CUCS-295-87

GENUS DISTRIBUTIONS FOR BOUQUETS OF CIRCLES

Jonathan L. Gross[1]
Department of Computer Science
Columbia University
New York, NY 10027

David P. Robbins
Institute for Defense Analyses
Princeton, NJ 08540

Thomas W. Tucker[2]
Department of Mathematics
Colgate University
Hamilton, NY 13346

ABSTRACT. The genus distribution of a graph G is defined to be the sequence $\{g_m\}$ such that g_m is the number of different imbeddings of G in the closed orientable surface of genus m. A counting formula of D.M. Jackson concerning the cycle structure of permutations is used to derive the genus distribution for any bouquet of circles B_n. It is proved that all these genus distributions for bouquets are strongly unimodal.

1. Supported by ONR Contract N00014-85-0768.

2. Supported by NSF Contract DMS-8601760.

1. Introduction

It is a reasonably straightforward matter to relate the number of imbeddings of a bouquet in the 2-sphere to Catalan numbers, but attempts at similarly direct approaches to counting the imbeddings of bouquets in the higher genus surfaces have not met with quick success. However, enumerative results by D. M. Jackson [11] concerning the cycle structure of permutations, and obtained by methods involving the representation theory of symmetric groups, permit us to make such counts. The principal task of this paper is to give a mathematical justification of a translation of our enumeration problem into Jackson's context, so that we may apply his formula.

By a bouquet of circles, or more briefly, a bouquet, we mean a graph with one vertex and some self-loops. In particular, the bouquet with n self-loops is denoted B_n. One sense in which bouquets are fundamental building blocks of topological graph theory is that any connected graph can be reduced to a bouquet by contracting a spanning tree to a point. Another is that Cayley graphs and many other regular graphs are covering spaces of bouquets, as has been demonstrated with voltage graphs by Gross and Tucker [7] and by Gross [4].

Every surface in this paper is closed and orientable. An oriented surface is one with a given fixed orientation. The closed orientable surface of genus m is denoted S_m; for instance, S_0 denotes the sphere and S_1 the torus.

As is customary in topological graph theory, a graph is permitted to have self-loops and/or multiple adjacencies. All imbeddings of graphs in surfaces are 2-cell imbeddings, that is, every complementary region or face of the imbedding is homeomorphic to an open disc. For general background in topological graph theory, see Gross and Tucker [8] or White [22].

By the genus distribution of a graph G, we mean the sequence $g_0, g_1, \ldots,$ where g_m is the number of different imbeddings of the graph G in the oriented surface S_m. If there is more than one graph at hand, we write $g_m(G)$ and $g_m(H)$ to distinguish their distributions. In particular, $g_m(B_n)$ denotes the number of imbeddings of the bouquet B_n in the surface S_m.

The outline of this paper is as follows. Section 2 discusses our notion of equivalence of imbeddings and its connection with rotation systems and permutations of the directed edge set of a graph. Section 3 relates the numbers $g_m(B_n)$ to the numbers computed by Jackson [11]. Section 4 shows that for every n, the sequence $g_m(B_n)$ is strongly unimodal. Section 5 poses various questions.

2. Equivalence of imbeddings and rotation systems

In order even to talk about the number of imbeddings of a given graph, we need to be more explicit about what we mean by graph, imbedding, and equivalence of imbeddings. A graph here is a topological space given in the following manner as a finite 1-dimensional CW-complex. The vertex set V of a graph G is a finite set of points. The edge set E is a finite number of copies of the unit interval $[0,1]$. For each edge e there is a function $f_e:\{0,1\}\rightarrow V$ telling where to attach the end-points of edge e. The graph G is then the identification space formed from the disjoint union of V and E by identifying end-points of edges with vertices via the functions f_e. Observe that this description in effect assigns a "direction" for every edge e with initial vertex $f_e(0)$ and terminal vertex $f_e(1)$, even if edge e is a self-loop.

If a graph has no multiple edges or loops, it can be given in a coded form simply be listing the symbols for the vertices, say integers $1, \ldots, n$, together with ordered pairs of vertices giving the initial and terminal vertices of each edge. A graph with multiple edges or loops can be coded in the following way. Symbols for edges are listed and to each edge e are associated two symbols, e^+ and e^-. The set $D(E)$ of all such symbols is called the <u>directed edge set</u> of the graph. Each vertex v is then specified by a subset of directed edge symbols: e^+ is in the subset if and only if vertex v is the initial vertex of edge e, and e^- is in the subset if and only if vertex v is the terminal vertex of edge e. Thus a graph can be given simply by a partition of its directed edge set.

