Growth rates of geometric grid classes of permutations

David Bevan

Department of Mathematics and Statistics
The Open University
Milton Keynes, England

David.Bevan@open.ac.uk

Submitted: Nov 17, 2014; Accepted: Nov 22, 2014; Published: Dec 4, 2014
Mathematics Subject Classifications: 05A05, 05A16, 05C31

Abstract

Geometric grid classes of permutations have proven to be key in investigations of classical permutation pattern classes. By considering the representation of gridded permutations as words in a trace monoid, we prove that every geometric grid class has a growth rate which is given by the square of the largest root of the matching polynomial of a related graph. As a consequence, we characterise the set of growth rates of geometric grid classes in terms of the spectral radii of trees, explore the influence of "cycle parity" on the growth rate, compare the growth rates of geometric grid classes against those of the corresponding monotone grid classes, and present new results concerning the effect of edge subdivision on the largest root of the matching polynomial.

1 Introduction

Following the proof by Marcus & Tardos [21] of the Stanley–Wilf conjecture, there has been particular interest in the growth rates of permutation classes. Kaiser & Klazar [18] determined the possible growth rates less than 2, and then Vatter [27] characterised all the (countably many) permutation classes with growth rates below $\kappa \approx 2.20557$ and established that there are uncountably many permutation classes with growth rate κ. Critical to these results has been the consideration of *grid classes* of permutations, and particularly of *geometric* grid classes. Geometric grid classes have also been used to achieve the enumeration of some specific permutation classes [1, 3, 24]. Following initial work on particular geometric grid classes by Waton [30], Vatter & Waton [29], and Elizalde [8], their general structural properties have been investigated in articles by Albert, Atkinson, Bouvel, Ruškuc & Vatter [2] and Albert, Ruškuc & Vatter [4]. We build on their work

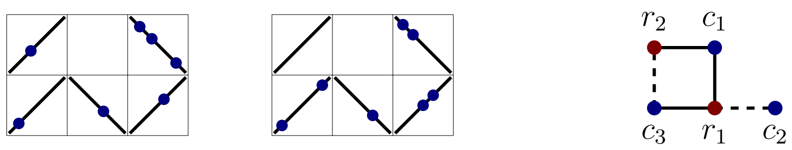

Figure 1: At left: The standard figure for $\mathsf{Geom}\left(\begin{smallmatrix}1 & 0 & -1\\ 1 & -1 & 1\end{smallmatrix}\right)$, showing two plots of the permutation 1527634 with distinct griddings. At right: Its row-column graph; positive edges are shown as solid lines, negative edges are dashed.

to establish the growth rate of any given geometric grid class. Before we can state our result, we need a number of definitions.

A geometric grid class is specified by a $0/\pm1$ matrix which represents the shape of plots of permutations in the class. 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* $\mathsf{Geom}(M)$ is then defined to be the set of permutations $\sigma_1\sigma_2\ldots\sigma_n$ that can be plotted as a subset of the standard figure, i.e. for which there exists a sequence of points $(x_1, y_1), \ldots, (x_n, y_n) \in \Lambda_M$ such that $x_1 < x_2 < \ldots < x_n$ and the sequence $y_1, \ldots, y_n$ is order-isomorphic to $\sigma_1, \ldots, \sigma_n$. See Figure 1 for an example.

If g_n is the number of permutations of length n in $\mathsf{Geom}(M)$, then the *growth rate* of the class is given by $\mathrm{gr}(\mathsf{Geom}(M)) = \lim_{n\to\infty} g_n^{1/n}$. We will demonstrate that this limit exists for geometric grid classes[1] and determine its value for any given $0/\pm1$ matrix M.

Much of the structure of a geometric grid class is reflected in a graph that we associate with the underlying matrix. If M is a $0/\pm1$ matrix of dimensions $t \times u$, the *row-column graph* $G(M)$ of M is the bipartite graph with vertices $r_1, \ldots, r_t, c_1, \ldots, c_u$ and an edge between r_i and c_j if and only if $M_{i,j} \neq 0$. We label each edge r_ic_j with the value of $M_{i,j}$. Edges labelled $+1$ are called *positive*; edges labelled -1 are called *negative*. See Figure 1 for an example.

