

On the Average Genus of a Graph

Jonathan L. Gross, E. Ward Klein, Robert G. Rieper

CUCS-482-89

On the Average Genus of a Graph

Jonathan L. Gross*, E. Ward Klein, Robert G. Rieper
Columbia University, New York, NY 10027

ABSTRACT. Not all rational numbers are possibilities for the average genus of an individual graph. The smallest such numbers are determined, and varied examples are constructed to demonstrate that a single value of average genus can be shared by arbitrarily many different graphs. It is proved that the number one is a limit point of the set of possible values for average genus and that the complete graph K_4 is the only 3-connected graph whose average genus is less than one. Several problems for future study are suggested.

version: 18 September 1989

* Partially supported by ONR Contract N00014-85-0768.

1. Introduction

By the average genus of a graph G , we mean the average value of the genus of the imbedding surface, taken over all orientable imbeddings of G . This value is evidently a rational number, and it is clearly an invariant of the homeomorphism type of a graph.

Studying the average genus of an individual graph was first suggested by Gross and Furst [1987], who placed it toward the bottom of an extensive hierarchy of invariants of isomorphism type, now known (Gross, Rieper, and Tucker [1989]) to contain complete invariants higher up. Other low-end invariants are the genus distribution and the face-size distribution.

A graph has average genus zero if and only if it has maximum genus zero. In Section 2, we establish that the smallest possible positive values of average genus are

$$1/3, \quad 1/2, \quad 5/9, \quad 2/3, \quad 19/27, \quad 3/4$$

It is easy to construct trivial examples of different graphs with the same average genus, as shown in Section 3. Our pursuit of non-trivial examples leads to a "necklace" construction in Section 4 that yields arbitrarily many 2-connected graphs with the same average genus, as well as a sequence of ascending values of

average genus whose limit is the number one. In Section 5, we establish that all the cutedge-free supergraphs of K_4 have average genus larger than one, from which it follows that K_4 is the only 3-connected graph whose average genus is less than one. Section 6 describes some related results and lists some open research problems.

We assume familiarity with the standard lore of topological graph theory, as described by Gross and Tucker [1987], or -- with minor terminological exceptions -- by White [1984]. It might also be helpful to review Gross and Furst [1987]. Here are a few of our definitions and notations.

A graph may have self-adjacencies or multiple adjacencies. It is taken to be connected unless one can infer otherwise from the immediate context. The orientable surface with j handles is denoted S_j . The graph imbeddings under consideration here are exclusively in orientable surfaces, and we observe that an analogous theory might be explored for non-orientable imbeddings.

A bar-amalgamation of two disjoint graphs G and H is obtained by running an edge from a vertex of G to a vertex of H . Proof of the following theorem of Gross and Furst [1987] is omitted. Its corollary is quite useful to our present investigation.

THEOREM 1.1 The genus distribution of a bar-amalgamation of two graphs is a scalar product of the convolution of their respective genus distributions.

COROLLARY 1.2 The average genus of a bar-amalgamation of two graphs equals the sum of their average genera.

2. Smallest positive values of average genus

A general question one might ask is, what positive rational numbers can be realized as the average genus of a graph? We begin our investigation by determining the smallest few numbers that can occur as the average genus of a graph. Zero is obviously the smallest, so we turn to the smallest positive numbers.

In what follows, we assume that the reader has sufficient familiarity with topological graph theory to calculate the genus distribution of a small graph. Various details of some of the calculations here are not given until later sections. We begin with a utility theorem.

THEOREM 2.1 The average genus of a graph is at least as large as the average genus of any of its subgraphs.

Proof. It suffices to consider the effect on average genus of adding an edge to a connected graph G . By Corollary 1.2, we may as well assume that the new edge is a self-loop or runs between two existing vertices of G . We denote the extension of G by $G+$.

If the new edge is a self-loop at a vertex v with valence d , then for each imbedding of G there are $d(d+1)$ imbeddings of $G+$, each in a surface of genus at least as large as the genus of

the imbedding surface for G from which it arose. If each imbedding of G_+ had exactly the same genus as the imbedding of G from which it arose, then the genus distribution of G_+ would be a scalar multiple of the genus distribution for G , and G_+ would have the same average genus as G . Shifting parts of the coordinate values higher can only raise the average genus.

