete Math.*, 270:333–334, 2003. Cited on pages [18](#), [20](#), [122](#), and [125](#).
- [117] Darla Kremer and Wai Chee Shiu. Finite transition matrices for permutations avoiding pairs of length four patterns. *Discrete Math.*, 268:171–183, 2003. Cited on page [17](#).
- [118] Ian Le. Wilf classes of pairs of permutations of length 4. *Electron. J. Combin.*, 12: Research paper 25, 26 pp., 2005. Cited on page [18](#).

- [119] P. W. H. Lemmens and J. J. Seidel. Equiangular lines. *J. Algebra*, 24:494–512, 1973. Cited on page 69.
- [120] Steve Linton, Nik Ruškuc, and Vincent Vatter, editors. *Permutation Patterns*, volume 376 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 2010. Papers from the 5th Conference held at the University of St. Andrews, St. Andrews, June 11–15, 2007. Cited on page 14.
- [121] L. Lipshitz. The diagonal of a D-finite power series is D-finite. *J. Algebra*, 113(2):373–378, 1988. Cited on page 45.
- [122] L. Lovász and J. Pelikán. On the eigenvalues of trees. *Period. Math. Hungar.*, 3:175–182, 1973. Cited on pages 70 and 88.
- [123] László Lovász. *Large networks and graph limits*. Amer. Math. Soc., 2012. Cited on page 79.
- [124] László Lovász and Michael D. Plummer. *Matching Theory*. AMS Chelsea Publishing, 2009. Corrected reprint of the 1986 original. Cited on pages 87 and 89.
- [125] László Lovász and Balázs Szegedy. Limits of dense graph sequences. *J. Combin. Theory Ser. B*, 96(6):933–957, 2006. Cited on page 79.
- [126] Percy A. MacMahon. *Combinatory analysis*, volume I. Cambridge University Press, 1915. Cited on page 17.
- [127] Neal Madras and Hailong Liu. Random pattern-avoiding permutations. In *Algorithmic Probability and Combinatorics*, volume 520 of *Contemp. Math.*, pages 173–194. Amer. Math. Soc., 2010. Cited on page 163.
- [128] Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Combin. Theory Ser. A*, 107(1):153–160, 2004. Cited on page 16.
- [129] Brendan McKay. Roots of matching polynomial of graph. MathOverflow. <http://mathoverflow.net/q/171721>, 2014. Cited on page 88.
- [130] Abbe Mowshowitz. The characteristic polynomial of a graph. *J. Combin. Theory Ser. B*, 12:177–193, 1972. Cited on page 88.
- [131] Maximillian M. Murphy and Vincent R. Vatter. Profile classes and partial well-order for permutations. *Electron. J. Combin.*, 9(2): Research paper 17, 30 pp, 2003. Cited on page 27.
- [132] John Noonan and Doron Zeilberger. The enumeration of permutations with a prescribed number of “forbidden” patterns. *Adv. in Appl. Math.*, 17(4):381–407, 1996. Cited on page 16.
- [133] The OEIS Foundation Inc. The On-Line Encyclopedia of Integer Sequences. Published electronically at <https://oeis.org>. Cited on pages 125, 133, 135, 141, 154, and 160.
- [134] Jay Pantone. The enumeration of permutations avoiding 3124 and 4312. *Preprint. arXiv:1309.0832*, 2013. Cited on page 18.
- [135] Blaise Pascal. *Pensées de M. Pascal sur la religion et sur quelques autres sujets*. 1670. Cited on page v.
- [136] Marco Pedicini. Greedy expansions and sets with deleted digits. *Theoret. Comput. Sci.*, 332(1-3):313–336, 2005. Cited on pages 101 and 102.
- [137] Alfred van der Poorten. A proof that Euler missed . . . Apéry’s proof of the irrationality of $\zeta(3)$. *Math. Intelligencer*, 1(4):195–203, 1979. Cited on page 141.