We need one final definition related to geometric grid classes. If M is a $0/\pm1$ matrix of dimensions $t \times u$, we define the *double refinement* $M^{\times 2}$ of M to be the $0/\pm1$ matrix of dimensions $2t \times 2u$ obtained from M by replacing each 0 with $\left(\begin{smallmatrix}0 & 0\\ 0 & 0\end{smallmatrix}\right)$, each 1 with $\left(\begin{smallmatrix}0 & 1\\ 1 & 0\end{smallmatrix}\right)$, and each -1 with $\left(\begin{smallmatrix}-1 & 0\\ 0 & -1\end{smallmatrix}\right)$. See Figure 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 1. $\mathsf{Geom}(M^{\times 2}) = \mathsf{Geom}(M)$ *for any* $0/\pm1$ *matrix* M.

We will demonstrate a connection between the growth rate of $\mathsf{Geom}(M)$ and the *matching polynomial* of the graph $G(M^{\times 2})$, the row-column graph of the double refinement of M.

[1] It is widely believed that all permutation classes have growth rates. The proof of the Stanley–Wilf conjecture by Marcus & Tardos [21] establishes only that each has an *upper* growth rate ($\limsup g_n^{1/n}$).

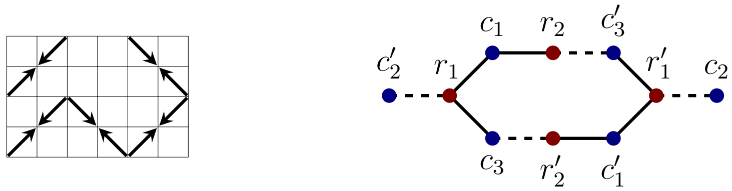

Figure 2: At left: The standard figure of $\begin{pmatrix} 1 & 0 & -1 \\ 1 & -1 & 1 \end{pmatrix}^{\times 2}$, with a consistent orientation marked. At right: Its row-column graph.

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 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

$$\mu_G(z) \;=\; \sum_{k \geqslant 0} (-1)^k m_k(G)\, z^{n-2k}. \tag{1}$$

Note that this *is* a polynomial since $m_k(G) = 0$ for $k > \lfloor n/2 \rfloor$. Observe also 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 2. *The growth rate of geometric grid class* $\mathsf{Geom}(M)$ *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 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.

2 Proof of Theorem 2

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 [28] to establish certain structural properties of grid classes, and was developed further in [2] from where we use a number of results. To begin with, we need to consider *griddings* of permutations.

If M has dimensions $t \times u$, then an *M-gridding* of a permutation $\sigma_1 \ldots \sigma_n$ of length n in $\mathsf{Geom}(M)$ consists of two sequences $c_1, \ldots, c_t$ and $r_1, \ldots, r_u$ such that there is some plot $(x_1, y_1), \ldots, (x_n, y_n)$ of σ for which c_i is the number of points (x_k, y_k) in column i (with $i-1 < x_k < i$), and r_j is the number of points in row j (with $j-1 < y_k < j$).[2] Note that a permutation may have multiple distinct griddings in a given geometric grid class; see Figure 1 for an example. We call a permutation together with one of its M-griddings an *M-gridded permutation*. We use $\mathsf{Geom}^{\#}(M)$ to denote the set of all M-gridded permutations.

From an enumerative perspective, it can be much easier working with M-gridded permutations than directly with the permutations themselves. The following observation means that we can, in fact, restrict our considerations to M-gridded permutations:

Lemma 3 (see Vatter [27] Proposition 2.1)**.** *If it exists, the growth rate of $\mathsf{Geom}(M)$ is equal to the growth rate of the corresponding class of M-gridded permutations $\mathsf{Geom}^{\#}(M)$.*