An imbedding of a graph G in a surface S is a continuous one-to-one function $f:G\to S$. Two imbeddings $f:G\to S$ and $g:G\to S$ of a graph G in the oriented surface S are equivalent if there is an orientation-preserving homeomorphism $h:S\to S$ such that $hf=g$. This means that the surface homeomorphism h must respect the labeling and directing of edges: for each edge e, h must take $f(e)$ to $g(e)$ and the plus direction of $f(e)$ to the plus direction of $g(e)$. Throughout this paper, "number of imbeddings" really means number of equivalence classes of imbeddings.

There are some weaker notions of "equivalence" of imbeddings. The most common is to say that imbeddings $f:G\to S$ and $g:G\to S$ are "equivalent" if there is a (not necessarily orientation-preserving) homeomorphism $h:S\to S$ such that h takes the image $f(G)$ onto the image $g(G)$ (equivalently, there is a graph automorphism $j:G\to G$ such that $hf=gj$.). As suggested in [8], we will call this weaker form of equivalence congruence. If the homeomorphism h preserves orientation, we call this oriented congruence. It is congruence that is found, for example, in Negami [16], Mull, Rieper, and White [15], and Bender and Canfield [1].

To illustrate the concept of equivalence, we consider the four imbeddings of the bouquet B_2 in the sphere S_0 given in Figure 1. The loops are labeled and directed because in order to know which imbedding $f:B_2 \to S_0$ is being illustrated, we must know one loop from another and, for each loop, which end is which. Clearly, all four imbeddings are congruent, since congruence ignores labels and directions. However, as the reader should be able to verify, these four imbeddings are mutually inequivalent.

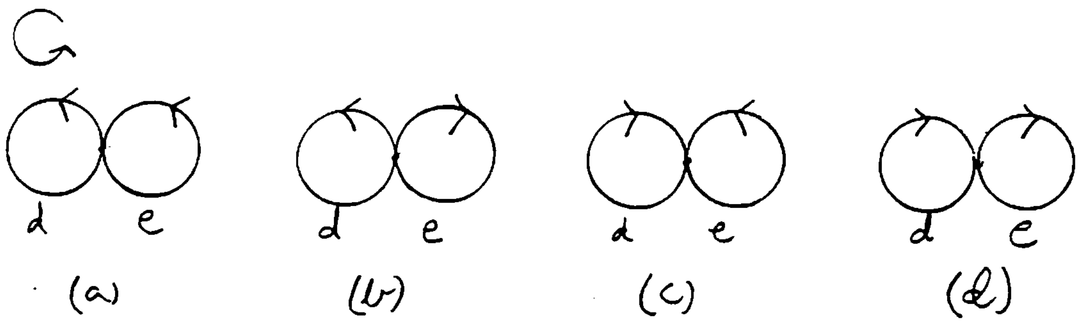

Figure 2.1 Four inequivalent imbeddings of the bouquet B_2 in the sphere.

Are there any other inequivalent imbeddings of the bouquet B_2 in the sphere? In Figure 2.2, there are illustrated some plausible candidates. We claim that all three are equivalent to imbedding (a) in Figure 2.1. For the "wild"-looking imbedding on the left, we appeal to the Schoenfliess Theorem to straighten out the closed curve representing loop e. For the middle imbedding, we remind the reader that the imbedding surface is the sphere, not the plane, and that if one chooses to put the point at infinity in the region between d and e, one has again imbedding (a) from Figure 2.1. Finally, for the imbedding on the right of Figure 2.2, a rotation by 180° about the vertex brings us back to imbedding (a). This imbedding could also be turned into imbedding (d) by a reflection in a vertical axis, but we require our homeomorphism h to be orientation-preserving; thus imbeddings (a) and (d) are not equivalent under our definition. At this point, we hope the reader believes that the four imbeddings in Figure 2.1 give all the equivalence classes of imbeddings of the bouquet B_2 in the sphere, even though we have not given a formal proof.

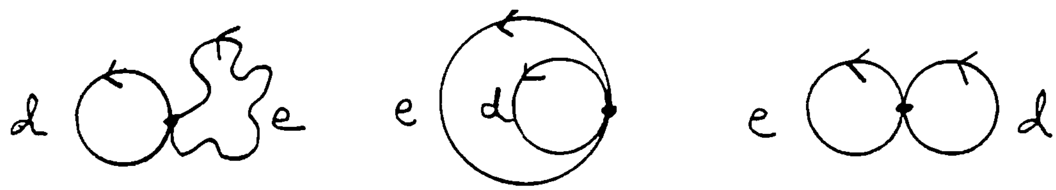

Figure 2.2 Three imbeddings of the bouquet B_2 in the 2-sphere.

The role of edge-directions and loops in equivalence of imbeddings deserves some mention. One must remember that the edge-directions of a given graph are fixed at the outset. Thus, if edge e in graph G is not a loop, we cannot obtain a "different" imbedding for G simply by reversing an arrow on edge e in some drawing of an imbedding of G; reversing the plus-direction of edge e changes the underlying representation of the graph G, which has already been agreed upon. On the other hand, if edge e is a loop, then reversing an arrow on edge e in a drawing of an imbedding of G is possible, since such a reversal can be obtained by maintaining the same agreed-upon direction for the loop e, but changing the imbedding by flipping the loop e over.