The case in which the new edge runs between two different vertices is quite similar. Let us assume that their respective valences are d_1 and d_2 . Then each imbedding of G leads to $d_1 d_2$ imbeddings of G_+ , none in a surface of genus lower than the genus from which it arises. As before, we observe that increasing the values of some members of a set of numbers cannot decrease the average. []

The bouquet B_n is the graph with one vertex and n self-loops. Since there are four imbeddings in the sphere, two in the torus, and no other imbeddings, the average genus of B_2 is $1/3$.

The dipole D_n is the graph with two vertices and n adjacencies between them. The dipole D_2 has two imbeddings in the sphere and two in the torus, for an average genus of $1/2$.

THEOREM 2.2 Let G be a graph with positive average genus. Then the average genus of G is at least $1/3$. Moreover, it cannot lie in the open interval $(1/3, 1/2)$.

Proof. In accordance with Corollary 1.2, we may as well assume that the graph G has no cut-edge. Of course, this precludes 1-valent vertices. Since subdividing an edge does not change the genus distribution, we may also assume that G has no 2-valent vertices.

Let C be a longest cycle in G . If there were a path in $G - C$ from any vertex of C to any other, then G would contain a homeomorph of the dipole D_3 , implying that the average genus of G would be at least $1/2$, by Theorem 2.1.

If there are no paths in $G - C$ from any vertex of the cycle C to any other vertex of C , then there is a path in $G - C$ from each vertex of C to itself, because the minimum valence is at least three. This implies that G contains a homeomorph of the bouquet B_3 , from which it follows that the average genus is at least $1/3$. Average genus $1/3$ is realized when the cycle C has only one vertex. If C had two vertices, then G would contain a dipole D_2 with a self-loop at each end, whose average genus is $5/9$. []

REMARK 2.3 The realizable values of average genus in the open interval $(1/2, 3/4)$ are

$5/9$, $2/3$, and $19/27$

Let n be the length of cycle C in the proof of Theorem 2.1. The cases $n = 1$ or 2 result in average genus at least $1/3$, $1/2$, or $5/9$, as established above. For $n = 3$, the graph G must have one of the following configurations: three self-loops, in which case its average genus is at least $19/27$; a "2-ended bridge" and a self-loop, in which case the average genus is at least $2/3$; a "3-ended bridge", in which case G has average genus at least $7/8$, since G contains the complete graph K_4 ; or two chords from the same vertex, in which case G has a subgraph with average genus $5/6$.

REMARK 2.4 The number $3/4$ is the average genus of a cycle with two chords.

3. Shared values of average genus

An easy way to construct an example of two non-isomorphic graphs with the same genus distribution is to subdivide an edge of any graph. The construction of an example of two non-homeomorphic graphs with the same average genus is a more interesting endeavor. With the aid of Corollary 1.2, we have rather rapid success, if cut-edges are to be permitted, as illustrated in Figure 3.1.

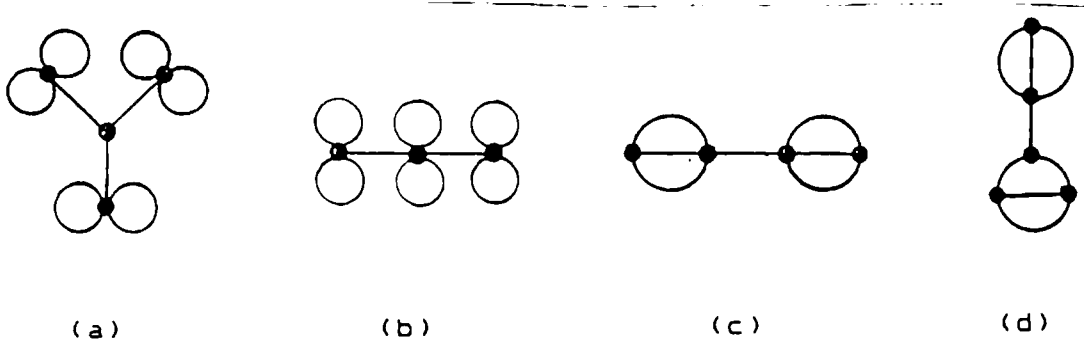


Figure 3.1 Four non-homeomorphic graphs with average genus equal to one.