Proof. Suppose that M has dimensions $t \times u$. Each permutation in $\mathsf{Geom}(M)$ has at least one gridding in $\mathsf{Geom}^{\#}(M)$, but no permutation of length n in $\mathsf{Geom}(M)$ can have more than $\binom{n+t-1}{t-1}\binom{n+u-1}{u-1}$ griddings in $\mathsf{Geom}^{\#}(M)$ because that is the number of ways of choosing the number of points in each column and row. Thus the number of M-gridded permutations of length n is no more than a polynomial multiple of the number of n-permutations in $\mathsf{Geom}(M)$; the result follows immediately from the definition of the growth rate. □

To determine the growth rate of $\mathsf{Geom}^{\#}(M)$, 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 $0/\pm1$ 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.[3] See Figures 2 and 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 4 (see Vatter & Waton [28] 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, $G\left(\begin{smallmatrix} 1 & 0 & -1 \\ 1 & -1 & 1 \end{smallmatrix}\right)$ contains a negative cycle so its standard figure has no consistent orientation (see Figure 1), whereas $G\left(\begin{smallmatrix} -1 & 0 & -1 \\ 1 & -1 & 1 \end{smallmatrix}\right)$ has no negative cycles so its standard figure has a consistent orientation (see Figure 3).

[2]This definition of an M-gridding is equivalent to the traditional one given in terms of the positions of the cell dividers relative to the points (k, σ_k).

[3]For ease of exposition, we use the concept of a consistent orientation rather than the approach used previously involving partial multiplication matrices; results from [2] follow *mutatis mutandis*.

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 2). So we have the following:

Lemma 5 (see [2] Proposition 4.1)**.** *If M is any $0/\pm1$ matrix, then $\Lambda_{M^{\times 2}}$ has a consistent orientation.*

Thus, by Lemma 4, the row-column graph of the double refinement of a matrix never contains a negative cycle. Figure 2 shows a consistent orientation of the standard figure of the double refinement of a matrix whose standard figure (shown in Figure 1) doesn't itself have a consistent orientation.

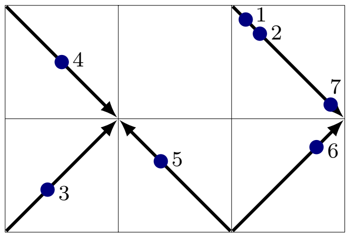

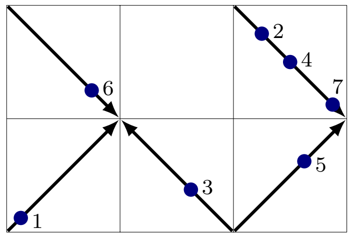

Figure 3: The plots of permutation 1527634 in $\mathsf{Geom}\begin{pmatrix} -1 & 0 & -1 \\ 1 & -1 & 1 \end{pmatrix}$ associated with the words $a_{32}a_{32}a_{11}a_{12}a_{21}a_{31}a_{32}$ and $a_{11}a_{32}a_{21}a_{32}a_{31}a_{12}a_{32}$. Both plots correspond to the same gridding.

We are now in a position to describe the association between words and M-gridded permutations. If M is a $0/\pm1$ matrix, then we let $\Sigma_M = \{a_{ij} : M_{i,j} \neq 0\}$ 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 $\mathsf{Geom}(M)$ 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 3 for two examples. Clearly, this induces a mapping from the set of all finite words over Σ_M to $\mathsf{Geom}^\#(M)$. In fact, it can readily be shown that this map is surjective, every M-gridded permutation corresponding to some word over Σ_M ([2] Proposition 5.3).

As can be seen in Figure 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 $a_{k\ell}$ commute (i.e. $a_{ij}a_{k\ell} = a_{k\ell}a_{ij}$) whenever $i \neq k$ and $j \neq \ell$. It is then relatively straightforward to show equivalence between gridded permutations and elements of the trace monoid:

Lemma 6 (see [2] Proposition 7.1)**.** *If the standard figure Λ_M has a consistent orientation, then gridded n-permutations in $\mathsf{Geom}^{\#}(M)$ are in bijection with equivalence classes of words of length n in $\mathbb{M}(M)$.*