- [138] Vaughan R. Pratt. Computing permutations with double-ended queues. Parallel stacks and parallel queues. In *Fifth Annual ACM Symposium on Theory of Computing (Austin, Tex., 1973)*, pages 268–277. Assoc. Comput. Mach., New York, 1973. Cited on page 15.
- [139] Helmut Prodinger. The kernel method: a collection of examples. *Sém. Lothar. Combin.*, 50: Art. B50f, 19 pp., 2004. Cited on page 8.
- [140] Lara Pudwell. Enumeration schemes for permutations avoiding barred patterns. *Electron. J. Combin.*, 17(1): Research paper 29, 27 pp, 2010. Cited on page 17.

- [141] Amitai Regev. Asymptotic values for degrees associated with strips of Young diagrams. *Adv. in Math.*, 41(2):115–136, 1981. Cited on pages 17 and 161.
- [142] A. Rényi. Representations for real numbers and their ergodic properties. *Acta Math. Acad. Sci. Hungar.*, 8:477–493, 1957. Cited on page 102.
- [143] John Riordan. *An introduction to combinatorial analysis*. Dover Publications Inc., 2002. Reprint of the 1958 original. Cited on page 86.
- [144] Horst Sachs. Beziehungen zwischen den in einem Graphen enthaltenen Kreisen und seinem charakteristischen Polynom. *Publ. Math. Debrecen*, 11:119–134, 1964. Cited on page 88.
- [145] James B. Shearer. On the distribution of the maximum eigenvalue of graphs. *Linear Algebra Appl.*, 114/115:17–20, 1989. Cited on page 71.
- [146] Slobodan K. Simić. On the largest eigenvalue of unicyclic graphs. *Publ. Inst. Math. (Beograd) (N.S.)*, 42(56):13–19, 1987. Cited on page 73.
- [147] Slobodan K. Simić. On the largest eigenvalue of bicyclic graphs. *Publ. Inst. Math. (Beograd) (N.S.)*, 46(60):1–6, 1989. Cited on page 73.
- [148] Rodica Simion and Frank W. Schmidt. Restricted permutations. *European J. Combin.*, 6(4):383–406, 1985. Cited on page 17.
- [149] John H. Smith. Some properties of the spectrum of a graph. In *Combinatorial Structures and their Applications*, pages 403–406. Gordon and Breach, 1970. Cited on page 69.
- [150] Zvezdelina Stankova and Julian West. A new class of Wilf-equivalent permutations. *J. Algebraic Combin.*, 15(3):271–290, 2002. Cited on page 18.
- [151] Zvezdelina E. Stankova. Forbidden subsequences. *Discrete Math.*, 132:291–316, 1994. Cited on pages 18 and 161.
- [152] Richard P. Stanley. *Enumerative Combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 1999. Cited on page 41.
- [153] Richard P. Stanley. Catalan addendum. <http://www-math.mit.edu/~rstan/ec/catadd.pdf>, 2013. Cited on page 9.
- [154] Einar Steingrímsson. Some open problems on permutation patterns. In *Surveys in Combinatorics 2013*, volume 409 of *London Math. Soc. Lecture Note Ser.*, pages 239–263. Cambridge Univ. Press, 2013. Cited on pages 14 and 162.
- [155] Robert Tarjan. Sorting using networks of queues and stacks. *J. Assoc. Comput. Mach.*, 19:341–346, 1972. Cited on page 15.
- [156] E. J. Janse van Rensburg. *The Statistical Mechanics of Interacting Walks, Polygons, Animals and Vesicles*, volume 18 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, 2000. Cited on page 120.
- [157] Vincent Vatter. Enumeration schemes for restricted permutations. *Combin. Probab. Comput.*, 17(1):137–159, 2008. Cited on page 17.
- [158] Vincent Vatter. Permutation classes of every growth rate above 2.48188. *Mathematika*, 56(1):182–192, 2010. Cited on pages 100, 105, 106, 107, 110, and 113.