Equivalence of imbeddings ignores small differences in the appearance of imbedding, and focuses instead on the underlying combinatorial structure of a labeled graph imbedding. One should expect, therefore, that a given equivalence class can be encoded in some finite way using only the symbols for the edges.

There is one obvious way to encode imbeddings. Simply list in cyclic order the directed edges encountered in a closed walk around the boundary of each face in the imbedding, oriented by the given orientation of the surface. Since any surface homeomorphism providing an equivalence between two imbeddings must preserve labels and directions of edges and must take each oriented face to a like-oriented face, equivalent imbeddings clearly generate the same cyclic lists of boundary walks. Conversely, given such cyclic lists of boundary walks, one can recover the imbedding surface, up to equivalence, simply by identifying, for each cyclic list, the boundary of a polygonal disk to the closed walk in the graph given by that list.

However, there is a difficulty with this method of encoding an imbedding. Not every partition of the edges of a graph into cyclic lists gives an imbedding: the cyclic lists must correspond to closed circuits in the graph and, even then, one may obtain a pseudo-surface rather than a surface, that is, a surface that is "pinched" by identifying some points together.

Instead, we adopt the dual viewpoint of specifying, for each vertex v, a cyclic list of the directed edges encountered in a trip around (following the given orientation of the surface) the boundary of a small disk centered at vertex v. In this list, an occurrence of the edge e is written as e^+ if v is its initial vertex and e^- if v is its terminal vertex. Such a partition of the directed edges of the graph G into cyclic lists is called a <u>rotation system</u> for the graph G. For example, the rotation systems defined by the four imbeddings in Figure 2.1 are as follows, if the given orientation of the sphere is counterclockwise:

(a) $(d^+ \ d^- \ e^+ \ e^-)$

(b) $(d^+ \ d^- \ e^- \ e^+)$

(c) $(d^- \ d^+ \ e^+ \ e^-)$

(d) $(d^- \ d^+ \ e^- \ e^+)$

Equivalent imbeddings of a graph define the same rotation system, since any equivalence between two imbeddings respects edge labels and directions and takes any small oriented disk centered at one vertex to another small like-oriented disk at the image vertex. Conversely, given any rotation system for a graph, one can recover the associated imbedding by determining the cyclic lists of edges giving the face boundary walks, as follows. Choose any directed edge e as the start of the face boundary walk. Then the next edge in that walk is the directed edge that immediately follows the opposite of e in the rotation at the vertex at which the opposite of e terminates. Continue generating successive directed edges until the walk closes. Then choose any unused directed edge as the start of the next closed walk. When all directed edges have been used, stop.

This one-to-one correspondence between rotation systems and equivalence classes of imbeddings was given by Edmonds [2] for the case of simplicial graphs. A less convenient dual form was previously known to Heffter [10]. The general form given here for graphs with possible loops and multiple edges is due to Gross and Alpert [5].

The cyclic lists given by a rotation system suggest that we consider each rotation system simply as a single permutation of the directed edge set of the graph given in cycle form. It turns out that this viewpoint has a substantial payoff: it allows an immediate algebraic method of determining the faces of the imbedding. Given a graph G and a rotation system viewed as a permutation ρ on the directed edge-set $D(E)$ of G, then the faces of the associated imbedding are given by the cycles of the permutation $\rho \circ \beta$, where β is the permutation of $D(E)$ that interchanges e^+ and e^-, for each edge e. To see this, simply think about how faces are recovered from a rotation system. A directed edge, say e^ϵ where $\epsilon = \pm$, is selected as the first edge in a face boundary. The next edge is obtained by first going to the vertex v where e^{ϵ} terminates. In the cyclic list given by the rotation at that vertex v, one finds $e^{-\epsilon}$ listed. The next edge in the face boundary is the one following $e^{-\epsilon}$ in the rotation at vertex v. Thus we first apply β to e^{ϵ} and then apply ρ to the result. For example, in imbedding (a) of Figure 1 we have:

$$\rho = (d^+ \, d^- \, e^+ \, e^-) \qquad \beta = (d^+ \, d^-)(e^+ \, e^-)$$

$$\rho \circ \beta = (d^+ \, e^+)(d^-)(e^-)$$

Observe that the faces are oriented by the opposite, clockwise orientation of the sphere (the edge directions d^+ and e^+ for the "outside" face of the imbedding really do run against the counterclockwise orientation given by the little circular arrow at the upper left of Figure 2.1.)