Hence, by combining Lemmas 3, 5 and 6 with Observation 1, we know that the growth rate of $\mathsf{Geom}(M)$ 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 [6]. Using extended Möbius inversion, they determined the general form of the generating function, as follows:

Lemma 7 ([6]; see also Flajolet & Sedgewick [10] Note V.10)**.** *The ordinary generating function for $\mathbb{M}(M)$ is given by*

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

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 [25]). 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 8. *For all $k \geqslant 0$, $r_k(M) = m_k(G(M))$.*

Now, by elementary analytic combinatorics, 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 (see [10] Theorem IV.7). The fact that this polynomial has a *unique* root of smallest modulus was proved by Goldwurm & Santini in [14]. It is real and positive by Pringsheim's Theorem.

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

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

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 8, 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 2: The growth rate of $\mathsf{Geom}(M)$ 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 4, 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 2:

Corollary 9. *If $G(M)$ contains no negative cycles, then the growth rate of $\mathsf{Geom}(M)$ is equal to the square of the largest root of $\mu_{G(M)}(z)$.*

3 Consequences

In this final section, we investigate some of the implications of Theorem 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 resulting from deleting the edge e from G. The number of connected components of G is represented by $\mathrm{comp}(G)$. 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 [16] when investigating monomer-dimer systems in statistical physics. It was first studied from a combinatorial perspective by Farrell [9] and Gutman [15]. The theory was then further developed by Godsil & Gutman [13] and Godsil [11]. An introduction can be found in the books by Godsil [12] and Lovász & Plummer [20].

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 10 (Farrell [9], Gutman [15])**.** *The matching polynomial of a graph is the product of the matching polynomials of its connected components.*

Thus, in particular:

Corollary 11. *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 12 (Godsil & Gutman [13])**.** *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*

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

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

As an immediate consequence, we have the following:

Corollary 13 (Sachs [26], Mowshowitz [23], Lovász & Pelikán [19])**.** *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 14. *If G is a forest, then $\lambda(G) = \rho(G)$.*

Thus, using Corollaries 9 and 11, we have the following alternative characterisation for the growth rates of acyclic geometric grid classes:

Corollary 15. *If $G(M)$ is a forest, then $\operatorname{gr}(\mathsf{Geom}(M)) = \rho(G(M))^2$.*

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 [11], 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*. 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.

Lemma 16 (Godsil [11]; see also [12] and [20])**.** *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 4). Then $\lambda(G) = \lambda(K)$.*

The process that is described in Lemma 16 we will 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 will thus eventually result in a forest F such that $\lambda(F) = \lambda(G)$. We shall say that F results from *fully expanding* G. Hence, by Corollaries 11 and 14, the

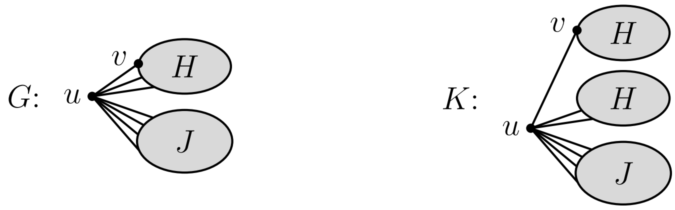

Figure 4: Expanding G at u along uv; H is the component of $G - u$ that contains v

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 17. *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 18 ([7] Proposition 1.3.10)**.** *If G is connected and H is a proper subgraph of G, then we have $\rho(H) < \rho(G)$.*

Lemma 16 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 19. *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 $u_1, \ldots, u_k$, say), then the result is a forest F such that $\lambda(H) = \rho(F)$. Now suppose that we repeatedly expand G analogously at $u_1, \ldots, u_k$, 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 18. □

Adding a non-zero cell to a $0/\pm1$ matrix M adds an edge to $G(M)$. Thus, geometric grid classes satisfy the following monotonicity condition:

Corollary 20. *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* $\mathrm{gr}(\mathsf{Geom}(M')) > \mathrm{gr}(\mathsf{Geom}(M))$.