Figures 3.1a and 3.1b illustrate two different iterated bar-amalgamations of three copies of the bouquet B_2 . By Corollary 1.2, both have average genus equal to

$$1/3 + 1/3 + 1/3$$

which equals one. Figure 3.1c shows a bar-amalgamation of two copies of the dipole D_3 , and Figure 3.1d shows a subdivision of a copy of D_3 bar-amalgamated to another copy of D_3 . By Corollary 1.2, both these graphs have average genus equal to

$$1/2 + 1/2$$

which also equals one.

By restricting our attention to graphs of minimum valence at least two and no cut-edges, we eliminate the simplest examples. We shall see in the next section that, even with these restrictions, arbitrarily many graphs can share the same average genus. For the time being, we consider the examples illustrated in Figure 3.2

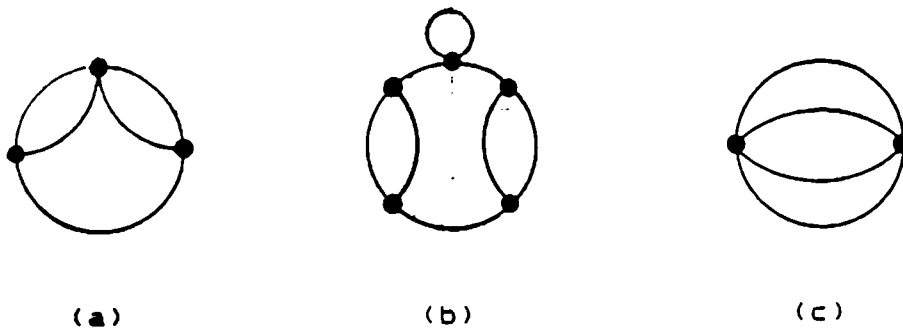


Figure 3.2 Three non-homeomorphic graphs with average valence equal to $5/6$.

The graph in Figure 3.2a has two 3-valent vertices and one 4-valent vertex. It follows that it has

$$(2!)^2(3!) = 24$$

imbeddings. Exactly four of them are in the sphere. Since its cycle rank is equal to three, its maximum genus is at most one. Thus, there are 20 toroidal imbeddings. It follows that the average genus is $5/6$.

The graph in Figure 3.2b has four 3-valent vertices and one 4-valent vertex. Thus, it has

$$(2!)^4(3!) = 96$$

imbeddings. Of these, 16 are spherical. Although its cycle rank is four, every spanning tree has at least two odd components in its cotree, so the theorem of Xuong [1979] implies that the maximum genus is one. It follows that there are 80 toroidal imbeddings. Thus the average genus is $5/6$.

The dipole D_4 in Figure 3.2c has maximum genus equal to one, by Xuong's theorem, and 36 imbeddings in all. Exactly six of its imbeddings are spherical. Thus, the average genus is $5/6$.

4. Necklaces

Examples in the previous section demonstrate that two different graphs can have the same average genus, even if there are no cut-edges or subdivisions. We now introduce a systematic method to construct arbitrarily many homeomorphism types of 2-connected graphs with the same average genus.

Suppose that r disjoint edges of a cycle are doubled and that a self-loop is added at each vertex which is not an endpoint of a doubled edge. Suppose this results in s self-loops. Then the resulting graph is called a necklace of type (r, s) . Figure 4.1 illustrates two necklaces of type $(2, 3)$.




Figure 4.1 Two non-homeomorphic necklaces with the same genus distribution.

The number of different necklaces of type (r, s) equals the coefficient of $x^r y^s$ when $x + y$ is substituted into the cycle index polynomial for the dihedral group D_{r+s} . A detailed explanation of this application of Polya's enumeration theorem is given, for example, by Harary [1969] or Tucker [1984]. In the present context, it is sufficient to realize that for $r = 2$, there are at least $s/2$ different necklaces of type (r, s) , corresponding to the minimum number of self-loop "beads" between the two doubled-edge beads encountered in a traversal of the necklace.

THEOREM 4.1 A necklace H of type (r, s) has the genus distribution

$$g_0 = 2^{r-4s} \quad g_1 = 4^{r-6s} - 2^{r-4s}$$

Proof. Since there are $2r$ vertices of valence 3 and s vertices of valence 4, it follows that the number of imbeddings is 4^{r-6s} . An induction argument can be used to demonstrate that no matter what spanning tree is chosen for the necklace H , the cotree has $r + s - 1$ odd components. Since the cycle rank of H is $r + s + 1$, it follows from Xuong's theorem [1979] that the maximum genus is one. Accordingly, our remaining task is to determine the number of imbeddings in S_0 .