This viewpoint of rotation systems as permutations of the directed edge set has been exploited by a number of authors, especially Stahl [20]. As Stahl [21] has observed, it makes some graph imbedding questions accessible to techniques from the representation theory of symmetric groups. We summarize this discussion in the following theorem.

THEOREM 2.1 Every equivalence class of imbeddings of a graph G in an oriented surface corresponds uniquely to a permutation ρ of the directed edge set of G such that each cycle of ρ gives the list of edges encountered in an oriented trip on the surface around a vertex of G. Moreover, the faces of the imbedding (oppositely oriented) are given by the cycles of $\rho \circ \beta$, where β is the involution on the directed edge set of G that takes each directed edge to its reverse.

The following obvious corollary to Theorem 2.1 indicates why it is infeasible to calculate the genus distribution of a graph simply by writing down each possible rotation system ρ and counting the number of cycles in $\rho \circ \beta$: there are of course too many rotation systems.

THEOREM 2.2 Let G be an n-vertex graph who vertices have valences $d_1, \ldots, d_n$. Then the number of equivalence classes of imbeddings of G is the product

$$\prod_{i=1}^{n} (d_i - 1)!$$

The problem of calculating genus distributions in a less costly manner than case-by-case has been raised by Gross and Furst [6]. Since knowing the genus distribution implies knowing the minimum genus, one might expect that quite powerful enumerative methods would be required to calculate the genus distributions for many standard classes of graphs. For instance, a calculation of the genus distribution of the complete graphs would yield a new proof of the Ringel-Youngs Theorem [18] that solved the Heawood map-color problem.

3. Applications of Jackson's Formula

Jackson [11] observes in his paper, indeed, in the title itself, that his results on counting certain kinds of permutations with a given number of cycles has application to topological problems. To apply Jackson's results to the genus distribution of the bouquet B_n, we first consider what amounts to a special case of Theorem 2.1.

THEOREM 3.1 The number $g_m(B_n)$ of imbeddings of the bouquet B_n in the oriented surface of genus m is equal to the number of permutations ρ on the $2n$ symbols $\{e_1^+, e_1^-, e_2^+, e_2^-, \ldots, e_n^+, e_n^-\}$ such that ρ has a single cycle of length $2n$ and $\rho \circ \beta_0$ has $n - 2m + 1$ cycles, where β_0 is the involution $(e_1^+ \, e_1^-)(e_2^+ \, e_2^-)\ldots(e_2^+ \, e_2^-)$.

Proof. Since the bouquet B_n has only one vertex, it follows that each rotation system ρ for B_n consists of a single 2n-long cycle of directed edges. The number k of cycles of $\rho \circ \beta_0$ equals the number of faces of the associated imbedding. Hence by Euler's equation for the surface of genus m:

$$2 - 2m = V - E + F = 1 - n + k.$$

Thus, the composition permutation $\rho \circ \beta_0$ has $k = n - 2m + 1$ cycles.

We can interpret $g_m(B_n)$ in terms of Σ_{2n}, the full symmetric group on $2n$ symbols. Call $\rho \in \Sigma_{2n}$ a <u>long cycle</u> if it consists of a single cycle of length $2n$, and call $\beta \in \Sigma_{2n}$ a <u>full involution</u> if it consists of n cycles of length 2. Fix a full involution $\beta_0 \in \Sigma_{2n}$. Then $g_m(B_n)$ is a number of elements in the set:

$$\{\pi \in \Sigma_{2n} \mid \pi \text{ has } n - 2m + 1 \text{ cycles and } \pi = \rho \circ \beta_0 \text{ for some long cycle } \rho\}.$$

By contrast, fix a long-cycle ρ_0. Then Jackson counts, among other things, the number of elements in the set:

$$\{\pi \in \Sigma_{2n} \mid \pi \text{ has } k \text{ cycles and } \pi = \rho_0 \circ \beta \text{ for some full involution } \beta\}.$$

Jackson denotes this number $e_k^{(2)}(n)$ which we abbreviate here to $e_k(n)$.

It is not difficult to find the numerical relationship between $g_m(B_n)$ and $e_k(n)$, but it is also instructive to see what Jackson's $e_k(n)$ counts in terms of graph imbeddings. We begin with a picture of what an imbedding of the bouquet B_n looks like near its only vertex v: there we see $2n$ "spokes" radiating out from vertex v as illustrated on the left of Figure 3 for $n = 3$. If Jackson's fixed $2n$-cycle ρ_0 is denoted $(1\ 2\ 3\ \ldots\ 2n)$, choose one spoke of the imbedding to label 1 and then label the other spokes around the vertex v by $2,\ldots,2n$ using the cyclic order determined by the given orientation of the imbedding surface. Then Jackson's varying involution β encodes which pairs of spokes are connected in order to complete the loops of the bouquet B_n, as illustrated in the middle of Figure 3.1 for $\beta = (13)(26)(45)$. The faces of the imbedding correspond to the cycles of $\rho_0 \circ \beta$, but unlike a usual rotation system, the edges are unlabeled and undirected. In fact, after agreeing to let spoke 1 be the initial end of the directed edge of e^+, there are $2! \cdot 2^2$ different ways to label and direct the remaining two loops of the imbedding. On the right of Figure 3.1 is shown one of the possible ways, which yields the rotation $(e_1^+ e_3^- e_1^- e_2^+ e_2^- e_3^+)$.