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 16, we have $\lambda(C_n) = \rho(P_{2n-1})$, where P_n is the path graph on n vertices. The spectral radius of a graph on n vertices is $2\cos\frac{\pi}{n+1}$. So,

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

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,

$$\mathrm{gr}\big(\mathsf{Geom}\left(\begin{smallmatrix}1 & 0 & -1\\ 1 & -1 & 1\end{smallmatrix}\right)\big) \;=\; 3+\sqrt{2} \;\approx\; 4.41421, \tag{4}$$

whereas

$$\mathrm{gr}\big(\mathsf{Geom}\left(\begin{smallmatrix}-1 & 0 & -1\\ 1 & -1 & 1\end{smallmatrix}\right)\big) \;=\; 4. \tag{5}$$

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

Corollary 21. *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* $\mathrm{gr}(\mathsf{Geom}(M_1)) > \mathrm{gr}(\mathsf{Geom}(M))$.

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 $\{v_1, \ldots, v_n\}$, then we let $G(M^{\times 2})$ have vertices $v_1, \ldots, v_n$ and $v'_1, \ldots, v'_n$. If v_iv_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_iv_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 2 against that in Figure 1.

Note that if $v_1, \ldots, v_k$ 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 $\{v_1, \ldots, v_k, v'_1, \ldots, v'_k\}$. In contrast, if $v_1, \ldots, v_\ell$ is a negative ℓ-cycle in $G(M)$, then $G(M^{\times 2})$ contains a (positive) 2ℓ-cycle on $\{v_1, \ldots, v_\ell, v'_1, \ldots, v'_\ell\}$ 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, $1 \leqslant i \leqslant \ell$. We make the following additional observations:

Observation 22. *If $G(M)$ has no odd cycles, then $G(M^{\times 2}) = G(M) + G(M)$.*

Observation 23. *If $G(M)$ is connected and has an odd cycle, then $G(M^{\times 2})$ is connected.*

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

Proof of Corollary 21. 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 22, $G(M^{\times 2}) = G + G$. Thus, since G is connected, it has the form at the left of Figure 5, in which H is the component of $G - u$ containing v. Moreover, we have $\mathrm{gr}(\mathsf{Geom}(M)) = \lambda(G)$. (This also follows from Corollary 9.) Now, if

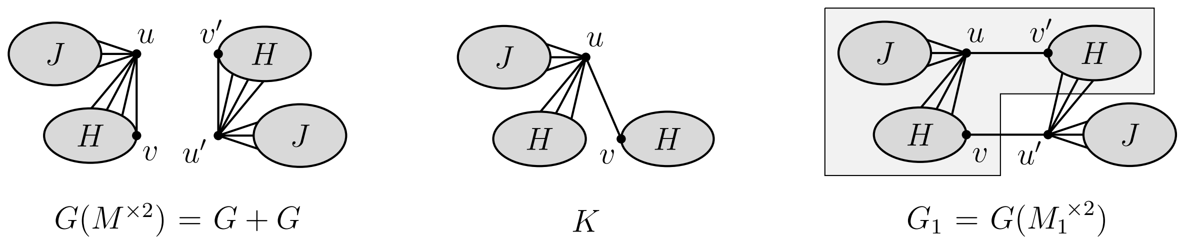

Figure 5: Graphs used in the proof of Corollary 21

we expand G at u along uv, by Lemma 16, $\lambda(G) = \lambda(K)$, where K is the graph in the centre of Figure 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 5. It is readily observed that K is a proper subgraph of G_1 (see the shaded box in Figure 5), and hence, by Corollary 19, $\lambda(K) < \lambda(G_1)$. Since $\mathrm{gr}(\mathsf{Geom}(M_1)) = \lambda(G_1)$, the result follows. □

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

Conjecture 24. 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 we have $\mathrm{gr}(\mathsf{Geom}(M_1)) > \mathrm{gr}(\mathsf{Geom}(M))$.

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 $G(M^{\times 2})$.