In order to count the spherical imbeddings, let us suppose that C is a maximum cycle in H . Thus, C is a Hamiltonian cycle that contains one edge of each doubled edge-pair and none of the self-loops. In the plane, we draw a rotation projection (see Gross and Tucker [1987]) for H so that C lies on a circle. The corresponding imbedding is spherical if and only if both ends of each of the r other edges of a doubled-edge pair lie on the same side of C and both ends of each self-loop lie on the same side of C . It follows that the number of spherical imbeddings is 2^{r-4} . []

COROLLARY 4.2 The average genus of any necklace of type (r, s) is

$$1 - (1/2)^r (2/3)^s \quad []$$

COROLLARY 4.3 Arbitrarily many mutually non-homeomorphic 2-connected graphs can have the same average genus. []

COROLLARY 4.4 The average genus of a graph with non-trivial genus range can lie arbitrarily close to the maximum genus. []

COROLLARY 4.5 The number one is an upper limit point of the set of possible values of average genus. []

Constructing examples of non-homeomorphic 3-connected graphs with the same average genus is a more difficult task. The earliest known pair (Furst and Gross [1985]), illustrated in Figure 4.2, comprises two non-simplicial graphs. McGeoch [1987] developed a general method for generating such pairs of non-simplicial graphs. Rieper [1988] has used methods from Jackson [1987] to generate arbitrarily many simplicial graphs with identical genus distribution and, hence, identical average genus.

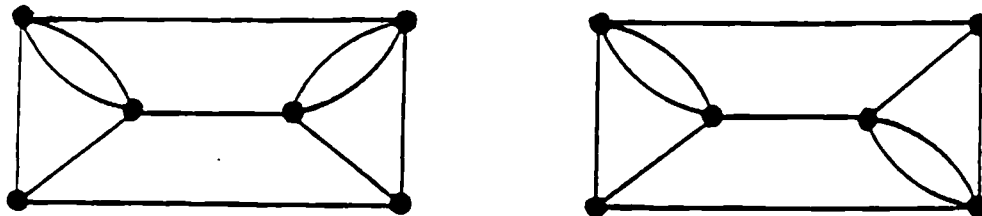


Figure 4.2 Two non-homeomorphic 3-connected graphs with the genus distribution 8, 536, 3416, 1224.

5. On the average genus of 3-connected graphs

Two easily derived properties of the complete graph K_4 are that its average genus is $7/8$ and that it is contained homeomorphically as a subgraph of every 3-connected graph. In view of Theorem 2.1, this implies that every 3-connected graph has genus at least $7/8$. We shall prove that except for K_4 itself, the average genus of a 3-connected graph is larger than one.

THEOREM 5.1 The average genus of the complete graph K_4 is equal to $7/8$.

Proof. As explained by Mull, Rieper, and White [1988], there are three conjugacy classes of imbeddings of the complete graph K_4 , which are illustrated in Figure 5.1. There are two imbeddings of class (a), six of class (b), and eight of class (c). Since class (a) is the only spherical class and classes (b) and (c) are toroidal, it follows that the average genus is $14/16$, which equals $7/8$. []

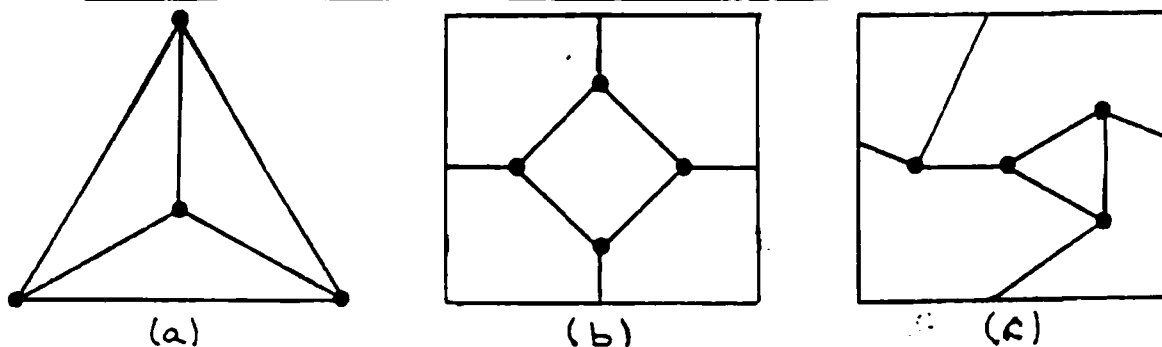


Figure 5.1 The three conjugacy classes of imbeddings of K_4 .