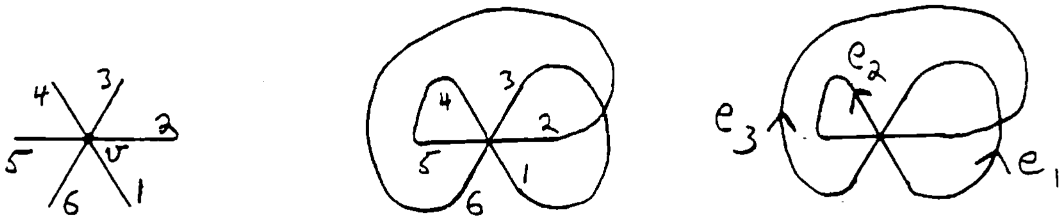

Figure 3.1 The imbedding of B_3 corresponding to the involution $\beta = (13)(26)(45)$.

Jackson's number $e_k(n)$ can be interpreted in terms of congruence classes of imbeddings in the following way. Attach to the bouquet B_n a single edge ("stem") leading to a new vertex of valence one. This stem can then be used to mark which spoke is to be labeled 1 as the first spoke encountered after the stem using the given orientation of the imbedding surface. Let B_n' denote such a <u>stemmed</u> <u>bouquet</u>. Figure 3.2 shows the five oriented congruence classes of imbeddings in the sphere for the stemmed bouquet B_3' together with the corresponding involution $\beta \in \Sigma_6$. The reader should check that for every one of the ten other full involutions in Σ_6, the product of such an involution with $\rho_0 = (1\ 2\ 3\ 4\ 5\ 6)$ has two cycles and hence represents an imbedding in the surface of genus one.

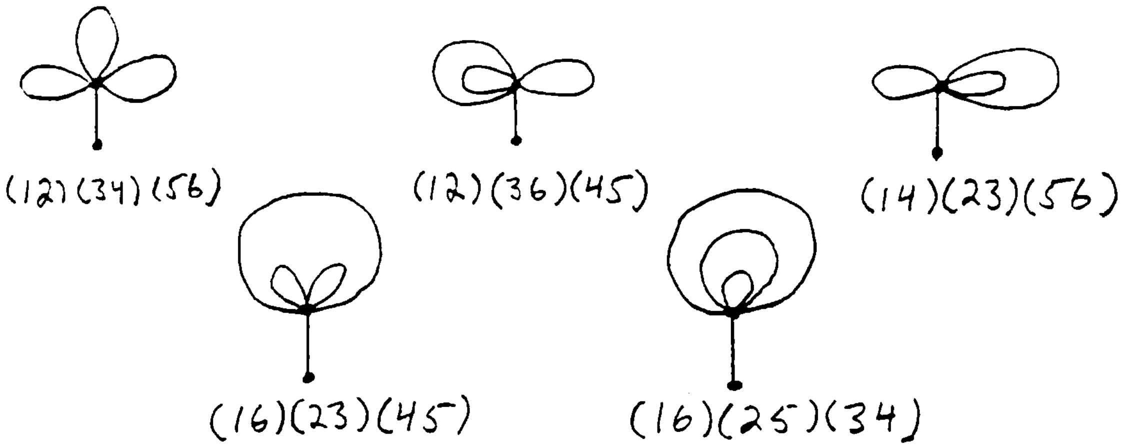

Figure 3.2 The five imbeddings in the sphere for the stemmed bouquet B_3'.

Let us abbreviate $g_m(B_n)$ by $g_m(n)$. Then this discussion can be summarized in the following theorem.

THEOREM 3.2 The number of oriented congruence classes of the stemmed bouquet B_n' in the oriented surface of genus m is $e_k(n)$, where $k = n - 2m + 1$. Moreover, $g_m(n) = (n - 1)! \cdot 2^{n-1} \cdot e_k(n)$.

Proof. The equation $k = n - 2m + 1$ has already been explained in Theorem 3.1. The one - to - $(n - 1)! \cdot 2^{n-1}$ correspondence between Jackson's involution β and our rotation ρ is given as follows. If $\beta(1) = j$, then positions 1 and j in the 2n-cycle ρ are occupied by e_1^+ and e_1^-, respectively. For each other 2-cycle (ij) in the involution β, choose an edge label e from $e_2, \ldots, e_n$ and put e^+ or e^- in position i of ρ and e^- or e^+, respectively, in position j. Conversely, given a rotation ρ for B_n, write ρ in cyclic form with e_1^+ listed first. Then each transposition in the corresponding involution β corresponds to two positions in ρ occupied by the same edge label. □

Jackson describes the numbers $e_k(n)$ in terms of closed formulas, generating functions, and recurrence relations. All of these descriptions can therefore be applied to $g_m(n)$. We give here the recurrence relation for $g_m(n)$.