3.2 Monotone grid classes

In a recent paper [5], we established the growth rates of *monotone* grid classes. If M is a $0/\pm 1$ matrix, then the *monotone grid class* $\mathsf{Grid}(M)$ 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 $\mathsf{Geom}(M)$ is a subset of $\mathsf{Grid}(M)$ and we have $\mathrm{gr}(\mathsf{Geom}(M)) \leqslant \mathrm{gr}(\mathsf{Grid}(M))$. In fact, the geometric grid class $\mathsf{Geom}(M)$ and the monotone grid class $\mathsf{Grid}(M)$ are identical if and only if $G(M)$ is acyclic (Theorem 3.2 in [2]). Hence, if $G(M)$ is a forest, $\mathrm{gr}(\mathsf{Geom}(M)) = \mathrm{gr}(\mathsf{Grid}(M))$. We determined in [5] that the growth rate of monotone grid class $\mathsf{Grid}(M)$ is equal to the square of the spectral radius of $G(M)$. For acyclic $G(M)$, this is consistent with the growth rate of the geometric grid class as given by Corollary 15.

Typically, the growth rate of a monotone grid class will be greater than that of the corresponding geometric grid class. For example, if $G(M)$ is a cycle then $\mathrm{gr}(\mathsf{Grid}(M)) = 4$, whereas from (3) we have $\mathrm{gr}(\mathsf{Geom}(M)) < 4$. And we have

$$\mathrm{gr}\big(\mathsf{Grid}\left(\begin{smallmatrix}1 & 0 & -1\\ 1 & -1 & 1\end{smallmatrix}\right)\big) \;=\; \mathrm{gr}\big(\mathsf{Grid}\left(\begin{smallmatrix}-1 & 0 & -1\\ 1 & -1 & 1\end{smallmatrix}\right)\big) \;=\; \tfrac{1}{2}(5+\sqrt{17}) \;\approx\; 4.56155,$$

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 25 (Godsil & Gutman [13])**.** *If G is connected and contains a cycle, then we have $\lambda(G) < \rho(G)$.*

Proof. By Lemma 18, if C is a nonempty subgraph of G, then $\rho(G - C) < \rho(G)$. So we have $\Phi_{G-C}(z) > 0$ for all $z \geqslant \rho(G)$. Moreover, $\Phi_G(z) \geqslant 0$ for $z \geqslant \rho(G)$. So, since G contains a cycle, from Lemma 12 we can deduce that $\mu_G(z) > 0$ if $z \geqslant \rho(G)$, and thus $\lambda(G) < \rho(G)$. □

Note that, analogously to Observation 1, $\mathsf{Grid}(M^{\times 2}) = \mathsf{Grid}(M)$. 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 25 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 26. *If $G(M)$ is connected, then* $\mathrm{gr}(\mathsf{Geom}(M)) < \mathrm{gr}(\mathsf{Grid}(M))$ *if and only if $G(M)$ contains a cycle.*

3.3 Subdivision of edges

One particularly surprising result in [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. 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 27 (Hoffman & Smith [17])**.** *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), $\rho(G') \leqslant \rho(G)$, with equality if and only if G is a cycle or has the following form (which we call an "H graph":*

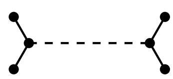

Thus for monotone grid classes, if $G(M)$ is connected, and $G(M')$ is obtained from $G(M)$ by the subdivision of one or more internal edges, then $\mathrm{gr}(\mathsf{Grid}(M')) \leqslant \mathrm{gr}(\mathsf{Grid}(M))$.

As we will see, 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 8–10 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 28. *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 $\lambda(G') \leqslant \lambda(G)$, with equality if and only if G is an H graph.*

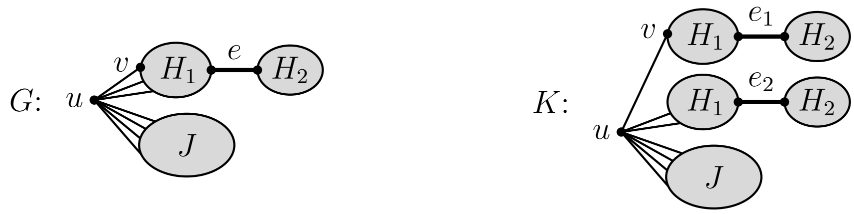

Figure 6: Graphs used in the proof of Lemma 28

Proof. If e lies on an endpath, then G is a proper subgraph of G' and so the result follows from Corollary 19. On the other hand, if e is internal and G is acyclic, the conclusion is a consequence of Corollary 14 and Lemma 27. 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 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. □