LEMMA 5.2 Every 3-connected graph G contains a homeomorphic subdivision of the complete graph K_4 .

Proof. Since adding self-loops or doubling edges does not change the connectivity of a graph, we may as well assume that the graph G is simplicial. Let u and v be two arbitrarily selected vertices of the graph G . By Menger's theorem (e.g., see Bondy and Murty [1976] or Harary [1969]), there exist three internally-disjoint paths in G between u and v . Since G is simplicial, at least two of these paths contain internal vertices. Therefore, we may choose an internal vertex w on one of the three paths from u to v , denoted P_1 , and another vertex x on another such path, denoted P_2 .

By Menger's theorem, there are three internally disjoint paths from w to x . Clearly, one of them, say path P , does not go through either of the points u or v . Without loss of generality, we assume that the path P goes directly from w to x without ever internally intersecting either of the paths P_1 or P_2 , for otherwise, we might replace w by the last vertex in which P intersects P_1 and x by the first vertex in which P intersects P_2 .

Let H be the subgraph of G formed by the three paths between u and v plus the path P from w to x . If path P does not intersect the third path between u and v , then H

is homeomorphic to K_4 . If the intersection is a single vertex, then H is homeomorphic to the wheel W_4 with four spokes, which contains K_4 . It is not difficult to verify that even if the intersection contains more than one vertex, the graph H still contains K_4 . \square

THEOREM 5.3 Every cut-edge-free proper supergraph of K_4 has average genus larger than one.

Proof. There are eight different ways to add an edge to K_4 so that the resulting graph has no cut-edges. It is sufficient to demonstrate that each of them has more imbeddings in the surface S_2 than in the sphere. (Since this result is a first cousin to a forbidden-subgraph theorem, some case-by-case analysis seems inevitable.) We repeatedly refer back to the conjugacy classes of imbeddings of K_4 described above.

Case 1: attach a self-loop at a vertex of K_4 . There are two class (a) imbeddings. Three different faces meet at the vertex with the self-loop. There are two directions in which to run the positive sense of the loop in each of those three faces. Thus, there are 12 imbeddings in S_0 . On the other hand, there are six class (b) imbeddings. In each of them a genus four imbedding can be obtained by placing one end of the self-loop in the 4-sided face and the other end in one of two "angles" of the 8-sided face. Thus, the number of imbeddings in S_2 that arise

from class (b) alone is 24.

Case 2: subdivide an edge, and attach a self-loop at the new vertex. There are four ways to extend a class (a) imbedding of K_4 to an imbedding of this supergraph, so there are 8 spherical imbeddings of the supergraph. The subdivided edge lies on two different faces in four of the class (b) imbeddings of K_4 and in four of the class (c) imbeddings. For each such imbedding, there are two ways to install the self-loop so that the resulting imbedding has an additional handle. Thus, the number of imbeddings in S_e is 16.

Case 3: run a new edge between two existing vertices, thereby creating a parallel adjacency to edge e . There are two ways to extend each class (a) imbedding to a spherical imbedding of the new graph, yielding a total of 4 spherical imbeddings. For each of the four class (b) imbeddings in which edge e lies on two faces, there are four ways to install the parallel edge so that its ends are not in the same face, yielding 16 imbeddings in S_e from class (b) alone.

Case 4: run a new edge from a vertex v to the midpoint of an edge e which is incident on v . As in case 3, there are only 4 spherical imbeddings. For each of the four class (b) imbeddings such that edge e lies on two faces, there are three ways to install the new edge on a new handle.

From here on, we do not bother to count the imbeddings in S_2 , since the details would repeat part of the previous cases.

Case 5: run a new edge from a vertex v to the midpoint of an edge e that is not incident on v . Then the resulting graph is isomorphic to the wheel W_4 with four spokes and has only two spherical imbeddings.

Case 6: run a new edge between the midpoints of two edges that meet. Then the resulting graph is simplicial and 3-connected, so it has only two spherical imbeddings.

Case 7: run a new edge between the midpoints of two edges that do not meet. Then the resulting graph is isomorphic to $K_{3,3}$ and has average genus $11/8$.