THEOREM 3.3 The numbers $g_m(n)$ satisfy the recurrence for $n > 2$:

$$(n+1)g_m(n) = 4(2n-1)(2n-3)(n-1)^2(n-2)g_{m-1}(n-2) + 4(2n-1)(n-1)g_m(n-1)$$

with the boundary conditions

$$g_m(n) = 0 \quad \text{if} \quad m < 0 \quad \text{or} \quad n < 0,$$
$$g_0(0) = g_0(1) = 1 \quad \text{and} \quad g_m(0) = g_m(1) = 0 \quad \text{for} \quad m > 0,$$
$$g_0(2) = 4,\ g_1(2) = 2,\ g_m(2) = 0 \quad \text{for} \quad m > 1.$$

Proof. By Lemma 6.1 of Jackson [11],

$$(n+1)e_k(n) = (2n-1)(2n-3)(n-1)e_k(n-2) + 2\,(2n-1)e_{k-1}(n-1),$$

where n and k have opposite parity and $n > 0$. The desired recurrence for $g_m(n)$, $n > 2$, is then obtained by substituting $n - 2m + 1$ for k, multiplying both sides by $(n-1)!\,2^{n-1}$, and using Theorem 3.2. The recurrence for $g_m(n)$ begins only at $n = 3$, because multiplying by $(n-1)!$ causes a spurious factor of 0 in front of the

$g_{m-1}(n-2)$ term when $n \le 2$. The boundary conditions for $n \le 2$ are easily verified.

Table 1 Some values of $g_m(n)$

n/m	0	1	2	total
0	1			1
1	1			1 = 1!
2	4	2		6 = 3!
3	40	80		120 = 5!
4	672	3360	1008	5040 = 7!
5	16128	161280	185472	362880 = 9!

4. Unimodality

A nonnegative sequence $\{a_m\}$ is said to be unimodal if there exists at least one integer M such that

$$a_{m-1} \le a_m \quad \text{for all} \quad m \le M, \text{ and}$$
$$a_m \ge a_{m+1} \quad \text{for all} \quad m \ge M.$$

Although this includes nondecreasing sequences that eventually level off and nonincreasing sequences that start out level, a typical unimodal sequence first rises and then falls.

A sequence $\{a_m\}$ is called strongly unimodal if its convolution with any unimodal sequence $\{b_m\}$ is unimodal. Keilson and Gerber [12] have proved that $\{a_m\}$ is strongly unimodal if and only if

$$a_m^2 \ge a_{m+1}\, a_{m-1} \qquad \text{for all} \quad m$$

or, equivalently, if and only if $\{a_m\}$ is unimodal and

$$\frac{a_{m+1}}{a_m} \le \frac{a_m}{a_{m-1}}$$

wherever these ratios are defined. This property is also called log concavity.

THEOREM 4.1 The genus distribution of the bouquet B_m is strongly unimodal.

Proof. The proof is by induction on n. The genus distributions for B_0, B_1, and B_2, are strongly unimodal, thus providing a basis for an induction on $n > 2$ using Theorem 3.3. The recurrence for $g_m(n)$ given by Theorem 3.3 can be written in the form:

$$g_m(n) = A(n)g_m(n-1) + B(n)g_{m-1}(n-2) \qquad (4.1.1)$$

where $A(n)$ and $B(n)$ are positive functions of n. In this form, the recurrence holds not only for $n > 2$, but also for $n > 0$, if we set $A(2) = 4$, $B(2) = 2$, $A(1) = 1$, $B(1) = 0$.

We wish to prove that

$$g_m(n) \geq g_{m+1}(n)g_{m-1}(n) . \tag{4.1.2}$$

We may assume that $n - 2m + 1 > 0$, since $g_m(n) = g_{m+1}(n) = 0$ by Theorem 3.1 whenever $n - 2m + 1 \leq 0$. By using recurrence (4.1.4), we expand the left side of inequality (4.1.2) to

$$A^2(n)g_m^2(n-1) + B^2(n)g_{m-1}^2(n-2)$$
$$+ 2A(n)B(n)g_m(n-1)g_{m-1}(n-2)$$

and the right side to

$$A^2(n)g_{m+1}(n-1)g_{m-1}(n-1) + B^2(n)g_m(n-2)g_{m-2}(n-2)$$
$$+ A(n)B(n)[g_{m+1}(n-1)g_{m-2}(n-2) + g_{m-1}(n-1)g_m(n-2)].$$