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 29. *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 $\mathrm{gr}(\mathsf{Geom}(M')) \leqslant \mathrm{gr}(\mathsf{Geom}(M))$, with equality if and only if $G(M)$ is an H graph.*

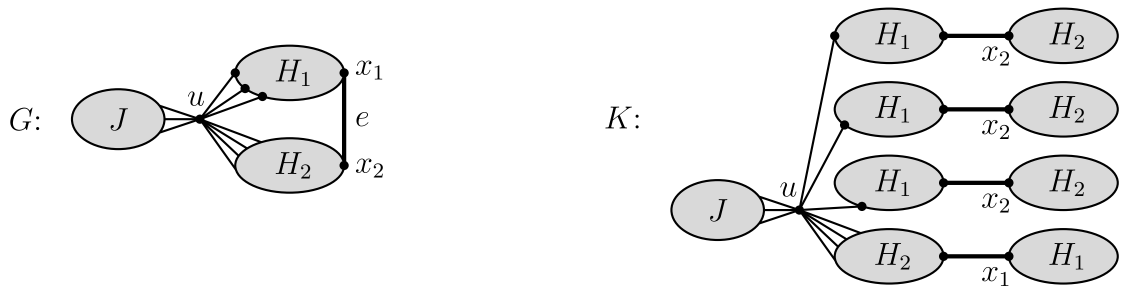

Figure 7: Graphs used in Lemma 30

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 for an illustration. We leave the consideration of multiply-connected graphs that fail to satisfy this condition for future study.

Lemma 30. *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. 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.

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 28, 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 28, 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 9, these effects balance exactly;

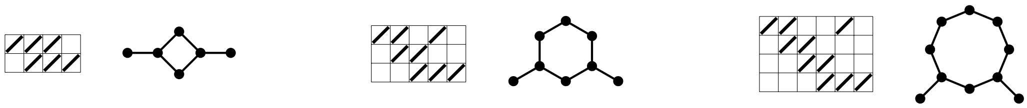

Figure 8: Standard figures and row-column graphs of geometric grid classes whose growth rates increase from left to right

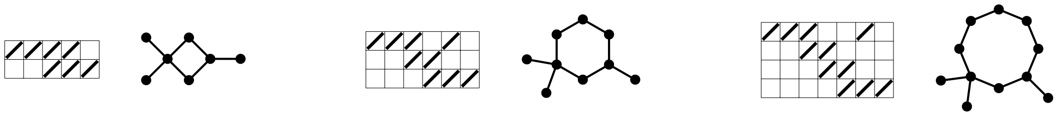

Figure 9: Standard figures and row-column graphs of geometric grid classes whose growth rates are all the same (equal to 5)

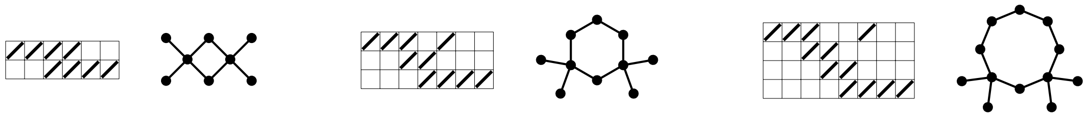

Figure 10: Standard figures and row-column graphs of geometric grid classes whose growth rates decrease from left to right

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 odd cycles and $G^{\times}(M)$ to be $G(M^{\times 2})$ otherwise.

Corollary 31. *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 30, then $\mathrm{gr}(\mathsf{Geom}(M')) > \mathrm{gr}(\mathsf{Geom}(M))$.

(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 30, then $\mathrm{gr}(\mathsf{Geom}(M')) < \mathrm{gr}(\mathsf{Geom}(M))$.

Figure 8 provides an illustration of part (a) and Figure 10 an illustration of part (b).