Case 8: run a new edge between two new subdivision points on the same edge of K_4 . Then there are only four spherical imbeddings. []

4

COROLLARY 5.4 Every 3-connected graph except K_4 has average genus larger than one.

Proof. This follows immediately from Theorem 2.1, Lemma 5.2, and Theorem 5.3. []

6. Research problems and related work

In the course of this research, it became clear that there are numerous immediate possibilities for continuation. We now formulate several of them as specific problems.

(6.1) First of all, the "necklaces" of Section 4 are an infinite family of cutedge-free graphs whose average genus is less than one. Find a concise way to characterize the other homeomorphism types of cutedge-free graphs with average genus less than one.

(6.2) Characterize the set of limit points of the values of the average genus of 2-connected graphs and of 3-connected graphs.

(6.3) The number one is an upper limit point of the set of values of average genus. Are there any lower limit points?

Rieper [1988] has proved that the average genus of a 3-regular graph is at least half the maximum genus.

Stahl [1989] has proved that the average genus of K_n is asymptotic to the maximum genus, and he has obtained upper bounds for the mean and variance of the genus distribution of an arbitrary graph. Determining the average genus of all imbeddings over the class of graphs with a fixed number of edges is a somewhat related problem that has also been explored by

Stahl [1983].

Furst, Gross, and Statman [1989] have calculated the genus distributions of two infinite classes of graphs, called closed-end ladders and cobblestone paths. McGeoch [1987] has calculated the genus distributions of circular ladders and of Mobius ladders.

Gross, Robbins, and Tucker [1989] have calculated the genus distributions of bouquets of circles. Stahl [1989] has subsequently elaborated upon this by demonstrating that the genus distribution (or "region distribution", if one prefers) of a bouquet is a close approximation to the distribution of the unsigned Stirling cycle numbers.

(6.4) All the known genus distributions are strongly unimodal. Decide whether the the genus distribution of every graph is strongly unimodal.

Lee and White [1989] and Schwenk and White [1989] have explored some of the variations that occur in studying imbedding distributions, depending on whether or not one prescribes labeling of the graph or orientability of the imbedding surface.

References

- J. A. Bondy and U. S. R. Murty [1976], Graph Theory with Applications, American Elsevier, New York.
- M. L. Furst and J. L. Gross [1985], unpublished example.
- M. L. Furst, J. L. Gross, and R. Statman [1989], Genus distributions for two classes of graphs, J. Combinatorial Theory 46, 22-36.
- J. L. Gross and M. L. Furst [1987], Hierarchy for imbedding-distribution invariants of a graph, J. Graph Theory 11, 205-220.
- J. L. Gross, R. G. Rieper, and T. W. Tucker [1989], Stratified graphs, manuscript.
- J. L. Gross, D. P. Robbins, and T. W. Tucker [1989], Genus distributions for bouquets of circles, J. Combinatorial Theory 47, to appear.
- J. L. Gross and T. W. Tucker [1987], Topological Graph Theory, Wiley-Interscience, New York.
- F. Harary [1969], Graph Theory, Addison-Wesley, Reading.

D. M. Jackson [1987], Counting cycles in permutations by group characters, with an application to a topological problem, Trans. Amer. Math Soc. 299, 785-801.

S. H. Lee and A. T. White [1989], Random topological graph theory, I, manuscript.

L. A. McGeoch [1989], Genus distributions for circular and Mobius ladders, manuscript. (originally in Ph.D. Thesis [1987], Carnegie-Mellon University)

B. P. Mull, R. G. Rieper, and A. T. White [1988], Enumerating 2-cell imbeddings of connected graphs, Proc. Amer. Math. Soc. 103, 321-330.

R. G. Rieper [1988a], unpublished.

R. G. Rieper [1988b], unpublished.

A. J. Schwenk and A. T. White [1989], Random topological graph theory, II, manuscript.

S. Stahl [1983], The average genus of classes of graph embeddings, Congressus Numerantium 40, 375-388.

S. Stahl [1989a], Region distributions of graph embeddings and

Stirling numbers, to appear.

S. Stahl [1989b], An upper bound for the average number of regions, manuscript.

A. Tucker [1984], Applied Combinatorics (2d ed.), John Wiley, New York.

A. T. White [1984], Graphs, Groups and Surfaces (rev. ed.), North-Holland, Amsterdam.

N. H. Xuong [1979], How to determine the maximum genus of a graph, J. Combin. Theory Ser. B 26, 217-225.