The induction hypothesis implies that the first and second terms of the expanded left side dominate the first and second terms, respectively, of the expanded right side. Accordingly, the theorem is proved if we can establish that

$$g_m(n-1)g_{m-1}(n-2) \geq g_{m+1}(n-1)g_{m-2}(n-2), \tag{4.1.3a}$$
$$g_m(n-1)g_{m-1}(n-2) \geq g_{m-1}(n-1)g_m(n-2) . \tag{4.1.3b}$$

Applying recurrence (4.1.1) to the $(n-1)$ terms, we expand the left side of inequality (4.1.3a) to the formula

$$[A(n-1)g_m(n-2) + B(n-1)g_{m-2}(n-2)]g_{m-1}(n-2)$$

and the right side to the formula

$$[A(n-1)g_{m+1}(n-2) + B(n-1)g_m(n-3)]g_{m-2}(n-2).$$

It is easy to show that a strongly unimodal sequence $\{a_m\}$ satisfies $a_m a_{m-1} \geq a_{m+1}a_{m-2}$. Thus by the induction hypothesis, we have

$$g_m(n-2)g_{m-1}(n-2) \geq g_{m+1}(n-2)g_{m-2}(n-2).$$

Therefore, in order to establish (4.1.3a), it suffices to prove that

$$g_{m-2}(n-2)g_{m-1}(n-2) \geq g_m(n-3)g_{m-2}(n-2). \qquad (4.1.4a)$$

If we once again apply the recurrence to the $(n-2)$ terms in (4.1.4a) and use the strong unimodality of $g_m(n-3)$, we find that it suffices to prove that

$$g_{m-1}(n-3)g_{m-2}(n-4) \geq g_m(n-3)g_{m-3}(n-2),$$

which is just (4.1.3a) with m reduced by 1 and n by 2. Recalling that the recurrence (4.1.1) holds for $n > 0$ and that $n-2m+1 > 0$, we can safely iterate this argument until we reach the inequality

$$g_1(n-2m+1)g_0(n-2m) \geq g_2(n-2m+1)g_{-1}(n-2m).$$

Since $g_{-1}(n-2m) = 0$, this last inequality is true and inequality (4.1.3a) is proved.

A similar argument confirms inequality (4.1.3b) by reducing it to the form

$$g_1(n-2m+1)g_0(n-2m) \geq g_0(n-2m+1)g_1(n-2m).$$

A final application of recurrence (4.1.4) to the $(n-2m+1)$ terms yields for the left side,

$$[A(n-2m+1)g_1(n-2m) + B(n-2m+1)g_0(n-2m-1)]g_0(n-2m),$$

which clearly dominates the right side.

$$A(n-2m+1)g_0(n-2m)g_1(n-2m).$$

Thus, both inequalities (4.1.3a) and (4.1.3b) have been established, thereby completing the proof. □

Clearly, Theorem 4.1 applies to any triangular array satisfying a recurrence of the form (4.1.1) with appropriate initial conditions. The only restriction on the coefficients $A(n)$ and $B(n)$ is that they are nonnegative. Sagan [19] has obtained a similar result for a slightly different recurrence, but there appears to be no direct connection between Sagan's result and Theorem 4.1.

5. Problems for further study

The determination of the genus distribution $g_m(B_n)$ raises a variety of questions. Some of the problems listed here may well be amenable to an attack like Jackson's using the representation theory for the symmetric group.

(5.1) Compute the genus distribution for some other interesting graphs.

Furst, Gross, and Statman [3] have computed genus distributions for "ladders" and "cobblestone paths". McGeogh [14] has done circular ladders and Mobius ladders, and Klein [13] has done Ringel ladders. A reasonable candidate for study is the dipole D_n, which consists of 2 vertices joined by n edges. A rotation system for the dipole D_n is a permutation ρ of the form $(e_1^+ \ldots)(e_1^- \ldots)$, where both cycles have length n and the first cycle uses only the symbols $e_1^+, \ldots, e_n^+$ and the second cycle uses only the symbols $e_1^- \ldots e_n^-$. To compute $g_m(D_n)$, one must count the number of such permutations ρ such that $\rho \circ \beta_0$ has $n - 2m$ cycles; as with bouquets $\beta_0 = (e_1^+ e_1^-)(e_2^+ e_2^-) \ldots (e_n^+ e_n^-)$. One could also enumerate oriented congruence classes analogous to Jackson's $e_k(n)$ by counting the number of full involutions $\beta \in \Sigma_{2n}$ such that $\beta(1) = 2$, $\beta(i)$ and i have opposite parity for all i, and $\beta \circ (1\ 3\ 6\ 7 \ldots)(2\ 4\ 6\ 8 \ldots)$ has k cycles. Rieper [17] has computed, both for the dipole D_n and the bouquet B_n, the number of equivalence classes of imbeddings having a specified structure of region sizes, that is, rotations ρ such that $\rho \circ \beta_0$ has a given cycle structure. Rieper uses these computations for asymptotic enumeration. He does not, however, derive a closed formula for $g_m(D_n)$ or recurrence relations for $g_m(B_n)$ or $g_m(D_n)$.