Acknowledgements

The author would like to thank Penny Bevan and Robert Brignall for reading earlier drafts of this paper. Their feedback led to significant improvements in its presentation. Thanks are also due to an anonymous referee whose comments resulted in a much expanded final section, and to Brendan McKay, whose helpful response to a question on *MathOverflow* [22] alerted the author to the significance of Godsil's Lemma 16. *S.D.G.*

References

[1] 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. (electronic), 2012.

[2] 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.

[3] 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.

[4] Michael H. Albert, Nik Ruškuc, and Vincent Vatter. Inflations of geometric grid classes of permutations. *Israel J. Math.*, in press; arXiv:1202.1833.

[5] David Bevan. Growth rates of permutation grid classes, tours on graphs, and the spectral radius. *Trans. Amer. Math. Soc.*, in press; arXiv:1302.2037.

[6] P. Cartier and D. Foata. *Problèmes combinatoires de commutation et réarrangements.* Lecture Notes in Mathematics, No. 85. Springer-Verlag, 1969.

[7] 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.

[8] Sergi Elizalde. The X-class and almost-increasing permutations. *Ann. Comb.*, 15(1):51–68, 2011.

[9] E. J. Farrell. An introduction to matching polynomials. *J. Combin. Theory Ser. B*, 27(1):75–86, 1979.

[10] Philippe Flajolet and Robert Sedgewick. *Analytic Combinatorics*. Cambridge University Press, 2009.

[11] C. D. Godsil. Matchings and walks in graphs. *J. Graph Theory*, 5(3):285–297, 1981.

[12] C. D. Godsil. *Algebraic combinatorics*. Chapman & Hall, 1993.

[13] C. D. Godsil and I. Gutman. On the theory of the matching polynomial. *J. Graph Theory*, 5(2):137–144, 1981.

[14] Massimiliano Goldwurm and Massimo Santini. Clique polynomials have a unique root of smallest modulus. *Inform. Process. Lett.*, 75(3):127–132, 2000.

[15] Ivan Gutman. The acyclic polynomial of a graph. *Publ. Inst. Math. (Beograd) (N.S.)*, 22(36):63–69, 1977.

[16] Ole J. Heilmann and Elliott H. Lieb. Monomers and dimers. *Phys. Rev. Lett.*, 24:1412–1414, 1970.

[17] 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.

[18] Tomáš Kaiser and Martin Klazar. On growth rates of closed permutation classes. *Electron. J. Combin.*, 9(2): Research paper 10, 20 pp. (electronic), 2003.

[19] L. Lovász and J. Pelikán. On the eigenvalues of trees. *Period. Math. Hungar.*, 3:175–182, 1973.

[20] László Lovász and Michael D. Plummer. *Matching Theory*. AMS Chelsea Publishing, 2009. Corrected reprint of the 1986 original.

[21] Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Combin. Theory Ser. A*, 107(1):153–160, 2004.

[22] Brendan McKay. Roots of matching polynomial of graph. MathOverflow, 2014; http://mathoverflow.net/q/171721.

[23] Abbe Mowshowitz. The characteristic polynomial of a graph. *J. Combin. Theory Ser. B*, 12:177–193, 1972.

[24] Jay Pantone. The enumeration of permutations avoiding 3124 and 4312. Preprint; arXiv:1309.0832.

[25] John Riordan. *An introduction to combinatorial analysis*. Dover Publications Inc., 2002. Reprint of the 1958 original.

[26] Horst Sachs. Beziehungen zwischen den in einem Graphen enthaltenen Kreisen und seinem charakteristischen Polynom. *Publ. Math. Debrecen*, 11:119–134, 1964.

[27] Vincent Vatter. Small permutation classes. *Proc. Lond. Math. Soc.*, 103(5):879–921, 2011.

[28] Vincent Vatter and Steve Waton. On partial well-order for monotone grid classes of permutations. *Order*, 28(2):193–199, 2011.

[29] Vincent Vatter and Steve Waton. On points drawn from a circle. *Electron. J. Combin.*, 18(1): Paper 223, 10 pp. (electronic), 2011.

[30] 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.