- [159] Vincent Vatter. Small permutation classes. *Proc. Lond. Math. Soc.*, 103(5):879–921, 2011. Cited on pages 26, 27, 48, 67, 69, 70, 75, and 99.
- [160] Vincent Vatter. Finding regular insertion encodings for permutation classes. *J. Symbolic Comput.*, 47(3):259–265, 2012. Cited on page 18.
- [161] Vincent Vatter. Permutation classes. In Miklós Bóna, editor, *The Handbook of Enumerative Combinatorics*. CRC Press, in preparation; [arXiv:1409.5159](https://arxiv.org/abs/1409.5159). Cited on pages 14 and 16.

- [162] Vincent Vatter and Steve Waton. On partial well-order for monotone grid classes of permutations. *Order*, 28(2):193–199, 2011. Cited on pages 27, 83, and 84.
- [163] Stephen D. Waton. *On permutation classes defined by token passing networks, gridding matrices and pictures: Three flavours of involvement*. PhD thesis, University of St Andrews, 2007. Cited on page 27.
- [164] Julian West. *Permutations with restricted subsequences and stack-sortable permutations*. PhD thesis, Massachusetts Institute of Technology, 1990. Cited on pages 15 and 18.
- [165] Julian West. Sorting twice through a stack. *Theoret. Comput. Sci.*, 117(1-2):303–313, 1993. Conference on Formal Power Series and Algebraic Combinatorics (Bordeaux, 1991). Cited on page 15.
- [166] Julian West. Generating trees and the Catalan and Schröder numbers. *Discrete Math.*, 146(1-3):247–262, 1995. Cited on page 17.
- [167] Julian West. Generating trees and forbidden subsequences. In *Proceedings of the 6th Conference on Formal Power Series and Algebraic Combinatorics (New Brunswick, NJ, 1994)*, volume 157, pages 363–374, 1996. Cited on page 17.
- [168] Eugene P. Wigner. The unreasonable effectiveness of mathematics in the natural sciences. Richard Courant lecture in mathematical sciences delivered at New York University, May 11, 1959. *Communications on Pure and Applied Mathematics*, 13(1):1–14, 1960. Cited on page vi.
- [169] Wikipedia. Enumerations of specific permutation classes. http://en.wikipedia.org/wiki/Enumerations_of_specific_permutation_classes. Cited on pages 17 and 129.
- [170] Herbert S. Wilf. What is an answer? *Amer. Math. Monthly*, 89(5):289–292, 1982. Cited on page 2.
- [171] Herbert S. Wilf. *generatingfunctionology*. A K Peters Ltd., third edition, 2006. Cited on page 2.
- [172] Wolfram Research, Inc. Mathematica. Version 9.0. www.wolfram.com/mathematica, 2012. Cited on pages vi, 133, 141, and 181.
- [173] Renee Woo and Arnold Neumaier. On graphs whose spectral radius is bounded by $\frac{3}{2}\sqrt{2}$. *Graphs Combin.*, 23(6):713–726, 2007. Cited on page 70.
- [174] J. Yarkony, C. Fowlkes, and A. Ihler. Covering trees and lower-bounds on quadratic assignment. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 887–894, 2010. Cited on page 53.
- [175] Doron Zeilberger. A proof of Julian West’s conjecture that the number of two-stack-sortable permutations of length n is $2(3n)!/((n+1)!(2n+1)!)$. *Discrete Math.*, 102(1):85–93, 1992. Cited on page 15.
- [176] Doron Zeilberger. Enumeration schemes and, more importantly, their automatic generation. *Ann. Comb.*, 2(2):185–195, 1998. Cited on page 17.
- [177] Doron Zeilberger. Enumerative and algebraic combinatorics. In Timothy Gowers, editor, *The Princeton Companion to Mathematics*. Princeton University Press, 2008. Cited on page 1.