Another candidate is the wheel graph W_n, which consists of a n vertices on an n-cycle together with an extra vertex or "hub" joined to the n vertices by n edges or "spokes". For each imbedding of W_n, there is an associated imbedding for the bouquet B_n obtained by contracting the spokes to the hub. To compute $g_m(W_n)$ it then suffices to count the number of full involutions $\beta \in \Sigma_{2n}$ such that $\beta \circ (12)(34)\ldots(2n-1\ 2n)$ has two n-cycles and $\beta \circ (1\ 2\ 3\ldots 2n)$ has k cycles.

(5.2) Find a way to exploit the fact that every graph contracts along a spanning tree to a bouquet of circles.

For example, computing the genus distribution for the wheel W_n reduces to a bouquet imbedding problem with an extra restriction. Perhaps the contracted spanning tree marks the bouquet in some manner analogous to the stem introduced in Section 3.

(5.3) Prove the recurrence equation of Theorem 3.3 using direct topological methods.

A topological derivation of the recurrence might generalize to other graphs while methods using the representation theory of the symmetric group might not. Bender and Canfield [1] enumerate asymptotically all "rooted" n-edge imbeddings of genus m. Can their methods be be applied to a restricted family of graphs like the bouquets B_n?

(5.4) Prove that $g_m(G)$ is strongly unimodal for any graph G.

All available evidence supports the conjecture, but very few genus distributions are known. It is conceivable that any graph G can be resolved into a sequence of graphs $G = G_N, G_{N-1}, \ldots, G_0$, where G_n is obtained by deleting an edge from G_{n+1} and G_0 is a spanning tree for G, and that in this context there is a recurrence of the form (4.1.1) for the graphs G_n.

References

1. E. A. Bender and E. R. Canfield, The asymptotic number of rooted maps on a surface, J. Combin. Theory Ser. A 43 (1986), 244-257.
2. J. Edmonds, A combinatorial representation for polyhedral surfaces (abstract), Notices Amer. Math. Soc 7 (1960), 646.
3. M. Furst, J. L. Gross, and R. Statman, Genus distributions for two classes of graphs, preprint (1985).
4. J. L. Gross, Every connected regular graph of even degree is a Schreier coset graph, J. Combinatorial Theory B 22 (1977), 227-232.
5. J. L. Gross and S. R. Alpert, The topological theory of current graphs, J. Combinatorial Theory B 17 (1974), 218-233.
6. J. L. Gross and M. Furst, Hierarchy for imbedding-distribution invariants of a graph, J. Graph Theory 11 (1987), to appear.
7. J. L. Gross and T. W. Tucker, Generating all graph coverings by permutation voltage assignments. Discrete Math. 18 (1977), 273-283.
8. J. L. Gross and T. W. Tucker, Topological Graph Theory, Wiley-Interscience (1987), to appear.
9. P. J. Heawood, Map-colour theorem, Quart. J. Math. 24 (1890), 332-338.
10. L. Heffter, Uber das Problem der Nachbargebiete, Math. Annalen 38 (1981), 477-508.
11. D. M. Jackson, Counting cycles in permutations by group characters, with an application to a topological problem, Trans. Amer. Math. Soc. 299 (1987), 785-801.

12. J. Keilson and H. Gerber, Some results for discrete unimodality, J. Amer. Stat. Asso. 66 (1971), 386-389.

13. E. W. Klein, Private communication.

14. L. McGeoch, Ph.D Thesis, Carnegie-Mellon, 1987.

15. B. P. Mull, R. G. Reiper, and A. T. White, Enumerating 2-cell imbeddings of complete graphs, Proc. Amer. Math. Soc., to appear.

16. S. Negami, Uniqueness and faithfulness of embedding of toroidal graphs, Discrete Math. 44 (1983), 161-180.

17. R. G. Rieper, Ph.D. Thesis, Western Michigan University, 1987.

18. G. Ringel and J. W. T. Youngs, Solution of the Heawood map-coloring problem, Proc. Nat. Acad. Sci. USA 60 (1968), 438-445.

19. B. Sagan, Inductive and injective proofs of log concavity results, preprint (1987).

20. S. Stahl, Permutation-partition pairs: a combinatorial generalization of graph embeddings, Trans. Amer. Math. Soc. 259 (1980).

21. S. Stahl, Average genus of classes of graph embeddings, Congressus Numerantium 40 (1983), 375-388.

22. A. T. White, Graphs, Groups and Surfaces, 2nd edition, North-Holland, (1984).